

ETHICAL DECISION-MAKING COMPASS (EDMC)

A FRAMEWORK FOR LEADERSHIP

IN A COMPLEX WORLD

by

Indira Bunic, MAES, MECP, MLCP, CASPGA

DISSERTATION

Presented to the Swiss School of Business and Management Geneva

In Partial Fulfillment

Of the Requirements

For the Degree

DOCTOR OF BUSINESS ADMINISTRATION

SWISS SCHOOL OF BUSINESS AND MANAGEMENT GENEVA

OCTOBER, 2025

ETHICAL DECISION-MAKING COMPASS (EDMC)
A FRAMEWORK FOR LEADERSHIP
IN A COMPLEX WORLD

by

Indira Bunic

APPROVED BY

Ava Buljibasic

Dissertation chair

RECEIVED/APPROVED BY:

Renee Goldstein Osmic

Admissions Director

Dedication

To my family,
whose love has been the steady star by which I set my course,
whose trust and support have carried me through every storm.

To my mentors and all the leaders I have learned from—
guides and fellow travelers on the winding path of ethical discovery—
who taught me that integrity, resilience, and curiosity are the truest bearings.

And to all who, knowingly or not, helped me find and hold my ethical compass,
so that I might help others navigate the complex seas of decision and leadership.

Acknowledgements

This dissertation is the culmination of a journey shaped by the wisdom, encouragement, and generosity of many remarkable people.

I am profoundly grateful to my family, whose patience, love, and unwavering belief in me have been my anchor and my inspiration. Your support has been the foundation upon which every achievement rests.

To my mentors and academic advisors, especially Dr. Anna Provodnikova, thank you for your rigorous guidance, insightful feedback, and steadfast encouragement. Your example has set the standard for integrity and excellence in scholarship.

I extend my heartfelt appreciation to the leaders, colleagues, and experts who shared their experiences and perspectives through interviews, surveys, and conversations. Your openness and candor enriched this research and brought the Ethical Decision-Making Compass to life.

To my friends and fellow travelers on the path of ethical discovery, thank you for your camaraderie, your challenges, and your belief in the value of this work.

Finally, I am indebted to all who, knowingly or not, helped me find and hold my ethical compass—reminding me that the truest bearings are found not in certainty, but in curiosity, resilience, and the courage to choose wisely.

This work is as much yours as it is mine. Thank you!

ABSTRACT

ETHICAL DECISION-MAKING COMPASS (EDMC)

A FRAMEWORK FOR LEADERSHIP

IN A COMPLEX WORLD

Indira Bunic
2025

Dissertation Chair: Dr. Anna Provodnikova

In today's fast, AI-intensive environments, leaders and teams must make principled choices under time pressure, complexity, and uncertainty. While many ethical frameworks illuminate parts of the journey, they often prove too siloed, slow, or abstract for high-velocity decisions—where ethical salience is easily missed and organizational learning is lost. This thesis develops and validates the Ethical Decision-Making Compass (EDMC): a first-tier, time-aware, auditable architecture structured by PASO—Principles (why), Actions (how), Skills (who), and Outcomes (what).

EDMC is doctrine-agnostic about which principles a community adopts, yet uncompromising that whatever is named must be enacted, owned, and evidenced. It functions as a Compass, not a checklist: orienting attention at the moment of choice, linking intent to safeguards and stewardship, and leaving a transparent record for review. Rather than cataloging gaps, this work synthesizes strengths from sixteen leading frameworks, integrating their most practical recurring elements into a unified, adaptable tool for diverse, AI-driven contexts.

Validation draws on two evidence streams: an international expert panel (N=30) and researcher-led applications across eight well-documented public-facts cases. Results show that EDMC clarifies criteria, counters ethical blindness with engineered salience cues and thresholds, and supports cautious, context-sensitive action. In AI-heavy settings, it translates principles into concrete guardrails, human-in-the-loop roles, bias audits, rollback conditions, and principle-linked KPIs, while remaining stable as priorities shift. For leaders and teams, EDMC turns intent into safeguarded action with capable ownership and accountable outcomes. For boards and policymakers, it provides credible, pre-emptive diligence and proportionate disclosure. Repeated use builds capability and culture: the doctrine-agnostic PASO DNA, state the principle, specify the action, assign a skilled owner, define the check, functions as an operating code that creates a learning loop rather than paperwork. While practice-focused and modest in scope, the evidence indicates EDMC is adaptable, teachable, and ready for digital embedding and cross-context deployment.

In sum, EDMC reframes ethical decision-making as a disciplined, time-aware, and auditable, learnable practice. Used consistently, its PASO DNA makes ethical noticing unavoidable and turns principled intent into traceable action—offering a Compass for navigating uncertainty at the speed of modern work.

Keywords:

ethical decision-making, responsible leadership, decision architecture, PASO (Principles–Actions–Skills–Outcomes), Ethical Decision-Making Compass (EDMC), doctrine-agnostic, behavioral ethics, ethical blindness, time-pressure, traceability, auditability, AI governance, cross-cultural adaptability, human-in-the-loop (HITL), digital transformation

TABLE OF CONTENTS

CHAPTER I: INTRODUCTION.....	1
1.1 Introduction.....	1
1.1.1 Background – Floor vs. Ceiling.....	1
1.1.2 Why and Where is Ethical Decision-Making Needed?.....	3
1.2 Research Problem.....	4
1.2.1 Problem Statement.....	5
1.3 Research Purpose, Focus Area, Vision.....	6
1.3.1 Research Purpose - The Practical Need for a New Framework.....	6
1.3.2 Research Focus Area.....	7
1.3.3 Research Vision - Bridging Theory and Practice.....	9
1.4 Significance of the Study.....	9
1.4.1 Theoretical Significance.....	10
1.4.2 Practical Significance.....	10
1.4.3 Broader Societal Significance.....	11
1.5 Research Questions.....	12
1.5.1 Primary Research Question.....	12
1.5.2 Three-Phase Research Approach.....	12
1.5.3 Alignment of Questions, Limitations, Objectives, and PASO.....	14
1.6 Objectives – Long-Term Research Goal.....	16
1.6.1 Primary Research Objective.....	17
1.6.2 Research Objectives and Sub-Objectives.....	17
1.6.3 Strategic Implementation Objectives (beyond Q1–Q6).....	20
1.6.4 Expected Contributions and Benefits.....	20
1.7 Theoretical Framework.....	22
1.7.1 Foundational Theoretical Pillars.....	22
1.7.2 Toward an Integrated Framework: The PASO Components.....	23
1.7.3 Bridging Theory and Practice.....	24
1.8 Methodology.....	24
1.8.1 Introduction.....	24
1.8.2 Methodological Foundation.....	25

1.8.3 Three-Phase Research Design	26
1.8.4 Analytical Methods	27
1.8.5 Quality Assurance	27
1.8.6 Validation Methods	27
1.8.7 Ethical Considerations and Reflexivity	28
1.8.8 Methodological Limitations and Innovations.....	28
1.8.9 Practical Application and Implementation	29
1.8.10 Implementation Roadmap	30
1.8.11 Future Research Directions	30
1.9 Limitations, Delimitations, and Assumptions	31
1.9.1 Limitations.....	31
1.9.2 Delimitations	33
1.9.3 Assumptions	34
1.9.4 Conclusions	35
<i>CHAPTER II: REVIEW OF LITERATURE</i>	<i>37</i>
2.1 Ethical Decision-Making - Theoretical Foundations and Review Strategy.....	37
2.1.1 Introduction – Purpose and Scope.....	37
2.1.2 Background and Evolution of Ethical Decision-Making Frameworks	38
2.1.3 Review Scope and Method.....	39
2.1.4 Literature Review Methodology and Framework Selection	41
2.2 Comparative Review of Ethical Decision-Making Frameworks	42
2.2.1 Foundational Ethical Theory	43
2.2.2 Structured Decision-Making Models	45
2.2.3 Leadership-Oriented Approaches.....	47
2.2.4 Contemporary Challenge-Specific Frameworks (e.g. Sustainability, AI, Governance)	48
2.3 Critical Patterns, Fragmentations, and Limitations in Frameworks	51
2.3.1 Key Dimensions of Comparative Analysis	51
2.3.2 Identified Limitations and Fragmentations in Frameworks	56
2.4 Defining Ethical Decision-Making: Research Definition and Foundations	61
2.4.1 Research Definition.....	61
2.4.2 Academic Foundation and Literature Grounding.....	62
2.4.3 Theoretical Integration and Practical Implications (PASO → EDMC)	63

2.4.4 Methodological Significance and Construct Validity	64
2.5 Emergence and Integration of EDMC’s PASO Components	65
2.5.1 Synthesis of Foundational Elements	65
2.5.2 Structuring the EDMC on the PASO Foundation	70
2.6 Advancing Ethical Decision-Making: the EDMC Framework.....	74
2.6.1 Addressing Fragmentation (L1)	75
2.6.2 Ethical Blindness Mitigation (L5).....	76
2.6.3 Cross-Cultural Adaptability (L1/L4).....	77
2.6.4 Time-Pressure Functionality (L4/L5).....	77
2.6.5 Practical Usability (L2/L3).....	78
2.7 From Theoretical Review to Methodological Foundation: Advancing the EDMC Framework	79
2.7.1 From Fragmentation to Foundation: Synthesizing the PASO Components.....	79
2.7.2 Justifying a Three-Phase Methodological Approach	80
2.7.3 Advancing the Field: EDMC’s Unique Contributions	81
2.7.4 From Theory to Empirical Application: Designing the Research.....	82
2.8 Summary and Implications for Method	83
<i>CHAPTER III: METHODOLOGY</i>.....	85
3.1 Introduction.....	85
3.2 Operationalization of Theoretical Constructs.....	86
3.3 Research Purpose and Questions	88
3.4 Research Design	90
3.5 Population and Sample.....	92
3.5.1 Case Corpus (C1–C8).....	92
3.5.2 Expert Panel (EDMC Instrument Validation)	93
3.6 Participant Selection.....	94
3.7 Instrumentation	96
3.8 Data Collection Procedures.....	99
3.9 Data Analysis.....	102
3.9.1 Quantitative Analysis (Expert-Panel Dataset).....	102
3.9.2 Quantitative Analysis (Researcher-Led Case Applications)	104
3.9.3 Qualitative Analysis (Free-Text Comments and Scoring Rationales)	105
3.9.4 Missing Data, N/A, and Exclusions	106

3.9.5 Robustness, Reliability, and Transparency	106
3.10 Research Design Limitations	107
3.11 Ethical Considerations and Researcher Reflexivity	109
3.12 Conclusion	110
CHAPTER IV: RESULTS.....	113
4.1 Introduction.....	113
4.2 Expert-Panel Survey Results Introduction.....	114
4.2.1 Sample Profile	116
4.2.2 Usability and PASO Content Validity.....	120
4.2.3 Decision Stance and Confidence	122
4.2.4 Global Items (Portability, Integration, Time-Pressure, Coherence, Tech/AI)	124
4.2.5 Qualitative Insights (Directed Content Analysis).....	125
4.2.6 Interim Summary and Synthesis of Expert-Panel Evidence.....	127
4.3 Researcher-Led Case Applications	128
4.3.1 Why a Second Evidence Stream?.....	129
4.3.2 Alignment to the Research Questions	130
4.3.3 Case Set, Scoring, and Weight Presets	131
4.3.4 Results Under Equal Weights.....	132
4.3.5 Sensitivity to Decision Priorities (Weight Presets).....	134
4.3.6 Cross-Case Synthesis (What the Metrics Say).....	137
4.3.7 Reproducibility Note and Provenance (RQ Alignment to Q4–Q6).....	139
4.3.8 What This Adds Beyond the Survey	141
4.4 Sensitivity and Robustness.....	142
4.4.1 Expert-Panel Data Integrity (SENS-1)	143
4.4.2 Alternative Denominator / N-Handling.....	144
4.4.3 Researcher Case-Application Robustness Beyond Weights	144
4.4.4 Threats to Validity and Mitigations.....	145
4.4.5 Bottom Line.....	145
4.5 Summary of Key Findings	146
4.6 Conclusions.....	147
CHAPTER V: DISCUSSION.....	150
5.1 Discussion — Purpose, Why Now, and Vision	150
5.2 Summary of the Study and Research Journey (From Problem to Evidence).....	152

5.3 Linking Findings to Objectives and Research Questions (Closing the Loop)	154
5.4 What These Results Mean (and Why they Matter)	159
5.5 Contributions	162
5.5.1 Positioning: EDMC as a First-Tier Filter in a Tiered Decision Stack.....	165
5.5.2 Distinctive Contribution: Engineered Mitigation of “Ethical Blindness” Under Time Pressure.....	168
5.5.3 Theoretical Contributions	170
5.5.4 Methodological Contributions	172
5.5.5 Practical Contributions (Leadership and Governance).....	174
5.5.6 Societal and Policy-Adjacent Contributions	176
5.5.7 Where This Contribution Stops (and Why That’s a Strength)	177
5.5.8 Novelty and Real Merit of This Study	179
5.6 Limitations and Boundary Conditions (and Why Claims are Modest)	180
5.7 Practice Implications (Preview; Elaborated in Chapter 6)	182
5.8 Directions for Further Research	183
5.9 Chapter Takeaway	185
<i>CHAPTER VI: SUMMARY, IMPLICATIONS, AND RECOMMENDATIONS</i>	<i>186</i>
6.1 Summary	186
6.2 Implications	187
6.2.1 Theoretical Implications	187
6.2.2 Practical and Leadership Implications.....	188
6.2.3 Policy and Governance Implications.....	189
6.2.4 Societal Implications	189
6.2.5 Capacity, Empowerment, Awareness, and Trust.....	190
6.3 Recommendations	190
6.3.1 Organization-Wide Implementation Roadmap.....	190
6.3.2 For AI, Data, and Socio-Technical Programs	192
6.3.3 For Boards, Audit, and Policy/Governance Bodies.....	193
6.3.4 For Researchers, Educators, and Partners	193
6.4 Implementation and Evaluation: How We’ll Know EDMC + PASO Is Working	194
6.5 Limitations and Boundary Conditions	196
6.6 Conclusions	197

APPENDIX A: RESEARCHER-LED CASE PACKETS (PUBLIC CASES) — ONE-LINE SCORING RATIONALES (C1–C8)	201
APPENDIX B: SCORING RUBRIC, WEIGHTS, AND COMPUTATION (0–3 anchors)	207
APPENDIX C: EXPERT-PANEL PROTOCOL AND SURVEY INSTRUMENT (INDEX AND VERBATIM)	212
APPENDIX D: EXPERT-PANEL CASE NOTES (HYPOTHETICAL VIGNETTES)	216
APPENDIX E: LITERATURE SYNTHESIS TABLES	220
APPENDIX F: ANALYSIS PLAN AND DERIVED METRICS	227
APPENDIX G: EXPERT-PANEL RESULTS — SUMMARY TABLES AND TIMELINE...	231
APPENDIX H: EXPERT-PANEL FIGURES AND METHODS NOTES	234
APPENDIX I: DATA AVAILABILITY AND RETENTION	241
APPENDIX J: GLOSSARY AND ABBREVIATIONS	242
APPENDIX K: PRIVACY AND DATA HANDLING (GDPR / SWISS FADP)	244
APPENDIX L: CONSENT / NDA WORDING AND PARTICIPANT-CODE PROCESS	247
APPENDIX M: EXCLUSION RULES AND SENSITIVITY CHECKS	248
APPENDIX N: VERSION LOCK AND FIELDWORK WINDOW	250
APPENDIX O: RESEARCH JOURNEY FIGURES	251
APPENDIX P: EXTENDED METHODS SCAFFOLDS	254
REFERENCES	255

LIST OF TABLES

<i>Table 1 Recent Ethical-Risk Cases Across Sectors</i>	3
<i>Table 2 Alignment of Research Questions with Addressed Limitations, Objectives, and PASO/EDMC Focus</i>	14
<i>Table 3 Objectives and Sub-Objectives with Linked Research Questions</i>	19
<i>Table 4 Sixteen Frameworks Included in the Review</i>	40
<i>Table 5 Foundational Ethical Theory: Contributions and Typical Challenges</i>	44
<i>Table 6 Structured Ethical Decision-Making Models</i>	46
<i>Table 7 Leadership-Oriented Ethical Decision-Making Frameworks and Typical Challenges</i>	48
<i>Table 8 Contemporary Ethical Challenges: Sustainability, AI, and Governance</i>	50
<i>Table 9 How PASO Components Address Critical Limitations</i>	73
<i>Table 10 How the EDMC Framework Addresses Critical Limitations</i>	74
<i>Table 11 Compact Mapping—Construct and Survey Indicator(s)</i>	87
<i>Table 12 Phases and Alignment to RQs</i>	91
<i>Table 13 Gender (N = 30)</i>	117
<i>Table 14 Age Band (N = 30)</i>	117
<i>Table 15 Highest Education (N = 30)</i>	117
<i>Table 16 Field of Expertise (Self-Described; N = 30)</i>	117
<i>Table 17 Years of Experience (N = 30)</i>	118
<i>Table 18 Current Role/Position (N = 30)</i>	118
<i>Table 19 Primary Domain (N = 30)</i>	118
<i>Table 20 Sector (N = 30)</i>	118
<i>Table 21 Agreement ($\geq 4/5$) for Q1–Q8 by Case (% , N)</i>	121
<i>Table 22 PASO Content Validity (I-CVI Per Item; S-CVI/Ave Per Case)</i>	121
<i>Table 23 Stance-Change Matrix (Counts; Pooled Across Cases, N = 49)</i>	123
<i>Table 24 Confidence Uplift (Post – Baseline, 0–10)</i>	123
<i>Table 25 Global Wrap-Up Items (Means, 1–5; N = 30)</i>	124
<i>Table 26 Themes and Exemplar Quotes (Anonymized, ≤ 20 Words)</i>	126
<i>Table 27 EDMC Scores by Case (Means 0–3; Composites 0–100)</i>	132
<i>Table 28 Sensitivity Summary (Δ Versus Equal, Percentage Points)</i>	136
<i>Table 29 Findings → RQs/Objectives Crosswalk (Discussion Focus)</i>	158
<i>Table 30 EDMC vs. Other Models—Distinctive Features</i>	169

<i>Table 31 EDMC Contributions at a Glance</i>	190
<i>Table 32 EDMC Implementation Phases</i>	191
<i>Table 33 Barriers and Mitigation</i>	192
<i>Table 34 Appendix B1: EDMC Effectiveness Metrics—0–3 Anchors (Overview)</i>	208
<i>Table 35 Appendix B2: Worked Example–Per-Metric Scoring Sheet (Blank Template)</i>	210
<i>Table 36 PASO Mapping Across Reviewed Frameworks (P/A/S/O)</i>	220
<i>Table 37 Comparative Analysis of Ethical Decision-Making Frameworks (Across Six Dimensions)</i>	221
<i>Table 38 Comprehensive Review Across 16 Definitions, Components, Strengths, Gaps, and EDMC Relevance</i>	222
<i>Table 39 Comparative Analysis — EDMC Versus Prior Ethical Decision-Making Frameworks (Closing the Loop)</i>	225
<i>Table 40 Appendix G1: Expert-Panel Key Indicators by Case</i>	231
<i>Table 41 Appendix G2: PASO Content Validity Coverage by Case</i>	231
<i>Table 42 Appendix G3: Global Wrap-Up Item Means</i>	232
<i>Table 43 Appendix G4a: Sufficiency Distribution Across the Three Cases</i>	232
<i>Table 44 Appendix G4b: PASO Consistency Distribution (All 3 / 2 / Not Present)</i>	232
<i>Table 45 Appendix G5: Post-Framework Confidence by Case</i>	232
<i>Table 46 Appendix G6: Q1–Q8 Means Pooled Across Cases 1–3</i>	233

LIST OF FIGURES

<i>Figure 1 Significance of the Study</i>	12
<i>Figure 2 Research-Question Flow Across Phases and Linkage to the EDMC/PASO Architecture</i>	15
<i>Figure 3 Validation and Refinement Loops Linking Outcomes to the EDMC/PASO Architecture</i>	16
<i>Figure 4 Historical Evolution of the Ethical Decision-Making Frameworks</i>	38
<i>Figure 5 Comparative Evolution of the Ethical Decision-Making Frameworks</i>	42
<i>Figure 6 Identified Gaps and Limitations</i>	61
<i>Figure 7 The Research Gap and EDMC Framework Contribution</i>	64
<i>Figure 8 Systematic Gap Analysis Matrix of Ethical Decision-Making Frameworks (n = 16) Across Five Gap Dimensions (L1–L5)</i>	67
<i>Figure 9 Mapping Foundational Components-From Reviewed Frameworks to PASO Anchors (Principles, Actions, Skills, Outcomes)</i>	69
<i>Figure 10 Bridging Gaps and Limitations: EDMC Mechanisms Mapped to L1–L5 and PASO (Critical Limitation)</i>	79
<i>Figure 11 From Framework Analysis to Approach: Flow from Synthesis and Gaps to the Three-Phase Research Design</i>	82
<i>Figure 12 Three-Phase EDMC Methodology—From Literature Synthesis → PASO → EDMC Build and Validation</i>	91
<i>Figure 13 Distribution of Primary Domain (N = 30)</i>	120
<i>Figure 14 EDMC Composite Scores (0–100), Equal Weights (n = 8)</i>	133
<i>Figure 15 EDMC Composite Scores (0–100), Stakeholder/Outcomes Emphasis (n = 8)</i>	135
<i>Figure 16 EDMC Composite Scores (0–100), Time-Pressure Blindness/Implementation Emphasis (n = 8)</i>	136
<i>Figure 17 Appendix H1: Heatmap of Agreement ($\geq 4/5$) by Item \times Case</i>	234
<i>Figure 18 Appendix H2: PASO Coverage by Case (I-CVI = $\% \geq 4$)</i>	235
<i>Figure 19 Appendix H3: Efficiency Impact by Case</i>	236
<i>Figure 20 Appendix H4: Confidence Uplift by Case</i>	237
<i>Figure 21 Appendix H5: Global Item Means (Q1–Q8 Pooled Across Cases)</i>	238
<i>Figure 22 Appendix H6: (a) PASO Consistency and Sufficiency (All 3 / 2 / Not Present)</i>	239
<i>Figure 23 Appendix H6: (b) PASO Sufficient and Defensible Decision (0-3 Cases)</i>	239
<i>Figure 24 Appendix H7: Global Wrap-Up Items (Means, 1–5)</i>	240
<i>Figure 25 Literature Review Journey—Seven Steps (Chapter II)</i>	253

CHAPTER I: INTRODUCTION

1.1 Introduction

1.1.1 Background – Floor vs. Ceiling

A dangerous illusion has taken hold in leadership: that ticking every compliance box is the same as making an ethical decision. Law and regulation set the floor—the minimum below which organizations must not fall. Ethical decision-making aims at the ceiling—the standard leaders should strive to reach. Most real decisions live in the wide space between floor and ceiling: ambiguous, cross-cultural, time-pressured, and consequential.

Leaders face increasingly complex ethical challenges in an era of unprecedented technological progress, global interconnectedness, and rapidly changing societal expectations. The incredible growth of artificial intelligence, the pressing needs of climate change, and shifting stakeholder demands have created a landscape where traditional methods of ethical decision-making are no longer sufficient.

Leaders are often tasked with making important decisions quickly, across diverse cultures, and within complex socio-technical systems. Research in behavioral ethics shows how limited awareness and context can distort judgment, creating “blind spots” that decision-makers often overlook (Bazerman and Tenbrunsel, 2011). Under time pressure, this can lead to “ethical blindness,” where individuals overlook the ethical aspects of their own decisions (Palazzo et al., 2012). Reviews of ethical decision-making and leadership highlight that culture, systems, and governance influence behavior, not just individual intent (Craft, 2013; Treviño et al., 2014). Recent research on digital-ethics governance emphasizes a practical point: organizations need operational mechanisms, such as transparency, accountability, and control tools, to implement ethics effectively on a large scale (Zhang and Wade, 2022).

Ethical decision-making is essential to leadership, corporate governance, sustainability, and emerging areas such as AI ethics. It emphasizes organizations' ability to navigate complex ethical dilemmas, ensuring accountability, trust, and long-term credibility. It fosters trust, enhances reputation, and contributes to long-term success in an increasingly complex and interconnected global environment (Forbes Business Council, 2023; Harvard Professional Development, 2023; Treviño and Nelson, 2016).

Ethical decision-making dilemmas are common, affecting leadership, governance, corporate responsibility, and the rapidly evolving field of AI ethics. They require simple, strong, and adaptable ethical decision-making frameworks—from addressing corporate scandals to mitigating biased algorithms.

Years of research have yielded numerous models, yet the field remains fragmented and slow to be applied; most frameworks are limited to specific areas, abstract, or difficult to implement quickly. As a result, organizations lack a practical approach for fast-paced, cross-cultural, complex situations (Craft, 2013; Jobin et al., 2019; Kaptein, 2023; Mittelstadt et al., 2016; Treviño et al., 2014; Treviño and Nelson, 2016). There is a need for a simple, robust, and adaptable ethical decision-making framework that is applicable across cultures, from addressing corporate scandals to mitigating biased algorithms.

This research introduces the Ethical Decision-Making Compass (EDMC), a practical framework for making informed decisions when rules are silent, conflicting, or absent. EDMC is organized around four components: Principles, Actions, Skills, Outcomes (PASO), so values, context, capability, and impact can be weighed together, transparently and repeatably. EDMC aims beyond mere compliance: it enables decisions that are timely, explainable, auditable, and learnable across sectors and cultures. In short, it helps leaders move from awareness to action when it matters most.

1.1.2 Why and Where is Ethical Decision-Making Needed?

Ethical decision-making dilemmas are widespread across leadership, governance, corporate responsibility, and the fast-moving domain of AI ethics. Robust, adaptable frameworks matter because ethics underpins accountability, trust, reputation, and long-term performance in complex, global environments (Forbes Business Council, 2023; Harvard Professional Development, 2023; Treviño and Nelson, 2016).

Ethical failures carry measurable costs. The average global cost of a data breach reached USD 4.88 million in 2024 (IBM, 2024). Capital markets also penalize firms after negative ESG/CSR incidents: event studies document statistically significant negative abnormal returns around such news (Flammer, 2013; Krüger, 2015). Beyond direct losses, failures erode trust and legitimacy, undermining sustainable (Floridi et al., 2018; Jobin et al., 2019; Martin et al., 2021) . In AI-enabled settings, irresponsible use and weak controls can scale harm and increase systemic risk (Deloitte, 2023; Zhang and Wade, 2022). *Table 1* summarizes recent cross-sector cases (environmental, labor, discrimination, AI ethics), illustrating the prevalence and severity of ethical risk. These examples motivate the need for transparent, auditable, and context-adaptive decision processes—needs the EDMC is designed to address.

Table 1 Recent Ethical-Risk Cases Across Sectors

Source: Author (2025). Original table created from study data and materials.

Category	Issue	Details	Source	Year
Environmental Issues	Fast Fashion Impact	Shein faced allegations of hazardous chemical violations exceeding EU REACH limits by 100 times	(Reuters, 2023)	2023
Environmental Issues	Oil Industry	Shell's Jackdaw gas field approval sparked environmental protests and legal challenges	(Ambrose, 2022)	2024
Labor Rights	Tech Industry Labor	Apple's main supplier Foxconn violated labor laws in China's Zhengzhou facility	(Bloomberg News, 2023)	2023
Labor Rights	Union Issues	Starbucks faced over 500 labor violations charges	(National Labor Relations Board, 2023)	2023
Discrimination	Gender	Google agreed to pay \$118 million in gender discrimination lawsuit settlement	(Bhattarai, 2022)	2022

Discrimination	Age Bias	Tech industry age discrimination claims increased by 41% during 2023	(U.S. Equal Employment Opportunity Commission, 2024)	2024
AI Ethics	Recruitment Bias	Amazon discontinued AI hiring tool due to gender bias against women	(Dastin, 2018)	2023
AI Ethics	Governance AI	OpenAI leadership crisis revealed gaps in AI governance structures	(Heaven, 2023)	2023

These cases illustrate only a few examples of the numerous sustainability and ethical issues that organizations face when striving to develop better corporate responsibility strategies and ethical decision-making. Organizations and businesses, including both for-profit and non-profit entities, NGOs, government agencies, and public services, face unprecedented ethical challenges regardless of context or borders. Failures in ethical decision-making result in irreparable losses, diminished trust, reputational harm, and unsustainable practices (Floridi et al., 2018; Treviño and Nelson, 2016).

1.2 Research Problem

Ethical decision-making remains difficult for leaders operating in complex, high-pressure environments. Despite decades of development, many frameworks remain siloed, theory-heavy, or difficult to apply under real-world constraints (Kaptein, 2023; Treviño and Nelson, 2016). Good intentions can be overwhelmed by situational pressures and time constraints, contributing to “ethical blindness” (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012).

A persistent gap between theory and practice shows up in recurring scandals and issues like algorithmic bias, which have real costs (see *Table 1*). Current approaches are often limited to specific contexts or cultures (Jobin et al., 2019; Kaptein, 2023).

Traditional theories—Deontology and Utilitarianism (Kant, 1785; Mill, 1863), struggle with modern dilemmas involving AI and sustainability (Floridi, 2013; Schwartz, 2023). Even organized models (Rest, 1986; Treviño, 1986) tend to focus only on thinking or context, often ignoring measurable results and key skills (Mittelstadt et al., 2016).

Contemporary, rule-based frameworks, such as GDPR (2016) and Floridi's AI Ethics Model (2013; 2020), prioritize regulatory compliance but lack global adaptability, cultural inclusivity, and practical implementation tools (2023).

The urgency goes beyond business into AI governance, climate policy, healthcare systems, and global development (Deloitte, 2023; Newman et al., 2023). Leaders also face ethical paralysis—the inability to make confident decisions amid ambiguity—caused by complexity and perceived risk (Bazerman and Tenbrunsel, 2011).

Many frameworks do not adequately address skill development needed for practical ethical reasoning (Harrison, 2020; Kim and Patel, 2018) and lack mechanisms to counter cognitive bias and ethical blindness (Moore and Gino, 2013; Palazzo et al., 2012). At the same time, employees, consumers, and investors are raising expectations for transparent, accountable leadership (Babalola et al., 2023; Harvard Professional Development, 2023a).

Beyond substantive limitations, a methodological gap exists. Research often tests predetermined models rather than allowing solutions to emerge from in-depth exploration of the problem (Alvesson and Sandberg, 2013; Davis, 1971). Emerging-methods scholarship advocates for inductive, problem-driven development when theorizing in complex fields (Edmondson and McManus, 2007; Eisenhardt et al., 2016; Gioia et al., 2013; Van de Ven, 2007). This has limited progress on practical, practice-oriented frameworks that connect theory with leadership actions (Treviño and Nelson, 2021; Woiceshyn, 2011).

1.2.1 Problem Statement

There is a need for a comprehensive, practical, and adaptable framework that: (i) integrates normative principles with situational constraints, skill requirements, and

outcome assessment; (ii) functions effectively under time pressure and across different cultures; and (iii) offers transparent, auditable reasoning paths that minimize blind spots and promote accountable leadership (Greenwood et al., 2022; Johnson, 2021). This thesis addresses that need by developing the Ethical Decision-Making Compass (EDMC) through an emergent discovery approach that identifies recurring patterns and consolidates them into foundational pillars.

1.3 Research Purpose, Focus Area, Vision

1.3.1 Research Purpose - The Practical Need for a New Framework

This study recognizes that an ethical decision-making framework is not merely an academic contribution but a vital real-world requirement. Drawing on leadership experience across sectors and countries, we observe that the most effective and sustainable decisions are grounded in ethics. Yet leaders routinely face interconnected constraints that hinder responsible action:

- Time pressure and cognitive constraints. In high-pressure contexts, time scarcity impairs deliberative reasoning and increases reliance on heuristics, elevating the risk of ethical compromise (Del Popolo Cristaldi et al., 2024; Emser et al., 2021).
- Ethical blindness. Even well-intentioned leaders may temporarily fail to see the ethical dimension of a choice under contextual pressures; this effect is amplified in complex stakeholder environments and cultures that normalize questionable practices (Bazerman and Tenbrunsel, 2011; Moore and Gino, 2013; Palazzo et al., 2012).
- Legal and regulatory complexity. National and international rules may be fragmented, rapidly evolving, or ill-suited to systemic inequities, complicating interpretation and application—especially for multinational organizations and AI-intensive settings

(Donaldson and Dunfee, 1999; Floridi, 2013; Kim and Patel, 2018; Mittelstadt et al., 2016).

- Cross-sector and cross-cultural demands. Modern leadership spans sectors, cultures, and stakeholder groups, requiring tools that remain actionable across contexts (Craft, 2013; Eisenbeiss, 2012).

The consequences of inadequate ethical decision-making extend far beyond individual firms to societal trust, environmental sustainability, and human welfare. Recent cases across technology, finance, labor, and the environment (see *Table 1*) underscore the urgency of a framework that works under real-world constraints.

Purpose statement. Accordingly, the purpose of this study is to design, articulate, and evaluate the Ethical Decision-Making Compass (EDMC). This practical, cross-domain framework integrates normative reasoning with situational constraints and skill requirements. EDMC is intended to reduce ethical blind spots, improve transparency and traceability, and support timely, context-sensitive decisions across sectors and cultures. Methodologically, we address the gap between theory and practice through an emergent discovery approach (outlined in Chapter 3) that synthesizes recurring patterns from scholarship and practitioner insight (Edmondson and McManus, 2007; Gioia et al., 2013; Van de Ven, 2007).

1.3.2 Research Focus Area

This study develops the Ethical Decision-Making Compass (EDMC) by starting with a structured review of existing ethical frameworks and, through iterative synthesis, distilling their strengths, limitations, gaps, and recurring patterns into four foundational components—Principles, Actions, Skills, Outcomes (PASO). The aim is to provide a

framework that leaders can use under real-world conditions (time pressure, cognitive load, and competing interests) and across various cultural and sectoral contexts.

Analogous to targeted interventions that restore systemic balance, EDMC is designed for precision and flexibility, helping decision-makers see the ethical dimension early, structure judgments transparently, and move from awareness to action. Chapter 2 (Section 2.4) specifies how this thesis defines *ethical decision-making* and situates EDMC within that definition. To keep the scope sharp, this study:

Does not:

- replicate existing frameworks or create another sector-specific model;
- argue for the necessity of ethics per se, as well established. (Babalola et al., 2023; Brown and Treviño, 2006; Ete et al., 2023; Mo et al., 2024; Newman et al., 2023)
- catalogue all causes of unethical behavior;
- replace existing laws, policies, or compliance systems.

Instead, it does:

- develop a practical, strategic, high-level framework that complements regulation and governance;
- provide tools intended for cross-sector and cross-cultural use;
- frame ethical decision-making as an opportunity for better performance, not a compliance burden;
- address ethical blindness under time constraints with simple, repeatable steps leaders can apply first, before engaging detailed compliance checks.

Positioned as the first-tier ethical filter, EDMC ensures that ethics guide decisions from the beginning; compliance then follows with clearer intent and traceability.

1.3.3 Research Vision - Bridging Theory and Practice

Organizations need tools that connect ethical awareness to action under real constraints. EDMC is intended to bridge the gap between normative ethics and everyday leadership practice by translating well-founded principles into usable steps and artifacts (Johnson, 2021). Consistent with problem-first, emergent approaches, this work prioritizes understanding the decision environment before prescribing solutions (Edmondson and McManus, 2007; Eisenhardt et al., 2016; Gioia et al., 2013; Van de Ven, 2007).

What EDMC offers

- Perception aids: prompts and structure to surface ethical salience during decision-making.
- Time-sensitive use: simple, repeatable moves that work under time and cognitive pressure.
- Cross-context portability: applicability across sectors and cultures without re-engineering.
- From awareness to action: guidance to move from diagnosis to justified, auditable decisions.

Accessible and scalable, EDMC is designed for leaders at all levels, from boardrooms to frontline operations, to apply ethical reasoning with clarity, confidence, and accountability.

1.4 Significance of the Study

This study matters on three levels. Theoretically, it combines insights from ethics, leadership, and governance to clarify how ethical decisions are made and put into action, and it formalizes the PASO components (Principles, Actions, Skills, Outcomes) as an organizing framework. Practically, it turns those insights into a usable framework

(EDMC) that enables quick, transparent, and consistent decisions under time constraints and across different cultures. Societally, it fosters trust and accountability by making ethical reasoning more transparent, auditable, and scalable across public, private, and nonprofit sectors. The following subsections detail the study's theoretical (1.4.1), practical (1.4.2), and broader societal (1.4.3) contributions.

1.4.1 Theoretical Significance

This study advances understanding of ethical decision-making by integrating insights across disciplines into a coherent, practice-ready architecture. It introduces the PASO components, Principles, Actions, Skills, Outcomes, as a conceptual scaffold for how ethical judgments are formed, enacted, and assessed across contexts. The framework addresses ethical blindness by linking psychological mechanisms that suppress ethical perception under pressure with concrete countermeasures embedded in process design (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012; Tenbrunsel and Smith-Crowe, 2008). More broadly, it positions ethical awareness as a first-order capability, moving beyond “knowing principles” toward reliably noticing the ethical salience of everyday decisions (Bazerman and Tenbrunsel, 2011). Finally, by avoiding Western-centric presuppositions and emphasizing portability, the study contributes to cross-cultural ethical decision theory (Jobin et al., 2019).

1.4.2 Practical Significance

For organizations and leaders, the contribution is a set of usable tools that make ethics actionable under real constraints (time pressure, ambiguity, competing interests). EDMC provides:

- a clear implementation pathway that can be embedded into existing decision, risk, and governance routines;
- a first-tier ethical filter that surfaces issues early, before detailed compliance checks, improving traceability and auditability;
- structured practices to raise ethical awareness at individual, team, and organizational levels (Brown and Treviño, 2006; Treviño and Nelson, 2016; Werhane, 2008).
- measurement and accountability hooks to evaluate outcomes and demonstrate business value.

The result is to turn ethical decision-making from a perceived constraint into a leadership strength—supporting faster, clearer, and more defensible choices in pressured environments.

1.4.3 Broader Societal Significance

At the societal level, EDMC supports the (re)building of trust through transparent reasoning, documented trade-offs, and accountable outcomes. By offering shared principles and process discipline, it enables cross-sector collaboration among public, private, and nonprofit actors. The framework equips leaders to confront complex challenges, such as AI governance and climate action, by making ethical dimensions visible and tractable, rather than incidental (Deloitte, 2023; Floridi, 2013). Because EDMC is simple and teachable, it helps democratize ethical competence beyond elite roles, spreading awareness and capability across organizations and communities. In sum, the study bridges theory and leadership practice with a framework that elevates perception, structures judgment, and improves outcomes—thereby enhancing organizational effectiveness and societal well-being (see *Figure 1*).

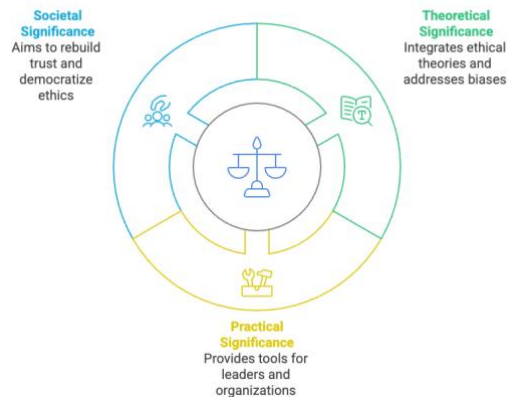


Figure 1 Significance of the Study

Source: Author (2025). Original figure created from study data and materials.

1.5 Research Questions

This study follows an emergent, three-phase design in which research questions are refined as insight develops, rather than fixed a priori—an approach suited to complex, practice-oriented inquiry (Edmondson and McManus, 2007; Eisenhardt et al., 2016; Van de Ven, 2007).

1.5.1 Primary Research Question

How can the strengths and gaps of existing ethical decision-making frameworks be synthesized into a practical, cross-cultural tool (EDMC) that enables timely, auditable decisions under real-world constraints?

1.5.2 Three-Phase Research Approach

Aligned with emergent, practice-oriented inquiry (Edmondson and McManus, 2007; Eisenhardt et al., 2016; Van de Ven, 2007), the study proceeds in three phases: Analysis and Assessment (Initial Phase), Integration and Synthesis (Bridge Phase), and

Development and Validation (Emergent Phase), each informing the next through iterative learning.

Phase I — Analysis and Assessment (Initial Phase)

- Q1. What are the strengths, limitations, and gaps in existing ethical decision-making frameworks?
- Q2. Which recurring patterns emerge across frameworks, and how do these patterns specify requirements for a more integrated approach?

Method note. Systematic comparative analysis of sixteen frameworks using directed content analysis and matrix mapping across six dimensions (foundation, adaptability, cultural inclusivity, implementation guidance, competency requirements, evaluation mechanisms) (Hsieh and Shannon, 2005).

Phase II — Integration and Synthesis (Bridge Phase)

- Q3. How can the identified strengths and patterns be synthesized into the PASO components (Principles, Actions, Skills, Outcomes) to form a coherent foundation?

Method note. Integrative/critical interpretive synthesis to derive PASO from cross-framework regularities (Dixon-Woods et al., 2006).

Phase III — Development and Validation (Emergent Phase)

- Q4. How should PASO components be organized and integrated into the EDMC framework?
- Q5. How can EDMC be translated into practical tools that leaders can deploy across organizational contexts?
- Q6. To what extent does EDMC enhance ethical clarity, confidence, and application in varied leadership settings?

Method note. Theory-building and applied validation via case applications (Lynham, 2002; Yin, 2018), structured expert-panel review for content validity (Grant and Davis, 1997; Lynn, 1986), and comparative assessment against established frameworks.

These phases reflect the iterative logic of the research project: review → synthesize (PASO) → build and test (EDMC). *Figure 2* and *Figure 3* visualize the flow and loops among questions and PASO components.

1.5.3 Alignment of Questions, Limitations, Objectives, and PASO

The research questions, along with their alignment with limitations, objectives, and framework components, provide a structured yet adaptable approach to addressing the research gap identified in the introduction (Craft, 2013; Lehnert et al., 2015). This prioritization ensures that each aspect of the research contributes directly to developing the EDMC framework and addresses specific limitations in existing approaches to ethical decision-making (MacDougall et al., 2015; Treviño and Brown, 2004).

The matrix below illustrates how each question addresses specific limitations, advances concrete objectives, and relates to PASO. *Figure 2* and *Figure 3* visualize the phased flow and feedback loops that refine PASO and the EDMC architecture.

Table 2 Alignment of Research Questions with Addressed Limitations, Objectives, and PASO/EDMC Focus

Source: Author (2025). Original table created from study data and materials.

Research question	Addressed limitation(s)	Primary objective(s) served (Section 1.6.2)	PASO / EDMC focus
Primary RQ	Fragmentation; adaptability failures; lack of tools; skills neglect	Build an integrated, practice-ready framework drawing on theory and practice	All (P, A, S, O)
Q1	Fragmentation; inconsistent scope	1.6.2.1 Analyze and assess existing frameworks	Patterns → draft PASO mapping
Q2	Silos; cultural narrowness	1.6.2.1 Analyze and assess existing frameworks	Cross-framework patterns → PASO inputs
Q3	Integration challenges	1.6.2.2 Synthesize into PASO foundations	Define P, A, S, O components and interactions

Q4	Framework coherence	1.6.2.3 Develop EDMC conceptual framework	EDMC architecture (how PASO composes the system)
Q5	Lack of usable tools; cross-context portability	1.6.2.4 Structure EDMC and develop toolkit	Toolkit, guidance, anti-blindness mechanisms (A, S)
Q6	Limited outcome evidence	1.6.2.5 Validate EDMC effectiveness	Outcomes evidence (O) with P, A, S inputs; feedback into PASO/EDMC

Note: PASO = Principles, Actions, Skills, Outcomes.

This alignment highlights how the primary and secondary research questions target limitations in existing ethical decision-making frameworks, connect to specific research objectives, and integrate PASO components. This structure ensures the full problem space is addressed systematically. *Figure 2* maps the discovery-to-validation flow across Q1–Q6 and its linkage to PASO/EDMC, while *Figure 3* shows the validation and refinement loops that feed outcomes back into PASO and the EDMC architecture.

Figure 2 sketches the discovery-to-validation flow across Q1–Q6 and its linkage to PASO/EDMC. *Figure 3* and *Figure 3* show the feedback loops that return validation insights to refine PASO and the EDMC architecture. Detailed crosswalks appear in the relevant chapters.

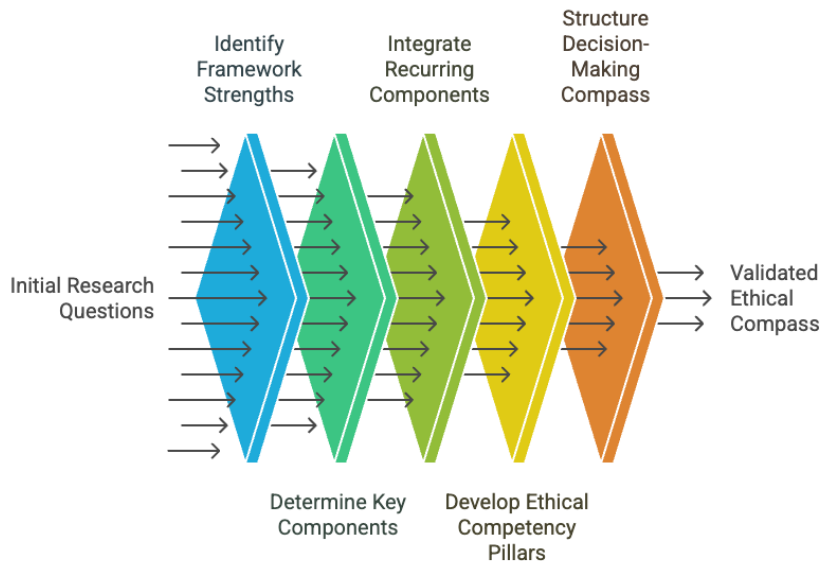


Figure 2 Research-Question Flow Across Phases and Linkage to the EDMC/PASO Architecture

Source: Author (2025). Original figure created from study data and materials.




Figure 3 Validation and Refinement Loops Linking Outcomes to the EDMC/PASO Architecture

Source: Author (2025). Original figure created from study data and materials.

Operational testing (Q6) produces evidence on clarity, confidence, and consistency; findings loop back to adjust PASO definitions/weights, update decision rules in Actions, refine Principles prioritization, and improve EDMC usability while preserving traceability.

1.6 Objectives – Long-Term Research Goal

This research aims to strengthen ethical decision-making in leadership practice by developing an integrated framework that connects theory with application. It leverages the strengths of existing models and addresses the gap between frameworks and practice by offering a simple, scalable approach that works under time pressure, mitigates ethical blindness, and operates across cultures.

By establishing a practical and widely applicable Ethical Decision-Making Compass (EDMC), the study seeks to equip leaders at all levels to navigate complex ethical dilemmas with greater confidence and consistency. The long-term goal is to help rebuild trust in leadership through enhanced ethical capability, improving governance across public, private, and nonprofit sectors in an increasingly interconnected world.

1.6.1 Primary Research Objective

Develop and validate a comprehensive ethical decision-making framework, the Ethical Decision-Making Compass (EDMC), that preserves proven elements of prior models while addressing their limitations, serving as a strategic, accessible “first-tier” tool that:

- provides clear, actionable guidance for leaders at all levels;
- functions effectively in both time-constrained and deliberative contexts;
- bridges theoretical rigor and practical application;
- accommodates diverse cultural perspectives on ethical reasoning;
- and forms a foundation on which more advanced ethical skills can be built.

Scope: a systematic analysis of sixteen frameworks; synthesis into PASO (Principles, Actions, Skills, Outcomes); development of the EDMC architecture and toolkit; and validation through comparative case applications and expert input. Organizational roll-out is noted but not the primary emphasis.

1.6.2 Research Objectives and Sub-Objectives

1.6.2.1 Analyze and Assess Existing Frameworks (Initial Phase: Q1, Q2)

- Systematically analyze and assess ethical decision-making frameworks using a six-dimensional lens (philosophical foundation, contextual adaptability, cultural

inclusivity, implementation guidance, competency requirements, evaluation mechanisms).

- Identify recurring patterns and insights that can inform the design of the framework.
- Offer a working definition of ethical decision-making suited to this research.
- Generate new knowledge on how fragmentation, limited adaptability, lack of tools, skills neglect, and weak responses to ethical blindness reduce the effectiveness of current approaches.

1.6.2.2 Synthesize into PASO Foundational Components (Bridge Phase: Q3)

- Map relationships among the identified patterns to clarify how effective ethical decisions are produced.
- Leverage, integrate, and synthesize strengths and insights into the four PASO components.
- Advance theory by showing how PASO components interact synergistically, retaining strengths while addressing known gaps.

1.6.2.3 Develop the EDMC Conceptual Framework (Emergent Phase: Q4)

- Specify how PASO components organize and integrate into the EDMC architecture.
- Articulate the interactions among PASO elements that produce a coherent decision process.

1.6.2.4 Structure EDMC and Develop the Practical Toolkit (Emergent Phase: Q5)

- Translate the EDMC into usable tools and guidance for diverse organizational contexts.
- Provide mechanisms that counteract ethical blindness and enhance awareness under time and cognitive constraints.
- Support cross-cultural governance by ensuring effectiveness across Western and non-Western paradigms.

1.6.2.5 Validate EDMC Effectiveness (Emergent Phase: Q6)

- Validate gains in ethical clarity, confidence, and practical application (e.g., stakeholder consideration, consistent reasoning).
- Demonstrate that EDMC preserves valuable elements of existing frameworks while addressing their limitations across governance, AI ethics, sustainability, and public administration.
- Show EDMC performance under time constraints—for both crisis response and strategic planning.

The objectives above are sequenced to answer the research questions: analyze → synthesize (PASO) → build EDMC → operationalize → validate. *Table 3* consolidates this alignment by mapping each research question to the objective(s) in Section 1.6.2 and the PASO/EDMC element engaged. For completeness, *Table 2* (Section 1.5.3) links the questions to the limitations they address, and *Figures 2–3* depict the phased flow and the validation feedback loops that feed results back into PASO and the EDMC architecture. The table below consolidates the objectives, the phase in which they are pursued, and the linked research questions; sub-objectives appear in the right-hand column.

Table 3 Objectives and Sub-Objectives with Linked Research Questions

Source: Author (2025). Original table created from study data and materials.

Objective	Sub-Objectives
Analyze and Assess Existing Ethical Decision-Making Frameworks (Initial Phase; RQs Q1–Q2)	<ul style="list-style-type: none"> • Systematically analyze ethical decision-making frameworks across disciplines using a structured six-dimension lens (philosophical foundation, contextual adaptability, cultural inclusivity, implementation guidance, competency requirements, evaluation mechanisms). • Identify recurring patterns and insights that can inform framework design. • Offer a working definition of ethical decision-making suited to this research. • Generate new knowledge on how fragmentation, limited adaptability, lack of tools, skills neglect, and weak responses to ethical blindness reduce the effectiveness of current approaches.
Synthesize into PASO Foundational Components (Bridge Phase; RQ Q3)	<ul style="list-style-type: none"> • Map relationships among the identified patterns to clarify how effective ethical decisions are produced. • Leverage, integrate, and synthesize strengths and insights into the four PASO components. • Advance theory by showing how PASO components interact synergistically, retaining strengths while addressing known gaps.

Develop the EDMC Conceptual Framework (Emergent Phase; RQ Q4)	<ul style="list-style-type: none"> Specify how PASO components organize and integrate into the EDMC architecture. Articulate the interactions among PASO elements that produce a coherent decision process.
Structure EDMC and Develop the Practical Toolkit (Emergent Phase; RQ Q5)	<ul style="list-style-type: none"> Translate the EDMC into usable tools and guidance for diverse organizational contexts. Provide mechanisms that counteract ethical blindness and enhance awareness under time and cognitive constraints. Support cross-cultural governance by ensuring effectiveness across Western and non-Western paradigms.
Validate EDMC Effectiveness (Emergent Phase; RQ Q6)	<ul style="list-style-type: none"> Validate gains in ethical clarity, confidence, and practical application (e.g., stakeholder consideration, consistent reasoning). Demonstrate how EDMC preserves valuable elements of existing frameworks while addressing their limitations across governance, AI ethics, sustainability, and public administration. Show EDMC performance under time constraints—for crisis response as well as strategic planning.
Enhance Global Ethical Leadership Capability (Implementation; beyond Q1–Q6)	<ul style="list-style-type: none"> Provide implementation guidelines for varied organizational settings and leadership levels. Increase accessibility of ethical tools across cultures and sectors. Support sector-spanning governance by aligning EDMC usage with existing compliance systems and controls.

This alignment keeps the development path coherent from analysis through validation while preserving a clear link to PASO and the EDMC architecture.

1.6.3 Strategic Implementation Objectives (beyond Q1–Q6)

- Provide implementation guidelines for varied organizational settings and leadership levels
- Increase accessibility of ethical tools across cultures and sectors
- Support cross-sector governance by aligning EDMC usage with existing compliance systems and controls

1.6.4 Expected Contributions and Benefits

This section summarizes the anticipated contributions if the research objectives are achieved. The contributions are organized in three layers, theoretical, practical, and societal, and are grounded in the PASO model and the EDMC architecture developed in this thesis. They indicate the intended impact of the work rather than pre-judging results; empirical support and qualification will be presented in Chapters 4–5. For context, *Table*

2 links research questions to objectives and PASO/EDMC focus, while *Figure 2*, and *Figure 3* visualize the phased build and feedback loops that underpin these contributions.

1.6.4.1 Theoretical Contributions

- Integrated understanding: a clear, structured account of ethical decision-making by combining insights from various disciplines.
- PASO model: a novel organizing framework (Principles, Actions, Skills, Outcomes) that specifies component interactions
- Conceptual clarity: a working definition of ethical decision-making aligned to PASO and practice needs
- Methodological value: an emergent, synthesis-driven approach that can be reused to study ethics under real-world constraints

1.6.4.2 Practical Contributions

- Usable first-tier tool: a practice-ready compass that works under time pressure and in deliberative settings
- Cross-cultural portability: guidance effective across Western and non-Western paradigms
- Anti-blindness mechanisms: structured prompts and checks that surface ethical salience when it is easiest to miss
- Implementation pathway: clear steps for piloting, adoption, and governance integration

1.6.4.3 Broader Societal Impact

- Trust rebuilding: more transparent, explainable decisions that can be audited and learned from

- Cross-sector collaboration: shared ethical language and mechanisms for public-private-nonprofit partnerships
- Capacity for complex challenges: tools that help leaders address AI governance, climate, and other systemic issues with ethical awareness beyond compliance

Note. Alignment of research questions with limitations, objectives, and PASO/EDMC focus appears in *Table 2* (Section 1.5.3). *Figure 2* and *Figure 3* visualize the phased flow and feedback loops used to refine PASO and the EDMC architecture.

1.7 Theoretical Framework

This study is grounded in a robust theoretical base that synthesizes insights across ethical decision-making theory, leadership studies, and applied ethics. Through a systematic comparative analysis of sixteen ethical decision-making frameworks, grouped into four categories (classical ethical theories, structured ethical models, leadership-oriented approaches, and contemporary applied-ethics frameworks), the research identifies recurring patterns, critical limitations, and essential components for an integrated approach. This theoretical grounding informs the development of the Ethical Decision-Making Compass (EDMC) and its PASO components: Principles, Actions, Skills, and Outcomes.

1.7.1 Foundational Theoretical Pillars

Each category contributes distinct strengths:

- Classical Ethical Theories (e.g., Deontology, Utilitarianism, Virtue Ethics) supply normative principles that guide moral reasoning and judgment (Kant, 1785; MacIntyre, 1984; Mill, 1863).

- Structured Models such as Rest’s Four-Component Model (1986) and Treviño’s Interactionist Model (1986) provide psychologically grounded process views of ethical cognition and action.
- Leadership-Oriented Frameworks (e.g., Transformational Leadership, Servant Leadership) align ethics with competencies, relationships, and influence (Bass, 1985; Greenleaf, 1977).
- Contemporary Applied-Ethics Frameworks, including GDPR (2016), AI ethics (L Floridi, 2020), and the Triple Bottom Line (Elkington, 1997), address current socio-technical challenges and offer implementation tools.

Chapter 2 details contributions and exposes five systemic limitations across the landscape: (1) disciplinary fragmentation, (2) weak adaptability under real-world constraints, (3) lack of actionable tools, (4) neglect of competency development, and (5) inadequate mechanisms to counter ethical blindness, especially under time pressure.

1.7.2 Toward an Integrated Framework: The PASO Components

Rather than discarding prior models, this research takes a constructive, integrative stance—synthesizing valuable elements into a framework that addresses the collective limitations above. The resulting PASO structure fuses:

- Normative grounding (from classical theories),
- Structured processes (from cognitive and organizational models),
- Competency-based insights (from leadership theories), and
- Pragmatic application (from contemporary frameworks).

This synthesis preserves academic rigor while emphasizing practical utility. It directly responds to identified methodological and theoretical gaps and follows best practices in integrative theory building (Alvesson and Sandberg, 2013; Gioia et al., 2013).

1.7.3 Bridging Theory and Practice

The EDMC's theoretical framework functions as a bridge between abstract ethical ideals and the lived complexity of leadership decisions by providing:

- Cross-disciplinary integration to reduce fragmentation,
- Adaptability mechanisms (e.g., tiered application protocols),
- Cognitive aids to mitigate ethical blindness (e.g., concise ethical heuristics),
- Competency-development pathways that treat ethics as a learnable leadership capability, and
- Implementation guidance that links awareness to action.

The development of this framework follows a three-phase design: Initial, Bridge, and Emergent, so that patterns, tools, and theory emerge from both scholarship and practice. This positions PASO and EDMC as theoretically grounded and operationally usable in real-world, time-pressured, and cross-cultural contexts.

1.8 Methodology

1.8.1 Introduction

This study follows a three-phase path. We begin with a comprehensive, comparative, matrix-based analysis of sixteen ethical decision-making frameworks, mapping strengths, limitations, fragmentation, and recurring design patterns across six analytic dimensions and against five systemic limitations (L1–L5) identified in the review. From this comparative mapping, four recurring strength-patterns consistently explain where decisions hold up in practice and indicate practical ways to close the five limitations. We consolidate and name these as the foundational components—Principles, Actions, Skills, Outcomes (PASO). These PASO components form the basis of the Ethical Decision-Making Compass (EDMC). Finally, we validate EDMC through

researcher-led analytic case applications and an expert-panel assessment, interpreting the two evidence streams together. The stance is pragmatic: we care about what works under real constraints and can be used by leaders in practice (Morgan, 2014). Methodologically, we adopt a Design Science Research approach (design–build–evaluate) with convergent mixed-methods, validation, and triangulation (Creswell and Plano Clark, 2018; Denzin, 1978; Hevner et al., 2004; Jick, 1979; Peffers et al., 2007).

1.8.2 Methodological Foundation

We base the study on pragmatism, practical knowledge for real decisions (Morgan, 2014). Our approach is discovery with discipline: we let patterns emerge and then test them, drawing on grounded-theory sensibilities and abductive reasoning to move between data and ideas (Charmaz, 2006; Timmermans and Tavory, 2012). We use design thinking to refine tools people can use, supported by reflective practice to check assumptions (Brown, 2008; Schön, 1983).

We apply a convergent mixed-methods design with explicit triangulation (Creswell and Plano Clark, 2018).

- Methodological triangulation. Two complementary validation streams—a researcher-led analytic case-application stream and an expert-panel assessment—are interpreted together (Denzin, 1978; Jick, 1979).

- Data/source triangulation. Literature-derived patterns inform scenario-based calibration and public-facts case applications, while expert judgments provide Likert $\geq 4/5$ and content-validity indices (I-CVI; S-CVI/Ave) (Lynn, 1986; Polit and Beck, 2006).

Phase-I empirical finding (foundational components). Comparative, matrix-based mapping across six analytic dimensions and against five systemic limitations (L1–L5) yields a stable, recurring four-part pattern. We consolidate and name these recurring

strengths as the foundational components—Principles, Actions, Skills, Outcomes (PASO)—and carry this empirically derived foundation into design and validation.

1.8.3 Three-Phase Research Design

This research consists of three phases:

Phase I — Initial: analysis and assessment (Q1, Q2)

- Systematically analyzes sixteen ethical decision-making frameworks.
- Identifies strengths, limitations, and recurring patterns.
- Uses a six-dimension lens: philosophical foundation, contextual adaptability, cultural inclusivity, implementation guidance, competency requirements, and evaluation mechanisms, explicitly examined against five systemic limitations (L1–L5).
- Methods: directed content analysis and comparative matrix mapping (Hsieh and Shannon, 2005).
- Output: an empirically derived PASO pattern that becomes EDMC’s foundational components.

Phase II — Bridge: integration and synthesis (Q3)

- Synthesizes Phase-I patterns into the PASO components.
- Methods: integrative/critical interpretive synthesis and visual summaries (Dixon-Woods et al., 2006).
- Documents component emergence via recursive pattern analysis.

Phase III — Emergent: development and validation (Q4, Q5, Q6)

- Integrates PASO into the EDMC framework.
- Translates EDMC into practical tools and guidance (e.g., a PASO-embedded rationale capture).

- Validates via researcher-led analytic case applications and an expert-panel assessment (N = 30) (Yin, 2018), plus comparative assessment against the 16 frameworks.

Figure 2 and *Figure 3* map the phased flow and the validation feedback loops; detailed crosswalks appear in the relevant chapters.

1.8.4 Analytical Methods

Directed content analysis and matrix-based, computer-assisted comparative analysis across six dimensions (Hsieh and Shannon, 2005), with visual summaries (heat-maps, comparison charts) to keep judgments transparent and traceable.

1.8.5 Quality Assurance

We support rigor through:

- Theoretical validity: explicit links to established ethical and leadership theories.
- Analytical reliability: anchored rubrics, structured matrices, version-controlled memos.
- Transparency: audit trails for coding, scoring rationales (≤ 15 words per metric), and design decisions.
- Triangulation: convergent interpretation of researcher-led applications and expert-panel evidence.
- Reflexivity: ongoing memoing and peer debriefs to surface assumptions (Schön, 1983).

1.8.6 Validation Methods

We validate EDMC through two streams interpreted together (triangulation):

- Stream A — Researcher-led analytic case applications.

- A1 Scenario-based calibration (H1–H8).
- A2 Public-facts case applications (C1–C8).
- Scoring: a 0–3 anchored rubric on six metrics (PASO linkage; blindness mitigation; cultural adaptability; time-pressure suitability; implementation guidance; outcomes evaluability) with transparent rationales; composites computed as means (0–3) and rescaled to 0–100; weighting sensitivity reported where relevant.
- Design note: single-rater scoring (no ICC/Kendall; not applicable).
- Stream B — Expert-panel assessment (N = 30).
- Evidence: Likert $\% \geq 4/5$ and content-validity indices (I-CVI; S-CVI/Ave).

Validation criteria. Theoretical robustness; practical applicability under constraint; cross-cultural relevance; technological alignment for digital ethics (Floridi, 2020).

1.8.7 Ethical Considerations and Reflexivity

We follow good-practice standards for informed consent, confidentiality, and data security (Israel, 2015). Reflexive practices include memoing, peer debriefing, and explicit assumption checks (Alvesson and Sköldbberg, 2017; Finlay, 2002).

1.8.8 Methodological Limitations and Innovations

This research has several limitations, but also brings innovations.

- Limitations.
 - Single-rater scoring (Stream A). No inter-rater statistics are applicable; anchors and sensitivity checks mitigate, but do not remove, subjectivity.
 - Public-facts dependence (C-cases). No internal data; findings are formative, not causal.
 - Scenario constraints (H-cases). Ecological realism is bounded.

- Expert-judgment replicability (Stream B). Content validity evidence reflects this expert sample at this time; replication with new samples would test stability and generalizability.

- Innovations.

- Strength-first constructive synthesis. Instead of a deficit-only review, we recombined recurring strengths across sixteen frameworks to derive PASO as foundational components—explicitly mapping against the six analytic dimensions and the five systemic limitations (L1–L5).

- Glass-box scoring. A six-metric, 0–3 anchored rubric with ≤ 15 -word rationales, stated N/A rules, and disclosed weighting/sensitivity produces auditable 0–100 composites.

- Dual-stream triangulation. Researcher-led analytic case applications + expert-panel content validity (Likert $\% \geq 4/5$; I-CVI; S-CVI/Ave) provide convergent, fit-for-purpose validation for a practice-oriented DBA study.

- PASO-embedded rationale capture. A brief, auditable instantiation of the architecture that links Why, How, Who, and What under time pressure.

1.8.9 Practical Application and Implementation

In this study, we apply the EDMC decision architecture (PASO foundational components) through three lightweight artifacts used in the two validation streams:

- a PASO-embedded rationale capture (architecture instantiation) that makes Why (Principles), How (Actions), Who (Skills/ownership), and What (Outcomes) visible and traceable—interpretive, not prescriptive;
- a six-metric, 0–3 anchored rubric with ≤ 15 -word rationales to score clarity, safeguards, capability, and evaluability;

- a brief reporting note (time-box, thresholds/KPIs, escalation path) to support auditability.

In Stream A (researcher-led analytic case applications), the rationale capture is completed within a ≤ 30 -minute time-box and scored with the rubric (composite 0–3, rescaled 0–100; N/A rules disclosed). In Stream B (expert panel), excerpts from the PASO-embedded capture serve as stimulus while experts judge clarity and coverage of PASO (Likert $\% \geq 4/5$; I-CVI; S-CVI/Ave). The two streams are interpreted together (triangulation) to assess fitness-for-use under constraint and traceability.

Moreover there is an intended use beyond this study. The primary asset is the EDMC decision architecture. Organizations can embed the architecture within existing decision, risk, and governance routines. Where documentation is required, a brief PASO-embedded rationale capture can provide proportional transparency for assurance and learning—again, architecture-instantiating, not checklist-prescribing.

1.8.10 Implementation Roadmap

- Phase I (Initial): completed Nov 2024
- Phase II (Bridge): completed Mar 2025
- Phase III (Emergent): Apr–Oct 2025

1.8.11 Future Research Directions

Field pilots and longitudinal studies; cross-cultural validation and localization heuristics; measurement science (including inter-rater reliability with multiple raters); lightweight digital deployment that preserves salience; applications in AI ethics and sustainability.

This methodology provides a transparent, phase-linked path from literature to a usable framework, aligning PASO and EDMC with both scholarly standards and real-world constraints.

1.9 Limitations, Delimitations, and Assumptions

No test–retest or multi-rater reliability. Researcher-led case scores were produced by a single assessor using disclosed anchors and one-line rationales; inter-rater reliability and test–retest checks were not performed and are left to future work.

Hypothetical applications, not field trials. Case demonstrations use public documentation to evidence analytic transferability; they do not estimate population effects or organizational impact.

1.9.1 Limitations

This study acknowledges several methodological and conceptual limits.

- Coverage and corpus
 - Framework selection and coverage. The analysis considers sixteen ethical decision-making frameworks; it cannot represent the entire range. Selection emphasized (i) theoretical diversity (e.g., deontological, utilitarian, virtue-ethics, stakeholder traditions (MacIntyre, 2007); (ii) professional context spread (business/leadership ethics, AI governance, sustainability) (O’Fallon and Butterfield, 2005); (iii) temporal range (classic to contemporary); and (iv) implementation specificity (from conceptual models to applied tools). This improves breadth and comparability but may exclude emerging or less-documented approaches.

- Cultural representation. Despite efforts to include diverse perspectives, the literature base is weighted toward Western traditions—reflecting publication patterns rather than intentional exclusion (Donaldson and Dunfee, 1999).
- Language and publication bias. Reliance on English-language, peer-reviewed and practitioner sources may underrepresent non-English or grey-literature contributions.
 - Design and validation limits
- Single-rater scoring (Stream A). Inter-rater statistics are not applicable; anchored scoring and sensitivity checks mitigate—but do not remove—subjectivity.
- Scenario constraints (H-cases). Hypothetical scenarios aid (Stream B) clarity and speed but limit ecological realism.
- Public-facts dependence (C-cases). Applications rely on verifiable public sources; no internal data were accessed.
 - Measurement and integration
- Measurement challenges. Assessing ethical decision quality is context-sensitive and multidimensional; the proposed measures acknowledge these constraints (Schwartz, 2016).
- Theoretical integration tensions. Synthesizing across philosophy, psychology, economics, leadership, coaching, and AI ethics requires reconciling field-specific meanings (Scherer and Palazzo, 2011).
 - Researcher stance and context
- Researcher positionality. Practitioner insight can introduce interpretive bias; reflexive memoing, audit trails, and peer debriefs mitigate but do not eliminate this risk (Alvesson and Sköldberg, 2017; Finlay, 2002).
- Context volatility. Rapid shifts in AI and regulation can outpace any static framework; EDMC will require periodic updates.

Future work (pointer). Subsequent phases should add multi-rater scoring with ICC, a short test–retest, and organizational field studies to assess external validity.

1.9.2 Delimitations

This study makes intentional scope choices to keep the inquiry focused and tractable:

- Domain focus. The primary unit of interest is leadership decision-making in organizational contexts (public, private, nonprofit). Interpersonal/household ethics and clinical bioethics are outside the scope.
- Corpus size and language. Sixteen frameworks are analyzed, drawn primarily from English-language literature. Classic sources and contemporary models are included; lesser-documented or non-English frameworks may be underrepresented (see 1.9.1).
- Level of analysis. The inquiry centers on framework constructs and processes, not on large-sample outcome effects.
- Validation modes. Early validation relies on expert panel review and hypothetical case applications, plus comparative mapping against existing frameworks; it does not include field experiments or randomized trials (Grant and Davis, 1997; Lynham, 2002; Lynn, 1986).
- Method orientation. Emphasis is on emergent, synthesis-driven theory-building and case-logic rather than statistical generalization (Dixon-Woods et al., 2006; Edmondson and McManus, 2007; Eisenhardt et al., 2016; Yin, 2018).
- Sector coverage. Examples span corporate governance, AI ethics, sustainability, and public administration; deep sectoral rulebooks and micro-policy specifics are out of scope.

- Implementation scope. The work delivers an architecture and toolkit with implementation guidelines; full organizational change programs (training at scale, IT integration, incentive redesign) are beyond scope.
- Cultural/legal scope. EDMC aims for cross-cultural portability and references major regimes (e.g., GDPR), but does not provide jurisdiction-by-jurisdiction legal analysis.
- Ethical traditions. Core emphasis is on deontology, consequentialism, virtue/care, and stakeholder perspectives; religious or communitarian traditions are acknowledged but not treated in depth.

These delimitations maintain focus on building and preliminarily validating a practical, portable decision framework (EDMC).

1.9.3 Assumptions

This study proceeds under the following working assumptions, which inform the design, validation, and interpretation of findings:

- Ethical universals amid cultural variation. Certain principles, e.g., human dignity, fairness, transparency, have broad cross-cultural salience, though interpretation and priority vary locally (Donaldson and Dunfee, 1999; Hofstede, 2001).
- Improvability of ethical decision-making. Decision quality can be improved through clear, simple, and structured processes (breakdown), prompts, and skill development—beyond character or culture alone (Rest et al., 1999; Treviño and Nelson, 2016).
- Value of integration (PASO). Integrating Principles, Actions, Skills, Outcomes provides a more complete decision structure than approaches focusing on a single dimension (Lynham, 2002).

- Cross-sector transferability. Core process elements of ethical decision-making are sufficiently common to support a cross-disciplinary framework with local adaptations (Kaptein, 2017; Newman et al., 2017).
- Time pressure as a pervasive constraint. Time pressure systematically affects ethical judgment across organizational contexts (Bazerman and Tenbrunsel, 2011).
- Ethical blindness as contextual. Ethical salience is often masked by situational forces rather than fixed dispositions; surfacing context is therefore essential (Palazzo et al., 2012).
- Validity of an emergent, iterative method. Allowing constructs to emerge through iterative analysis is appropriate for complex, practice-oriented problems (Edmondson and McManus, 2007; Gioia et al., 2013; Jacobson et al., 1999).

By articulating these limitations, delimitations, and assumptions, the research maintains methodological transparency while reinforcing its commitment to both theoretical rigor and practical relevance.

1.9.4 Conclusions

Chapter 1 set the stage for this thesis. It framed the leadership challenge of acting ethically beyond compliance and explained why organizations need a practical, portable approach that works under time pressure and across cultures. It stated the research problem and purpose, outlined the research questions, and clarified the theoretical, practical, and societal significance of developing the Ethical Decision-Making Compass (EDMC).

The chapter also introduced the study's design: a three-phase path in which a comprehensive, comparative, matrix-based analysis of existing frameworks (viewed across six analytic dimensions and against five systemic limitations, L1–L5) leads to the

emergence of PASO—the foundational components (Principles, Actions, Skills, Outcomes) that structure EDMC. It summarized how EDMC will be validated through two convergent evidence streams—researcher-led analytic case applications and an expert-panel assessment—and it set appropriate expectations by stating the study’s limitations, delimitations, and assumptions.

Together, these elements provide the rationale, aims, and roadmap for the work ahead: to translate well-grounded ethical insight into a framework leaders can use; where documentation is needed, a brief PASO-embedded rationale captures links to Why, How, Who, and What in a traceable way.

Next: Chapter 2 surveys the major ethical decision-making frameworks, identifies recurrent gaps and strengths, and details the comparative analysis that motivates the PASO components on which EDMC is built.

CHAPTER II: REVIEW OF LITERATURE

2.1 Ethical Decision-Making - Theoretical Foundations and Review Strategy

2.1.1 Introduction – Purpose and Scope

This chapter thoroughly reviews sixteen ethical decision-making frameworks across four streams, highlighting their complementary strengths and common limitations. Instead of dismissing previous work, the analysis clarifies which aspects should be maintained and where innovation is necessary. From this synthesis, the PASO foundational components—Principles, Actions, Skills, Outcomes—stand out as an integration of valuable elements from various traditions and as targeted responses to gaps.

The analytical journey unfolded through seven interconnected steps. The analytical journey unfolded through seven interconnected steps (*Figure 25*). Note: *Figure 25* retains the Research Proposal's Chapter IV step numbering for provenance; chapter numbering in this thesis differs.

We conduct a comprehensive, comparative, matrix-based analysis of sixteen frameworks across six analytic dimensions and read findings against five systemic limitations (L1–L5), allowing recurring strengths to surface and later synthesize into PASO. This comparison highlights complementary strengths within the four streams: theoretical depth in classical ethical theories, systematic process in structured models, competency insights in leadership frameworks, and practical tooling in contemporary applied approaches.

The synthesis also reveals consistent patterns of limitations that restrict effectiveness. Using a structured gap-analysis matrix, we show how each framework addresses distinct aspects of ethical decision-making and where limitations persist. The goal is integrative: keep what works, and specify where innovation is required.

A key insight is the identification of foundational components to preserve and integrate in any comprehensive approach. By tracing essentials across frameworks, the analysis shows how PASO emerged through comparative mapping—not as a prior assumption but as a grounded consolidation of recurring strengths. This prepares the way for the Ethical Decision-Making Compass (EDMC), which builds on PASO to bridge long-standing gaps between theory and practice, philosophical depth and practical use, and universal principles and contextual needs. *Figure 25* provides a visual roadmap of the journey’s te, showing how each step builds on previous insights to map the landscape and its opportunities for improvement.

2.1.2 Background and Evolution of Ethical Decision-Making Frameworks

Numerous ethical decision-making models exist. A systematic, comparative review *Table 37* addresses the field’s fragmentation by tracing the evolution from normative theory to applied mechanisms (*Figure 4*), surfacing enduring strengths and boundary conditions that inform an integrated approach (Treviño and Nelson, 2016). Fragmentation leaves leaders less equipped to navigate complex ethical dilemmas across diverse cultural and organizational contexts—precisely when ethical guidance is most needed (Kaptein, 2023; Palazzo et al., 2012).

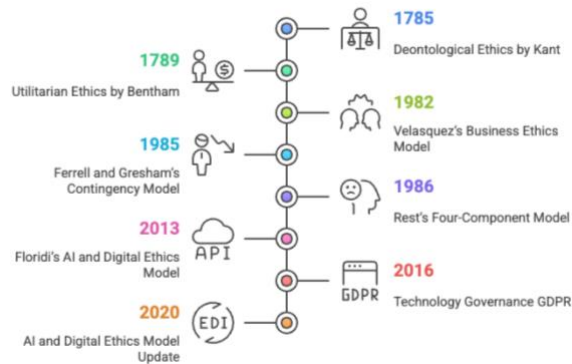


Figure 4 Historical Evolution of the Ethical Decision-Making Frameworks
 Source: Author (2025). Original figure created from study data and materials.

This review analyzed sixteen frameworks—listed in *Table 4*, and compared in detail in *Table 38*, to identify strengths, limitations, recurring patterns, and foundational components through comprehensive, matrix-based analysis. Grouping frameworks into four streams enabled us to trace the field’s movement from abstract principles to practical applications in complex environments (Kaptein, 2023). That historical perspective helped locate critical gaps while highlighting essential elements worth preserving.

The review served three purposes:

1. Categorize and critically analyze ethical frameworks across disciplines;
2. Identify strengths, essential components, recurring patterns, and limitations via systematic comparative analysis;
3. Establish the theoretical foundation for developing the Ethical Decision-Making Compass (EDMC).

This comparative groundwork sets the stage for an integrated approach that transcends disciplinary silos while maintaining theoretical rigor and practical applicability.

2.1.3 Review Scope and Method

We focus on peer-reviewed scholarship and policy/standards sources relevant to ethical decision-making in organizational and technology-mediated contexts.

Databases and sources: Scopus; Web of Science Core Collection; Google Scholar; targeted searches in IEEE Xplore and ACM Digital Library for AI-related ethics; and official repositories for governance instruments (EUR-Lex for EU law, OECD iLibrary, and the ISO/IEC Online Browsing Platform).

Date window: Contemporary/organizational literature: 1977–2025. Foundational philosophical texts: no lower bound (to include canonical works).

Search strategy: Keyword strings combined controlled and free terms such as: ethical decision-making; deontology; utilitarianism; virtue ethics; moral psychology; ethical leadership; servant leadership; transformational leadership; authentic leadership; governance; GDPR; EU AI Act; AI ethics; ISO/IEC; algorithmic accountability; risk management.

Inclusion criteria: (i) Conceptual clarity and practical relevance to decisions/governance; (ii) methodological transparency for empirical works; (iii) explicit treatment of decision processes, mechanisms, or safeguards.

Exclusion criteria: Narrowly technical papers lacking ethical content; commentary without identifiable argument/evidence.

Comparative method: Sixteen frameworks were selected using transparent inclusion criteria and analyzed via directed content analysis (Hsieh and Shannon, 2005) and integrative/critical interpretive synthesis (Dixon-Woods et al., 2006). A comparative matrix mapped contributions/limits to PASO anchors (Principles, Actions, Skills, Outcomes) to identify where EDMC requires multi-perspective synthesis. The definitive set of frameworks appears in *Table 4*, and the full matrix in *Table 36*.

Clarification. “Technology governance” is an umbrella for institutional, regulatory, and standards-based instruments—e.g., GDPR (2016), OECD AI Principles (2019), ISO/IEC standards, and the EU AI Act—mechanisms that codify principles and controls beyond organization-specific policies.

Table 4 Sixteen Frameworks Included in the Review

Source: Author (2025). Original table created from study data and materials.

Category	Framework and Author(s)/Year
Foundational ethical theory	Deontological Ethics (Kant, 1785)
	Utilitarian Ethics (Bentham, 1789; Mill, 1863)
	Virtue Ethics (Hursthouse, 1999; MacIntyre, 1984)
	Stakeholder Theory (Freeman, 1984)
Structured decision models	Rest’s Four-Component Model (Rest, 1986)
	Treviño’s Interactionist Model (Treviño, 1986)

	Ferrell and Gresham’s Contingency Model (Ferrell and Gresham, 1985)
	Velasquez’s Business Ethics Model (Velasquez, 1982)
	Hunt and Vitell’s Theory of Marketing Ethics (Hunt and Vitell, 1986)
Leadership-oriented approaches	Transformational Leadership (Bass, 1985)
	Servant Leadership (Greenleaf, 1977)
	Ethical Leadership in Public Administration (Denhardt and Denhardt, 2003)
Technology governance / digital ethics	Triple Bottom Line (Elkington, 1997)
	AI and Digital Ethics (Floridi, 2013; 2020)
	General Data Protection Regulation (GDPR) (2016)
	Big Data and Smart Cities Ethics (Kitchin, 2014)

2.1.4 Literature Review Methodology and Framework Selection

This review followed a discovery-oriented, iterative approach aligned with the overall research design (Chapter 3). Rather than a static catalogue, the review functioned as an active site of exploration, supported by constructivist grounded-theory sensibilities (Charmaz, 2006), abductive analysis (Timmermans and Tavory, 2012), and problematization (Alvesson and Sandberg, 2013). Guided by a problem-definition-first logic (Dorst, 2015; Wedell-Wedellsborg, 2020), it advanced through recursive cycles of comparison → synthesis → refinement, consistent with the three-phase design (Initial, Bridge, Emergent; Section 1.5.2).

Through this process, PASO—the foundational components (Principles, Actions, Skills, Outcomes)—emerged through comparative mapping, not as a prior assumption. This is a methodological contribution: rigorous literature synthesis can generate novel constructs that subsequently guide methodological choices.

The sixteen frameworks were examined with this triangulated analytical stance and then grouped into four streams—(1) foundational ethical theory, (2) structured decision models, (3) leadership-oriented approaches, and (4) technology governance/digital ethics—to surface best practices and gaps in each stream. This organization also allowed us to trace the field’s evolution from abstract theory to practical mechanisms (see *Figure 5*).

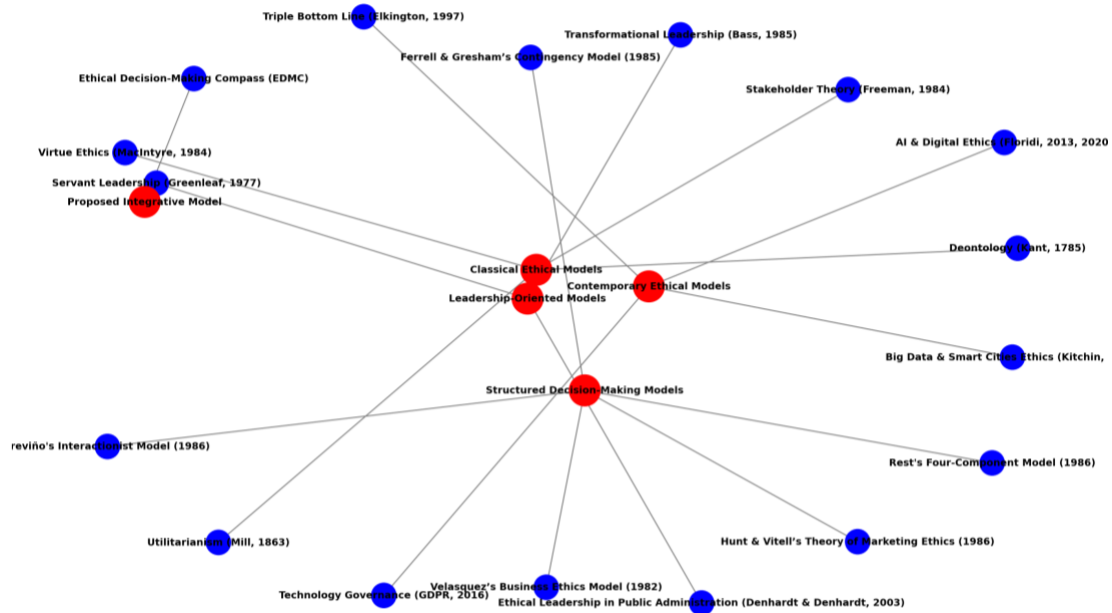


Figure 5 Comparative Evolution of the Ethical Decision-Making Frameworks
 Source: Author (2025). Original figure created from study data and materials.

Selection rationale: Frameworks were chosen for historical significance, academic impact, professional adoption, and theoretical diversity; they reflect multiple traditions (e.g., deontology, utilitarianism), leadership relevance, coverage of diverse domains (business, AI, sustainability), and both historical and contemporary approaches. To assess responsiveness to five critical limitations identified in the literature, a three-tier legend was used in the matrix: ✓ = addressed; ● = partially addressed; ✗ = not addressed. For procedural detail, see Chapter 3 (Framework Selection Criteria).

2.2 Comparative Review of Ethical Decision-Making Frameworks

This comparative analysis provides the foundation for the Ethical Decision-Making Compass (EDMC), a practice-oriented tool designed to navigate complex ethical dilemmas with clarity and confidence. The sections that follow examine sixteen frameworks grouped into four streams: foundational ethical theory, structured decision

models, leadership-oriented approaches, and technology governance/digital ethics. For each stream, the review considers theoretical principles, core components, and practical applications.

2.2.1 Foundational Ethical Theory

Key focus. Philosophical underpinnings that shape subsequent structured models and inform what counts as legitimate reasons, constraints, and virtues in ethical decision-making.

These traditions supply the philosophical groundwork on which later, more structured decision models were developed. Their primary contribution is normative depth, clarifying duties, outcomes, or character, while their primary limitation is operational guidance in real-world, multi-stakeholder contexts (see *Table 5*).

2.2.1.1 Deontological and Utilitarian Perspectives

1. Deontological ethics (Kant, 1785). Duty-based ethics emphasize moral obligations and rule-guided action, offering moral certainty and rights protection; however, they often struggle with conflicting duties and context-sensitive trade-offs (Alexander and Moore, 2016).
2. Utilitarian ethics (Bentham, 1789; Mill, 1863). Outcome-focused reasoning prioritizes the maximization of overall benefit, providing pragmatic metrics for evaluation; yet it faces challenges in quantifying long-term and cross-stakeholder impacts and in guarding minority interests (Driver, 2012).

2.2.1.2 Virtue Ethics and Stakeholder Theory

- Virtue ethics (Hursthouse, 1999; MacIntyre, 1984). Centers ethical character and moral motivation, enriching the *why* behind ethical action; its limitation is less

concrete guidance for specific decisions and a weaker procedural structure (Kristjánsson, 2020).

- Stakeholder theory (Freeman, 1984). Broadens moral scope beyond shareholders and integrates corporate responsibilities across stakeholder groups; typical limitations involve prioritizing among conflicting stakeholder interests and specifying actionable resolution mechanisms (Harrison et al., 2019).

These foundational traditions shaped the emergence of structured frameworks that formalized moral reasoning. They offer indispensable moral justifications and constraints, but require contextualization and operationalization to support decisions in dynamic environments, see *Figure 5*.

Table 5 Foundational Ethical Theory: Contributions and Typical Challenges
Source: Author (2025). Original table created from study data and materials.

Framework	Core emphasis	Typical challenge
Deontological ethics (Kant, 1785)	Duty, rules, rights	Conflicting duties; limited flexibility (Alexander and Moore, 2016)
Utilitarian ethics (Bentham, 1789; Mill, 1863)	Maximization of overall benefit	Quantifying long-term, cross-stakeholder impacts (Driver, 2012)
Virtue ethics (Hursthouse, 1999; MacIntyre, 1984)	Moral character, dispositions	Limited procedural guidance for specific decisions (Kristjánsson, 2020)
Stakeholder theory (Freeman, 1984)	Multi-stakeholder responsibility	Prioritization under conflict; actionable mechanisms (Harrison et al., 2019)

Foundational care-ethics sources (e.g., Gilligan, 1982; Held, 2006) are cited for context but not included in the set of 16 frameworks analyzed. Care ethics (Gilligan, 1982; Held, 2006) emphasizes relational responsibilities and context-aware responsiveness. It complements the four main traditions but is cited here as an interpretive lens rather than a primary framework in the comparative set.

Implications for PASO/EDMC:

- Contribution: Supplies Principles (norms/constraints), the justificatory logic for Outcomes, and the motivational basis for Skills (virtues).
- Limitation: Under-specifies Actions (repeatable steps) and conflict resolution when principles and stakeholder interests collide.
- Integration: PASO anchors these necessities; EDMC brings a multi-perspective synthesis that makes trade-offs explicit and supports auditable recommendations.

2.2.2 Structured Decision-Making Models

Key focus. Development of operational models that incorporate ethical reasoning in business and governance while specifying steps, roles, and evaluative criteria.

These models provide systematic procedures for ethical decision-making, yet their effectiveness in complex, real-world contexts often varies by organizational conditions, stakeholder diversity, and evidence requirements (see *Table 6*).

2.2.2.1 Cognitive and Psychological Models

- Rest's Four-Component Model (1986). Identifies moral sensitivity, judgment, motivation, and character, offering a comprehensive psychological account of why ethical action succeeds or fails; however, it tends to underweight organizational influences on perception and behavior (Sonenshein, 2007).
- Treviño's Interactionist Model (1986). Integrates individual and situational factors (e.g., ethical climate, reinforcement systems), providing richer contextualization; limitations include limited standardized assessment tools for consistent application across settings (Kaptein, 2008).

2.2.2.2 Corporate and Business Ethics Models

- Ferrell and Gresham’s Contingency Model (1985). Applies ethical analysis within marketing/corporate decision contexts, highlighting environmental contingencies; critics note an insufficient depth of moral analysis beyond compliance-oriented reasoning (Jones, 1991).
- Velasquez’s Business Ethics Model (1982). Aligns corporate responsibility with stakeholder engagement and organizational policy guidance; the framework does not explicitly incorporate individual moral cognition, limiting micro-level diagnostic power (Donaldson and Dunfee, 1999).
- Hunt and Vitell’s Theory of Marketing Ethics (1986). Bridges deontological and teleological evaluation within decision processes and acknowledges cultural influences; empirical testing remains concentrated in marketing domains (McClaren, 2015; O’Fallon and Butterfield, 2005).
- Overall, these structured models contribute clear steps, role definitions, and evaluative logics. Their typical limitations are disciplinary scope, measurement standardization, and the integration of organizational conditions with individual cognition—constraints that reduce cross-sector portability in global, high-stakes environments.

Table 6 Structured Ethical Decision-Making Models

Source: Author (2025). Original table created from study data and materials.

Framework	Core emphasis	Typical challenge
Rest’s Four-Component Model (1986)	Moral sensitivity, judgment, motivation, character	Underweights organizational influences (Sonenshein, 2007)
Treviño’s Interactionist Model (1986)	Person–situation interaction; ethical climate	Lacks standardized assessment tools (Kaptein, 2008)
Ferrell and Gresham’s Contingency Model (1985)	Environmental/marketing contingencies in decisions	Limited moral depth beyond compliance (Jones, 1991)
Velasquez’s Business Ethics Model (1982)	Corporate responsibility and stakeholder alignment	Omits individual moral cognition (Donaldson and Dunfee, 1999)
Hunt and Vitell’s Theory of Marketing Ethics (1986)	Deontological + teleological evaluation; culture	Limited testing beyond marketing (Hunt and Vitell, 1986; Hunter et al., 2013)

Implications for PASO/EDMC:

- **Contribution:** Specifies Actions (repeatable steps) and clarifies evaluation logic for Outcomes; surfaces organizational levers relevant to Skills (e.g., climate, reinforcement).
- **Limitation:** Partial integration of Principles across conflicting stakeholder claims; limited standardization for cross-sector use.
- **Integration:** PASO anchors stepwise practice (Actions) with principled constraints (Principles) and capability needs (Skills), while EDMC provides a multi-perspective synthesis to reconcile trade-offs and render recommendations auditable.

2.2.3 Leadership-Oriented Approaches

Key focus. The role of ethical influence and organizational culture in shaping norms, behaviors, and climates supportive of ethical decision-making.

Leadership-oriented frameworks contribute valuable perspectives on ethical influence, role modeling, and organizational culture, yet their practical implementation often varies across contexts, authority structures, and accountability mechanisms (see *Table 7*).

- **Transformational leadership** (Bass, 1985). Highlights ethically oriented change driven by vision, values, and inspiration; critiques note equity-related challenges and potential over-concentration of ethical agency in formal leaders (Eagly, 2007).
- **Servant leadership** (Greenleaf, 1977). Centers the well-being of stakeholders and service-first motivation; limitations include less-defined governance and accountability mechanisms for scalable implementation (Eva et al., 2019).
- **Ethical leadership in public administration** (Denhardt and Denhardt, 2003). Encourages transparency, accountability, and stewardship within public service; faces

institutional constraints and bureaucratic resistance that can impede sustained practice (Van Wart, 2013).

These approaches enrich ethical climates and competencies (Skills), yet benefit from coupling with repeatable safeguards (Actions) and principled constraints (Principles) for cross-context consistency.

Table 7 Leadership-Oriented Ethical Decision-Making Frameworks and Typical Challenges

Source: Author (2025). Original table created from study data and materials.

Framework	Core emphasis	Typical challenge
Transformational leadership (Bass, 1985)	Vision-led ethical change; inspiration; values alignment	Equity-related challenges; concentrated ethical agency (Eagly, 2007)
Servant leadership (Greenleaf, 1977)	Stakeholder well-being; service-first motivation	Limited governance/accountability mechanisms (Eva et al., 2019)
Ethical leadership in public administration (Denhardt and Denhardt, 2003)	Transparency; accountability; stewardship	Institutional constraints; bureaucratic resistance (Van Wart, 2013)

Implications for PASO/EDMC.

- **Contribution:** Strengthens Skills (relational/ethical competencies) and supports Outcomes through climate-building and accountability norms.
- **Limitation:** Under-specifies Actions (repeatable steps/safeguards) and may dilute Principles when values conflict across stakeholders or levels.
- **Integration:** PASO links leadership behaviors to codified practice (Actions) under principled constraints (Principles); EDMC provides multi-perspective synthesis and auditable recommendations, distributing ethical agency beyond individual leaders and embedding it organization-wide.

2.2.4 Contemporary Challenge-Specific Frameworks (e.g. Sustainability, AI, Governance)

Key focus. Emerging instruments and models for technological, sustainability, and governance ethics, with attention to implementation consistency, standardized metrics, and cross-jurisdictional applicability.

These frameworks address contemporary ethical frontiers in sustainability, AI/digital systems, and governance, providing safeguards, documentation, and enforceable standards. Yet they often operate in isolation from broader ethical traditions, contributing to the fragmentation identified in this review (see *Table 8*).

- Triple Bottom Line (Elkington, 1997). Integrates social, environmental, and economic responsibilities, expanding ethical consideration beyond financial outcomes; common issues include inconsistent implementation, measurement standardization, and prioritization across the three pillars (Slaper and Hall, 2011).
- AI and Digital Ethics (Floridi, 2013; 2020). Articulates principles such as fairness and transparency for digital ecosystems; limitations include a lack of standardized ethical metrics and difficulties evaluating trade-offs at scale (Mittelstadt et al., 2016).
- General Data Protection Regulation GDPR (2016). Provides rights-based data protection and accountability mechanisms (e.g., documentation, DPIAs, lawful bases), but faces cross-jurisdictional implementation challenges and rapid technological change (Floridi, 2013).
- Big Data and Smart Cities Ethics (Kitchin, 2014). Addresses ethical issues at the intersection of datafication, urban governance, and public space; regulatory frameworks remain fragmented and unevenly enforced (Mökander et al., 2023).

Taken together, these contemporary models sharpen attention to evidence, documentation, and safeguards, but they remain siloed within their domains. To support complex, cross-sector decisions, they benefit from integration with philosophically grounded justifications, leadership skills, and structured decision steps.

Table 8 Contemporary Ethical Challenges: Sustainability, AI, and Governance
Source: Author (2025). Original table created from study data and materials.

Framework	Core emphasis	Typical challenge
Triple Bottom Line (Elkington, 1997)	Social, environmental, and economic responsibility	Implementation inconsistencies; measurement/prioritization issues (Slaper and Hall, 2011)
AI and Digital Ethics (Floridi, 2013; 2020)	Fairness, transparency, and digital ecosystem principles	Lack of standardized ethical metrics; evaluation at scale (Mittelstadt et al., 2016)
GDPR (2016)	Rights-based data protection; accountability mechanisms	Cross-jurisdictional implementation; rapid tech change (Floridi, 2013)
Big Data and Smart Cities Ethics (Kitchin, 2014)	AI/data ethics in urban planning and governance	Fragmented regulatory frameworks (Mökander et al., 2023)

Related instruments (not part of the 16 frameworks). Widely used comparators such as the OECD AI Principles (2019) and the NIST AI Risk Management Framework (2023) inform practice and are referenced where relevant, but they are not included in the comparative set to preserve scope

Implications for PASO/EDMC.

- Contribution: Codifies Principles (rights/safeguards), prescribes Actions (documentation, DPIAs, controls), and sets evidence demands for Outcomes.
- Limitation: Siloed domain scope; limited integration with foundational justifications and leadership skills.
- Integration: PASO aligns safeguards and metrics with principled constraints and capability needs; EDMC provides multi-perspective synthesis to balance competing values and render recommendations auditable across jurisdictions and sectors.

Taken together, the four streams provide the raw materials for a coherent, portable practice: PASO consolidates its strengths while EDMC operationalizes them under real-world constraints—linking principles, structured actions, skills development, and outcomes we can audit and learn from.

2.3 Critical Patterns, Fragmentations, and Limitations in Frameworks

2.3.1 Key Dimensions of Comparative Analysis

The sixteen frameworks were evaluated across six dimensions—philosophical foundations, implementation guidance, cultural inclusivity, adaptability under time constraints, competency requirements, and evaluation mechanisms *Table 37*.

The comparative analysis surfaced commonalities and differences across traditions, sectors, and contexts—highlighting strengths, recurring limitations, and critical gaps.

2.3.1.1 Philosophical Foundations and Theoretical Underpinnings

Classical philosophical frameworks (Bentham, 1789; Hursthouse, 1999; Kant, 1785; MacIntyre, 1984; Mill, 1863) offer robust normative principles yet provide limited mechanisms for practical implementation. Stakeholder theory (Freeman, 1984) partially bridges this gap by connecting principles to organizational contexts.

Structured decision models (Ferrell and Gresham, 1985; Hunt and Vitell, 1986; Rest, 1986; Treviño, 1986; Velasquez, 1982) integrate philosophical lenses with practice but often prioritize organizational concerns over broader ethical justification. Notably, Hunt and Vitell (1986) synthesize deontological and teleological evaluation more explicitly than most structured models.

Leadership-oriented approaches (Bass, 1985; Denhardt and Denhardt, 2003; Greenleaf, 1977) draw on service, transformation, and public-interest principles but are less systematically integrated with established ethical theories.

Contemporary frameworks (Elkington, 1997; GDPR, 2016; Floridi, 2013; 2020; Kitchin, 2014) introduce domain-specific principles, sustainability, information ethics, rights-based protections, and urban technology ethics, yet frequently remain siloed from foundational traditions, limiting conceptual continuity.

Overall, a gap persists between theoretical depth and actionable guidance, underscoring the need for an integrated approach that preserves philosophical rigor while enabling practical decision-making in organizational contexts.

2.3.1.2 Implementation Guidance and Contextual Adaptability

Structured frameworks such as Rest (1986) and Treviño (1986) specify systematic steps but often presuppose deliberative conditions that are rare in organizational practice.

Ferrell and Gresham (1985) offer a contingency-oriented structure, while Velasquez (1982) presents a balanced decision process integrating multiple perspectives. Hunt and Vitell (1986) outline a theory-guided process but offer limited application tools.

Leadership-oriented frameworks (Bass, 1985; Denhardt and Denhardt, 2003; Greenleaf, 1977) contribute relational insight yet frequently lack repeatable decision protocols; Ethical Leadership in Public Administration (Denhardt and Denhardt, 2003) remains largely governance-centered. Contemporary approaches, Triple Bottom Line (Elkington, 1997), AI ethics (Floridi, 2013; 2020), GDPR (2016), and smart cities ethics (Kitchin, 2014), provide domain-specific guidance but show limited cross-context portability.

Few frameworks address time pressure explicitly—even though empirical work shows that time constraints impair deliberative ethical reasoning (Del Popolo Cristaldi et al., 2024; Emser et al., 2021).

2.3.1.3 Cultural Inclusivity and Cross-Contextual Application

Classical Western frameworks (Bentham, 1789; Kant, 1785; Mill, 1863) often presume universal applicability without adequately engaging diverse cultural perspectives. Virtue ethics (Hursthouse, 1999; MacIntyre, 1984) and stakeholder theory (Freeman, 1984) show variable cultural applicability and moderate adaptation, respectively. Among structured models, Hunt and Vitell (1986) explicitly acknowledge

cultural influence, and Treviño (1986) demonstrates moderate cultural sensitivity; few frameworks provide robust mechanisms for cross-cultural adaptation.

Within leadership models, servant leadership (Greenleaf, 1977) exhibits moderate cultural adaptability, while Ethical Leadership in Public Administration (Denhardt and Denhardt, 2003) shows context-dependent adaptability across governance systems. Contemporary frameworks present mixed results: Triple Bottom Line (Elkington, 1997) has moderate global applicability; GDPR (2016) faces cross-jurisdiction challenges tied to its European-centered design; big data and smart cities ethics (Kitchin, 2014) vary by local urban context.

This dimension reveals a persistent Western-centric orientation that limits global applicability. Few frameworks specify mechanisms for cultural adaptation, reinforcing the need to balance normative guidance with local legitimacy (see also Integrative Social Contracts Theory; Donaldson and Dunfee, 1999).

2.3.1.4 Competency Requirements and Skills Development

Classical frameworks often assume sophisticated reasoning without offering developmental pathways. Psychological accounts like Rest (1986) identify cognitive components yet provide limited guidance for skills development. Within structured models, Treviño (1986) requires individual and situational awareness, and Hunt and Vitell (1986) presume normative evaluation skills; both offer limited development pathways.

Leadership approaches emphasize distinct competency sets. Transformational leadership (Bass, 1985) emphasizes inspirational leadership; servant leadership (Greenleaf, 1977) emphasizes service orientation; public-administration ethics (Denhardt and Denhardt, 2003) emphasizes public-service competence, yet frequently under-specifies analytical reasoning skills.

Contemporary frameworks require specialized competencies: Triple Bottom Line (Elkington, 1997) demands multi-dimensional analysis; AI/digital ethics (Floridi, 2013; 2020) requires technical and ethical expertise; GDPR (2016) requires legal and technical competence; smart cities ethics (Kitchin, 2014) requires urban planning and technology-ethics expertise. The lack of structured capability development approaches remains a major limitation at both individual and organizational levels.

2.3.1.5 Evaluation Mechanisms and Outcome Assessment

Utilitarian frameworks (Bentham, 1789; Mill, 1863) foreground outcome assessment but struggle to incorporate diverse stakeholder considerations. Stakeholder theory (Freeman, 1984) supports multi-perspective evaluation yet faces implementation challenges. Among structured models, Rest (1986) emphasizes process-oriented assessment, and Hunt and Vitell (1986) offer a dual-evaluation logic (deontological and teleological).

Leadership approaches vary: transformational leadership (Bass, 1985) uses vision-alignment indicators; servant leadership (Greenleaf, 1977) emphasizes stakeholder benefit; Ethical Leadership in Public Administration (Denhardt and Denhardt, 2003) relies on accountability-based evaluation. Contemporary frameworks introduce specialized mechanisms: Triple Bottom Line (Elkington, 1997) leverages comprehensive metrics; AI/digital ethics (Floridi, 2013; 2020) uses technology-impact assessment; GDPR (2016) applies compliance-focused evaluation; big data/smart cities ethics (Kitchin, 2014) deploys urban-impact assessment.

Despite this diversity, many frameworks lack closed-loop evaluation, that is, systematic feedback and continuous-improvement cycles tied to the original decision criteria. Stakeholders increasingly demand measurable ethical performance and auditable

accountability, yet few frameworks operationalize outcome assessment alongside decision steps.

2.3.1.6 Synthesis of Comparative Patterns

The cross-framework comparison yields five patterns:

1. Complementary strengths across categories. No single framework covers all dimensions; classical theory offers depth; structured models add process; leadership contributes competencies; contemporary instruments add safeguards and documentation, suggesting value in integration.
2. Consistent limitations. Comprehensive, cross-cultural implementation guidance is rare, and accommodation of time pressure is uncommon, a critical practical vulnerability.
3. Disciplinary fragmentation. Business-ethics models (Ferrell and Gresham, 1985; Hunt and Vitell, 1986; Velasquez, 1982), individual moral cognition; leadership ethics (Bass; Denhardt and Denhardt; Greenleaf) under-specify structured protocols; contemporary governance models (GDPR; Floridi; Kitchin) often lack connection to classical justification.
4. Theory–practice disconnection. Theoretical sophistication (Bentham, 1789; Kant, 1785; Mill, 1863) often arrives without operational tools, while practice-forward approaches (Triple Bottom Line; GDPR) can under-specify the depth of justification.
5. Evolutionary progression. The field has moved from abstract principles toward context-specific practice, acknowledging complexity, but has not yet produced a fully integrated, portable framework.

Implications for PASO/EDMC.

- Design requirements. The patterns above justify three requirements carried into the methodology: Transparency (make value assumptions/constraints explicit),

- Transferability (link principles to repeatable steps and skills development), and Accountability (tie decisions to evidence and stakeholder-relevant outcomes).
- PASO alignment. PASO anchors the necessities surfaced here: Principles (norms/constraints), Actions (steps/safeguards), Skills (capabilities/relational practices), and Outcomes (stakeholder-visible results).
 - EDMC orchestration. EDMC provides a multi-perspective synthesis that makes trade-offs explicit (rights, utility, virtue/care) and records rationale for audit and learning, improving performance under time pressure and across cultural contexts.

2.3.2 Identified Limitations and Fragmentations in Frameworks

Building on the thematic categorization, this section examines the critical fragmentations and limitations across the sixteen reviewed frameworks. Seven recurrent limitations materially affect effectiveness across contexts and cultures. Collectively, they create barriers to ethical decision-making in contemporary leadership settings and justify a more integrated approach.

2.3.2.1 Disciplinary Silos and Theoretical Isolation

Ethical decision-making frameworks have emerged within distinct academic and professional domains, generating fragmentation across knowledge areas. Business-ethics models often underuse insights from individual moral cognition; leadership-ethics models overlook structured decision protocols; and AI governance models can under-integrate broader ethical considerations. This isolation hinders cross-fertilization and obscures the multidimensional nature of contemporary dilemmas.

Cognitive models (Rest, 1986; Treviño, 1986) explain individual moral processes yet underweight organizational factors. Conversely, corporate ethics frameworks (Ferrell and Gresham, 1985) emphasize organizational structures while under-addressing

individual cognition. Theory–practice disconnection. Theoretical sophistication (Bentham, 1789; Kant, 1785; Mill, 1863) rarely arrives with tools; practice-forward approaches (Triple Bottom Line; GDPR) can lack justificatory depth. The result is a coordination gap across individual, organizational, and societal levels.

2.3.2.2 Cultural and Contextual Fragmentation

The field remains dominated by Western ethical philosophy, with limited incorporation of diverse cultural perspectives. Classical frameworks (Bentham, 1789; Kant, 1785; Mill, 1863) and many contemporary models presume universal applicability without specifying mechanisms for cultural adaptation.

Non-Western traditions, e.g., Confucian ethics and Ubuntu, are underrepresented despite their contributions to communal and relational ethics. Even frameworks that acknowledge culture, such as Hunt and Vitell’s Theory of Marketing Ethics (1986), provide limited operational guidance for cross-cultural application. Where adaptation is discussed, guidance tends to remain conceptual rather than procedural.

2.3.2.3 Jurisdictional Inconsistencies in Ethical Standards

Divergent regulatory regimes complicate ethical practice for multinationals, creating uneven obligations and enforcement risk. For example, the EU’s GDPR (2016) offers comprehensive rights-based data protection, whereas U.S. approaches are more fragmented across states and sectors. These inconsistencies make global, uniform standards difficult to implement.

Contemporary instruments such as GDPR (2016) and the Triple Bottom Line (Elkington, 1997) seek cross-boundary coherence but face implementation friction where regulatory regimes diverge. Frameworks need to operate effectively across jurisdictions while maintaining ethical integrity.

2.3.2.4 Inflexibility Under Real-World Constraints

Many frameworks assume ideal deliberative conditions rarely found in organizations. Three constraint clusters are especially salient:

1. Artificial intelligence and automation pressures. Rapid technological change raises dilemmas of bias, privacy, and governance that traditional frameworks do not fully address.
2. Global sustainability challenges. Balancing environmental, social, and economic objectives introduces complex trade-offs that many models do not integrate coherently.
3. Geopolitical and regulatory shifts. Divergent digital-governance regimes complicate global implementation, with limited guidance for navigating conflicts.

While AI/digital ethics (Floridi, 2013; 2020) targets technological contexts, it often requires stronger integration with broader ethical principles and clearer guidance for balancing competing considerations under real-world constraints. Guidance for prioritizing among competing constraints under time pressure remains limited.

2.3.2.5 Absence of Operational and Actionable Tools

A persistent gap separates ethical principles from practical implementation. Three dimensions are notable:

1. Ethics-washing risk. Adoption of ethical language without meaningful change creates reputational cover rather than substantive practice.
2. Lack of standardization. Few commonly accepted metrics exist to assess ethical performance and compliance, weakening accountability.
3. Corporate ineffectiveness. Without clear KPIs and responsibility structures, guidelines remain symbolic rather than operational.

Although instruments such as the Triple Bottom Line (Elkington, 1997) and GDPR (2016) offer more specific guidance, many frameworks remain principle-heavy and tool-light.

2.3.2.6 Neglect of Competency Development and Ethical Capabilities

Many frameworks identify needed competencies yet lack structured development pathways. Three gaps recur:

1. Missing pathways. Skills are named but not cultivated through explicit learning designs.
2. Over-emphasis on analytical cognition. Emotional, interpersonal, and systems competencies receive less structured attention.
3. Lack of organizational capability models. Few frameworks specify how to build collective skills beyond individual training.

Leadership-oriented frameworks (Bass, 1985; Denhardt and Denhardt, 2003; Greenleaf, 1977) engage competencies more directly but often lack systematic, scalable development models.

2.3.2.7 Inadequate Response to Ethical Blindness and Time Pressure

A critical cross-cutting limitation is the insufficient response to ethical blindness under time pressure. Across the sixteen frameworks, few provide a comprehensive mechanism for this problem. Three facets stand out:

1. Assumption of ideal deliberation. Most models presume slow, reflective decision windows, despite many leadership decisions occurring under severe time constraints.
2. Ethical blindness under pressure. Time pressure, culture, and competing priorities can suppress recognition of ethical salience (Palazzo et al., 2012).
3. Lack of tiered application. Few frameworks provide graded/tiered approaches that function at different time/resource levels.

Empirical work shows that time constraints can impair deliberative reasoning (Del Popolo Cristaldi et al., 2024; Emser et al., 2021), yet guidance on maintaining ethical awareness in such conditions is limited.

2.3.2.8 Synthesis of Limitations and Justification for the EDMC Framework

The analysis reveals seven limitations that, together, constrain effectiveness: disciplinary fragmentation; cultural/contextual gaps; jurisdictional inconsistency; inflexibility under real-world constraints; lack of actionable tools; weak competency development; and inadequate response to time-pressured ethical blindness.

The time-pressure gap is particularly acute and under-addressed, despite its prevalence in leadership contexts. These limitations directly motivate the Ethical Decision-Making Compass (EDMC), which integrates across disciplines, embeds competency-development pathways, and provides adaptable, portable tools—explicitly addressing decision-making under time pressure.

Implications for PASO/EDMC.

- PASO targets the gaps: Principles (rights/constraints across jurisdictions and cultures), Actions (tiered, time-sensitive steps and safeguards), Skills (individual and organizational skills development), and Outcomes (evidence and metrics to prevent ethics-washing).
- EDMC orchestrates a multi-perspective synthesis (rights/utility/virtue-care) under real-world constraints, documents trade-offs for audit and learning, and distributes ethical agency beyond individuals to organizations and systems.

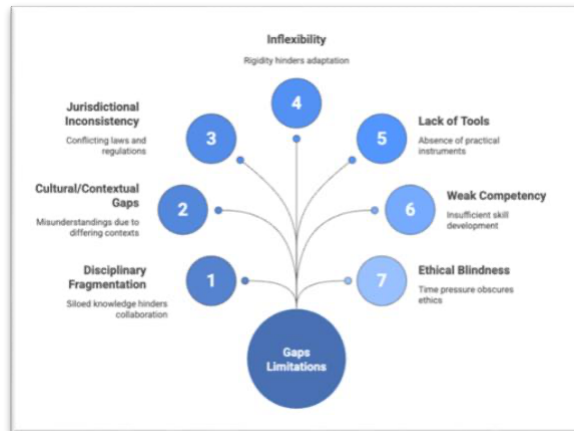


Figure 6 Identified Gaps and Limitations

Source: Author (2025). Original figure created from study data and materials.

2.4 Defining Ethical Decision-Making: Research Definition and Foundations

Key focus — working research definition; academic grounding; PASO integration; practical implications for EDMC.

2.4.1 Research Definition

This section presents the study’s working definition, derived from the comparative analysis in Chapter 2 and used to anchor PASO and the EDMC methodology.

Definition (this study): Ethical decision-making is a multidimensional process that integrates principles, actions, skills, and outcomes. It enables, inspires, and empowers individuals and organizations to navigate ethical dilemmas by balancing accountability, adaptability, and stakeholder considerations while ensuring sustainable, transparent, and inclusive practices.

This conceptualization builds on prior frameworks by incorporating cognitive moral development (Rest, 1986), person–situation and organizational influences (Treviño, 1986), stakeholder engagement, and digital-ethics considerations (Florida,

2013; 2020), yielding a holistic and adaptable approach. (Section 2.4.2 provides the academic grounding.)

2.4.2 Academic Foundation and Literature Grounding

Key focus — situate the definition in the comparative review; justify PASO anchors with evidence from the sixteen frameworks.

This definition arises from the systematic comparative analysis in Section 2.3, which identified strengths and limitations across the literature. Explicitly integrating Principles, Actions, Skills, and Outcomes (PASO) addresses four theoretical and practical gaps:

1. **Fragmentation.** The definition integrates philosophical foundations, psychological processes, organizational context, and practical application (Jobin et al., 2019; Kaptein, 2023). While prior work often emphasizes single dimensions, e.g., cognitive-developmental approaches (Rest, 1986) or person–situation models (Treviño, 1986), this definition synthesizes these perspectives into a cohesive whole. Evidence of persistent isolation limiting cross-sector applicability (e.g., Schwartz, 2023) motivates explicit integration across traditions.
2. **Implementation gap.** Theoretical accounts often fail to translate into practice (Schwartz, 2023; Weaver et al., 1999). By making actions and skills core components, the definition bridges the knowing–doing gap that limits practical utility (Treviño and Nelson, 2016). Ethical awareness without execution has constrained organizational impact; hence, the emphasis on skills development and concrete steps.
3. **Pluralism and context.** Balancing accountability, adaptability, and stakeholder considerations responds to the limits of universal principles in diverse settings (Donaldson and Dunfee, 1999; Kim and Patel, 2018). Integrative Social Contracts

Theory highlights calibrating “hypernorms” with contextual adaptation (Donaldson and Dunfee, 1999); the definition operationalizes this balance through PASO.

4. Forward-looking outcomes. By embedding sustainable, transparent, and inclusive practices as salient outcomes, the definition engages challenges tied to technological change, global interconnectedness, and environmental concerns (Elkington, 1997; Floridi, 2020; Newman et al., 2023), while staying linked to established traditions.

2.4.3 Theoretical Integration and Practical Implications (PASO → EDMC)

Key focus — show how the definition integrates theory streams and addresses ethical blindness; explain how PASO operationalizes into EDMC for real-world use.

The definition moves beyond fragmented approaches to support both academic rigor and practical application. By combining strategic dimensions (principles and outcomes), process elements (actions), and skills development (skills), it establishes the foundation for EDMC built on PASO (see *Figure 7*).

It also targets ethical blindness (Palazzo et al., 2012) by foregrounding skills development—especially ethical awareness and recognition—as a core component (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012). Ethical blindness arises when decision-makers fail to perceive the ethical dimensions of choices (Tenbrunsel and Smith-Crowe, 2008). Explicitly integrating skills with principles, actions, and outcomes equips leaders to counter blindness through structured practice.

In short, the definition provides the platform for the Ethical Decision-Making Compass (EDMC), which operationalizes PASO into a practical guide for ethical leadership under real-world constraints.

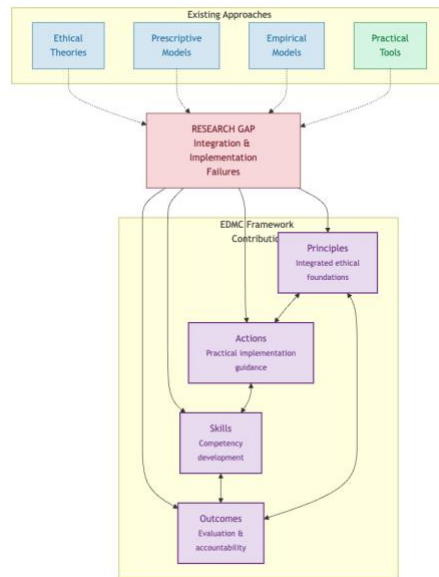


Figure 7 The Research Gap and EDMC Framework Contribution
Source: Author (2025). Original figure created from study data and materials.

2.4.4 Methodological Significance and Construct Validity

Key focus — explain how the definition underpins the research design, measures, and analysis; clarify validity boundaries and assessment criteria.

Methodologically, the definition anchors the study’s conceptual framework and research design. Identifying the foundational PASO components supplies analytic categories for reviewing existing frameworks and for developing the integrative EDMC approach, strengthening construct validity across the research process. This definitional clarity sets parameters for assessing practical effectiveness.

Implications for PASO/EDMC

- Principles — articulate non-negotiable constraints and justificatory standards.
- Actions — specify repeatable, auditable steps (tiered for time pressure).
- Skills — develop individual and organizational capabilities to prevent ethical blindness.
- Outcomes — define evidence, metrics, and learning loops to avoid ethics-washing.

2.5 Emergence and Integration of EDMC’s PASO Components

Key focus — from comparative patterns to PASO synthesis; gap-matrix evidence; how PASO integrates into EDMC.

2.5.1 Synthesis of Foundational Elements

The systematic identification of limitations across existing ethical decision-making frameworks clarifies where current theory and practice fall short. To move beyond descriptive critique, a structured comparative analysis examined how these limitations manifest across framework categories and surfaced patterns that inform an integrated approach.

While Section 2.3 detailed patterns, limitations, and fragmentations, the gap-matrix analysis (*Figure 8*) mapped the extent to which each limitation appears in specific frameworks across categories. This visualization highlights where gaps persist and where complementary strengths can be leveraged.

The gap-matrix serves three purposes:

1. Pattern recognition. Rating each framework against each limitation makes recurrent patterns visible—what excels where, and which blind spots repeat.
2. Strength identification. Alongside limits, the matrix surfaces complementary strengths to synthesize into a more complete design.
3. Development justification. The comparative evidence grounds the move to an integrated framework that addresses gaps while preserving foundations.

2.5.1.1 Assessment Methodology

The gap-matrix assessed how sixteen frameworks address five gap dimensions (L1–L5), derived by clustering the seven limitations in Section 2.3.2 for analytic parsimony:

- L1 — Theory–practice fragmentation (combines disciplinary silos + theory–practice disconnection).
- L2 — Absence of actionable tools.
- L3 — Neglect of competency development.
- L4 — Insufficient adaptability to context (combines cultural/contextual fragmentation + jurisdictional inconsistency).
- L5 — Inadequate response to time pressure (ethical blindness under time pressure).

Frameworks were selected for historical significance, citation influence, representation across philosophical traditions (O’Fallon and Butterfield, 2005), and relevance to professional decision-making (Craft, 2013). Following comparative-framework precedents (Ferrell et al., 2013; Whittier et al., 2006), we grouped them as foundational ethical theory, structured decision models, leadership-oriented approaches, and technology governance/digital ethics, and applied a three-tier legend:

Legend: ✓ well addressed; ● partly addressed; X not addressed.

Each framework–limitation cell was evaluated using primary texts, secondary reviews, empirical studies, and scholarly critiques (e.g., Tenbrunsel and Smith-Crowe, 2008). We triangulated sources, applied consistent criteria, and documented support for each rating (Gioia et al., 2013).

2.5.1.2 Summary of Findings From the Gap Matrix Analysis

The gap matrix (*Figure 8*) compares sixteen frameworks across four categories against the five gap dimensions. The legend (✓, ●, X) provides an at-a-glance view.

Critical Limitations	Classical Ethical Foundations				Structured Decision-Making Models					Leadership-Oriented Models			Technology governance/digital ethics			
	Deontology	Utilitarianism	Virtue Ethics	Stakeholder Theory	Rest's Model	Treviño's Model	Ferrell and Gresham	Velasquez	Hunt and Vitell	Transformational Leadership	Servant Leadership	Public Admin Ethics	Triple Bottom Line	AI Ethics	GDPR Framework	Smart Cities Ethics
Theory-Practice Fragmentation	X	X	●	●	●	✓	●	●	●	●	●	●	●	X	X	X
Absence of Actionable Tools	X	X	●	●	●	●	●	●	●	●	●	●	✓	●	✓	●
Neglect of Skill Development	X	X	●	●	✓	●	X	X	●	✓	✓	✓	X	X	X	X
Insufficient Adaptability to Context	X	●	✓	✓	X	●	●	●	✓	●	✓	●	●	X	X	●
Inadequate Response to Time Pressure	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X

Figure 8 Systematic Gap Analysis Matrix of Ethical Decision-Making Frameworks (n = 16) Across Five Gap Dimensions (L1–L5)

Source: Author (2025). Original figure created from study data and materials.

Legend: ✓ well addressed; ● partly addressed; X not addressed. Categories:

foundational ethical theory; structured decision models; leadership-oriented approaches; technology governance/digital ethics.

Key insights:

1. Foundational theory supplies strong normative grounding yet consistently lacks implementation and competency development; all scored X on time pressure.
2. Structured models balance coverage (e.g., Rest on competencies; Treviño on person–situation) but uniformly overlook time pressure.
3. Leadership approaches are strong on competency development yet inconsistent on tools and cross-context adaptability.
4. Tech governance/digital ethics excels in actionable tools and documentation/accountability but often lacks competency pathways and philosophical integration.

5. Universal L5 gap (time pressure). In this set (n = 16), all frameworks score X on time-pressured decision-making—precisely where failures are more likely (Palazzo et al., 2012).

2.5.1.3 Recognition and Synthesis of Foundational Components

Despite gaps, the analysis identifies foundational components to preserve and integrate. PASO, Principles, Actions, Skills, Outcomes, is not merely reactive; it synthesizes recurring necessities evident across the sixteen frameworks:

- Principles — foundational elements
 - Normative foundations. Classical theory (Bentham, 1789; Kant, 1785; Mill, 1863) supplies durable criteria for right/wrong.
 - Stakeholder salience. Stakeholder theory (Freeman, 1984) provides a framework for affected-party analysis.
 - Universal value recognition. AI/digital ethics and governance instruments (Floridi, 2013; 2020; GDPR, 2016) underscore transparency, explainability, privacy, and accountability.
 - Virtue orientation. Virtue ethics (Hursthouse, 1999; MacIntyre, 1984) and servant leadership (Greenleaf, 1977) foreground character and intent.
- Actions — foundational elements
 - Structured process. Rest (1986) and Treviño (1986) provide stepwise approaches.
 - Context appraisal. Hunt and Vitell (1986) enable situational assessment and dual evaluation.
 - Implementation protocols. Triple Bottom Line (Elkington, 1997) and GDPR (2016) translate principles into procedures.
 - Compliance mechanisms. DPIAs/PIAs and similar controls support verifiable accountability.
- Skills — foundational elements

- Ethical awareness. Recognizing ethical salience (Rest, 1986).
- Perspective-taking. Leadership models (Bass, 1985; Denhardt and Denhardt, 2003).
- Systems thinking. Interdependencies and long-term effects (Elkington, 1997).
- Cultural intelligence. Cross-cultural competence for diverse contexts.
- Outcomes — foundational elements
 - Stakeholder impact measurement. Multi-stakeholder effects (Elkington, 1997; Freeman, 1984).
 - Process quality. Evaluate the decision process, not only outputs (Denhardt and Denhardt, 2003).
 - Implementation verification. Compliance and control checks (GDPR, 2016).
 - Long-term effects. Intergenerational/systemic consequences.

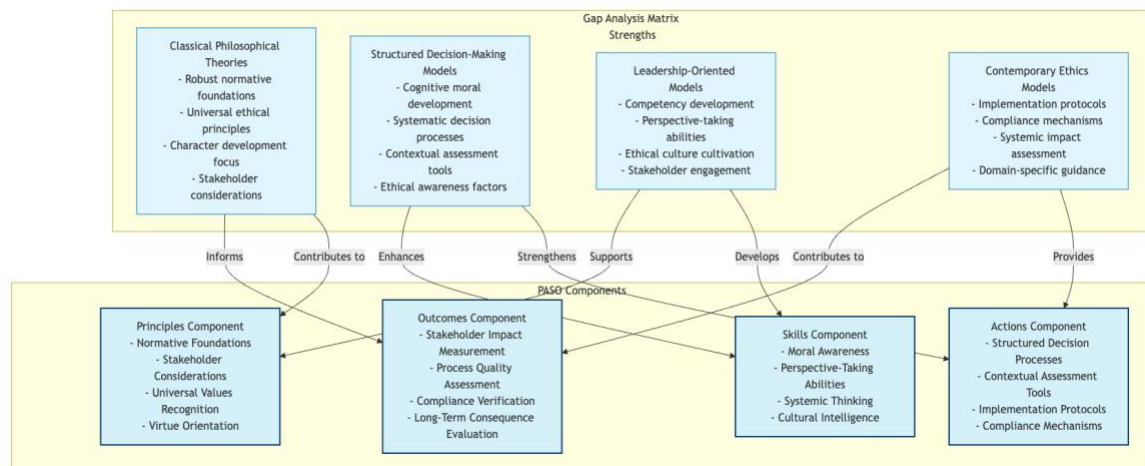


Figure 9 Mapping Foundational Components-From Reviewed Frameworks to PASO Anchors (Principles, Actions, Skills, Outcomes)

Source: Author (2025). Original figure created from study data and materials.

Implications for PASO/EDMC

- PASO consolidates necessities repeatedly observed across the literature: Principles (normative constraints/values), Actions (procedures, protocols, documentation), Skills

(capabilities from awareness to systems thinking), and Outcomes (stakeholder-visible evidence and effects).

- EDMC orchestrates these components under real-world conditions, explicitly addressing time pressure through tiered application, bridging disciplines, and recording trade-offs for audit and learning.

2.5.2 Structuring the EDMC on the PASO Foundation

The gap analysis clarifies limitations and reveals complementary strengths. PASO emerged through a methodical analytical process as a synthesis of recognized foundations while simultaneously addressing the identified gaps. No single framework covers all five gap dimensions, but each category contributes elements that, when synthesized, form the basis of an integrated approach (Kaptein, 2020; Treviño and Nelson, 2021):

- Classical foundations provide robust ethical principles and philosophical depth (Whetstone, 2001).
- Structured models offer systematic decision processes and cognitive frameworks (Rest et al., 1999).
- Leadership models contribute competency insights and skill-development approaches (Brown and Treviño, 2006).
- Contemporary models introduce implementation guidance and practical tools (Floridi, 2013).

The universal time-pressure gap signals a critical design requirement. EDMC therefore incorporates tiered, time-aware application and prompts to mitigate ethical blindness.

2.5.2.1 Emergence of PASO foundational components

Recurring components coalesced, through iterative synthesis, into four interconnected elements: Principles, Actions, Skills, Outcomes (PASO)—derived from the literature rather than predetermined. PASO structures ethical decision-making around four essential questions:

- Principles (Why?) — What ethical principles and values should guide decisions?
- Actions (How?) — What specific procedures translate principles into practice?
- Skills (Who?) — What competencies must decision-makers develop to implement actions effectively—especially under constraints?
- Outcomes (What?) — How do we measure and evaluate the impact of ethical decisions?

These components are cyclical and interconnected, reflecting the dynamic nature of ethical decision-making.

2.5.2.2 Principles: Theoretical Foundations

The analysis identified a need for integrated principles that maintain normative clarity while accommodating contextual diversity:

- Universal principles. Core ethical commitments synthesized from classical approaches (Bentham, 1789; Kant, 1785; Mill, 1863) and contemporary extensions (Floridi, 2013; 2020).
- Stakeholder considerations. Systematic identification and integration of affected parties (Freeman, 1984) and corporate social responsibility (Carroll, 1991).
- Cultural adaptation mechanisms. Ways to adapt principles across cultures while retaining core values (Donaldson and Dunfee, 1999).
- Regulatory-ethics integration. Alignment of regulatory requirements (e.g., General Data Protection Regulation, 2016) with broader ethical principles.

This component addresses fragmentation by integrating diverse philosophical traditions and supporting cross-cultural/jurisdictional adaptation.

2.5.2.3 Actions: Operational Mechanisms

The review highlighted the need for specific mechanisms that translate principles into practice:

- Decision-process protocols. Structured steps adapted from cognitive models (Rest, 1986; Treviño, 1986) for different time/resource scenarios.
- Assessment tools. Practical instruments for evaluating ethical dimensions, adapted to account for time pressure.
- Implementation pathways. Procedures for turning decisions into organizational action, informed by corporate ethics and governance frameworks.
- Communication strategies. Approaches for explaining decisions to diverse stakeholders.

This component addresses the lack of actionable tools and adds tiered application so EDMC works when time is scarce.

2.5.2.4 Skills: Competency Requirements

Effective decision-making requires explicit capability building:

- Ethical awareness. Recognizing ethical salience—especially under time pressure (Rest, 1986; Treviño, 1986).
- Stakeholder analysis. Systematically identifying and balancing perspectives (Freeman, 1984).
- Ethical reasoning. Applying principles to complex cases with practical learning designs.
- Cultural intelligence. Navigating diverse norms (Donaldson and Dunfee, 1999).

- Systems thinking. Understanding interdependencies and long-term effects (Elkington, 1997).

This component remedies the common neglect of competency development by embedding structured pathways for practice.

2.5.2.5 Outcomes: Impact Assessment

Frameworks should include clear, auditable outcome measures:

- Stakeholder impact metrics. Effects across stakeholder groups (Carroll, 1991; Freeman, 1984).
- Process-quality indicators. Evaluating the quality of the decision process, not only outputs.
- Implementation effectiveness. How well decisions are enacted, including compliance and controls (GDPR, 2016).
- Long-term consequences. Intergenerational and systemic effects (Elkington, 1997).

This component links principles and actions to results that matter and can be learned from.

2.5.2.6 Integrating PASO Components Into a Coherent EDMC Framework

The PASO components are interconnected, not modular. *Table 9* shows how each component systematically addresses the limitations identified in existing frameworks.

Table 9 How PASO Components Address Critical Limitations

Source: Author (2025). Original table created from study data and materials.

PASO component	Critical limitation addressed	Integration mechanism
Principles	Fragmentation across disciplines and cultures	Integrates diverse philosophical traditions; provides cross-cultural adaptation mechanisms
	Lack of adaptability under real-world constraints	Balances universal principles with contextual flexibility
Actions	Absence of actionable tools	Provides operational mechanisms, assessment tools, and implementation pathways
	Inadequate response to ethical blindness under time pressure	Uses tiered application protocols designed for varying time constraints

Skills	Neglect of essential ethical competencies	Embeds explicit competency-development pathways
	Ethical blindness under time pressure	Builds awareness and recognition that function even under pressure
Outcomes	Theory–practice disconnection	Links principles to measurable, stakeholder-relevant results
	Inconsistent evaluation standards	Provides a comprehensive assessment framework (stakeholder impacts, process quality, long-term consequences)

This integrated approach shows how the PASO components collectively address the limitations in existing frameworks, providing a comprehensive foundation for ethical decision-making that bridges disciplines while maintaining both theoretical rigour, and practical applicability. PASO closes the specific gaps identified in Section 2.3 while preserving the strengths of prior work, providing the foundation EDMC operationalizes under real-world constraints.

2.6 Advancing Ethical Decision-Making: the EDMC Framework

Key focus — how EDMC (built on PASO) closes L1–L5: integration across disciplines/cultures, actionable tools, capability building, adaptability, and time-pressure resilience.

The EDMC framework, built on the PASO foundational components, advances beyond existing approaches by directly addressing the five critical limitations identified in Section 2.3.2 (L1–L5). *Table 10* summarizes how each EDMC component responds to these limitations, demonstrating the integrated nature of the solution.

Table 10 How the EDMC Framework Addresses Critical Limitations
Source: Author (2025). Original table created from study data and materials.

Framework limitation	Principles component	Actions component	Skills component	Outcomes component
Fragmentation across disciplines and cultures (L1)	Integrates universal ethical principles with contextual adaptation while acknowledging cultural variation	Provides standardized decision protocols adaptable across sectors/disciplines	Develops cross-cultural competencies and interdisciplinary understanding	Ensures comprehensive stakeholder-impact assessment across diverse contexts
Lack of adaptability under	Offers prioritization guidelines for	Supplies tiered protocols for	Builds rapid ethical reasoning skills and	Includes efficiency and adaptability

real-world constraints (L4)	applying principles under constraints and time pressure	comprehensive / abbreviated / rapid timeframes	mental models for complex environments	indicators alongside ethical quality measures
Absence of actionable tools (L2)	Translates abstract principles into operational guidelines	Provides templates, checklists, decision matrices, and implementation pathways	Emphasizes practical application skills across contexts	Tracks implementation success, process quality, and effectiveness
Neglect of essential ethical competencies (L3)	Identifies required competencies for principled reasoning	Links specific actions to the skills needed to execute them	Offers progressive development pathways with assessment tools and learning mechanisms	Measures competency development as a core outcome with feedback loops
Inadequate response to ethical blindness under time pressure (L5)	Embeds stakeholder perspective-taking and awareness triggers	Includes red-flag prompts and time-sensitive protocols	Develops metacognition, bias recognition, and awareness under pressure	Provides feedback mechanisms that strengthen vigilance in time-constrained contexts

Note. Abbreviations: L1–L5 = five gap dimensions defined in Section 2.5.1.2; GDPR = General Data Protection Regulation.

This table demonstrates how EDMC provides a comprehensive, integrated response to the limitations in existing frameworks. Each PASO component contributes distinct yet complementary mechanisms, maintaining ethical integrity while functioning across contexts and time constraints.

2.6.1 Addressing Fragmentation (L1)

EDMC transcends disciplinary boundaries by integrating insights from ethical theory, decision science, leadership studies, and contemporary challenge domains. It addresses fragmentation in three ways:

1. Cross-disciplinary integration. By synthesizing complementary strengths from psychological, philosophical, leadership, and domain-specific approaches, EDMC enables holistic treatment of multi-domain problems (e.g., AI ethics requiring technical, philosophical, and governance synthesis).
2. Cultural-integration mechanisms. EDMC combines universal principles (e.g., respect for dignity; prevention of harm) with cultural adaptation mechanisms informed by

Integrative Social Contracts Theory (Donaldson and Dunfee, 1999), enabling global application while respecting legitimate diversity.

3. Jurisdictional navigation. EDMC includes methods for navigating inconsistent regulatory environments so decisions satisfy core compliance across jurisdictions while maintaining ethical integrity beyond minimum legal standards.

By synthesizing insights from stakeholder theory (Freeman, 1984), virtue ethics (Hursthouse, 1999; MacIntyre, 1984), leadership ethics (Bass, 1985; Greenleaf, 1977), and contemporary frameworks (sustainability and digital governance), EDMC responds directly to the disciplinary silos identified in the gap analysis.

2.6.2 Ethical Blindness Mitigation (L5)

EDMC explicitly targets ethical blindness—the failure to recognize ethical dimensions, especially under pressure (Palazzo et al., 2012), through three mechanisms:

1. Ethical-awareness triggers. Structured prompts (stakeholder impacts, principle tests, unintended consequences) act as ethical alarm bells that still fire when time is scarce.
2. Organizational-culture integration. Drawing on Treviño's (1986) person–situation lens, EDMC embeds climate-building practices that sustain ethical sensitivity amid competing pressures.
3. Contextual adaptation tools. Domain-specific checklists and prompts for technology, sustainability, and governance (Elkington, 1997; GDPR, 2016; Floridi, 2013; 2020) surface typical blind spots in each environment.

Focusing on awareness mechanisms under time pressure addresses a key vulnerability documented by research (Del Popolo Cristaldi et al., 2024; Palazzo et al., 2012).

2.6.3 Cross-Cultural Adaptability (L1/L4)

EDMC maintains normative guidance while enabling cultural fit:

1. Hypernorm integration. Following Integrative Social Contracts Theory (Donaldson and Dunfee, 1999), EDMC anchors decisions to hypernorms (e.g., dignity, non-maleficence) while allowing local adaptation in practices.
2. Cultural-intelligence pathways. Skills development pathways strengthen leaders' ability to read cultural cues, anticipate value tensions, and adapt decision routes without diluting core commitments.
3. Intercultural dialogue mechanisms. Structured dialogue practices help negotiate ethical differences, identify shared ground, and document justified deviations.

This principled adaptability advances beyond rigid universalism or unconstrained relativism.

2.6.4 Time-Pressure Functionality (L4/L5)

A distinctive contribution of EDMC is designing for time scarcity:

1. Tiered application protocols. Three tiers keep guidance alive under different time windows:
 - Comprehensive — full analysis (stakeholders, principles, consequences).
 - Abbreviated — streamlined focus on essentials and key stakeholders.
 - Rapid — pre-agreed ethical heuristics and mental models for crises.
2. Ethical heuristics. Informed by heuristic decision research (Gigerenzer and Gaissmaier, 2011), EDMC includes tools like the reversibility test (“Would I accept this if I were the affected party?”) and the transparency test (“Could I explain this publicly?”) to protect integrity when cognition is taxed.

3. Pre-decision preparation. Scenario work, value clarification, and practice drills establish readiness so ethical scripts are available when seconds matter.

Additionally, EDMC embeds time-aware prompts within all tools, templates, and decision matrices so that essential ethical checks surface under any time budget.

2.6.5 Practical Usability (L2/L3)

EDMC bridges the theory–practice gap, see *Figure 10* with concrete implementation supports:

1. Application tools. Decision protocols; stakeholder-impact templates; implementation checklists; context-specific assessment guides (technology, sustainability, governance).
2. Skill-development pathways. Progressive development of awareness, reasoning, motivation, and implementation skills, drawing on Rest’s (1986) four components and contemporary competency research.
3. Organizational integration. Methods to embed ethics into existing processes and culture: integration points in governance, culture-building tools, measurement systems for decision quality, and feedback loops for continuous improvement.

These mechanisms move beyond principle-heavy, tool-light models while preserving philosophical depth (Bentham, 1789; Kant, 1785; Mill, 1863).



Figure 10 Bridging Gaps and Limitations: EDMC Mechanisms Mapped to L1–L5 and PASO (Critical Limitation)

Source: Author (2025). Original figure created from study data and materials.

2.7 From Theoretical Review to Methodological Foundation: Advancing the EDMC Framework

Key focus — synthesizing streams → PASO; why a three-phase method; EDMC’s distinct contributions; bridge to empirical design.

2.7.1 From Fragmentation to Foundation: Synthesizing the PASO Components

Building on the preceding analysis, this section synthesizes the complementary strengths and core deficiencies of classical, structured, leadership-oriented, and contemporary frameworks for ethical decision-making.

Classical ethical theories—Deontology (Kant, 1785), Utilitarianism (Bentham, 1789; Mill, 1863), and Virtue ethics (Hursthouse, 1999; MacIntyre, 1984), provide robust normative foundations essential for ethical reasoning. Structured frameworks, including Rest’s Four-Component Model (1986), and Treviño’s Interactionist Model (1986), contribute psychological insight and systematic decision processes. Leadership-oriented

approaches such as Transformational leadership (Bass, 1985), and Servant leadership (Greenleaf, 1977), add competency-development pathways and relational perspectives. Contemporary frameworks, AI and Digital Ethics (Floridi, 2013; 2020), the General Data Protection Regulation (2016) as a technology-governance instrument, and the Triple Bottom Line (Elkington, 1997), supply domain-specific tools and guidance for emerging challenges.

Collectively, these insights lay the conceptual groundwork for the Ethical Decision-Making Compass (EDMC), a holistic and operational response to persistent shortcomings across models. The gap analysis shows that although many frameworks offer singular virtues, philosophical depth, psychological insight, or practical tools, most falter amid real-world complexity: fragmentation across disciplines, limited adaptability, insufficient actionable tools, and neglect of essential competencies recur. A shared vulnerability also emerges: ethical blindness and decision fatigue under time pressure.

In response, the PASO components—Principles, Actions, Skills, Outcomes—emerged through cross-disciplinary synthesis. Each addresses critical limitations:

- 1) Principles anchor normative reasoning and consistency.
- 2) Actions ensure pragmatic application and operational clarity.
- 3) Skills foster awareness, reasoning, motivation, and implementation.
- 4) Outcomes provide measurement and feedback for continuous improvement.

This architecture preserves the strongest elements of existing frameworks while systematically addressing their shared limitations.

2.7.2 Justifying a Three-Phase Methodological Approach

To combine theoretical depth with practical usability, the study adopts a three-phase methodology:

- 1) Phase I — Initial: analysis and assessment. A rigorous evaluation of existing theories and frameworks to identify gaps and opportunities for integration.
- 2) Phase II — Bridge: integration and synthesis. Bridging disciplinary silos and consolidating strengths into the PASO structure, forming the basis for EDMC.
- 3) Phase III — Emergent: development and validation. Iterative prototyping, simulation, and applied testing to ensure EDMC’s adaptability and effectiveness across contexts.

This design not only connects theory and practice; it allows the framework to evolve through empirical engagement with real ethical dilemmas (Section 1.5.2).

2.7.3 Advancing the Field: EDMC’s Unique Contributions

The EDMC framework is distinctive in:

- 1) Integrating philosophical depth with actionable tools and measurement systems.
- 2) Counteracting ethical blindness and enabling decisions under time constraints via tiered protocols, time-aware prompts, and preparedness.
- 3) Supporting cross-cultural adaptability and organizational scalability through hypernorm anchoring and dialogue mechanisms.
- 4) Fostering a feed-forward learning orientation, anticipating challenges, not only reacting, through documentation, metrics, and feedback loops.

Rather than replacing established models, EDMC synthesizes their strengths and resolves collective limitations through a unified, matrix-based approach. Its tiered protocols and built-in learning loops help leaders navigate complexity, build trust, and enact decisions that are both principled and practical. Treating ethical decision-making as a complex, time-sensitive, multi-contextual process, EDMC offers a holistic and operational framework.

2.7.4 From Theory to Empirical Application: Designing the Research

This chapter culminates by mapping the progression from critical literature review to methodological design. Insights from the comparative analysis directly shape the empirical research, which unfolds in three phases:

- 1) Initial phase: grounding the methodology in theoretical synthesis and gap analysis.
- 2) Bridge phase: developing and iteratively refining PASO components through practical engagement and stakeholder feedback.
- 3) Emergent phase: validating the integrated EDMC in varied settings to ensure relevance and usability.

This structured transition ensures that EDMC evolves from conceptual rigour and lived experience, positioning it as a robust tool for future research and application. The subsequent chapter details the research design, showing how each phase operationalizes insights distilled here and ushers in the practical development of EDMC.

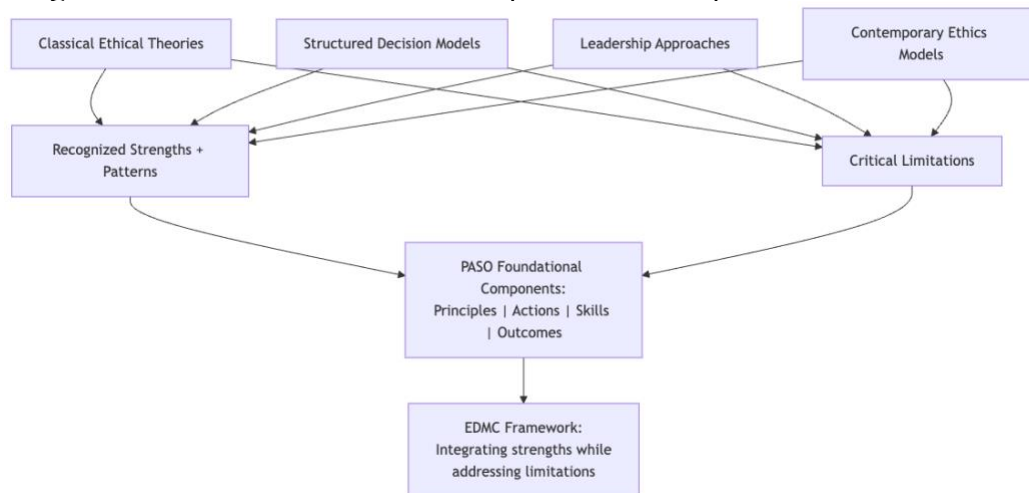


Figure 11 From Framework Analysis to Approach: Flow from Synthesis and Gaps to the Three-Phase Research Design

Source: Author (2025). Original figure created from study data and materials.

Phases: Initial (analysis and assessment); Bridge (integration and synthesis); Emergent (development and validation). PASO anchors: P = Principles; A = Actions; S = Skills; O = Outcomes.

2.8 Summary and Implications for Method

Key focus — what the review established; what PASO/EDMC operationalizes; three design requirements that carry into Chapter 3.

The comparative analysis shows that each stream contributes essential—yet partial—skills for ethical decision-making. Foundational theories supply normative constraints and justificatory depth; structured models add process discipline and cognitive steps; leadership-oriented approaches develop relational competencies and ethical-climate mechanisms; and technology-governance instruments provide safeguards, documentation, and enforceable accountability. Considered in isolation, these strengths leave recurrent tensions unresolved—most notably value trade-offs across stakeholders, the portability of abstract principles into practice, demonstrable accountability in high-risk, AI-enabled settings, and performance under time pressure.

The EDMC, built on PASO (Principles, Actions, Skills, Outcomes), consolidates recurring necessities across literatures and operationalizes a multi-perspective synthesis that makes trade-offs explicit and recommendations auditable. PASO anchors link why (principles) to how (actions), who (skills), and what (outcomes), closing the theory–practice gap while supporting cross-cultural and cross-jurisdictional use.

For the methodology that follows, three design requirements arise:

- Transparency — make value assumptions, constraints, and trade-offs visible (documented rationale, stakeholder salience, principle tests).
- Transferability — link theory to repeatable steps and skills development through time-aware prompts embedded in instruments (matrices, templates, decision records) and reusable decision matrices that travel across contexts.
- Accountability — tie decisions to stakeholder-relevant outcomes and evidence (impact metrics, process-quality indicators, feedback loops).

These requirements justify the three-phase research design and the planned validation through hypothetical cases and expert review, ensuring continuity from literature to method and from claims to evaluation. They also position EDMC to address the universal vulnerability identified in the review, ethical blindness under time pressure, by combining tiered application protocols with time-aware prompts and preparedness mechanisms embedded in the tools themselves.

With this foundation in place, the next chapter sets out the research design in detail, showing how each phase operationalizes PASO and tests EDMC's practicality, adaptability, and auditability across diverse leadership contexts.

CHAPTER III: METHODOLOGY

3.1 Introduction

Leaders routinely make ethically charged calls under time pressure, cultural diversity, and shifting rules. Existing frameworks help, but only in parts: some offer philosophical depth without telling people how to act; others list steps but ignore the skills and ownership needed to carry them out; many travel poorly across contexts; and few make ethical stakes unmistakable when minutes are scarce (Kaptein, 2023; Palazzo et al., 2012; Treviño and Nelson, 2016). Building on the synthesis in Chapter 2, this study develops and validates the Ethical Decision-Making Compass (EDMC), a practical architecture organized around PASO (Principles, Actions, Skills, Outcomes) to support decisions that are both usable and auditable in real settings.

This chapter explains how the design moved from literature synthesis into PASO and then into working EDMC instruments (Sections 3.2–3.4); describes the data, participants, materials, and procedures used in the two evidence streams (Sections 3.5–3.8); and sets out the analysis approach—quantitative and qualitative— including the six-metric rubric and a brief codebook explainer (Section 3.9). It then addresses quality and governance: limitations, reliability/validity notes, research ethics, and version control (Sections 3.10–3.12). Full survey questions appear in Appendix C; computation rules and denominator conventions are in Appendix F; and technical anchors, formulas, and item indices are consolidated in Appendices B, C, and F. Unless otherwise noted, all tables and figures are the author’s own, created from study data or study instruments, with any external items explicitly cited in their captions.

3.2 Operationalization of Theoretical Constructs

The core outcome is ethical decision-making effectiveness as enabled by EDMC under real constraints. We operationalize this outcome with six PASO-aligned metrics, each scored on shared 0–3 anchors to keep evidence standards consistent and auditable across cases (Hsieh and Shannon, 2005; Miles et al., 2014). The six metrics are: (1) PASO linkage, i.e., coherence from Principles → Actions → Skills → Outcomes; (2) blindness mitigation, i.e., prompts and guards that surface ethical salience under pressure; (3) cultural adaptability, i.e., portability across cultures and jurisdictions while preserving hypernorms; (4) time-pressure suitability, i.e., cues and heuristics usable in short decision windows; (5) implementation guidance, i.e., concrete steps, roles, and artifacts/templates; and (6) outcomes evaluability, i.e., KPIs, go/stop conditions, and feedback/audit trail.

All six metrics share the same 0–3 anchor scale: 0 = not evidenced; 1 = minimally evidenced; 2 = adequately evidenced; 3 = clearly and consistently evidenced. Use N/A only when a metric truly does not apply, and exclude N/A from denominators when averaging. Detailed anchors and weight presets appear in Appendix B (*Table 34*), and one-line score rationales are logged in the C1–C8 case packets (Appendix A) (Hsieh and Shannon, 2005; Miles et al., 2014).

Score	Evidence standard (applies to any metric)
0	Not evidenced
1	Minimally evidenced
2	Adequately evidenced
3	Clearly and consistently evidenced

Note. Use N/A only where a metric truly does not apply; exclude N/A from denominators when averaging.

- We apply a two-lens measurement model.

Rubric lens (0–3 anchors). For each public case, the researcher (single rater in this phase) scores the six metrics using the Appendix B rubric and records ≤15-word rationales in the case packet (Appendix A). The blank scoring sheet is in Appendix B,

Table 35. Per-metric means are computed with N/A excluded; the case composite is the mean of the six per-metric means (Equal weights by default), with alternative presets available in Appendix B.

Expert-panel survey lens (Likert indicators). The expert-panel instrument maps observable items to the same constructs (plus adjacent outcomes). We report % ≥ 4 on 1–5 items and compute PASO I-CVI and S-CVI/Ave for the PASO micro-block following content-validity conventions (Held, 2006; Lynn, 1986; Polit et al., 2007). The compact item index is in Appendix C.2 and the verbatim, version-locked instrument PDF in Appendix C.5.

A construct→indicator crosswalk is presented in *Table 11* (Compact mapping – construct and survey indicator[s]). In brief: time-pressure suitability maps to Q4; defensibility/implementability to Q5 and Q7; cultural adaptability (portability) to Q6; actionability (PASO linkage + implementation) to Q1 plus the PASO micro-block; intention to use to Q8; and outcomes evaluability to the PASO “O” item together with the open “conditions to proceed.” Convergent effects on judgment are captured via stance change (Baseline→Post), confidence uplift (0–10), and the efficiency item. For clarity, Integratability denotes the ease of fitting EDMC into an organization’s existing decision/approval workflows with only minor tweaks (survey Q7).

Table 11 Compact Mapping–Construct and Survey Indicator(s)
Source: Author (2025). Original table created from study data and materials.

Construct	Survey indicator(s)
Time-pressure suitability	Q4 “time-pressure usability”
Defensibility / Implementability	Q5 “withstand scrutiny”; Q7 “integratability”
Cultural adaptability (portability)	Q6 (per-case) + Global wrap-up “Portability” (post-Case-3) + qualitative notes on local-adaptation hooks / hypernorms
Actionability (PASO linkage + implementation)	Q1 “actionable pathway” + PASO micro-block (P, A, S, O)
Intention to use	Q8 (behavioral proxy of perceived usefulness)
Outcomes evaluability	PASO O item + open “conditions to proceed”
Effect on judgement (convergent outcome)	Stance change (Baseline→Post), confidence uplift (0–10), efficiency item

Composites and reporting follow a simple rule set. Per-metric scores (0–3) are averaged to two decimals with N/A excluded; the case composite (0–3) is the mean of the six per-metric means (Equal by default); survey indicators include $\% \geq 4$, confidence uplift (Post – Baseline, 0–10), stance-change distributions, and I-CVI/S-CVI/Ave. Operational definitions and formulas are consolidated in Appendix F; summary tables and figures appear in Appendix G; the survey’s verbatim wording is archived in Appendix C.5.

Terminology note. PASO uses Skills; any occurrence of “Capabilities” in cases/materials is treated as the same construct.

3.3 Research Purpose and Questions

To synthesize strengths and close gaps across established ethics frameworks by building, operationalizing, and validating a practical, cross-cultural, time-aware decision tool, the Ethical Decision-Making Compass (EDMC), organized around PASO (Principles, Actions, Skills, Outcomes).

Research questions (see Section 1.5 for full statement and wording):

- Q1–Q2 (analyze/assess). What strengths and limitations recur across established frameworks, and how do these specify requirements for an integrated approach?
- Q3 (synthesize). How can recurring necessities be synthesized into the PASO scaffold?
- Q4–Q5 (design/translate). How should PASO be organized into EDMC and translated into usable instruments?
- Q6 (validate/effects). To what extent does EDMC enhance ethical clarity, confidence, and defensibility under time pressure across varied contexts?

Methodological alignment.

- Q1–Q2 (review/comparison): directed content analysis and cross-framework matrix mapping (Hsieh and Shannon, 2005; Miles et al., 2014).
- Q3 (synthesis): integrative/critical interpretive synthesis to derive PASO from cross-framework regularities (Dixon-Woods et al., 2006a).
- Q4–Q5 (design/translation): applied theory-building and mixed-methods guidance for developing and operationalizing EDMC instruments (Creswell and Plano Clark, 2018; Lynham, 2002).
- Q6 (validation): case-based analytic logic (Yin, 2018) and expert-panel content-validity procedures (Grant and Davis, 1997; Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006). Inter-rater reliability statistics (e.g., ICC, Kendall's W) were not computed in this phase due to the single-rater design of the analytical case applications; multi-rater reliability is reserved for future work (Chapter 6).

Cross-references. Scoring rubric and computation → Appendix B; instrument verbatim/PDF → Appendix C.5; analysis plan and derived metrics → Appendix F; dashboard evidence → Appendix G; version-lock and fieldwork window → Appendix N.

Operationalization of Q6 in Chapter 4.

Two complementary evidence streams:

1. Expert-panel survey (N = 30; 90 case-blocks): usability ($\% \geq 4/5$), PASO content validity (I-CVI/S-CVI/Ave), stance change, confidence uplift, efficiency.
2. Researcher-led public-case applications: six 0–3 rubric metrics → composite (0–100) with weight-preset sensitivity (Equal / Stakeholder / Time-pressure) to test robustness and traceability.

These streams directly test whether EDMC preserves valued elements of prior frameworks while addressing their shared constraints (PASO coverage, usability under time pressure, cross-cultural portability).

3.4 Research Design

The study uses a pragmatic, problem-first design to discover, build, and validate EDMC: first surfacing patterns in the literature, then synthesizing PASO, and finally translating it into auditable tools tested under realistic constraints. The logic is sequential and mixed-methods so that each phase informs the next and the evidence triangulates.

Methodological orientation: The study adopts a pragmatic stance that prioritizes problem-solving and fitness-for-use in real contexts (Morgan, 2014). The approach is emergent and discovery-oriented, combining constructivist grounded techniques for pattern surfacing (Charmaz, 2006), abductive cycling between data and theory (Timmermans and Tavory, 2012), and a sequential mixed-methods logic for triangulation (Creswell and Plano Clark, 2018; Morse, 1991). The design follows a problem-definition-first ethos—specifying the problem space before “freezing” solutions—consistent with practice-oriented theory building (Edmondson and McManus, 2007; Eisenhardt, 1989; Van de Ven, 2007).

Phases and alignment to RQs (*Table 12*). The design mirrors the thesis logic—review → synthesis (PASO) → build and validate (EDMC):

- Phase 1 - Initial (Q1–Q2): directed content analysis and cross-framework matrix comparison of 16 frameworks across core dimensions (Hsieh and Shannon, 2005; Miles et al., 2014).
- Phase 2 - Bridge (Q3): integrative/critical interpretive synthesis to derive PASO from recurring necessities (Dixon-Woods et al., 2006a).
- Phase 3 - Emergent (Q4–Q6): EDMC instrument design; hypothetical case applications; expert-panel content-validity review; and sensitivity checks (Grant and Davis, 1997; Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006).

Table 12 Phases and Alignment to RQs

Source: Author (2025). Original table created from study data and materials.

Phase and Timeline	Activities	References
Initial (Q1–Q2)	Directed content analysis and cross-framework matrix comparison of 16 frameworks across core dimensions	(Hsieh and Shannon, 2005; Miles, Huberman and Saldaña, 2014)
Bridge (Q3)	Integrative/critical interpretive synthesis to derive PASO from recurring necessities	(Dixon-Woods et al., 2006)
Emergent (Q4–Q6)	EDMC instrument design; hypothetical case applications; expert-panel content-validity review; sensitivity checks	(Lynn, 1986; Grant and Davis, 1997; Polit and Beck, 2006; Polit, Beck and Owen, 2007)

Time-aware design choice. EDMC supports tiered application (comprehensive / abbreviated / rapid) without extra paperwork: the tiers are implemented as embedded time-aware prompts inside the same tools and decision matrices, so core ethical checks surface under any time budget while preserving auditability.




Figure 12 Three-Phase EDMC Methodology—From Literature Synthesis → PASO → EDMC Build and Validation

Source: Author (2025). Original figure created from study data and materials.

Legend (for Figure 12).

- Phase 1 – Initial (Q1–Q2): Directed content analysis of 16 frameworks → recurring strengths + five gaps (fragmentation; limited adaptability under constraints; weak tooling; thin skills/ownership; time-pressure blindness).
- Phase 2 – Bridge (Q3): Integrative synthesis → PASO scaffold (Principles, Actions, Skills, Outcomes) and component interactions.
- Phase 3 – Emergent (Q4–Q6): EDMC tools (concise EDMC decision record i.e. one-page case note; 0–3 rubric; survey) and validation via (i) expert panel (I-CVI/S-CVI, $\% \geq 4$, stance and confidence) and (ii) researcher-led case applications (six metrics → composite; Equal/Stakeholder/Time-pressure presets).

- Governance lanes: version-lock and fieldwork window; predefined exclusion/handling rules; ethics/consent; full audit trail.
- Traceability: archived in Appendices B (rubric/weights), C (instrument & verbatim), F (analysis plan), G (summary tables), H (figures), and M–N (exclusions; version-lock/fieldwork).

3.5 Population and Sample

This study draws on two purposively sampled sources: (i) a public case corpus used for hypothetical EDMC applications, and (ii) an expert panel used for content validity and usability evidence. The aim is transferability and instrument validation, not statistical generalization (Miles et al., 2014; Yin, 2018).

3.5.1 Case Corpus (C1–C8)

Eight public, well-documented cases (see *Table 1*) were selected to span sectors and decision types (e.g., environment, labor, discrimination, AI/tech governance). Illustrative anchors include Shein/Shell (environment), Foxconn/Starbucks (labor), Google/sector age-bias (discrimination), and Amazon hiring/OpenAI governance (AI ethics). Selection followed maximum-variation logic to ensure (i) sectoral and topical diversity, (ii) recency, and (iii) accessible evidence sufficient for transparent scoring and replication (Miles et al., 2014; Yin, 2018). These cases are used as hypothetical EDMC applications, not organizational field trials.

Inclusion criteria. Publicly sourced, citable documentation; decision salience (rights/harms, fairness/equity, duties/power); sufficient detail to specify PASO elements (principles, actions/roles/artifacts, skills, outcomes/KPIs).

Exclusion criteria. Proprietary or unverifiable accounts; insufficient detail to identify PASO elements.

Case-selection bias note. Cases were purposively chosen to span sectors, decision types, and recency, using public documentation sufficient for transparent scoring. As hypothetical applications, they evidence analytic transferability rather than population generalization (Miles et al., 2014; Yin, 2018).

3.5.2 Expert Panel (EDMC Instrument Validation)

A multidisciplinary expert panel reviewed three EDMC case notes and completed the structured instrument. Achieved sample: N = 30 complete responses (privacy-safe demographics reported in Chapter 4). Sampling was purposive (with light snowballing) to capture senior practitioners and scholar-practitioners with cross-context experience in governance/ethics/compliance/AI/sustainability.

Inclusion criteria. ≥ 10 years' domain experience; exposure to cross-cultural/multi-jurisdiction decisions and oversight; ability to evaluate frameworks/tools under time pressure; no direct conflicts with the focal cases.

Role mix (targeted). Governance/compliance; ethics/sustainability; AI/data/privacy; labor/human rights; senior management/board advisory.

Rationale for N. Content-validity guidance accepts small panels ($\approx 6-10$) (Grant and Davis, 1997; Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006); N = 30 provides more stable item-level estimates (e.g., I-CVI, $\% \geq 4$) and richer qualitative comments.

Administration. Participation was pseudonymous (NDA codes). Mis-coded/duplicate entries were excluded; corrected re-submissions were included per pre-specified rules (Appendix M). The instrument and fieldwork window were version-

locked (Appendix N). Consent and data protection followed the participant-facing notices (Appendix L; Appendix K).

Purpose of this sample. To assess EDMC’s content validity, time-pressure usability, portability and defensibility—not to estimate population parameters. Aggregated results, not identities, are reported.

3.6 Participant Selection

The expert panel was assembled through purposive sampling, with light snowballing where appropriate, to capture senior, cross-context decision experience relevant to evaluating EDMC (Miles et al., 2014). The aim was to generate content-validity and usability evidence rather than to estimate population parameters (Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006).

Eligible experts typically had ten or more years of professional practice in domains such as governance/compliance, ethics/sustainability, AI/data/privacy, labor/human rights, or senior management/board advisory roles. We sought participants with cross-cultural or multi-jurisdiction exposure who could review three EDMC case notes and complete the structured instrument under realistic time pressure (~15–20 minutes). Experts self-declared the absence of any direct conflict with organizations featured in the case summaries. Individuals were excluded if they held current employment, consulting engagements, or material financial interests in any focal organization, if they were unable to complete the instrument independently within the time window, or if they declined the NDA terms.

Recruitment occurred through professional networks and associations. Invitations described the study purpose, time burden, anonymity, and data handling. Access to the case notes required agreement to an NDA, and informed consent was recorded at the start

of the survey. Participation was voluntary, with the option to withdraw any time before anonymization (Israel, 2015). Consent/privacy wording appears in Appendix L; the GDPR/Swiss FADP notice is provided in Appendix K.

To protect identities, each invitee received a unique NDA-XX code used solely to verify consent at entry. The code↔person mapping was stored separately (PI-only) and destroyed at anonymization. The survey captured no names or emails—only privacy-safe demographics and responses.

Version control applied to the survey dataset throughout. The instrument was version-locked on 23 Sep 2025, 12:30 CEST; the fieldwork window closed on 29 Sep 2025 at 14:40 CEST (Appendix N). A single Case-3 link issue identified on 23 Sep 2025 at 12:15 CEST was corrected without altering any item wording. Three submissions used incorrect or duplicate NDA codes; these rows were excluded. Where an expert re-submitted with the correct code, only the latest valid row was retained. All inclusion/exclusion decisions and sensitivity checks are documented in Appendix M.

The achieved sample comprised $N = 30$ complete expert responses meeting all criteria. Demographic summaries (privacy-safe bins), sector/role mix, and regional spread appear in Chapter 4. The separate stream of researcher-led case applications drew solely on public sources and did not involve human participants.

Panel size was chosen for methodological fit: in content-validity work, small expert panels are acceptable and common) (Grant and Davis, 1997; Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006). With $N = 30$, item-level indicators such as $\% \geq 4$ and I-CVI/S-CVI/Ave are sufficiently stable while keeping respondent burden feasible under NDA and time-pressure constraints.

Cross-references. Rubric/weights → Appendix B; instrument overview + verbatim → Appendix C; summary tables → Appendix G; figures → Appendix H;

exclusions/sensitivity → Appendix M; version-lock & fieldwork window → Appendix N; consent/privacy → Appendices K–L.

3.7 Instrumentation

The study’s instruments were designed to test EDMC for clarity, usability under time pressure, content validity, and practical applicability, drawing on established guidance for scale development, survey design, and mixed-methods evaluation (Creswell and Plano Clark, 2018; DeVellis, 2017; Krosnick and Presser, 2010; Lynn, 1986; Polit and Beck, 2006). EDMC’s PASO-DNA (Principles–Actions–Skills–Outcomes) is operationalized through three coordinated artifact families: (i) concise EDMC decision records (“case notes”) scaffolded by a PASO plan, (ii) a six-metric 0–3 scoring rubric, and (iii) a short expert-panel survey mapped to those same metrics. Together these produce quantitative signals (agreement rates, CVI, stance/confidence shifts, rubric composites) and qualitative insights (directed/thematic analysis), with a transparent audit trail throughout (Braun and Clarke, 2006; Hsieh and Shannon, 2005; Miles et al., 2014).

Case notes (concise EDMC decision records). Three anonymized case notes were used as stimulus materials for the expert panel. Each note makes ethically salient elements explicit while remaining portable across contexts, following the PASO chain (Principles → Actions with roles/artifacts → Skills/ownership → Outcomes/KPIs with thresholds and review cadence). For study purposes, “concise EDMC decision record” and “case note” are used interchangeably. The case-note texts are archived for examiners under NDA; the panel-facing versions and format appear in Appendix D. Versioning and access controls are documented in Appendix N.

PASO plan scaffold. A compact internal scaffold ensured every case note and its assessment shared the same logic and traceability: Principles (why), Actions with

roles/artifacts (how), Skills (who), and Outcomes/KPIs and mitigations (what).

Terminology: “Skills” and “Capabilities” are treated as the same PASO construct in this thesis.

EDMC scoring rubric (0–3). The rubric contains six metrics—PASO linkage, blindness mitigation, cultural adaptability, time-pressure suitability, implementation guidance, and outcomes evaluability—each scored on anchors from 0 (not evident) to 3 (clearly and consistently evidenced). Rubric anchors, weighting options, and computation rules are consolidated in Appendix B (*Table 34* and Section B.3). A blank worked example (per-metric scoring sheet) appears in Appendix B2 (*Table 35*). For the researcher-led public cases, a ≤ 15 -word rationale accompanies each metric score; those one-liners are recorded in Appendix A (Grant and Davis, 1997; Lynn, 1986; Polit and Beck, 2006).

Expert-panel survey. The online instrument presents, for each of the three cases, a fixed sequence: brief baseline judgments (stance: Proceed / Proceed with conditions / Pause; confidence 0–10, with a short rationale) → exposure to the EDMC case note → post-EDMC items, including Q1–Q8 (1–5 Likert on actionability, salience, defensibility, portability, integratability, time-pressure usability, and intention to use), a PASO micro-block (P/A/S/O, 1–5), post-stance, post-confidence (0–10), an efficiency item (Less / Same / More / Uncertain), and an optional brief comment. A global wrap-up follows once (overall effectiveness; sufficiency count across cases; learning/efficiency from Case 1→3; portability; integratability; time-pressure usability; conceptual coherence; optional tech/AI relevance; two short open prompts). The item index is in Appendix C.2 and the verbatim, version-locked PDF is in Appendix C.5. Order and timing were fixed and designed for ~15–20 minutes (Krosnick and Presser, 2010). Versioning/admin handling are documented in Section 3.8 and Appendices M–N.

Presets and composites. Where relevant, case composites were computed as the mean of the six metrics (0–3), with optional weight presets (Equal; Stakeholder-emphasis; Time-pressure-emphasis). N/A values are excluded from denominators and, if any metric is N/A, remaining weights are renormalized. Definitions and formulas appear in Appendix B.3 and are applied in the Results chapter.

PASO-embedded rationale capture under time pressure. To reflect real constraints, brief “under-pressure” prompts (red flags, stakeholder sweeps, go/stop thresholds) are embedded alongside the case notes and rubric sheets. These cues support succinct justification of stance changes and “proceed with conditions” decisions.

Expert packet (delivery bundle). Participants received consent wording, an NDA reminder, the three case notes, the survey link, and a one-page rubric overview—enough common ground to ensure consistent interpretation without adding burden.

Derived variables and reporting. From the survey we report % ≥ 4 agreement on 1–5 Likert items (with N and N_eff), PASO content-validity indices (I-CVI per P/A/S/O; S-CVI/Ave as the mean of the four), stance-change rates, confidence uplift (Post – Baseline on 0–10), and efficiency shifts. From the rubric we report per-metric and composite scores (0–3), optionally rescaled to 0–100 for plotting. Operational definitions and formulas are consolidated in Appendix F; summary tables and figures appear in Appendices G–H.

Quality safeguards. Content-validity conventions (Grant and Davis, 1997; Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006) guided the PASO micro-block. Instrument and case notes are version-locked (23 Sep 2025), with a single pre-lock link correction and no wording change (Appendix N). Anonymity and privacy safeguards include pseudonymous participant codes, privacy-safe demographics, and GDPR/FADP-aligned handling (Appendix K). Inclusion/exclusion rules were pre-specified: duplicate/mis-

coded rows were excluded; corrected resubmissions included; any post-closure entries were marked LATE/EXCLUDE (Appendix M).

Why this configuration? Pairing a practical decision artifact (the concise EDMC decision record with PASO-embedded prompts) with short, validated judgment items produces evidence on usability, portability, and defensibility under time constraints—and triangulates that evidence with a rubric that renders the PASO chain auditable. This balances rigor and burden, consistent with best practice in expert-panel content validation (DeVellis, 2017; Lynn, 1986; Polit and Beck, 2006).

Cross-references. Item index and verbatim instrument → Appendices C.2 and C.5; rubric anchors, weights, and formulas → Appendix B; analysis plan and derived metrics → Appendix F; results summaries and visuals → Appendices G–H; privacy/consent → Appendices K–L; exclusions/versioning → Appendices M–N.

3.8 Data Collection Procedures

To evidence EDMC’s usability, validity, and robustness without burdening organizations, the study used two coordinated streams: (A) researcher-led public case applications of EDMC to eight documented cases, and (B) an online expert-panel survey using three concise EDMC decision records (“case notes”) as stimuli. The streams are complementary: the public case applications demonstrate coverage and transparent scoring, while the panel provides content-validity indices, stance/confidence shifts, portability/integratability judgments, and targeted qualitative feedback. All study artifacts were version-locked on 23 Sep 2025, 12:30 CEST; the fieldwork window ran 22–29 Sep 2025 and closed at 14:40 CEST (Appendix N).

The survey was administered in Google Forms with automatic capture to Google Sheets. No custom scripting was used; all summaries were computed with spreadsheet

formulas. The compact item index appears in Appendix C.2, and the verbatim, delivered instrument (PDF) is archived in Appendix C.5 (Dillman et al., 2014; Krosnick and Presser, 2010).

Researcher-led public case applications (C1–C8). For each case, public facts and sources were assembled (no proprietary data). A short PASO plan was drafted (Principles; Actions/roles/artifacts; Skills; Outcomes/KPIs and mitigations). The six EDMC metrics were then scored on the 0–3 rubric (Appendix B), with a ≤ 15 -word rationale logged per metric (Appendix A). Case composites were computed under the disclosed weight presets (Equal; Stakeholder; Time-pressure), and all calculations were archived to preserve transparency.

Expert-panel survey (three vignettes). Invitees received consent wording, an NDA reminder, and a unique NDA-XX participant code. Proceeding to the form indicated consent; codes verified NDA status and were removed at anonymization (Appendices K–L). For each of three anonymized, cross-sector vignettes, participants provided a baseline stance (Proceed / Proceed with conditions / Pause) and baseline confidence (0–10) with a brief rationale, reviewed the EDMC case note, and then completed the post-EDMC block: Q1–Q8 (1–5 Likert), a PASO micro-block (P/A/S/O, 1–5), post-stance, post-confidence (0–10), an efficiency comparison, and an optional brief comment. A single global wrap-up followed after Case 3 (overall effectiveness; sufficiency count; portability; integratability; time-pressure usability; conceptual coherence; optional tech/AI controls; two short open prompts). The fixed sequence minimized order effects and reflected realistic decision flow; completion time was ~15–20 minutes (Krosnick and Presser, 2010). Public-facing vignette texts and the case-note format appear in Appendix D; versioning and access control are documented in Appendix N.

Why hypothetical vignettes for the panel? Using generic, broadly structured scenarios reduces prior-knowledge and media-salience biases and keeps attention on EDMC’s clarity, usability, and portability rather than on the particulars of any incident. Vignette designs are widely used to study expert judgment because they balance control with realism (Aguinis and Bradley, 2014; Evans et al., 2015; Tremblay et al., 2022). The separate stream of public-facts cases (C1–C8) then tests robustness and portability on documented scenarios. The two datasets are analyzed and reported separately.

Pre-specified inclusion/exclusion rules. Only responses on instrument v1.0 (post-fix, version-locked) were included. If the same participant code appeared more than once, the latest valid submission was analyzed and any earlier rows were excluded. Mis-coded or duplicate codes were excluded. Submissions time-stamped after 29 Sep 2025, 14:40 CEST were marked LATE/EXCLUDE. Actions taken are documented in Appendix M, along with the sensitivity check outcomes.

Cleaning, derivations, and reporting. The analysis used the live Form→Sheet export with the above rules applied. Derived indicators include % \geq 4 agreement on 1–5 items (with N and effective N), PASO content-validity indices (I-CVI per item and S-CVI/Ave per case), confidence uplift (Post – Baseline on a 0–10 scale), stance change (Baseline→Post), and efficiency shift categories, plus rubric per-metric and composite scores (0–3), optionally rescaled to 0–100 for visualization. Operational definitions and formulas are consolidated in Appendix F; summary tables appear in Appendix G and figures in Appendix H. Means are reported to two decimals; whole percentages are used for % \geq 4; item-level Ns are shown.

Ethics and data protection. Responses were stored pseudonymously. No names/emails were collected; demographics used privacy-safe bins. The code↔person mapping was held separately (admin-only) and destroyed at anonymization; participant

rights and retention are described in Appendix K. NDA governed access to stimulus case notes; only public-facts summaries are reproduced in Appendix D.

Cross-references. Instrument index and verbatim: Appendices C.2/C.5. Rubric anchors, weights, and computation: Appendix B. Analysis plan and derived metrics: Appendix F. Results summaries and visuals: Appendices G–H. Privacy/consent: Appendices K–L. Exclusions and versioning: Appendices M–N.

3.9 Data Analysis

We combine numbers and narratives to judge how well EDMC works in practice—and why. Data analysis in this study was designed to rigorously evaluate the effectiveness, usability, and robustness of the Ethical Decision-Making Compass (EDMC) using both quantitative and qualitative methods. Integrating these strands provides a comprehensive view of how EDMC performs across diverse scenarios and user perspectives, consistent with mixed-methods guidance for practice-oriented inquiry (Creswell and Plano Clark, 2018). We analyze two quantitative datasets—(i) the expert-panel survey and (ii) the researcher-led case tallies—and a qualitative corpus of free-text comments and scoring rationales. Each strand is specified separately and then triangulated in the Results (Chapter IV).

3.9.1 Quantitative Analysis (Expert-Panel Dataset)

The aim is to quantify EDMC’s usability, content validity, decision effects, and perceived efficiency under time pressure. The unit of analysis is participant \times case (three repeated case blocks per participant); a small set of global items is analyzed once per participant.

Before analysis, we apply the pre-specified inclusion/exclusion rules (Appendix M) and keep only version-locked rows (v1.0; Appendix N). If a participant code appears more than once, only the latest valid submission is analyzed; earlier rows are marked EXCLUDE. Submissions after the close timestamp are EXCLUDE. “N/A” responses are retained in the dataset but are excluded from denominators when computing proportions and means; every statistic reports the N it uses.

Agreement for each Likert item is reported as the whole-percent share of responses rated 4 or 5 on the 1–5 scale: Agreement (%) = $[\text{count}(4-5) \div N \text{ valid (excluding N/A)}] \times 100$. Content validity for the PASO micro-block follows standard conventions: the item-level I-CVI is the % ≥ 4 for that PASO item, and the per-case S-CVI/Ave is the mean of the four I-CVIs (P, A, S, O) (Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006). Confidence uplift is computed per respondent as Post – Baseline on the 0–10 scale; we report the mean uplift (and, where helpful, the median/IQR). Stance change is summarized by cross-tabulating Baseline vs. Post categories (Proceed; Proceed with conditions; Pause) and reporting the share whose category changed, optionally split by direction. Perceived efficiency is shown as the proportions in the four response categories (Significantly less; Moderately less; Same; More; Uncertain). When aggregating across a participant’s three cases, we use simple means with equal case weights unless noted; we always display the effective N for each aggregate (available-case basis).

Descriptive outputs use whole-percent % ≥ 4 and means to two decimals. Key tables and figures are located as follows: headline indicators by case (Appendix G1); PASO coverage and S-CVI/Ave (Appendix G2); the Q1–Q8 % ≥ 4 heatmap (Appendix H1); PASO coverage plot (Appendix H2); efficiency distribution (Appendix H3); confidence uplift (Appendix H4); pooled Q1–Q8 means (Appendix H5); PASO

consistency and “sufficient for a defensible decision” distributions (Appendix H6); and global wrap-up means (Portability, Integratability, Time-pressure usability, Conceptual coherence, Tech/AI) in Appendix H7. A stance-change matrix with row percentages and valid Ns appears in Chapter 4.

Analytically, this phase is descriptive by design—agreement rates, CVI, stance/confidence shifts, and efficiency shares—chosen to align with content-validation goals and the study’s sample design. No inferential tests are reported here; non-parametric checks (e.g., Wilcoxon, Friedman) are reserved for future extensions if warranted by design and power.

3.9.2 Quantitative Analysis (Researcher-Led Case Applications)

This strand quantifies EDMC on eight public-facts cases (C1–C8) using a transparent, rubric-based procedure. For each case, the assessor scores the six PASO-aligned metrics on a 0–3 scale using the anchors in Appendix B, Table B1, recording a ≤15-word rationale per metric in Appendix A. Per-metric means are then averaged to a case composite under equal weights by default (N/A values are allowed only when truly inapplicable and are excluded from denominators). Where helpful for plotting, composites on a 0–3 scale are linearly rescaled to 0–100 as:

$$\text{Composite}(0-100) = [\text{Composite}(0-3) / 3] \times 100 = [(\text{Score}_1 + \dots + \text{Score}_6) / 18] \times 100.$$

To check robustness to decision priorities, we recompute the same composites under the disclosed weight presets (Equal; Stakeholder-emphasis; Time-pressure-emphasis; see Appendix B.3), report the range and median across presets, and flag any rank-order changes (Aguinis and Solarino, 2019). All means are rounded to two decimals; the effective N is disclosed wherever N/A occurs; computation rules are consolidated in Appendix F.

3.9.3 Qualitative Analysis (Free-Text Comments and Scoring Rationales)

The qualitative strand explains why the numbers move by analyzing brief open comments and the assessor's one-line score rationales through a PASO-anchored lens. The purpose is to interpret quantitative signals (e.g., $\% \geq 4$ agreement, stance tightening, confidence uplift) and surface concrete delivery-layer refinements—such as clearer summary strips, checklists, RACI prompts, KPI exemplars, or jurisdiction hooks.

The approach follows directed content analysis (Hsieh and Shannon, 2005) with an initial, deductive codebook aligned to PASO and the six rubric metrics, while allowing inductive additions where needed; themes are then distilled using focused/thematic synthesis (Braun and Clarke, 2006; Miles et al., 2014). The codebook explainer in Section 3.9.3 outlines each code's name, definition, inclusion/exclusion notes, and a short example; shells and quick examples are provided in Appendix H. As an illustration, the Blindness-mitigation code captures mentions of prompts, go/stop gates, red flags, or checks intended to surface ethical salience under time pressure (e.g., “Add a go/no-go threshold for model bias before launch”), while excluding generic exhortations with no concrete trigger.

The corpus includes (i) per-case brief comments, (ii) two global prompts (limits; one improvement), and (iii) the researcher's one-line rationales. Text is segmented at the sentence or short-paragraph level. The frozen codebook is applied first; ambiguous assignments receive a brief in-line note, and genuinely new patterns may be captured via provisional inductive codes.

To support trustworthiness, we keep dated codebook snapshots and analytic memos, re-review ambiguous units after ~48 hours, and—where feasible—have a second coder review a subset, recording simple agreement and any discrepancy resolutions

(Saldaña, 2021). Cross-references: rubric anchors → Appendix B; computation rules → Appendix F; codebook shells/examples → Appendix H.

3.9.4 Missing Data, N/A, and Exclusions

We handle missingness conservatively and apply pre-specified rules so summaries aren't distorted. "N/A" responses are retained in the dataset but excluded from denominators for both $\% \geq 4$ and means; for every statistic we report the effective N (N_{eff}). We do not impute missing values and always show the item-level N used. Exclusions follow Appendix M (mis-coded/duplicate participant codes, late/closed-window rows); included vs. excluded counts are documented there. Version integrity is enforced by analyzing only v1.0 rows after the pre-lock fix; the Case-3 link correction did not change any wording (Appendix N).

3.9.5 Robustness, Reliability, and Transparency

To test stability, we recompute headline indicators after applying the exclusion rules in Appendix M and under the disclosed weight presets (Equal, Stakeholder, Time-pressure). We flag a material shift if S-CVI/Ave moves by ≥ 0.05 , if median confidence changes by ≥ 1 point (0–10), or if qualitative themes change in substance; otherwise we note "immaterial impact" in Results. Because a single assessor scored the rubric in this phase, inter-rater and test-retest statistics were not computed; instead, reproducibility rests on disclosed anchors, one-line rationales, and version-locked materials. All statistics were produced with formula-only spreadsheets (no scripts); summary screenshots and a minimal data dictionary are provided in Appendix G, and operational definitions and formulas are consolidated in Appendix F. No inferential tests were run in this phase;

future hypothesis testing, if undertaken, should be pre-registered and emphasize effect sizes over p-values.

3.10 Research Design Limitations

The design prioritizes rigor and fitness-for-use. To keep interpretation proportionate, we state key limitations alongside the concrete mitigations employed. Administrative details and change logs are documented in Appendices M–N.

- Case-based applications (analytic generalization). Findings from researcher-led applications to publicly documented cases support transferability and fitness-for-use rather than statistical generalization (Shadish et al., 2002; Yin, 2018). Mitigation: transparent PASO plans, explicit scoring criteria, and sensitivity checks to support analytic generalization.
- Vignette design for the expert panel. Using hypothetical cases enhances objectivity and comparability but cannot capture every nuance of real-world pressure (Aguinis and Bradley, 2014; Evans et al., 2015). Mitigation: we triangulate with researcher-led public-facts cases (C1–C8) and propose field pilots/longitudinal outcomes as next steps (Tremblay et al., 2022).
- Expert-panel scope and sampling. The panel is purposive, not probabilistic; results reflect this group’s contexts. $N \approx 30$ yields stable descriptive statistics but is not designed for population inference. Mitigation: multi-sector experts (≥ 10 years’ experience); item-level N and variation reported.
- Measurement subjectivity (rubric scoring). 0–3 judgments can be subjective. Mitigation: shared anchors (Appendix B), one-line rationales (Appendix A), and—where feasible—a second-rater check or simple agreement statistic. For PASO micro-items, we report I-CVI and S-CVI/Ave (Lynn, 1986; Polit and Beck, 2006).

- Common-method and desirability bias. Usability perceptions and post-stance are same-session self-reports; social desirability/halo may inflate ratings (Shadish et al., 2002). Mitigation: anonymity, NDA context (reduced performance signaling), neutral wording, and emphasis on practical critique.
- Version integrity and administrative issues. A minor Case-3 link fix occurred; wording did not change. Three duplicate “07” codes and one corrected re-submission were handled per protocol (Appendix M). Mitigation: version lock (Appendix N), rule-based exclusion/supersedence, and sensitivity analysis indicating no material impact on headline conclusions.
- Weights and value choices. Any weighting encodes priorities (e.g., stakeholder emphasis, time-pressure emphasis). Mitigation: disclose formulae, report equal-weight results by default, and show robustness across presets (Aguinis and Solarino, 2019).
- Instrument breadth vs. depth. EDMC targets decision usability under constraints; it does not replace legal/compliance audits or deep technical risk assessments. Outcomes are proximal (clarity, defensibility, conditions to proceed), not long-term societal effects.
- Data limitations. N/A responses reduce effective denominators; no imputation was performed. Mitigation: report effective N per statistic and exclude N/A by design.
- Researcher reflexivity. Practitioner-scholar priors are possible. Mitigation: explicit criteria, documented rationales, and an audit trail enhance transparency (Miles et al., 2014; Saldaña, 2021).
- Scope of comparison. This phase prioritizes EDMC fitness-for-use; quantitative head-to-head scoring of other frameworks was out of scope and treated qualitatively in Chapter 2.

- Generalizability / boundary conditions. Findings reflect hypothetical applications and an expert sample; results are transferable but not population-representative. Effects depend on local capability, data quality, and governance maturity (Shadish et al., 2002).

3.11 Ethical Considerations and Researcher Reflexivity

This study followed established ethical principles for human research and qualitative inquiry (Alvesson and Sköldbberg, 2017; Israel, 2015; Lincoln and Guba, 1985). The expert-panel survey captured only non-identifying demographic bins and pseudonymous participant codes; no directly identifying data were collected. Participation was voluntary, with the right to withdraw prior to anonymization; an information sheet and explicit consent were presented at instrument start (Appendix L). To reduce social-desirability and common-method bias, responses were anonymous, an N/A option was available, and item wording was neutral (Podsakoff et al., 2003).

Data were collected in Google Forms/Sheets within a restricted Google Workspace; access was limited to the researcher. Aggregate reporting avoids small-cell disclosure. Data handling complies with SSBM guidance and EU GDPR/Swiss FADP norms (Appendix K). The instrument and case notes were version-locked; a minor link fix did not change wording (Appendix N). Case notes used as stimuli were shared under NDA; only the anonymized stimuli text is reproduced in Appendix D, while public-facts case rationales are in Appendix A and source notes in Appendix O.

Researcher reflexivity acknowledges the dual role as practitioner-scholar. Assumptions were surfaced and bracketed via an audit trail (codebook snapshots, memos, and formula-only summary evidence) and through shared rubrics and disclosed criteria

(Kolb, 1984; Saldaña, 2021; Schön, 1983). Where feasible, a secondary check on a subset was used to gauge scoring consistency.

These safeguards are proportionate to an applied, practice-oriented design. The mitigations and audit trail described here and in Appendices K–N support credible, ethical inference.

3.12 Conclusion

This chapter established the methodological foundation for developing and validating the Ethical Decision-Making Compass (EDMC). Using a sequential mixed-methods design, the study moved from a comparative analysis of existing frameworks, through the synthesis and operationalization of PASO, to empirical validation via an expert-panel survey and researcher-led applications to public cases (Creswell and Plano Clark, 2018).

The methods were chosen to fit real decision conditions. One-page EDMC case notes and a PASO plan mirror how leaders actually decide—under time pressure, with limited attention, and a need for traceability. Where dense, multi-page toolkits tend to fail, compact artifacts keep ethical salience visible while preserving an audit trail. A six-metric 0–3 rubric then prioritizes fitness-for-use and auditability over fine-grained psychometrics: at an early, practice-oriented stage, clear anchors, transparent rationales, and replicable scoring are more valuable than complex latent models (DeVellis, 2017; Miles et al., 2014). For validation, a content-validity-oriented expert panel provides the right kind of evidence for an instrument at this maturity. Do knowledgeable reviewers agree that EDMC’s elements are present, usable, and defensible? - using established I-CVI/S-CVI conventions (Lynn, 1986; Polit and Beck, 2006). By contrast, large probabilistic surveys, EFA/CFA, or RCTs address different questions (population

inference, latent structure, causal impact), and would impose burdens disproportionate to the present aim of practice-ready usability under constraints (Shadish et al., 2002).

Finally, applying EDMC to publicly documented cases substituted for proprietary field trials, avoiding confidentiality barriers while still testing coverage, scoring transparency, and robustness (Yin, 2018).

This dual approach, hypothetical vignettes for the expert panel and real-world public cases for the researcher-led applications, balances objectivity with practical relevance. The vignettes focus expert judgment on EDMC's core features; the public cases probe robustness in realistic, complex scenarios. Together, they enhance the credibility and transferability of the findings.

In reporting, the chapter specified who participated; which instruments were used; how data were collected, cleaned, and analyzed; how robustness was checked; and how ethics and reflexivity were handled. Quantitatively, the expert-panel strand reports agreement rates, content-validity indices, and stance and confidence shifts; the case-application strand uses a transparent 0–3 rubric, a disclosed composite, and weight-sensitivity checks (Aguinis and Solarino, 2019). Qualitatively, directed and thematic analyses surface facilitators, barriers, and actionable improvements (Braun and Clarke, 2006; Hsieh and Shannon, 2005). These choices balance rigor with fitness-for-use and make the analysis reproducible via version-lock and a formula-only dashboard (Appendices F–G).

Methodologically, the constructive, pattern-first synthesis and the pairing of expert review with public-facts applications advance ethical decision-making research by combining transparent, auditable scoring with both expert and field-proximate perspectives. This bridges theoretical rigor and practical usability—essential for frameworks meant for high-stakes, time-pressured leadership contexts.

Limitations were acknowledged, including the use of hypothetical cases, a non-probabilistic expert sample, and the absence of longitudinal field trials. These boundaries keep interpretation proportionate and point to next steps such as site-based pilots, longitudinal adoption studies, and experimental comparisons (Shadish et al., 2002).

With this methodological foundation in place, the next chapter presents the empirical results—integrating quantitative and qualitative evidence—to assess EDMC’s effectiveness, usability, and decision impact against the research questions and objectives, showing not only whether EDMC works, but how and why it can be adopted across diverse leadership settings.

CHAPTER IV: RESULTS

4.1 Introduction

This chapter reports results from two coordinated evidence streams used to evaluate the Ethical Decision-Making Compass (EDMC). The first stream is an expert-panel survey (N = 30; 90 case blocks) in which specialists reviewed three one-page EDMC case notes (hypothetical vignettes, Case 1–3) and completed a structured instrument. The second stream comprises researcher-led applications of EDMC to eight publicly documented cases (C1–C8), each scored with a transparent 0–3 rubric aligned to PASO. Read together, these streams offer a practice-oriented view of EDMC’s effectiveness, usability under time pressure, and robustness (Morgan, 2014; Yin, 2018).

In the expert-panel strand, we summarize item-level agreement rates ($\% \geq 4$ on 1–5 items), PASO content-validity indices (I-CVI and S-CVI/Ave), stance-change matrices, confidence uplift (Post – Baseline on a 0–10 scale), and perceived efficiency (Less / Same / More / Uncertain). We also report global utility judgments—portability, integratability, time-pressure usability, conceptual coherence, and tech/AI controls. These indicators follow established conventions for content validation and mixed-methods survey evaluation (Braun and Clarke, 2006; Creswell and Plano Clark, 2018; Grant and Davis, 1997; Hsieh and Shannon, 2005; Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006).

In the researcher-led case applications, each public-facts case is evaluated on six PASO-aligned metrics—PASO linkage, blindness mitigation, cultural adaptability, time-pressure suitability, implementation guidance, and outcomes evaluability—scored 0–3 against shared anchors. We compute case composites (0–3; optionally rescaled to 0–100) and probe robustness by recalculating under disclosed weight presets (Equal; Stakeholder

emphasis; Time-pressure emphasis), following best practice for sensitivity analysis (Aguinis and Solarino, 2019).

Across both streams we apply a small set of standard computations, detailed in Appendix F. Agreement is the share of responses rated 4 or 5 on a 1–5 scale ($\% \geq 4$), excluding N/A. PASO content validity uses I-CVI per item and S-CVI/Ave per case (Lynn, 1986; Polit and Beck, 2006). Confidence uplift is the mean difference between post- and baseline confidence (0–10); stance change compares baseline vs. post decision categories; efficiency summarizes the four response categories. For rubric results, composites are the mean of the six 0–3 metrics (with N/A excluded), optionally rescaled to 0–100; sensitivity is shown across the preset weight schemes (Aguinis and Solarino, 2019).

Only version-locked instrument rows (v1.0, 23 Sep 2025) are included. The fieldwork window ran 22–29 Sep 2025 (form closed 29 Sep 2025, 14:40 CEST). Administrative handling and timestamps are documented in Appendices M and N. The chapter is organized as follows: Section 4.2 presents expert-panel results (demographics, quantitative indicators, brief qualitative themes); Section 4.3 presents the researcher-led case applications (metric scores, 0–100 composites, weight sensitivity); Section 4.4 summarizes pre-specified robustness checks; Section 4.5 synthesizes the evidence. Summary tables appear in Appendix G and figures with concise computation notes in Appendix H; operational definitions are consolidated in Appendix F.

4.2 Expert-Panel Survey Results Introduction

This section addresses RQ6 (primary) and supports RQ4–RQ5. It summarizes usability, PASO coverage, decision effects, and global utility from $N = 30$ experts who completed 90 case blocks across three hypothetical vignettes.

To assess EDMC's practical utility under real constraints, we ran an anonymous expert-panel survey using three one-page EDMC case notes (AI hiring bias; supplier labor compliance; data-privacy breach). For each vignette, participants recorded a baseline stance and baseline confidence (0–10), reviewed the EDMC case note, and then completed post-EDMC items covering clarity/actionability, ethical salience, defensibility, portability, time-pressure usability, integratability, and intention to use. They then provided a revised confidence score and a decision-change judgment. A short global block captured portability, integratability, time-pressure usability, conceptual coherence (principles → actions → outcomes), and tech/AI controls. PASO content coverage was evaluated via I-CVI (per item) and S-CVI/Ave (scale average) using established content-validity conventions (Grant and Davis, 1997; Lynn, 1986; Polit and Beck, 2006). Open responses were analyzed using directed content analysis with thematic coding (Braun and Clarke, 2006; Hsieh and Shannon, 2005).

Agreement items used a 1–5 Likert scale; we report % ≥ 4 with N/A excluded and N shown. Confidence was captured on 0–10 and summarized as uplift (Post – Baseline). PASO content validity is reported as I-CVI (proportion ≥ 4 per PASO item) and S-CVI/Ave (mean of the four PASO items per case). Participation was anonymous; demographics were optional and are reported only in aggregate. Computation rules follow Appendix F.

Overview of results. EDMC performed strongly on actionability, ethical salience, and defensibility across all three cases: Actionable pathway = 83–93%, Ethical dimensions explicit = 87–93%, and Defensible to stakeholders = 87–90% (see heatmap in Appendix H1; detail in Appendix G1). PASO content validity was high: S-CVI/Ave = 0.85, 0.82, 0.87 across Cases 1–3, with strong I-CVI values for P/A/S/O (Appendices H2, G2). On decision efficiency, the share reporting less time to a decision was 63%, 67%,

73% (Cases 1–3, respectively); more time was 23%, 20%, 13%; same \approx 0%; uncertain \approx 3% in Cases 1 and 3 (Appendix H3). Confidence uplift averaged +0.37, +0.03, +0.53 on the 0–10 scale (Cases 1–3; Appendix H4). Regarding behavioral sufficiency, 60% found EDMC sufficient in all three cases (2 cases: 27%; 1 case: 13%; Appendix H6, sufficiency panel). PASO consistency was also strong: 73% reported all four PASO elements present in all three cases (two cases: 23%; not present: 3%; Appendix H6, PASO-consistency panel). In the global wrap-up (means, 1–5), experts rated Portability = 4.09, Integratability = 3.96, Time-pressure usability = 4.20, Conceptual coherence = 4.23, and Tech/AI controls = 4.19 (Appendix G3; visualized in Appendix H7).

Cross-references: Figures H1–H7 (Appendix H) and Tables G1–G3/G4 (Appendix G). All figures/tables reflect the recomputed, version-locked dataset; N/A was excluded by design (Appendix F).

4.2.1 Sample Profile

To ground RQ6 with context for external validity, we summarize the makeup of the expert panel. All 30 invitees completed the three vignette blocks (90 case blocks total). The cohort skews senior and practice-oriented: 80.0% report 21+ years of experience (*Table 17*), and roles cluster in leadership (Executive/Board = 43.3%) and industry management (26.7%), with additional specialist and academic representation (*Table 18*). Education is correspondingly high, with Master’s (56.7%) and Doctorate (20.0%) the most common credentials (*Table 15*). Most respondents are male (73.3%) and aged 50+ (66.7%)—46.7% in 50–59 and 20.0% in 60+; one age response was missing (3.3%) (*Table 13*, *Table 14*). Sector experience is primarily private (70.0%), complemented by nonprofit/NGO (13.3%), public (6.7%), academia (3.3%), and other (6.7%) (*Table 20*). Domain coverage is broad, with notable concentrations in

Technology/Data/AI (16.7%) and Sustainability/ESG (13.3%), alongside governance/legal/compliance, finance/banking, HR/people, operations, and healthcare/public-sector roles (*Table 19; Figure 13 Distribution of Primary Domain (N = 30)*). The expert panel was internationally distributed at the point of residence: 43.3% of participants were based in Europe, 23.3% in North America, 16.7% in Asia-Pacific, 10.0% in the Middle East, and 6.7% in Africa; no respondents reported residence in Latin America & Caribbean (N = 30). This supports the claim that the assessment reflects multiple regulatory and cultural environments rather than a single jurisdiction.

Table 13 Gender (N = 30)

Source: Author (2025). Original table created from study data and materials.

Level	n	%
Male	22	73.3
Female	8	26.7

Table 14 Age Band (N = 30)

Source: Author (2025). Original table created from study data and materials.

Level	n	%
50-59	14	46.7
60 and above	6	20.0
40-49	6	20.0
30-39	3	10.0
Missing	1	3.3

Table 15 Highest Education (N = 30)

Source: Author (2025). Original table created from study data and materials.

Level	n	%
Master's degree	17	56.7
Doctorate (PhD, DBA, etc.)	6	20.0
Bachelor's degree	3	10.0
Doctorate Honoris Causa	1	3.3
Post-Graduate Diploma	1	3.3
Missing	1	3.3
MBA	1	3.3

Note: "MBA" appears as a separate self-reported category; summary statistics treat it separately for transparency.

Table 16 Field of Expertise (Self-Described; N = 30)

Source: Author (2025). Original table created from study data and materials.

Level	n	%
Leadership	18	60.0
Technology	5	16.7
Consulting	1	3.3

Humanitarian Principles (GBV and Human Rights)	1	3.3
Nonprofit management	1	3.3
Program Management, Projects and Change	1	3.3
Strategy and Finance	1	3.3
Auditing and Finance	1	3.3
Health and Wellness	1	3.3

Table 17 Years of Experience (N = 30)

Source: Author (2025). Original table created from study data and materials.

Level	n	%
21 years and above	24	80.0
16-20 years	3	10.0
11-15 years	2	6.7
6-10 years	1	3.3

Table 18 Current Role/Position (N = 30)

Source: Author (2025). Original table created from study data and materials.

Level	n	%
Executive/Board	13	43.3
Industry Professional (Manager, Executive, etc.)	8	26.7
Specialist/Advisor	4	13.3
Senior management	4	13.3
Academic (Professor, Researcher, etc.)	1	3.3

Table 19 Primary Domain (N = 30)

Source: Author (2025). Original table created from study data and materials.

Level	n	%
Technology/Data/AI	5	16.7
Sustainability/ESG	4	13.3
Other	4	13.3
Governance/Legal/Compliance	3	10.0
Finance/Banking	3	10.0
HR/People	3	10.0
Operations	3	10.0
Healthcare/Public sector	2	6.7
Risk/Audit	2	6.7
Education	1	3.3

Table 20 Sector (N = 30)

Source: Author (2025). Original table created from study data and materials.

Level	n	%
Private	21	70.0
Nonprofit/NGO	4	13.3
Other	2	6.7
Public	2	6.7
Academia	1	3.3

This purposive, senior panel fits the study's aim: to evaluate fitness-for-use and content coverage rather than estimate population parameters. Content validation requires knowledgeable judges who can assess whether the framework's elements are present and

usable against explicit criteria (Grant and Davis, 1997; Lynn, 1986; Polit and Beck, 2006). Using hypothetical one-page case notes (AI, labor, privacy) mirrors the time-pressured, information-light conditions under which leaders actually decide, while spanning cross-domain contexts to test portability—aligning with the study’s pragmatic stance (Morgan, 2014). Anonymous participation, brief open comments, and three within-person case blocks support candid feedback and comparability; disclosed coding and reporting conventions ($\% \geq 4$, I-CVI/S-CVI, confidence uplift) make judgments auditable (Braun and Clarke, 2006; Hsieh and Shannon, 2005). Because the panel is non-probabilistic and skews senior, findings support analytic generalization (not population inference); we therefore report item-level Ns, exclude N/A from denominators by design, and triangulate with the researcher-led public-facts cases to strengthen credibility (Miles et al., 2014; Yin, 2018).

One-line profile. A senior, practice-oriented cohort (80% with ≥ 21 years), leadership-heavy roles, broad domain spread (notably Technology/Data/AI and ESG), and mostly private-sector experience—appropriate for evaluating an executive decision tool.

The distribution details appear in Tables 13–20 (Gender, Age, Education, Field of Expertise, Years of Experience, Role/Position, Primary Domain, Sector) and *Figure 13 Distribution of Primary Domain (N = 30)*.

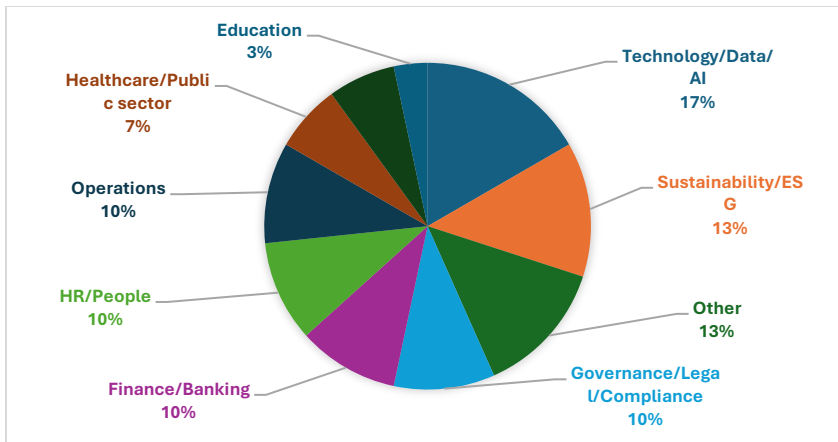


Figure 13 Distribution of Primary Domain (N = 30)

Source: Author (2025). Original figure created from study data and materials

4.2.2 Usability and PASO Content Validity

This section addresses RQ6 (usability under constraints) and supports RQ4 (coverage via PASO). We summarize item-level agreement on the eight usability/utility statements (Q1–Q8) and the PASO content-validity indices (I-CVI; S-CVI/Ave). As in Chapter 3, agreement is reported as $\% \geq 4$ on a 1–5 scale with N/A excluded and N shown, a top-two-box convention that keeps interpretation simple and comparable across items. For reference, the Q1–Q8 stems were: Q1 Actionable pathway, Q2 Ethical dimensions explicit, Q3 Prompted additional considerations, Q4 Time-pressure usability, Q5 Defensible to stakeholders, Q6 Portability, Q7 Integratability, and Q8 Intention to use.

Usability (Q1–Q8, $\% \geq 4$). Agreement was consistently high across cases (see Table 21). Across the 24 item×case cells, the range was 63.3–93.3%, and 22/24 cells were $\geq 80\%$. Case-level medians clustered in the mid-80s (C1 $\approx 85.0\%$, C2 83.3%, C3 86.7%). The soft spots were Q3 (Prompted additional considerations) in Case 1 (73.3%) and Case 2 (63.3%); Case 3 was uniformly $\geq 83.3\%$. The heatmap in Figure 17 (Appendix H) visually echoes these percentages. Computation note: $\% \geq 4 = \text{count}(4-5) \div N$ valid; N/A excluded; N = 30 per case.

Table 21 Agreement ($\geq 4/5$) for Q1–Q8 by Case (% , N)

Source: Author (2025). Original table created from study data and materials.

Item	Case 1 % ≥ 4	Case 2 % ≥ 4	Case 3 % ≥ 4	Case 1 N	Case 2 N	Case 3 N
Q1	90.0	83.3	93.3	30	30	30
Q2	86.7	86.7	93.3	30	30	30
Q3	73.3	63.3	83.3	30	30	30
Q4	86.7	83.3	86.7	30	30	30
Q5	86.7	86.7	90.0	30	30	30
Q6	83.3	83.3	83.3	30	30	30
Q7	83.3	80.0	86.7	30	30	30
Q8	80.0	83.3	83.3	30	30	30

(Computation: % ≥ 4 = count(4–5) \div N valid; N/A excluded; N per case = 30.)

PASO content validity (I-CVI; S-CVI/Ave). The PASO micro-block asked whether the note made Principles explicit, stated Actions/roles, identified Skills/resources, and specified Outcomes/KPIs and mitigations. We compute I-CVI as % ≥ 4 for each PASO element and S-CVI/Ave as the mean of the four I-CVIs per case, following content-validity conventions (Lynn, 1986; Polit and Beck, 2006). Scale-level coverage was good: S-CVI/Ave = 0.850 (C1), 0.817 (C2), 0.867 (C3) (see Table 22); visualized in Figure 18. By item, Principles and Outcomes were consistently strong (≈ 0.83 – 0.90). Skills dipped in Case 2 (0.733) while meeting ≥ 0.83 in Cases 1 and 3, suggesting a targeted refinement opportunity around clarifying required skills/resources and ownership.

PASO content validity by case: I-CVI (per item) and S-CVI/Ave (mean of P/A/S/O). Metric = % ≥ 4 on 1–5 (N/A excluded; N=30 per case).

Table 22 PASO Content Validity (I-CVI Per Item; S-CVI/Ave Per Case)

Source: Author (2025). Original table created from study data and materials.

PASO item	C1 I-CVI	C1 N	C2 I-CVI	C2 N	C3 I-CVI	C3 N
Principles	0.867	30	0.867	30	0.833	30
Actions	0.867	30	0.833	30	0.900	30
Skills	0.800	30	0.733	30	0.833	30
Outcomes	0.867	30	0.833	30	0.900	30
S-CVI/Ave	0.850	—	0.817	—	0.867	—

(Computation: I-CVI = %≥4 per PASO item; S-CVI/Ave = mean of P/A/S/O per case; N per case = 30; N/A excluded.)

How to read these findings. In this study, “usability” means a one-page EDMC note can be read quickly, applied under time pressure, and defended to stakeholders without extra training or tools. High %≥4 on Q1, Q4, and Q5 indicates that the notes were generally actionable, time-usable, and defensible. The modest troughs on Q3 highlight where brief, additional prompts (e.g., “what did we miss?” cues) could further strengthen salience under tight time budgets. The PASO indices indicate that experts judged the Principles→Actions→Skills→Outcomes chain to be present and coherent overall, with the Skills element in Case 2 calling for crisper articulation of resources and ownership.

These results establish that EDMC is perceived as usable and well-covered on PASO by a senior, cross-sector panel, with a clear, actionable cue for improvement (make Skills even more explicit in some contexts). Section 4.2.3 turns to decision effects—stance shifts and confidence uplift (0–10) after reviewing the EDMC note.

4.2.3 Decision Stance and Confidence

This section speaks to RQ6 by examining how EDMC affected experts’ policy stance (Proceed / Proceed with conditions / Pause) and their confidence (0–10) after reviewing each one-page case note. We treat any off-diagonal move in the Baseline × Post matrix as a change, and we define tightening as a shift toward greater caution (Proceed → Conditions/Pause; Conditions → Pause). Confidence uplift is computed per respondent × case as (Post – Baseline).

Stance change (pooled across cases). Among the 49 valid paired decisions, 12.2% changed, and all moves were tightenings toward stricter safeguards. Most entries

remained on the diagonal; when stances shifted, they did so in the cautious direction—consistent with EDMC surfacing conditions and guardrails without prompting wholesale reversals. See *Table 23*.

Table 23 Stance-Change Matrix (Counts; Pooled Across Cases, N = 49)

Source: Author (2025). Original table created from study data and materials.

Baseline ↓ / Post →	Proceed	Proceed w/ conditions	Pause	Row total
Proceed	0	3	1	4
Proceed w/ conditions	0	43	2	45
Pause	0	0	0	0
Column total	0	46	3	49

Notes. *N* valid = 49 (*N/A* excluded). *Off-diagonal* = change; “*Tightened*” = move toward greater caution (*Proceed* → *Conditions/Pause*; *Conditions* → *Pause*). Any change = 6/49 (12.2%), all tightening. See *Appendix G (Table 40)*.

Confidence uplift (0–10). Confidence increased slightly overall (across-case mean per respondent +0.31, *N* = 30), with case-level means of +0.37 (Case 1), +0.03 (Case 2), and +0.53 (Case 3). Medians of 0.00 with IQR = 1.00 indicate small, selective nudges—typical when baseline confidence is already high. See *Table 24* and *Figure 20* (*Appendix H*) for the visual.

Table 24 Confidence Uplift (Post – Baseline, 0–10)

Source: Author (2025). Original table created from study data and materials.

Case	Mean	Median	IQR	<i>N</i>
C1	0.37	0.00	1.00	30
C2	0.03	0.00	1.00	30
C3	0.53	0.00	1.00	30

Metric: mean/median/IQR per case; *N* = 30; *N/A* excluded from each statistic.

Visual: *Appendix H, Figure 20* shows the same means as bars.

How to read these effects: EDMC did not induce broad stance volatility; instead, it tightened a minority of decisions and delivered modest confidence lifts, especially in Case 3. Together, this pattern suggests the notes helped experts sharpen safeguards and justify “proceed with conditions” thresholds while avoiding overcorrection.

4.2.4 Global Items (Portability, Integration, Time-Pressure, Coherence, Tech/AI)

RQ linkage. To address RQ6—EDMC’s fit-for-use under real constraints—the survey closed with five global judgments asked once per participant.

Experts rated whether the PASO logic “hangs together” conceptually, how easily EDMC could slot into existing governance/approval flows (integratability), whether it remains workable under time pressure, how portable it feels across sectors and regions, and how relevant it is to contemporary tech/AI risk. On a 1–5 scale, means near or above 4.0 indicate strong endorsement; all five items met that bar, led by conceptual coherence (4.23) and time-pressure usability (4.20), with tech/AI relevance close behind (4.19); portability (4.09) and integratability (3.96) were also favorable.

All five items clustered at $\approx 4/5$. Conceptual coherence (4.23) and time-pressure usability (4.20) were highest, followed closely by tech/AI controls (4.19). Portability was strong (4.09). Integratability was slightly lower but still favorable (3.96), suggesting minor workflow alignment rather than structural barriers. In plain terms, experts viewed EDMC as coherent, usable under time constraints, broadly portable, and straightforward to integrate.

Global items (asked once per participant): conceptual coherence (principles → actions → outcomes), integratability (fit with existing governance/approval flows), time-pressure usability, portability across sectors/regions/teams, and tech/AI controls relevance.

Table 25 Global Wrap-Up Items (Means, 1–5; N = 30)

Source: Author (2025). Original table created from study data and materials.

Global item	Mean	N
Portability	4.09	30
Integratability	3.96	30
Time-pressure usability	4.20	30
Conceptual coherence	4.23	30
Tech/AI controls	4.19	30

Note. Means on a 1–5 scale; N = 30. Values are version-locked and match Appendix G (Table 42) and Appendix H (Figure 24). “Integratability” = ease of fitting EDMC into existing governance/approval flows (see Glossary).

4.2.5 Qualitative Insights (Directed Content Analysis)

RQ focus. Primarily Q6 (practical usability under constraints), with direct implications for Q5 (refinement of the EDMC tool).

What we analyzed and how: To understand why experts rated EDMC as they did and how to make it faster to use, we coded the brief per-case comments and two global prompts using the PASO-anchored codebook described in Section 3.9.3. We then grouped codes into themes and selected short, anonymized exemplars (≤ 20 words) to illustrate each point. This qualitative strand complements the quantitative signals by explaining which delivery-layer elements most affect usability and defensibility under time pressure.

What we found: Comments broadly affirmed the one-page traceability of the EDMC case note while asking for packaging refinements that accelerate use “in the moment.” The most recurrent suggestions converged on (i) a front-of-page “decision strip” and succinct checklist so a reviewer can scan the recommended stance and conditions at a glance, “Add a clear summary: shutdown vs proceed with conditions.”; (ii) sharper accountability cues (RACI) so roles and consequences are explicit, “Make clear who owns each step and consequences if missed.”; (iii) workflow alignment with governance/legal templates to ease adoption, “Align with governance templates; provide board/audit language.”; and (iv) concrete KPI and mitigation examples to standardize follow-through, “Standardized checklist with KPI/mitigation examples would increase consistency.”

Experts also urged modest jurisdiction hooks to improve portability (“Validate across regions and regulations; consider legal norms of countries.”) and plainer language for non-specialists (“Use more commonly understood terminology; clearly define concepts.”). Several comments pointed to scalability—including guidance for smaller organizations that may lack specialists (“Assumes specialists smaller organizations may lack; include training/awareness.”). One respondent flagged a faster entry path for high-tempo situations (“Too long—hard to do in 15–20 minutes.”), reinforcing the value of a quick-reference layer. For data-intensive cases, experts wanted system visibility (“Add a system diagram of data flows and stakeholders.”). Finally, a minority highlighted a business-continuity lens—explicitly weighing jobs or continuity when safeguards impose operational costs (“Consider jobs and continuity when ethics affect livelihoods.”).

Taken together, these themes translate directly into Q5-aligned refinements: a compact summary strip + checklist, embedded RACI prompts, template-ready phrasing for board/audit contexts, exemplar KPIs/thresholds, optional jurisdiction anchors, a quick-reference path for time pressure, and short guidance notes for small-team settings. They also help explain the quantitative pattern reported in Sections 4.2.2–4.2.4: high agreement on actionability/defensibility, strong time-pressure usability, good PASO coverage (S-CVI/Ave 0.82–0.87), selective stance-tightening (12.2%), and small but positive confidence nudges (mean +0.31).

Themes and exemplars are seen in *Table 26* for transparency.

Table 26 Themes and Exemplar Quotes (Anonymized, ≤20 Words)
Source: Author (2025). Original table created from study data and materials.

Theme	Summary	Exemplar (≤20 words)
Quick summary and checklist	One-look “decision strip”; steps at a glance	“Add a clear summary: shutdown vs proceed with conditions.”
Accountability and ownership (RACI)	Who owns what; consequences if missed	“Make clear who owns each step and consequences if missed.”
Templates and integration	Fit with governance/legal workflows	“Align with governance templates; provide board/audit language.”

KPIs and examples	Concrete KPI/mitigation examples	“Standardized checklist with KPI/mitigation examples would increase consistency.”
Legal/regional fit	Multi-jurisdiction hooks	“Validate across regions and regulations; consider legal norms of countries.”
Simpler language	Plainer terms; define key concepts	“Use more commonly understood terminology; clearly define concepts.”
Scalability/resources	Guidance for smaller organizations	“Assumes specialists smaller organizations may lack; include training/awareness.”
Time-pressure path	Even faster entry path requested	“Too long—hard to do in 15–20 minutes.”
Data/tech clarity	Show data flows and stakeholders	“Add a system diagram of data flows and stakeholders.”
Business continuity lens	Jobs/continuity noted in impacts	“Consider jobs and continuity when ethics affect livelihoods.”

Note. Exemplars are lightly edited for brevity and stripped of any identifying details. If desired, add a fourth column (“n respondents”) using your coded counts for each theme.

So what for the RQs? For Q6, experts confirm the case note is usable under realistic constraints, and they specify precisely which micro-features make it faster and more defensible. For Q5, the requested enhancements map cleanly to implementable changes (summary strip, checklist, RACI, KPIs, jurisdiction hooks, quick-path), giving a concrete roadmap for the next iteration of EDMC’s delivery layer.

4.2.6 Interim Summary and Synthesis of Expert-Panel Evidence

Integrated reading for RQ6 (practical usability under constraints), with implications for RQ5 (tool refinement).

Across three hypothetical cases (N = 30; 90 case blocks), experts judged EDMC highly usable and defensible: 22 of 24 Q1–Q8 item×case cells were ≥80% agreement (range 63.3–93.3), indicating consistent clarity, actionability, ethical salience, and stakeholder defensibility. PASO content coverage met conventional “good” thresholds, with S-CVI/Ave of 0.85, 0.82, and 0.87 by case—especially strong on Principles and Outcomes, with Skills the relative bottleneck in Case 2. Decision effects were selective and conservative: among paired baseline→post choices, 12.2% changed and all shifts tightened (from Proceed to Conditions/Pause, or from Conditions to Pause), suggesting

EDMC sharpened safeguards without wholesale reversals. Confidence moved modestly upward overall (mean +0.31), with the largest gain in Case 3 (+0.53). Global wrap-up means clustered around 4/5—Conceptual coherence 4.23; Time-pressure usability 4.20; Tech/AI controls 4.19; Portability 4.09; Integratability 3.96—signaling strong perceived fit-for-use. Efficiency distributions (reported in H3) aligned with these judgments, with most respondents indicating less time to decision after using EDMC.

Qualitative comments (Section 4.2.5) explain these numbers and point to concrete, low-burden refinements for RQ5: a front-of-page summary strip and quick checklist, embedded RACI prompts, template-ready phrasing for board/audit contexts, exemplar KPIs/thresholds, light jurisdiction hooks, a quick-reference path for high-tempo decisions, and brief guidance for small teams. Together, the panel’s quantitative and qualitative signals indicate that EDMC is coherent, portable, usable under time pressure, and readily integrable—with clear guidance on where packaging tweaks can make it even faster in practice.

Key indicators → G1; PASO CVI → G2/H2; Q1–Q8 heatmap → H1; stance/confidence → *Table 24* and *Figure 20*; efficiency → *Figure 19*; global items → *Table 42/Figure 24*; qualitative themes → *Table 26*.

4.3 Researcher-Led Case Applications

RQ linkage. This section primarily addresses RQ6 (practical fitness-for-use and robustness), while also supporting RQ4–RQ5 (how PASO is organized into EDMC and translated into working tools).

The expert-panel survey established that EDMC is usable under time pressure and tends to tighten safeguards. What it could not show is how EDMC behaves when a single analyst applies the full rubric end-to-end across diverse, real-world contexts. To close

that gap, a second evidence stream applied EDMC to eight publicly documented cases (C1–C8), scoring the six PASO-aligned metrics on the shared 0–3 anchors, recording ≤15-word rationales, and computing transparent composites with prespecified weight presets. In doing so, this strand demonstrates coverage and consistency on the metrics, robustness to plausible weighting priorities (Equal / Stakeholder / Time-pressure), and an auditable trail of “conditions to proceed,” owners, and KPIs (Appendices A–B). The goal is analytic generalization and fitness-for-use—not population inference—consistent with pragmatic, problem-first designs (Lynham, 2002; Miles et al., 2014; Morgan, 2014; Shadish et al., 2002; Yin, 2018).

4.3.1 Why a Second Evidence Stream?

The panel results answer whether knowledgeable practitioners perceive EDMC as clear, portable, and defensible under time pressure. The researcher-led applications answer a different, complementary question: what happens when EDMC is actually run through complex, public-facts cases by a single assessor using the disclosed rubric and computation rules? Applying EDMC to eight cross-sector cases shows (i) how the six metrics behave together on concrete material, (ii) whether composites remain stable when priorities shift (Stakeholder-emphasis vs. Time-pressure-emphasis vs. Equal), and (iii) how the rationale lines, thresholds, and KPIs create a glass-box audit trail. This is appropriate for practice-oriented validation: it evidences repeatable scoring and defensible summaries under realistic decision constraints without claiming causal effects or population-level parameters (Lynham, 2002; Morgan, 2014; Shadish et al., 2002; Yin, 2018).

Scope is intentionally bounded: full effectiveness claims would require longitudinal, in-situ field studies, which are proposed as future work. Here, transparency

(shared anchors; one-line rationales), reproducibility (disclosed formulas/weights), and robustness checks reflect the realities of executive decision-making under constraints (Appendix B).

4.3.2 Alignment to the Research Questions

This strand treats EDMC as a time-aware decision architecture that links Principles → Actions → Skills → Outcomes (PASO) and then tests that linkage end-to-end on real, documented scenarios. In terms of RQ4, the analyses inspect PASO linkage explicitly—whether principles, actions, skills/ownership, and outcomes/KPIs form a traceable chain—and observe how that linkage interacts with the other five rubric dimensions across cases. For RQ5, the strand evidences the working tools themselves: the one-page EDMC case note structure, the six-metric 0–3 rubric with anchors and computation rules, and the decision-record elements that make “proceed / proceed-with-conditions / pause” auditable (owners, thresholds, KPIs) (Appendices A–B).

For RQ6, three practical lenses guide interpretation through the rubric: clarity (PASO linkage plus outcomes evaluability—are outcomes and KPIs explicit and checkable?), confidence under constraints (blindness-mitigation prompts and time-pressure suitability—can EDMC be used decisively when time is tight?), and application in practice (implementation guidance—concrete steps, owners, and artifacts).

Operationally, each case is scored on the six metrics (0–3) with a ≤15-word justification per metric (Appendix A), then summarized via a case composite (equal-weighted by default). Robustness is checked by recomputing composites under the disclosed presets—Equal, Stakeholder-emphasis, and Time-pressure-emphasis—with renormalization if any metric is N/A (Appendix B). All anchors, formulas, and weight

presets are disclosed, making the results replicable while keeping the routine fit-for-use in real decision contexts (Lynham, 2002; Morgan, 2014; Yin, 2018).

4.3.3 Case Set, Scoring, and Weight Presets

We applied EDMC to eight publicly documented, cross-sector cases chosen for diversity of dilemma and domain (following *Table 1*): C1 Shein hazardous chemicals (EU REACH), C2 Shell Jackdaw protests, C3 Foxconn labor violations, C4 Starbucks union violations, C5 Google gender-discrimination settlement, C6 tech-sector age-bias trend, C7 Amazon AI hiring bias, and C8 OpenAI leadership/governance gap. Each case packet includes a PASO plan, stakeholder lens, and candidate KPIs (Appendix A).

Each case was scored on the six EDMC rubric metrics (Appendix B): PASO linkage, blindness mitigation, cultural adaptability, time-pressure suitability, implementation guidance, and outcomes evaluability. Scores use shared 0–3 anchors with a ≤ 15 -word rationale per metric (logged in Appendix A). For readability, we summarize each case with a 0–100 composite by rescaling the 0–3 profile; by default, we report Equal weights. When priorities plausibly differ (e.g., stakeholder emphasis vs. time-pressure), we recompute the composite under the pre-specified Stakeholder and Time-pressure presets to check robustness. Full computation rules—including handling of N/A metrics and weight renormalization—are consolidated in Appendix F, and preset definitions in Appendix B.3 (Aguinis and Solarino, 2019).

Roadmap. Section 4.3.4 reports Equal-weight results; Section 4.3.5 shows sensitivity under the Stakeholder and Time-pressure presets. Auditability is maintained end-to-end via shared anchors (Appendix B), per-metric rationales (Appendix A), and transparent formulas (Appendix F).

4.3.4 Results Under Equal Weights

Interpreting RQ6 (primary) and supporting RQ4–RQ5, this subsection reports the EDMC results when all six rubric metrics are weighted equally. Equal weighting is the neutral baseline for “fitness-for-use”: it rewards cases that make the PASO chain explicit and executable under constraints. *Table 27* (case profiles) and *Figure 14* (composites) summarize the outcomes; computation conventions are in Appendix F, and preset definitions in Appendix B.

Under Equal weights, composites span 75.0–91.7 with a median of ≈ 81.9 . Leaders—C1 Shein (91.7), C3 Foxconn (88.9), and C7 Amazon AI (86.1), pair strong PASO linkage with concrete implementation and checkable outcomes (e.g., testing regimes, named roles, KPIs/thresholds). Lower scores, C6 age-bias (75.0) and C8 AI governance (77.8), flag headroom in cultural adaptability (how the plan travels across contexts) and time-pressure suitability (prompts, run-books, escalation gates). In terms of RQ6, the equal-weight baseline indicates that EDMC yields a generally strong, traceable plan across diverse contexts, with improvements concentrated in “use under pressure” and “fit across settings”, precisely the packaging refinements surfaced by the expert panel (Section 4.2.5).

Table 27 EDMC Scores by Case (Means 0–3; Composites 0–100)

Source: Author (2025). Original table created from study data and materials.

CaseID	Short title	PASO	Blind	Cultural	Time	Implement	Outcomes	Equal	Stakeholder	Time-pressure
C1	Shein hazardous chemicals (EU REACH)	3.0	2.5	2.5	2.5	3.0	3.0	91.67	93.33	89.17
C2	Shell Jackdaw gas-field protests	2.5	2.5	2.0	2.5	2.5	2.5	80.56	80.00	81.67
C3	Foxconn labor violations	3.0	2.5	2.5	2.5	3.0	2.5	88.89	89.17	87.50
C4	Starbucks union violations	2.5	2.5	2.5	2.5	2.5	2.5	83.33	83.33	83.33
C5	Google gender discrimination settlement	2.5	2.0	2.5	2.5	2.5	2.5	80.56	81.67	79.17
C6	Tech sector age-bias trend	2.5	2.0	2.5	2.5	2.0	2.0	75.00	75.00	75.00
C7	Amazon AI hiring tool bias	3.0	2.5	2.0	2.5	3.0	2.5	86.11	85.83	85.83

C8	OpenAI leadership and governance gap	2.5	2.5	2.0	2.0	2.5	2.5	77.78	78.33	76.67
----	--------------------------------------	-----	-----	-----	-----	-----	-----	-------	-------	-------

Notes. Equal-weight composite is the focal result in this subsection; the Stakeholder and Time-pressure columns are reweighted results discussed in Section 4.3.5. Metrics are scored on 0–3 and composites are rescaled to 0–100. Computation rules (including N/A handling and renormalization) are in Appendix F; preset weights are in Appendix B.3.

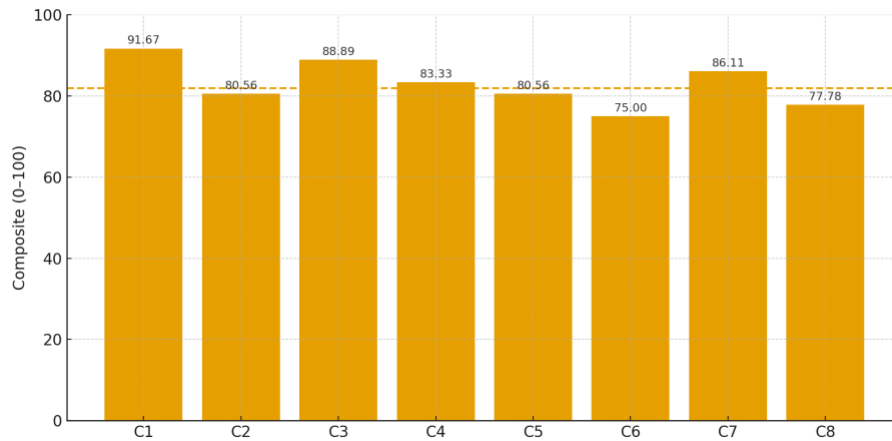


Figure 14 EDMC Composite Scores (0–100), Equal Weights ($n = 8$)
Source: Author (2025). Original figure created from study data and materials.

Figure 14. Equal-weight composites (0–100) across C1–C8; dashed line indicates the median (≈ 81.9). Reweighted results appear in Section 4.3.5.

C1 Shein (91.7): strong REACH-anchored stance, supplier testing cadence, remediation triggers → high PASO/Implementation/Outcomes. C3 Foxconn (88.9): specific actions (audits; worker-voice), named oversight skills, monitoring KPIs → coherent, auditable plan. C7 Amazon AI (86.1): bias-audit + human-in-the-loop; cultural portability is harder in tech but overall net-strong. C6 Age-bias (75.0): diffuse ownership and fewer concrete KPIs depress Implementation/Outcomes; time-pressure prompts less explicit. C8 AI governance (77.8): governance under time stress is thinner, reducing Time-pressure/Blindness despite some remedial actions.

Contextual read-through to RQs. For RQ4, these patterns show EDMC’s PASO logic operationalized consistently across heterogeneous dilemmas. For RQ5, the one-page case note + rubric + rationale trail produce auditable, comparable scores. For RQ6, equal-weight performance is strong overall, with actionable gaps aligned to the panel’s qualitative asks (summary strip, quick checklist, RACI cues, KPI exemplars, light jurisdiction hooks), which are addressed in the proposed v1.1 refinements.

4.3.5 Sensitivity to Decision Priorities (Weight Presets)

RQ focus (Q6): robustness when reasonable priorities shift (stakeholder impact vs. speed under pressure). Finding in one line: percentage-point movements were small (max $\approx \pm 2.50$ pp) and ranks stayed stable; only a C2/C5 tie was broken under the Stakeholder preset.

To test whether EDMC’s case composites are sensitive to plausible decision priorities, we re-weighted the six 0–3 metrics using two presets: a Stakeholder emphasis (higher weight on Cultural adaptability and Outcomes) and a Time-pressure emphasis (higher weight on Time, Blindness, and Implementation). Equal weights (reported in Section 4.3.4) serve as the baseline. Full weight vectors appear in Appendix B.3; composite conventions and N/A renormalization are in Appendix F.

How to read Figures 14–16. Each bar is a 0–100 composite derived from the six rubric metrics (PASO linkage, Blindness mitigation, Cultural adaptability, Time-pressure suitability, Implementation guidance, Outcomes evaluability). *Figure 14* (prior section) shows Equal weights; *Figure 15* re-weights toward Stakeholder/Outcomes; *Figure 16* re-weights toward Time/Blindness/Implementation. The dashed line marks the median. If case ordering is unchanged (or nearly so) across the three figures, that indicates

robustness to reasonable shifts in priorities. Case codes (C1–C8) correspond to the packets in Appendix A.

What changed (and what didn't). Reweighting produced small, interpretable nudges while leaving the overall pattern intact. Under the Stakeholder preset, cases with stronger Cultural hooks and Outcomes instrumentation rose slightly (e.g., C1 to 93.33, +1.66 pp); the only rank change was breaking the Equal-weights tie between C2 and C5 (placing C5 marginally higher). Under the Time-pressure preset, cases with explicit quick-use prompts and ownership cues edged up (e.g., C2 +1.11 pp, C7 unchanged to 85.83), while plans with leaner quick-reference elements slipped modestly (e.g., C1 –2.50 pp, C5 –1.39 pp). Two cases (C4, C6) were unchanged (0.00 pp). *Table 28* reports the exact percentage-point deltas relative to Equal weights.

Interpretation (RQ6). Because shifts were small and ranks stable, the composites appear robust to reasonable priority changes, supporting EDMC's fitness-for-use across stakeholder-oriented and time-critical contexts. Practically, this means organizations can dial weights to fit context without overturning which cases (or case features) look strongest—useful for governance discussions where priorities legitimately differ.




Figure 15 EDMC Composite Scores (0–100), Stakeholder/Outcomes Emphasis (n = 8)
Source: Author (2025). Original figure created from study data and materials.

Figure 15. Reweighted composites (0–100) under the Stakeholder emphasis; bars are weighted toward Cultural and Outcomes; dashed line = median (see *Table 28* for percentage-point deltas versus Equal).




Figure 16 EDMC Composite Scores (0–100), Time-Pressure Blindness/Implementation Emphasis ($n = 8$)

Source: Author (2025). Original figure created from study data and materials.

Figure 16. Reweighted composites (0–100) under the Time-pressure emphasis; bars are weighted toward Time, Blindness, and Implementation; dashed line = median. See *Table 28* for percentage-point deltas versus Equal.

Table 28 Sensitivity Summary (Δ Versus Equal, Percentage Points)

Source: Author (2025). Original table created from study data and materials.

Case	Δ Stakeholder	Δ Time-pressure
C1	+1.66	-2.50
C2	-0.56	+1.11
C3	+0.28	-1.39
C4	0.00	0.00
C5	+1.11	-1.39
C6	0.00	0.00
C7	-0.28	-0.28
C8	+0.55	-1.11

Notes. Δ values show Stakeholder–Equal and Time-pressure–Equal (in percentage points); positive = higher than Equal. Maximum absolute shift = 2.50 pp; rank order remained stable, with the Stakeholder preset only breaking the Equal-weights tie between

C2 and C5 (placing C5 marginally higher). Preset weights are in Appendix B.3; composite formulas and N/A renormalization are in Appendix F.

In sum, reweighting toward stakeholder impact and toward time-pressure produced only minor percentage-point movements (max $\approx \pm 2.50$ pp) and left the case ordering essentially intact—only the Stakeholder preset broke the Equal-weights tie between C2 and C5. For RQ6, this indicates that EDMC’s composites are robust to reasonable shifts in decision priorities: practitioners can tune weights to context without changing substantive conclusions. We therefore present Equal-weight results as the baseline and use the presets as a robustness check; the next section reports the pre-specified sensitivity and reproducibility checks across both evidence streams.

4.3.6 Cross-Case Synthesis (What the Metrics Say)

Framed against RQ6 (practical effectiveness) and supporting RQ5 (improvement targets), the eight researcher-led applications reveal a consistent pattern. Cases with the highest composites—C1 Shein ≈ 91.7 , C3 Foxconn ≈ 88.9 , and C7 Amazon ≈ 86.1 —made the PASO chain explicit and executable: principles were traced to concrete actions, owners/skills were named, and outcomes/KPIs were measurable. That coherence shows up as stronger Implementation and Outcomes scores, confirming that EDMC performs best when values are coupled to a workable plan.

Where headroom remained, it clustered in Cultural adaptability and Time-pressure/Blindness-mitigation. Variance on Cultural was widest in contested or multi-jurisdiction settings (e.g., C2 Shell, C8 OpenAI), and lower Time/Blindness scores appeared most in discrimination/governance contexts (C5, C6, C8). In practice, those dips co-occurred with fewer explicit prompts, run-book steps, and bias checks, signalling packaging rather than conceptual gaps.

These quantitative patterns align closely with the panel’s qualitative suggestions (Section 4.2.5): add a one-look decision summary strip (“pause” vs. “proceed with conditions”), a quick checklist, RACI ownership cues, a KPI/mitigation library, and light jurisdiction hooks. Importantly, the reweighting analysis (Section 4.3.5) showed stable ranks under both Stakeholder and Time-pressure presets, so these refinements are likely to raise absolute scores on Cultural/Time without changing the substantive ordering—useful for RQ6 robustness and RQ5 design guidance.

- Mini case reads (1–2 lines).

C1 Shein — Strong hypernorm stance (REACH), testing cadence, remediation triggers → high PASO/Implementation/Outcomes.

C2 Shell Jackdaw — Balanced overall; Cultural adaptability constrained amid contested mandates.

C3 Foxconn — Specific actions and oversight; worker-voice mechanisms; KPIs enable tracking.

C4 Starbucks — Even profile; protocols present; Cultural/Time elements adequate.

C5 Google gender pay — Clear actions/outcomes; slight dip on Blindness-mitigation relative to legacy practices.

C6 Age-bias trend — Thinner Implementation/Outcomes instrumentation; Time-pressure prompts less explicit.

C7 Amazon AI hiring — Strong mitigation (audits, HITL); Cultural adaptability more mixed in tech context.

C8 OpenAI governance — Governance/Time-pressure elements weaker; actions/outcomes present but less mature.

Takeaway for RQ5–RQ6. EDMC’s strengths are PASO linkage, implementation guidance, and outcomes evaluability; the main upgrade path is tighter time-aware

prompts and clearer portability hooks. Those are delivery-layer changes, summary strip, checklist, RACI, KPI exemplars, jurisdiction notes—rather than theoretical revisions, and they directly mirror expert feedback. The next section turns to the pre-specified sensitivity and reproducibility checks across both evidence streams.

4.3.7 Reproducibility Note and Provenance (RQ Alignment to Q4–Q6)

RQ focus (Q5–Q6). This strand emphasizes transparency and credibility: the scoring logic, inputs, and outputs are fully disclosed and version-locked so results can be reproduced.

Provenance and version lock. All outputs in Section 4.3 were generated from the EDMC scoring workbook (version 2025-09-09). Version control, administrative handling, and timestamps are documented in Appendices M–N (Yin, 2018).

Computation pipeline. Each case was scored on six 0–3 metrics—PASO linkage, blindness mitigation, cultural adaptability, time-pressure suitability, implementation guidance, and outcomes evaluability—using the anchors in Appendix B. Case-level summaries were then rescaled to a 0–100 composite via:

$$\text{Composite (0–100)} = (\text{Score}_1 + \text{Score}_2 + \text{Score}_3 + \text{Score}_4 + \text{Score}_5 + \text{Score}_6) \div 18 \times 100.$$

(Each metric is 0–3, so the maximum summed score is 18. Where a metric was truly N/A, it was excluded from the mean per Appendix B/F.) Weights and anchors were pre-specified to support transparent, practice-oriented validation (Lynham, 2002; Morgan, 2014).

Weight presets (pre-specified). To reflect realistic priority shifts, we reported composites under three disclosed presets (Appendix B), with renormalization if any metric was N/A (Aguinis and Solarino, 2019): Equal (1/6 each); Stakeholder emphasis (0.20, 0.10, 0.20, 0.10, 0.15, 0.25 for PASO, Blindness, Cultural, Time, Implementation, Outcomes); and Time-pressure emphasis (0.10, 0.25, 0.10, 0.30, 0.15, 0.10).

Audit trail and generation of figures. Metric anchors, weights, and formulas appear in Appendix B; computation rules and N/A handling are summarized in Appendix F. For each case, Appendix A provides the PASO plan, KPI sketch, and a ≤ 15 -word rationale per metric. Figures 14–16 were produced directly from the same version-locked workbook (no manual edits), enabling exact reproduction (Yin, 2018). In all figures, bars display 0–100 composites and the dashed line marks the median; Figure 14 uses Equal weights, Figure 15 re-weights toward Stakeholder/Outcomes, and Figure 16 re-weights toward Time-pressure/Blindness/Implementation (see Section 4.3.5 for a brief reading guide).

Consistency with the survey strand. Small preset shifts (maximum ≤ 2.50 percentage points) and stable rank order in this strand are consistent with the expert-panel signals: good PASO coverage (S-CVI/Ave 0.82–0.87), selective stance tightening (12.2%), modest confidence uplift (+0.31), and strong global means (≈ 3.96 –4.23). Together, the disclosed rubric and reproducible outputs provide an auditable chain aligned with case-based validation and pragmatic theory building (Lynham, 2002; Morgan, 2014; Yin, 2018).

Explicit RQ alignment. For Q4, PASO is operationalized as six rubric metrics with disclosed anchors and scoring logic. For Q5, those constructs are translated into concrete tools: case packets (PASO plans, KPIs, decision-record elements) and the scoring workbook (Appendices A–B). For Q6, clarity, constraints, and application are evidenced by the six metrics/composites and robustness under re-weighting (Aguinis and Solarino, 2019).

Appendix pointers (for replication). A—Case Packets (one-line scoring rationales); B—Scoring Rubric, Weights, and Computation; F—Analysis Plan and

Derived Metrics; M–N—Exclusion Rules, Sensitivity Checks, Version Lock and Fieldwork Window.

4.3.8 What This Adds Beyond the Survey

RQ focus (Q6). This subsection integrates the two evidence streams to assess EDMC’s practice-use under constraints.

The expert panel established that EDMC is usable under time pressure, selectively tightens safeguards (12.2% of paired decisions moved toward more caution), yields a modest confidence gain on average (+0.31), and covers PASO well (S-CVI/Ave 0.82–0.87). The researcher-led applications extend those signals by showing that a single analyst can apply the rubric end-to-end across diverse public-facts cases and obtain consistent, auditable results. Under Equal weights, composites ranged from 75.0 to 91.7 (median \approx 81.9); reweighting toward Stakeholder or Time-pressure priorities shifted results by at most \leq 2.50 percentage points, with the rank order essentially unchanged (the Stakeholder preset only broke the C2/C5 tie). Together, these patterns indicate that EDMC “travels” across contexts and remains decision-useful under realistic priority shifts—consistent with a pragmatic, problem-first stance (Aguinis and Solarino, 2019; Lynham, 2002; Morgan, 2014; Yin, 2018).

Importantly, the same analyses pinpoint delivery-layer refinements that speed use without altering the PASO-linked core: a one-look summary strip, a quick checklist, RACI ownership prompts, a small KPI/mitigation library, and light jurisdiction hooks. These upgrades address the lower-scoring dimensions (time-pressure prompts and cultural adaptability) and align with expert comments, making them strong candidates for v1.1 packaging.

Where to look. Case scores and preset shifts: Tables 27–28; Figures 14–16.
Computation rules and presets: Appendices B and F. Audit trail and versioning:
Appendices A and N.

4.4 Sensitivity and Robustness

RQ focus (Q6; supports Q4–Q5). We tested whether headline findings depend on administrative handling or reasonable modeling choices.

Pre-specification and purpose. Following a pragmatic validation logic (Aguinis and Solarino, 2019; Morgan, 2014; Yin, 2018), we pre-registered a small set of checks (Appendix M: SENS-1; Appendix N: version lock) to confirm that results are not artifacts of exclusions, denominators, or weighting choices.

Expert-panel strand (survey). We re-ran the summaries after applying the rule-based exclusions (mis-coded/duplicate IDs; late rows), keeping only version-locked responses. Agreement rates (≥ 4), PASO CVI (I-CVI and S-CVI/Ave), confidence uplift, and stance-change matrices were materially unchanged. Shifts, where present, were below the thresholds specified in the analysis plan (Appendix F/F.8), and no qualitative themes were overturned.

Researcher-led strand (case applications). Composites proved stable under the two disclosed presets. The maximum percentage-point movement versus Equal weights was ≤ 2.50 , and the rank order remained stable (the only change was breaking the Equal-weights tie between C2 and C5 under the Stakeholder preset). A simple leave-one-metric-out scan produced the same qualitative ordering, indicating no single metric drives the composites.

Reproducibility safeguards. All figures in Sections 4.3.4–4.3.6 were generated directly from the version-locked workbook; anchors, formulas, and N/A handling are

disclosed (Appendices B, F, N). These safeguards support exact reproduction and transparent audit.

Conclusion of checks. Across both strands, administrative exclusions and denominator rules did not alter conclusions, and preset reweighting produced only small, interpretable nudges. The findings are therefore robust to the reasonable variations anticipated in practice.

Where to look. Preset charts: Figures 14–16. Deltas: *Table 28*. Exclusion protocol and versioning: Appendices M–N. Computation rules and thresholds: Appendix F.

4.4.1 Expert-Panel Data Integrity (SENS-1)

RQ focus (Q6; supports Q4–Q5). This check asked whether headline patterns in the expert-panel strand depend on administrative handling or denominator choices. As pre-specified, we re-ran all indicators after (a) excluding any rows submitted before the Case-3 link-fix timestamp and any rows with invalid participant codes, (b) confirming that ≥ 4 excluded N/A responses and that each cell displayed the effective N, and (c) scanning for extreme or partial patterns that could distort percentages or means (Appendices M–N; computation rules in Appendix F).

What changed. Interpretations did not change. Usability remained high (most item \times case cells $\geq 80\%$); PASO content validity stayed in the “good” range (S-CVI/Ave 0.82–0.87); stance-tightening remained 12.2% of paired decisions; and the small, positive confidence-uplift pattern held. Item-level Ns varied slightly—as expected when excluding N/A—but the cross-item, cross-case pattern was stable (cf. Tables 21–25). In short, administrative tidy-ups and sensible denominator choices did not overturn the conclusion that a one-page EDMC note is usable under time pressure, coherently covers PASO, and nudges decisions toward safeguards without inflating confidence.

4.4.2 Alternative Denominator / N-Handling

RQ focus (Q6). Because N/A rates can vary by item, we confirmed the primary rule—compute $\% \geq 4$ with N/A excluded and display the effective N per cell—then inspected whether re-weighting items with more N/A or forcing equal Ns would alter conclusions. It did not. The “most items $\geq 80\%$ ” pattern persisted, and relative differences between items and cases were unchanged. We therefore retained the transparent approach (exclude N/A from denominators; show the effective N alongside each percentage), exactly as specified in Appendix F.

4.4.3 Researcher Case-Application Robustness Beyond Weights

RQ focus (Q6; supports Q4–Q5). For the case-application strand, we tested whether composites are sensitive to reasonable modeling choices beyond the Equal-weights baseline.

Preset reweighting. Reweighting toward Stakeholder/Outcomes and toward Time-pressure/Blindness/Implementation produced only small percentage-point shifts (maximum ≤ 2.50) and left the rank order essentially unchanged; the only difference was that the Stakeholder preset broke the Equal-weights tie between C2 and C5 (*Table 28*; *Figure 15-16*), consistent with the Equal-weight pattern (*Table 27*; *Figure 14*).

Leave-one-metric-out scan. Dropping each of the six 0–3 metrics in turn and recomputing the composite (no separate table shown) left the qualitative ordering stable or minimally affected, indicating that no single metric dominates the results. Together, these checks support the claim that the composite reflects balanced coverage rather than a fragile mix (rubric and formulas in Appendix B; computation notes in Appendix F).

4.4.4 Threats to Validity and Mitigations

RQ focus (Q6; supports Q4–Q5). We acknowledge standard threats and document proportionate mitigations. The expert sample is small and heterogeneous by design, so we emphasize descriptive indicators and triangulate with independent, researcher-led case applications rather than claim population inference. Because self-report can invite social desirability, we collected data anonymously, included an objective tightening indicator (stance moves toward “conditions” or “pause”), and cross-checked perceptions against rubric scores. Single-analyst scoring for cases can introduce subjectivity; we mitigated this with shared anchors, ≤ 15 -word rationales for every score, and disclosed formulas that enable reproduction or critique (Appendices A–B). To avoid construct conflation, global items (portability, integratability, time-pressure, coherence, tech/AI) are reported separately from Q1–Q8 and PASO CVI. Finally, we controlled versioning/admin drift with a version-locked instrument and dated change log (Appendix N). These safeguards, taken together, align with the thesis’s pragmatic aim: credible, transparent evidence of fitness-for-use under real-world constraints.

4.4.5 Bottom Line

RQ6 verdict. Across all pre-specified checks, the substantive conclusions held. On the expert-panel side, high usability ($\% \geq 4$ commonly $\geq 80\%$), good PASO coverage (S-CVI/Ave 0.82–0.87), cautious tightening of stance (12.2%), and a modest confidence uplift (+0.31) were unchanged by administrative exclusions or alternative denominator choices (Tables 21–25). On the researcher-led side, equal-weight composites (75.0–91.7; median ≈ 81.9) remained stable under the Stakeholder and Time-pressure presets (maximum shift ≤ 2.50 percentage points; rank order essentially unchanged) and under a leave-one-metric-out scan (Tables 27–28; Figures 14–16). In short, the signals supporting

EDMC's fitness-for-use under constraints are robust rather than artifacts of data handling or modeling (Aguinis and Solarino, 2019; Lynham, 2002; Morgan, 2014; Yin, 2018).

4.5 Summary of Key Findings

This section integrates the expert-panel evidence with the researcher-led applications to answer RQ6 and to indicate where RQ4–RQ5 point to further refinement.

What the expert panel showed (RQ6). Reviewing three one-page EDMC notes, experts consistently endorsed practical usability and coverage. Agreement on Q1–Q8 items was typically high (22/24 item×case cells $\geq 80\%$; range 63.3–93.3), and the global wrap-up means clustered between ~ 3.9 and 4.3—Conceptual coherence 4.23, Time-pressure usability 4.20, Tech/AI relevance 4.19, Portability 4.09, and Integratability 3.96 (Table 25; Figures H1–H7). PASO content validity for each vignette fell in the “good” range (S-CVI/Ave = 0.82–0.87; Table 22). Decision effects were selective and directionally conservative: 12.2% of paired stances tightened (toward “conditions” or “pause”), while mean confidence rose modestly (+0.31 on a 0–10 scale; Tables 23–24). Taken together, the panel read EDMC as coherent, usable under time pressure, and defensible to stakeholders—precisely the constructs RQ6 targets.

What the case applications added (RQ6; supports RQ4–RQ5). Applying the rubric end-to-end to eight public-facts cases produced transparent, auditable profiles with equal-weight composites spanning 75.0–91.7 (median ≈ 81.9 ; Table 27; Figure 14). Reweighting toward Stakeholder or Time-pressure priorities shifted values by at most ≤ 2.50 percentage points and left the rank order essentially intact (only the C2/C5 tie broke under the Stakeholder preset), indicating robust scores rather than fragile mixtures (Table 28; Figures 15–16). High performers shared a clear PASO chain and executable plans (principles \rightarrow actions \rightarrow named owners/skills \rightarrow measurable outcomes), while lower

scores reflected thinner cultural adaptability and fewer explicit time-aware prompts. This pattern aligns with EDMC’s intent (RQ4) and with how the tools were translated into practice (RQ5).

Where to improve (RQ5). The qualitative comments point to delivery-layer refinements that speed use without changing the core: a one-look summary strip (“proceed / proceed with conditions / pause”), a quick checklist, crisp RACI ownership cues, a small KPI/mitigation library, and light jurisdiction hooks. These map directly onto the dimensions with headroom (Cultural adaptability; Time-pressure suitability) and are tractable for a v1.1 update (Section 4.2.5).

So what for practice. For leaders working under compressed timelines, EDMC offers a compact, auditable way to surface ethical salience, make conditional decisions, and defend them. The evidence shows EDMC can tighten safeguards without inflating confidence, travels across contexts, and remains stable under reasonable shifts in decision priorities—features consistent with pragmatic, problem-first designs (Aguinis and Solarino, 2019; Lynham, 2002; Morgan, 2014; Yin, 2018).

Where this leaves us. Chapter 4 has shown that EDMC is usable under time pressure, principled and auditable via PASO, and robust across weighting schemes. The few, specific packaging upgrades suggested by experts are additive rather than foundational. Chapter 5 turns to interpretation and implications—what these results mean for executive decision practice, how v1.1 should be packaged, and where future field pilots and longitudinal studies can extend the evidence.

4.6 Conclusions

Two convergent evidence streams validate the Ethical Decision-Making Compass (EDMC) as a time-aware, PASO-linked decision architecture that is usable, auditable,

and robust under realistic constraints. In the expert-panel stream, three one-page notes were rated readable, portable, and defensible under pressure (global means $\approx 3.9\text{--}4.3$), PASO content coverage met “good” thresholds (S-CVI/Ave 0.82–0.87), and 12.2% of paired stances tightened while confidence rose modestly (+0.31/10). In the researcher-led stream, eight public-facts cases produced consistent, transparent scoring: equal-weight composites ranged 75.0–91.7 (median ≈ 81.9), and reweighting toward Stakeholder or Time-pressure priorities moved results by at most 2.50 percentage points with ranks essentially unchanged (Tables 27–28; Figures 14–16). Taken together, these patterns support EDMC’s fitness-for-use in pragmatic, problem-first settings (Aguinis and Solarino, 2019; Lynham, 2002; Morgan, 2014; Yin, 2018).

Read against the research questions: EDMC organizes PASO into a coherent architecture in use (RQ4), translates that architecture into practical, transparent tools (six 0–3 metrics with shared anchors and a reproducible composite: $\text{Composite (0–100)} = (\text{Score}_1 + \dots + \text{Score}_6) \div 6 \times 100$) (RQ5), and remains decision-useful under constraints—high usability, selective tightening without over-confidence, and score stability under reasonable priority shifts (RQ6).

Beyond the state of practice, the contribution is threefold: a compact, one-page artifact that keeps ethical salience visible when attention is scarce while preserving traceability via PASO; a validation approach that pairs content-validity indices with a practice-oriented 0–3 rubric to yield audit-friendly evidence at the right maturity; and simple weight-preset checks that probe robustness under plausible executive priorities. The workflow is reproducible end-to-end: anchors, formulas, and versions are disclosed, and tables/figures are generated directly from a formula-only workbook.

Limits and boundary conditions remain. The panel was purposive and small, case scoring used a single analyst, and public-facts materials cannot capture organizational

politics or data quality in full. Pre-specified checks, administrative exclusions, denominator rules, preset reweighting, and leave-one-metric-out—did not alter headline interpretations, suggesting findings are not artifacts of handling choices. EDMC does not replace legal/technical reviews; effects will vary with local capability, governance maturity, and data quality.

What to improve is the delivery layer, not the core logic. The evidence points to v1.1 packaging that speeds first-mile use: a one-look summary strip (“pause / proceed with safeguards”), a quick-reference cue card, crisp RACI prompts, a small KPI/mitigation library, and jurisdiction hooks for cross-region fit—precisely where headroom appeared (cultural adaptability; time-pressure suitability).

Implications for adoption are immediate: EDMC can be piloted inside existing governance steps, supports explicit “conditions to proceed” thresholds, and leaves an audit trail that withstands stakeholder scrutiny, including in AI/data contexts. Light digital embedding (forms/workflows) can trigger under-pressure prompts and auto-capture rationales for later review.

Chapter 5 develops these implications—v1.1 packaging, pilot guidance, longitudinal follow-ups, and digital choices that preserve PASO traceability without burden—while outlining next studies (multi-rater agreement, cross-cultural deployments, and outcome tracking).

CHAPTER V: DISCUSSION

5.1 Discussion — Purpose, Why Now, and Vision

This thesis addresses a persistent gap in ethical leadership: the distance between rich ethical theory and real decisions made under pressure. Many frameworks are siloed, abstract, or too cumbersome for short decision windows—conditions that foster ethical blindness even among well-intentioned leaders (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012). To close that gap, we designed and validated the Ethical Decision-Making Compass (EDMC), a practical, cross-context framework organized around PASO—Principles → Actions → Skills → Outcomes. EDMC exemplifies “ethical decision-making by design”: it makes salience explicit, names roles and escalation, and sets shared standards for progress so judgments are both principled and actionable (Treviño and Nelson, 2016).

The objective was to design and validate a first-tier ethical filter that leaders can deploy—actionable (clear steps and ownership), time-aware (usable in urgent and deliberative settings), bridging (linking rigor to application), culturally adaptable (across sectors and jurisdictions), and scalable (a base for capability building). Methodologically, the work followed an emergent, practice-oriented path—surfacing patterns from scholarship and practitioner insight, then iterating toward a design that fits real decision constraints (Edmondson and McManus, 2007; Gioia et al., 2013; Morgan, 2014; Van de Ven, 2007).

Why now. Decisions are faster, impacts are larger, and accountability is stricter. Short decision windows and data-intensive systems mean a single choice can trigger legal exposure, product reversals, reputational damage, talent loss, and missed value for communities. In global firms, those effects cross borders and jurisdictions. Under pressure, even well-intentioned teams miss what matters, ethical salience drops out of

view due to cognitive and organizational blind spots (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012). AI raises the stakes: models move at machine speed, scale errors instantly, and create opacity about who is responsible. Stakeholders and regulators now expect decisions that can be traced and defended.

What EDMC adds at this moment. EDMC makes the essentials explicit in the flow of work: which principles apply, what actions to take, who has the skills/ownership, and which outcomes/thresholds define “good enough to proceed.” Think of PASO as the decision DNA of the organization: four linked strands—Principles → Actions → Skills → Outcomes—that stay traceable from intent to impact and can be reused from a one-page note to a playbook. The “base pairs” are stable hypernorms; the “adapters” are local steps, roles, and KPIs that vary by context. EDMC encodes this PASO DNA into a compact artifact with under-pressure prompts, clear ownership and escalation, and go/stop conditions so teams can move quickly without dropping safeguards. In AI-heavy settings, the same scaffold becomes concrete guardrails—human-in-the-loop roles, bias audits, rollback procedures, and KPI thresholds grounded in principles—so leaders define the conditions under which a system may run (Bazerman and Tenbrunsel, 2011; Floridi, 2020; Palazzo et al., 2012; Zhang and Wade, 2022).

Vision. With PASO, EDMC offers a comprehensive view of ethical leadership—Why (Principles), How (Actions), Who (Skills), and What (Outcomes). It couples normative foundations with disciplined implementation and learning loops (KPIs and post-decision reviews), enabling organizations to build capability over time, not merely justify a single choice (Lynham, 2002; Treviño and Nelson, 2016). The sections that follow interpret the empirical results through this lens and develop implications for leadership and governance.

5.2 Summary of the Study and Research Journey (From Problem to Evidence)

Organizations lack a framework that (i) integrates principles with situational constraints, skills/ownership, and outcome assessment, (ii) works under time pressure and across cultures, and (iii) leaves an auditable reasoning path that reduces blind spots and supports accountable leadership. This study therefore built and evaluated an ethical decision system that is principled and auditable, usable under time and cognitive pressure, portable across sectors and cultures, and oriented to learning through outcomes and feedback.

How the journey unfolded. The work proceeded in five steps across three phases. First, a comparative review of sixteen frameworks surfaced strengths, limits, and recurrent patterns (Hsieh and Shannon, 2005). Second, those regularities were synthesized into PASO via integrative/critical interpretive synthesis (Dixon-Woods et al., 2006). Third, PASO was architected into the EDMC design. Fourth, we operationalized EDMC into usable instruments (one-page case notes, a six-metric rubric, and scoring logic) to test “ethical decision-making by design” under constraint. Fifth, we validated EDMC through two complementary evidence streams with pre-specified sensitivity checks (Lynham, 2002; Yin, 2018): the expert-panel stream and the researcher-led stream.

From review to PASO, the synthesis identified five systemic gaps—fragmentation, limited adaptability under real constraints, scarce actionable tools, neglected skills/ownership, and insufficient countermeasures to ethical blindness (Bazerman and Tenbrunsel, 2011; Craft, 2013; Kaptein, 2023; Mittelstadt et al., 2016; Palazzo et al., 2012). Rather than discard prior work, EDMC integrates its strengths into PASO: Principles (normative grounding), Actions (safeguards and escalation steps), Skills (roles, competencies, ownership), and Outcomes (KPIs and learning loops). The

design keeps values, execution, capability, and impact in view—under time pressure and across contexts.

To evaluate EDMC, we used two complementary evidence streams: an expert-panel survey (N = 30; three EDMC vignettes) and researcher-led applications to eight public-facts cases (C1–C8). For brevity, we refer to these as the expert-panel stream and the researcher-led stream. The expert-panel stream (N = 30; 90 case-blocks) assessed usability (Q1–Q8), PASO content validity via I-CVI and S-CVI/Ave, stance-tightening and confidence uplift, and global judgments (portability, integratability, time-pressure usability, conceptual coherence, tech/AI relevance) (Grant and Davis, 1997; Lynn, 1986; Polit et al., 2007; Polit and Beck, 2006). The researcher-led stream applied the disclosed 0–3 rubric to eight public-facts cases, computed 0–100 composites under Equal, Stakeholder, and Time-pressure presets, and logged ≤ 15 -word scoring rationales and KPIs to preserve auditability. Together, these streams address RQ6 directly and provide supporting evidence for RQ4–RQ5.

Why these cases. The researcher-led stream used public-facts cases, cross-sector and time-sensitive, including AI-relevant dilemmas, to probe EDMC where ethical salience, stakeholder impact, and speed co-exist. This selection supports external face validity while enabling transparent, reproducible scoring. The expert-panel stream, by contrast, used generic vignettes to minimize prior-knowledge bias and focus experts on the EDMC artifact itself. Together, the streams align with the thesis’s pragmatic intent: RQ4 (organizing PASO into EDMC), RQ5 (translating into usable tools), and RQ6 (validating clarity, time-aware usability, and application under constraints).

5.3 Linking Findings to Objectives and Research Questions (Closing the Loop)

To anchor the Discussion, we interpret the evidence against the study's objectives (Section 1.6) and research questions (Section 1.5). The aim is to move from what happened to so what—what these patterns imply for ethical leadership, how they extend or challenge prior work, and where they set boundaries for practice and future research.

- RQ1 — What gaps persist in existing ethical-decision frameworks, and why do they matter?

Interpretation. The review showed a durable disconnect between ethical aspiration and operational execution: many frameworks offer strong principles with weak delivery, or rich procedures with thin normative grounding. Under time pressure, this design gap enables ethical blindness—not for lack of intent, but because values, capabilities, and outcomes are not bound together at the moment of choice. The core issue is architectural, not merely educational: additional training cannot substitute for a decision system that keeps principles, actions, ownership, and outcomes aligned in real time.

Relation to prior work. Prior syntheses mapped blind spots (Craft, 2013; Kaptein, 2023; Mittelstadt et al., 2016). This study specifies which missing links matter most in use—ownership/skills, time-pressure readiness, and outcome hooks—and treats them as design constraints rather than afterthoughts.

- RQ2 — Which recurring patterns emerge across frameworks, and why do they matter?

Across the sixteen frameworks, three strengths reappeared with striking regularity: a stable ethical core (duties/rights, consequences, character, and stakeholder salience) that anchors fairness, harm minimization, dignity, and accountability; a familiar decision sequence from issue recognition through review; and calls for institutional accountability via codes, oversight, and audits (Treviño and Nelson, 2016).

Alongside these assets, five failure modes also repeated. Frameworks fragmented along single lenses, making judgments hard to carry across contexts; they rarely addressed time-pressure blindness with prompts, red flags, or go/stop thresholds (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012); they offered little decision-grade tooling leaders could actually run; they left skills and ownership thin—no RACI, no capability requirements; and they omitted outcome hooks such as KPIs, conditions to proceed, and feedback loops (Craft, 2013; Kaptein, 2023; Mittelstadt et al., 2016). A consistent co-occurrence sat behind these gaps: where principle statements were strong, ownership and outcomes were weak; where procedural checklists appeared, normative grounding thinned. Mentions of culture were common, but mechanisms for adaptation (hypernorms plus local fit) were scarce. Read together, the pattern is less about intent than architecture: ethics falters when values, action, ownership, and evidence are not coupled at the moment of decision (Hsieh and Shannon, 2005).

- RQ3 — How can these strengths and patterns be synthesized into PASO to form a coherent foundation?

The synthesis treated the recurring strengths as non-negotiables and the failure modes as design requirements, yielding PASO—a lens-agnostic but binding scaffold that turns ethical intent into owned action and measurable effects (Dixon-Woods et al., 2006). Principles remain explicit and portable, stating the governing norms (or hypernorms) and how local adaptation will work. Actions translate those principles into a small set of preventive, detective, and corrective steps with escalation paths and under-pressure prompts, so “what to do next” is operable even in short windows. Skills/Ownership names who does what, with which competencies and tools, closing the common “who/when” gap that derails otherwise principled plans. Outcomes specify what “good enough to proceed” means—KPIs and thresholds with a review cadence—so learning and

accountability are routine rather than aspirational. Crucially, the interfaces bind the parts: every principle maps to at least one action; every action has a named owner/skill; every outcome traces back to the principle it evidences; and the set includes a time-pressure prompt and a culture-fit note. These invariants convert a list of good ideas into a coherent decision architecture that travels across sectors and jurisdictions. This synthesis also generated the six EDMC rubric metrics used in Chapter 4 (PASO linkage, blindness-mitigation, cultural adaptability, time-pressure suitability, implementation guidance, outcomes evaluability). Expert judges recognized the PASO elements in practice (I-CVI high; S-CVI/Ave 0.82–0.87), and the strongest case composites appeared when these interfaces were explicit and executable. In short, PASO is neutral about which principles a context adopts but insistent that whatever is claimed ethically must be enacted (A), owned (S), and evidenced (O)—the combination that made EDMC portable, auditable, and usable under pressure in our results (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012; Treviño and Nelson, 2016). In this sense, PASO functions as the decision DNA: it encodes what matters (Principles), how to act (Actions), who owns it (Skills), and what counts as “good enough to proceed” (Outcomes).

- RQ4 — How should PASO be organized and integrated into EDMC as a coherent decision architecture?

Organized this way, EDMC is the expression of PASO’s decision DNA in use—a single, traceable system rather than a stack of disconnected tools. The expert-panel results rated conceptual coherence highly (mean 4.23/5), and the researcher-led applications scored highest where the PASO chain was explicit and executable: clear principles and guardrails, concrete steps with named owners, and outcomes/thresholds that define “good enough to proceed.” This convergence indicates EDMC operates as one coherent system

in practice, delivering the traceability from principle to KPI that boards and regulators increasingly expect (Yin, 2018).

- RQ5 — How can EDMC be translated into practical tools that leaders can deploy across organizational contexts?

Translation succeeded on both usability and portability. Across three EDMC vignettes, agreement on the eight usability/defensibility items was typically $\geq 80\%$, and global means for Portability, Integrability, Time-pressure usability, Conceptual coherence, and Tech/AI relevance clustered around 3.9–4.3/5 (Grant and Davis, 1997; Lynn, 1986). Qualitative comments did not seek a conceptual overhaul; instead, they asked for packaging refinements that speed use under pressure—an executive summary strip, a quick checklist, clearer RACI cues, KPI/mitigation exemplars, and light jurisdiction hooks. That pattern points to implementation maturity: the architecture is doing its job, and the v1.1 gains lie in how it is delivered, not in what it is. This translation was validated in both streams: high expert agreement and strong global means in the panel (usability, portability, time-pressure, coherence, tech/AI), and stable, auditable scoring across eight public-facts cases with composites robust to stakeholder and time-pressure reweighting.

- RQ6 — To what extent does EDMC enhance ethical clarity, confidence, and application in varied leadership settings?

Under real constraints, EDMC produced a disciplined pattern rather than bravado. After reviewing an EDMC note, 12.2% of paired stances changed and all moves tightened toward more cautious positions (Proceed → Conditions/Pause; Conditions → Pause), while mean confidence rose modestly (+0.31 on a 0–10 scale). Time-pressure usability averaged 4.20/5 and PASO coverage remained “good” (S-CVI/Ave 0.82–0.87). In the researcher-led stream, equal-weight composites spanned 75.0–91.7 (median ≈ 81.9) and

remained stable under Stakeholder and Time-pressure presets (maximum shift ≤ 2.50 percentage points; rank order essentially unchanged), indicating robustness to reasonable priority changes (Aguinis and Solarino, 2019; Morgan, 2014; Yin, 2018). Taken together, EDMC sharpened safeguards without inducing overconfidence and travelled across domains—including AI-relevant settings where scale magnifies consequences.

Bottom line. Across RQ1–RQ6, the evidence converges: ethical decision-making by design is feasible. EDMC’s PASO decision DNA—Principles, Actions, Skills, Outcomes—encodes why we act, how we act, who owns the work, and what counts as “good enough to proceed,” binding these elements at the moment of choice. In practice, that DNA turned ethics from an episodic reflection into an organizational capability: principled (norms made explicit), usable under time pressure (prompts, thresholds, escalation), auditable (traceability from principle to KPI), and portable across sectors and jurisdictions. The expert-panel stream showed high usability and “good” PASO coverage with cautious stance-tightening and modest confidence gains; the researcher-led stream showed stable, transparent scoring across diverse public-facts cases, robust to reasonable reweighting of priorities.

The contribution is not a new ethical doctrine; it is a portable architecture that makes principled decisions implementable, reviewable, and defensible—including in AI-intensive contexts where speed and scale amplify consequences. In short, PASO’s decision DNA gives leaders a shared, time-aware scaffold that reduces blind spots and raises the floor on decision quality without inducing overconfidence.

Table 29 Findings → RQs/Objectives Crosswalk (Discussion Focus)

Source: Author (2025). Original table created from study data and materials.

RQ	Objective focus	Empirical evidence (Ch. 4)	Interpretation (so-what)
Q1 (diagnose gaps)	Identify where existing frameworks	Ch. 2–3 synthesis* of 16 frameworks: fragmentation; thin time-pressure readiness; weak skills/ownership; missing outcome hooks; blindness unaddressed.	The problem is architectural, not motivational: without aligned principles→actions→skills→outcomes, time-pressured decisions invite ethical

	fail in practice		blindness; fixes must be design-led, not just more training.
Q2 (preserve strengths)	Retain proven normative foundations	Ch. 2–3 synthesis* shows enduring value of duty/consequence/virtue/stakeholder lenses; Ch. 4 conceptual coherence mean = 4.23/5 confirms appetite for principled traceability.	Don't reinvent ethical behaviour—standardise translation: keep normative clarity and institutionalise stewardship and review for portability.
Q3 (derive PASO)	Specify an integrative scaffold	PASO emerges in Ch. 2–3*; in Ch. 4 it functions as the coding/measurement spine: PASO I-CVI high; S-CVI/Ave 0.82–0.87 across cases.	PASO acts as decision DNA—neutral on which principles, insistent on how they are enacted and evidenced—enabling cross-sector, cross-culture use.
Q4 (organise PASO into EDMC)	Principled, auditable architecture	Conceptual coherence 4.23/5; top composites where PASO linkage + implementation + outcomes are explicit (C1, C3, C7).	EDMC integrates P–A–S–O into one operable unit; architecture holds in use from principle to KPI.
Q5 (translate into tools)	Usable, time-aware toolkit	Q1–Q8 mostly ≥80% agree (22/24 cells ≥80%); global means: Portability 4.09, Integratability 3.96, Time-pressure 4.20, Conceptual coherence 4.23, Tech/AI controls 4.19; comments request summary strip, checklist, RACI, KPI exemplars, jurisdiction hooks.	Tools are usable under pressure; requested upgrades target speed of use (delivery layer), not the conceptual core.
Q6 (clarity, confidence, application)	Decision-usefulness under constraints	12.2% stance-tightening; mean confidence +0.31/10; S-CVI/Ave 0.82–0.87; Time-pressure 4.20/5.	Anti-blindness pattern: clearer criteria → selectively more cautious decisions, modest confidence lifts—discipline without whiplash.
Q6 (robustness)	Stability across priorities	Composites 75.0–91.7 (median ≈ 81.9); preset shifts ≤ 2.50 pp; ranks essentially stable (Stakeholder preset only breaks the C2/C5 tie).	Conclusions don't flip under stakeholder/time-pressure emphasis → practical robustness.
Primary RQ (practical, cross-context)	Portability and AI relevance	Cross-sector cases (environment, labor, discrimination, AI); global means: Portability 4.09, Integratability 3.96, Time-pressure 4.20, Tech/AI 4.19.	EDMC travels across contexts; AI guardrails (HITL, bias audits, rollback) are naturally embedded.

Table values reference Chapter 4 exhibits (Figure 17-Figure 24; Table 21-Table 28) and align with conventions in Appendix F (denominators, N/A handling).

5.4 What These Results Mean (and Why they Matter)

The combination of a modest confidence shift (+0.31 on a 0–10 scale) and selective stance-tightening (12.2% of paired decisions, all toward more caution) is exactly the signature you want from a first-tier ethical filter. The goal is not wholesale reversals but clearer criteria and firmer safeguards at the margin. Seeing more “proceed with conditions/pause” paired with small confidence gains indicates EDMC raised the bar without whiplash—precisely as designed (Section 4.2.3).

From theory fragments to a working system. Content validity in the expert panel (S-CVI/Ave 0.82–0.87) and stable composites in the case applications show that EDMC integrates Principles → Actions → Skills → Outcomes into one operable unit. In the highest-scoring cases, that PASO chain was made explicit and executable: principles were tied to named actions and owners, capability was specified, and outcomes were pegged to KPIs and thresholds (Sections 4.3.4–4.3.6). This shifts ethics from aspiration to governable practice (Treviño and Nelson, 2016).

Confidence without complacency. A small average uplift (+0.31) alongside tightening is the right outcome for a gatekeeping tool: sharper judgment and stronger conditions to proceed, without rubber-stamping or dramatic flips (Section 4.2.3). The pattern matches the anti-blindness intent—notice earlier, tighten where warranted, and proceed with explicit conditions when needed (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012).

Time pressure as crucible. The 4.20/5 rating for time-pressure usability is both encouraging and diagnostic. Experts asked for delivery-layer refinements—an executive-summary strip and a quick checklist—to speed entry without changing the PASO core (Section 4.2.5). That is packaging, not a conceptual rework.

Portability and AI relevance. Means from 3.96 to 4.23 on portability, integratability, conceptual coherence, and tech/AI relevance suggest the framework travels well. In AI-relevant cases, performance was strongest when bias-audit hooks, human-in-the-loop roles, and rollback conditions were concrete—exactly the move from principles to guardrails (Floridi, 2020); Sections 4.2.4, 4.3.6.

Robustness to real-world priorities. Reweighting toward Stakeholder/Outcomes or toward Time-pressure/Blindness/Implementation moved composites by at most ≤ 2.50 percentage points, with near-identical ordering (Section 4.3.5). Translation: EDMC is not

“gameable” by modest preference changes. Organizations can tune emphasis without overturning conclusions; improvement targets remain obvious (sharpen time-aware cues; enrich stakeholder-impact instrumentation) (Aguinis and Solarino, 2019).

Where the gaps are—actionable and fixable. Lower composites clustered where cultural adaptability (how plans travel across jurisdictions/teams) and time-aware cues were thinner. Those are exactly the upgrades experts requested—summary strip, RACI cues, KPI exemplars, jurisdiction hooks (Section 4.2.5). In short, v1.1 is faster packaging, not a new theory.

Why this matters now (especially for AI). Under speed and cognitive load, teams often fail to notice the ethical dimension at all—ethical blindness (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012). By aligning attention, stewardship, and evidence in the same frame, EDMC turns ethics from an ideal into operational guardrails—the difference between hoping a model behaves and specifying the conditions under which it may run (Floridi, 2020).

Mechanism: how PASO’s decision-DNA counters ethical blindness

- Salience → Action (P→A). State the governing principle(s) up front, then translate them into immediate actions and escalation paths so ethical salience arrives before momentum locks in.
- Stakeholder sweep and harm checks. Brief attention prompts (“who is affected,” “who would disagree,” “what could go wrong fastest?”) interrupt tunnel vision and widen perspective beyond the immediate KPI (Moore and Gino, 2013).
- Go/stop conditions. Clear thresholds for “proceed with conditions” vs “pause” draw a visible line between acceptable risk and unacceptable harm.
- Named stewardship (S). RACI-style cues convert intent into accountable capability—who does what, by when.

- Outcome checkpoints (O). KPIs/feedback loops tied to the invoked principle (e.g., fairness \rightarrow bias delta \leq X; privacy \rightarrow exposure \leq Y) make ethical intent testable and learnable.
- Time-aware scaffolding. A concise decision-summary strip and quick checklist support first-pass use under pressure; deeper annexes (jurisdiction hooks, KPI exemplars) support considered review.

Net effect. EDMC makes certain moves unavoidable—state a principle, define an action, assign an owner, peg an outcome—so ethical salience is surfaced before speed or habit suppress it.

Glass-box, not black-box. Each decision record shows (i) the principle invoked, (ii) the action chosen (and considered alternatives), (iii) the accountable skill/owner, (iv) the outcome metric/threshold, (v) the go/stop conditions, and (vi) the review cadence. This glass-box traceability makes reasoning auditable and teachable; it curbs ethical blindness by replacing tacit assumptions with explicit, checkable fields (Treviño and Nelson, 2016; Yin, 2018).

Taken together, these mechanisms explain why RQ6 yielded selective stance-tightening with modest confidence gains and stable composites under reweighting: EDMC’s PASO “decision-DNA” makes the attention \rightarrow ownership \rightarrow evidence path explicit at the moment of choice (Sections 4.2.3, 4.3.5).

5.5 Contributions

Positioning (first-tier, glass-box): EDMC is a first-tier, time-aware decision architecture organized around PASO—Principles \rightarrow Actions \rightarrow Skills \rightarrow Outcomes. It sits upstream of legal/compliance review, risk and controls, assurance, and communications. Its job is to (i) surface ethical salience early, (ii) connect principles to

executable safeguards with named stewardship, (iii) attach outcome thresholds and learning loops, and (iv) leave an auditable reasoning trail. EDMC is not a code of ethics, a sectoral checklist, a maturity model, or an algorithm that decides for you; it structures judgment and accountability so choices are explainable, auditable, and improvable across contexts. For validation we instantiated EDMC as a compact, auditable decision record; in practice it can live inside governance workflows (risk registers, DPIA/AIIA templates, incident run-books) rather than as a standalone artifact.

Distinctive mechanism (engineered noticing under time pressure): A core advance over prior models is EDMC's explicit, operational response to ethical blindness under speed and cognitive load (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012). Where classic and structured models often assume reflective conditions (Ferrell and Gresham, 1985; Kant, 1785; Mill, 1863; Rest, 1986; Treviño, 1986), EDMC builds in attentional cues, a stakeholder sweep, go/stop thresholds, named stewardship (RACI-style), and outcome checkpoints tied to first principles. The observed signature—12.2% stance-tightening with a +0.31/10 confidence gain—is consistent with anti-blindness theory when noticing and thresholds are designed-in, not left to chance (Section 4.2.3).

Contribution to prior literature (what we extend, not replace): We do not displace normative grounding; we operationalize it. EDMC requires naming the governing principle(s) and tying them to concrete actions and measurable outcomes—addressing the critique that principles remain aspirational when unlinked to implementation (Treviño and Nelson, 2016). We extend process accounts (Rest; Treviño) by making two often-implicit elements explicit and auditable: Skills/ownership (S) and Outcome instrumentation (O), plus time-aware go/stop thresholds. This aligns with behavioral ethics evidence that salience collapses under pressure (Bazerman and Tenbrunsel, 2011; Moore and Gino, 2013; Palazzo et al., 2012). For applied domains (GDPR/DPIA, AI

ethics, ESG), integrative reviews note principle inflation without operational anchors (Jobin et al., 2019; Mittelstadt et al., 2016). EDMC serves as a glass-box first layer that can feed DPIAs/AI impact assessments/ESG controls with stated principles, safeguards, owners, KPIs, and proceed-with-conditions vs pause logic—interoperable scaffolding, not a competing checklist. Finally, we answer field-level critiques (fragmentation, weak adaptability, limited tools) by keeping values, execution capability, and outcomes in the same line of sight and demonstrating robustness to priority shifts via preset sensitivity (Craft, 2013; Kaptein, 2023), Section 4.3.5.

Methodological contribution: We model a constructive, pattern-first synthesis to derive PASO; pair a transparent six-metric 0–3 rubric with disclosed weights/presets; and use a reproducible workbook with version-lock so independent examiners can recreate every number from the appendices (Lynham, 2002; Morgan, 2014; Yin, 2018). This is a practical template for building usable ethics instruments without sacrificing auditability.

Practical contribution (leadership and governance): For leaders, EDMC provides a first-tier filter that works under time pressure and leaves a glass-box record. The expert-requested v1.1 delivery upgrades, decision summary strip, quick checklist, RACI cues, KPI/mitigation library, jurisdiction hooks, target speed-of-use, not the conceptual core (Section 4.2.5). That makes adoption feasible in existing workflows (risk gates, board packs, incident playbooks) with minimal retraining.

Societal contribution: By making reasoning visible and outcomes testable, EDMC reduces space for ethics-washing, supports cross-sector interoperability, and strengthens trust in tech-intensive contexts, e.g., HITL roles, bias audits, and rollback conditions tied to principles in AI scenarios (Floridi, 2020).

Beyond compliance: turning values into verifiable action: Compliance asks “is it legal?” EDMC adds “is it right, who owns it, and how will we know it worked?”

Evidence from Chapter 4—stance-tightening 12.2%, confidence +0.31/10, and global utility 3.9–4.3—shows a first-tier filter that raises the bar without theatrics and leaves a traceable record (Sections 4.2.3–4.2.4).

What is new, in one line: EDMC converts fragmented ethical theory into a glass-box, first-tier decision system that engineers ethical salience under time pressure—producing auditable, culturally portable decisions that remain stable under realistic priority shifts.

5.5.1 Positioning: EDMC as a First-Tier Filter in a Tiered Decision Stack

- What EDMC is for.

EDMC sits upstream of specialist regimes (legal/compliance review, risk & controls, assurance, communications). Its role is to surface ethical salience early, link principles to executable moves with named stewardship, attach outcome thresholds and learning loops, and leave a compact, auditable trail. After that first pass, domain processes (e.g., DPIAs, safety cases, labor-law procedures) “carry the baton.” EDMC does not replace domain governance; it prepares and disciplines it.

- What EDMC is—and is not.

EDMC is a time-aware decision architecture organized by PASO (Principles → Actions → Skills → Outcomes). It functions as a first-tier ethical filter in both fast and deliberative settings; it is also a traceability mechanism—decision record, RACI/ownership cues, KPI hooks—that enables audit and learning across sectors and cultures. EDMC is not a code of ethics, a legal rulebook, a sectoral checklist, a maturity model, or an algorithm that decides for you. It does not guarantee “the correct” answer; it reduces error and opacity by structuring judgment and accountability.

- Beyond compliance—when “legal” isn’t “ethical enough.”

Compliance sets the floor; EDMC aims at the ceiling. By insisting that principles be tied to actions, owners, and measurable outcomes, it helps leaders recognize when “legal” still isn’t “right”—and then make principled, auditable moves under pressure. In practice, EDMC turns values into verifiable action: who owns what, which safeguards are triggered, what thresholds define “good enough to proceed,” and how results will be reviewed.

- How this shows up in the evidence.

In the expert-panel stream, 12.2% of paired decisions tightened toward “proceed with conditions”/“pause,” and mean confidence rose by +0.31/10—a signature of clearer criteria without bravado (Section 4.2.3). Global wrap-up means clustered between 3.9 and 4.3 for portability, Integratability, time-pressure usability, conceptual coherence, and tech/AI relevance, indicating EDMC travels beyond a single regime or rulebook (Section 4.2.4). In the researcher-led stream, equal-weight composites spanned 75.0–91.7 with maximum ≤ 2.50 pp movement under Stakeholder/Time-pressure presets and essentially stable ranks (Sections 4.3.4–4.3.5). Together, this pattern fits EDMC’s anti-blindness intent under speed and load (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012).

- Mindset, not just mechanics.

EDMC is not “another checklist”; it is a discipline of clarity. The PASO chain forces traceability—from principle to KPI—with named ownership in between. That discipline explains why higher-scoring cases (e.g., C1, C3, C7) pair strong principles with executable safeguards, while lower scores flag where packaging must tighten (cultural adaptability; time-aware prompts). The system teaches as it evaluates.

- A portable language (from three vignettes to many domains).

PASO operates as a grammar for turning values into action across contexts. It does not pre-decide your values; it structures how whatever you claim ethically must be enacted

(A), owned (S), and evidenced (O). The panel’s global means (≈ 4.0 – 4.3), the case composites’ stability (75.0 – 91.7), and the small preset shifts (≤ 2.50 pp) show that the grammar holds when priorities shift (stakeholder emphasis vs time-pressure) and when domains change (environment, labor, discrimination, AI).

- Learning in action (why small confidence gains matter).

A modest confidence uplift alongside selective tightening is precisely what a first-tier filter should produce. It clarifies judgment and documents conditions for “go” without inflating certainty. Because outputs are glass-boxed—owners, triggers, KPIs—they feed a learning loop: outcomes \rightarrow review \rightarrow updated safeguards. Capability grows from case to case, not just compliance for a single decision.

- Practitioner voice (anonymized, Section 4.2.5).

“Add a clear summary: shutdown vs proceed with conditions.”

“Make clear who owns each step and the consequences if missed.”

“Align with governance templates; provide board/audit language.”

“Standardized checklist with KPI/mitigation examples would increase consistency.”

These are delivery-layer requests—summary strip, checklist, RACI cues, KPI exemplars, jurisdiction hooks—not challenges to the core. They point to a v1.1 packaging pass that speeds the first mile while keeping PASO intact.

Takeaway. EDMC is a first-tier, time-aware glass-box architecture that moves leaders beyond “Is it legal?” to “Is it right, who owns it, and how will we know it worked?” The evidence—high usability, good PASO coverage, small but real confidence gains, and stance-tightening—shows a system that clarifies judgment without theatrics, spans domains, and leaves a trail that regulators, boards, and teams can trust. Its engineered noticing under pressure directly addresses ethical blindness, the very failure

mode most likely when speed and complexity are highest (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012).

5.5.2 Distinctive Contribution: Engineered Mitigation of “Ethical Blindness” Under Time Pressure

Many respected models of ethical decision-making excel at principles, process, or stakeholder mapping, but they largely presume reflective conditions (Ferrell and Gresham, 1985; Rest, 1986; Treviño, 1986). In practice, failures often occur earlier—when speed, stress, or cognitive load prevent people from noticing the ethical dimension at all (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012). EDMC’s distinctive contribution is an explicit, operational answer to that problem. It engineers salience and accountability in the flow of work: short attention cues to what matters, explicit go/stop gates, named stewardship and escalation (RACI), and outcome thresholds—captured in a compact, auditable decision record (Sections 4.2–4.3). Put simply: EDMC moves beyond “know the right thing” toward “notice it, act on it, and evidence it—fast.”

- Why the evidence fits the mechanism.

In the expert-panel stream, decisions tightened in 12.2% of paired cases and confidence rose +0.31/10—a signature of clearer criteria without overconfidence (Section 4.2.3). In the researcher-led stream, case composites were stable under Stakeholder and Time-pressure reweighting (maximum shift ≤ 2.50 pp; ranks essentially unchanged), indicating that the conclusions do not flip when priorities shift toward speed/containment (Section 4.3.5). This is exactly what we expect if salience is engineered and thresholds/owners are explicit: ethics surfaces when minutes are scarce, and safeguards are tied to measurable outcomes.

Table 30 EDMC vs. Other Models—Distinctive Features

Source: Author (2025). Original table created from study data and materials.

Feature / aspect	Typical coverage in other models	EDMC approach and distinctiveness
Integration of diverse traditions	Sometimes (e.g., stakeholder theory; ISCT)	Yes — explicit synthesis into PASO
Adaptability under real constraints	Rare (some contingency/applied models)	Yes — tiered, time-aware scaffolding
Actionable tools and guidance	Sometimes (compliance/ESG checklists)	Yes — compact record, quick checklist, KPI hooks
Development of competencies	Sometimes (leadership/virtue)	Yes — Skills/ownership embedded (RACI)
Mitigation of ethical blindness under time pressure	Almost never (major gap)	Core design — attentional cues, go/stop gates, escalation
Structured decision process	Often (Rest; Treviño; Ferrell and Gresham)	Yes — PASO chain with traceable steps
Cross-cultural adaptation	Rare (ISCT; Hunt and Vitell)	Yes — hypernorms + local hooks
Stakeholder analysis	Often	Yes — systematic + assigned owners
Outcome measures	Sometimes (TBL/ESG)	Yes — KPIs, thresholds, review cadence

Table note. “Typical coverage” reflects patterns from the Chapter 2 comparative review, see Table 36 (PASO mapping across frameworks) and Table 37 (six-dimension comparison).

What is unique here:

Distinctive advance: an operational response to ethical blindness under time pressure.

Unique value: not “more principles,” but engineered noticing—so ethics is surfaced, owned, and evidenced when it matters most.

Auditability: a glass-box trail that regulators, boards, and teams can review.

- Position and mechanism in one view.

As a first-tier filter in a tiered decision stack, EDMC frames the choice with named principles, executable safeguards, accountable owners, and outcome thresholds—creating a defensible record before downstream specialization. Mechanistically, it combats ethical blindness (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012) through attentional cues, go/stop gates, RACI stewardship, and KPI thresholds tied back to principles, making ethical salience unavoidable—and auditable—when minutes are scarce.

- One-line capsule.

Not more principles—better plumbing. EDMC makes noticing unavoidable, operationalizes it, assigns ownership, and sets thresholds—fast enough for real work.

These patterns—selective tightening, modest confidence uplift, and robustness under realistic priority shifts—are the desired RQ6 signature of a first-tier filter that improves decision quality under constraint without theatrical reversals (Sections 4.2.3, 4.3.5).

5.5.3 Theoretical Contributions

This study advances ethical decision-making theory by turning dispersed insights into a compact, operable architecture and by specifying mechanisms that connect behavioral-ethics findings to concrete process controls (Edmondson and McManus, 2007; Gioia et al., 2013).

- PASO as an integrative “DNA” architecture.

We formalize PASO as an integrative lens and a DNA-like architecture that binds normative intent (*Principles*), designed safeguards/escalations (*Actions*), accountable capability (*Skills/ownership*), and evaluation/learning (*Outcomes*) into one operable unit. Treating PASO as “DNA” highlights its invariants (the four components) and interfaces (each principle maps to at least one action; each action has a named owner/skill; each outcome traces back to a principle; time-pressure and culture-fit prompts are explicit). Co-locating these elements addresses fragmentation by keeping values, execution, capability, and effects in the same line of sight, extending process views beyond cognition alone (Craft, 2013; Kaptein, 2023; Lynham, 2002; Rest, 1986; Treviño, 1986). We use “DNA” metaphorically to denote a minimal, portable architecture rather than a biological analogy

- From bounded ethicality to designed mechanisms.

EDMC operationalizes insights on ethical blindness by engineering salience and control points: attentional cues and stakeholder sweeps to counter narrow framing, explicit stewardship and escalation (RACI) so responsibility is owned, and go/stop thresholds tied to outcomes so “proceed with conditions” has testable meaning. In doing so, it connects well-documented blind-spot mechanisms to testable process requirements—moving from post-hoc reflection to designed-in practice (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012; Tenbrunsel and Smith-Crowe, 2008).

- Construct definition and scope conditions.

The thesis clarifies ethical decision-making as a multidimensional construct organized by PASO to enable timely, transparent, and accountable choices across cultures. This definition supports content-validity assessment (I-CVI / S-CVI/Ave in Section 4.2) and aligns with integrative, practice-oriented theory building (Dixon-Woods et al., 2006; Lynham, 2002; Morgan, 2014). Scope conditions are explicit: EDMC is a first-tier, time-aware filter that disciplines judgment before downstream legal/technical regimes engage.

- A practice-oriented mid-range theory.

By specifying a neutral, portable architecture (PASO “DNA”) and its mechanism (engineered noticing + ownership + thresholds), the study contributes a mid-range theory that travels via analytic generalization rather than statistical inference (Edmondson and McManus, 2007; Yin, 2018). It yields testable propositions—for example, strengthening the *Skills* or *Outcomes* links should reduce blind-spot errors and tighten “proceed with conditions” thresholds.

- Traceability as a theoretical bridge.

Finally, EDMC’s principle-to-outcome traceability links normative ethics with assurance and governance: ethical claims are anchored to designed safeguards, accountable roles, and measurable effects. This glass-box stance reframes ethics from episodic compliance

to verifiable capability, offering a unifying construct for future theory and measurement (Morgan, 2014; Treviño and Nelson, 2016).

EDMC's PASO DNA makes ethical intent operable under time pressure—
Principles → Actions → Skills/ownership → Outcomes, with explicit go/stop thresholds.

5.5.4 Methodological Contributions

Methodologically, this study shows a pragmatic route from literature to an operable architecture, pairs transparent scoring with pre-specified sensitivity checks, and triangulates evidence in a way that fits a DBA validation phase (Dixon-Woods et al., 2006; Edmondson and McManus, 2007; Gioia et al., 2013; Lynham, 2002; Yin, 2018). Rather than running a deficit-only review, we recombined strengths from sixteen frameworks and treated recurring failure modes as design requirements. Using integrative/critical interpretive synthesis, those patterns were distilled into PASO—a lens-agnostic but binding scaffold that ties Principles → Actions → Skills/ownership → Outcomes into a single, operable unit (Dixon-Woods et al., 2006; Gioia et al., 2013; Lynham, 2002; Morgan, 2014). This pattern-first, problem-first synthesis is explicitly pragmatic: it begins with the work decisions must accomplish under constraints and builds a theory that fits that work.

To test the architecture, we operationalized PASO with a six-metric, 0–3 rubric and disclosed every anchor, weight, and computation step. Case-level composites were reported on 0–100 and probed for robustness with two pre-specified reweightings—Stakeholder and Time-pressure—against an Equal baseline (Aguinis and Solarino, 2019). Shifts were small (≤ 2.50 percentage points) and rank ordering was essentially stable, illustrating how to check sensitivity without over-claiming generality. Reproducibility was engineered, not asserted: a version-locked workbook generated all figures; formulas

and denominator rules for $\% \geq 4$ (with N/A excluded and effective Ns shown) were documented; and a simple leave-one-metric-out scan confirmed that no single dimension dominated the composite (Appendices B, F, M–N; Yin, 2018).

Evidence was triangulated through two complementary streams. The expert-panel stream assessed usability and content coverage using established content-validity conventions (I-CVI, S-CVI/Ave), stance change, confidence uplift, and global fit-for-use judgments (Grant and Davis, 1997; Lynn, 1986). In parallel, the researcher-led stream applied the rubric to eight public-facts cases to examine cross-context consistency, transparency of scoring, and robustness under realistic priority shifts. This sequencing is appropriate for an applied DBA phase: it shows how the mechanism behaves before costly field pilots while leaving an auditable trail (Lynham, 2002; Yin, 2018).

The methodological stance is glass-box, not black-box. EDMC is not “another checklist”; it is a discipline of clarity. By requiring a traceable chain—principle → action → owner/skill → outcome—we measured ethical judgment under constraint in a way others can reproduce, critique, and extend. In Chapter 4, higher-scoring cases paired principled stances with executable safeguards, named stewardship, and outcome thresholds; mid/lower profiles pinpointed where packaging should be sped up (cultural adaptability; time-aware prompts). That is the method doing its job: the system teaches as it evaluates.

In sum, the contribution is a replicable protocol for turning dispersed theory into an auditable decision architecture and testing it transparently: constructive synthesis to derive PASO DNA, disclosed anchors/weights and denominator rules, pre-specified sensitivity checks, and dual-stream triangulation suited to practice-oriented validation (Aguinis and Solarino, 2019; Dixon-Woods et al., 2006; Edmondson and McManus, 2007; Lynham, 2002; Yin, 2018).

5.5.5 Practical Contributions (Leadership and Governance)

What leaders get in practice. The evidence supports three concrete gains: deployability as a first-tier filter, cross-context portability, and systematic auditability. In governance terms, EDMC turns ethical intent into operational control by clarifying the choice, assigning ownership, and tying actions to measurable outcomes in a traceable record.

A first-tier ethical filter you can ship. Instantiated as a compact decision record, EDMC works at the speed of executive decision-making: it couples RACI-style stewardship with KPI/threshold hooks so teams can act under time pressure without dropping safeguards. The expert-panel results match that intent—selective stance-tightening (12.2%) with a modest confidence gain (+0.31 on a 0–10 scale) and strong time-pressure usability (mean 4.20/5)—showing a filter that sharpens caution without whiplash (Sections 4.2.3–4.2.4). Importantly, the panel’s asks were delivery-layer accelerators, not conceptual rework: an executive summary strip, quick checklist, KPI exemplars, and light jurisdiction hooks (Section 4.2.5). In short, the system is deployable now and accelerable in a v1.1 pack.

Cross-context portability. Performance held across environment, labor, discrimination, and AI-governance scenarios, and the global wrap-up means—Portability 4.09, Integratability 3.96, Tech/AI relevance 4.19, Conceptual coherence 4.23, Time-pressure usability 4.20—indicate the architecture travels beyond a single regime or rulebook (Section 4.2.4). Researcher-led case composites ranged 75.0–91.7 (median \approx 81.9) and remained stable when priorities shifted toward Stakeholders or Time-pressure (maximum change \leq 2.50 percentage points; rank order essentially unchanged), which is the kind of robustness organizations need when legitimate emphases differ (Sections 4.3.4–4.3.5).

Auditability by design. EDMC leaves a “glass-box” trail: the invoked principles, chosen safeguards, named owners, KPIs/thresholds, and the proceed-with-conditions/pause logic are recorded up front. That reduces ethics-washing risk and strengthens board reporting, internal audit, and learning reviews, while handing clean inputs to downstream regimes (e.g., DPIAs/AI impact assessments, risk/controls, assurance) rather than competing with them (see Appendices A–B).

A universal language leaders can reuse. Experts from varied domains engaged three hypothetical EDMC vignettes—and could have applied the same scaffold to many more. PASO functions as a portable grammar: it does not pre-decide values; it structures how values become actions, ownership, and outcomes across contexts. The portability signal appears both in the global means above and in the stability of case composites under reweighting (Section 4.3.5).

Monday-morning readiness. The immediate adoption path is straightforward: deploy EDMC with the v1.1 delivery pack—an executive summary strip, a one-page quick checklist, explicit RACI prompts, KPI/mitigation exemplars, and jurisdiction hooks—embedded in existing governance templates (e.g., investment papers, change approvals, incident run-books). These speed the first pass without altering the PASO-linked core.

Takeaway. EDMC gives leadership teams a practical first tier that moves the conversation from “Is it legal?” to “Is it right, who owns it, and how will we know it worked?” The Chapter 4 evidence—high usability, good PASO coverage, selective tightening with small confidence gains, and stable composites under priority shifts—shows a system that clarifies judgment under pressure, travels across domains, and leaves an auditable trail regulators, boards, and teams can trust.

5.5.6 Societal and Policy-Adjacent Contributions

Why this matters beyond the firm. The same properties that make EDMC workable inside organizations—glass-box reasoning, named stewardship, and outcome hooks—also raise the quality of public accountability. Decisions stop being assertions (“trust us”) and become auditable narratives (“trace us”), with clear lines from the invoked principle to the safeguard chosen, who owns it, and how success will be checked.

Trust and accountability. Because EDMC records the why, what, who, and how we’ll know, organizations can explain controversial choices to boards, employees, communities, and regulators with evidence rather than rhetoric. In practical disputes (e.g., labor standards, AI fairness, environmental trade-offs), the record shows whether ethical salience was noticed in time, which alternatives were considered, and what thresholds gated a “proceed with conditions” versus a “pause.” This reduces ethics-washing risk and supports credible oversight and independent review.

Interoperability with governance regimes. EDMC outputs slot cleanly into existing policy tools as a front-end ethical trace: privacy or AI impact assessments (DPIA/AIIA), safety cases, risk registers, ESG disclosures, and assurance workpapers. Because the PASO DNA is principle-agnostic but structure-strict, it travels across regimes and jurisdictions—mirroring the portability and integratability signals reported in Chapter 4 (global means ≈ 4.0 – 4.3 ; stable composites under re-weighting). In practice, EDMC prepares downstream compliance rather than competing with it.

Capability building (from one-off fixes to learning). The Outcomes link (KPIs, thresholds, review cadence) turns each decision into a teachable case. Post-decision reviews can check whether the intended guardrails worked, feeding updates to templates, training, and governance artefacts. Over time, this creates an organizational (and sector-

level) memory for ethically salient patterns under pressure—exactly where many failures recur.

Public value in high-velocity domains (especially AI). Where speed and scale magnify harm, EDMC’s engineered noticing (attentional cues, go/stop gates, rollback plans, HITL assignments) offers a policy-adjacent safety layer. It gives regulators and civil society a consistent way to ask the same questions organizations should already be answering: Which principle governed? What safeguard ran? Who owned it? What threshold stopped the system?

Boundaries and good-faith use. EDMC does not replace law, technical audits, or sectoral standards; it disciplines the first pass and hands clean inputs to those regimes. Transparency must be balanced with confidentiality (commercial or personal data), but even redacted records preserve traceability for audit and oversight.

Takeaway. EDMC raises the public bar from “trust us” to “trace us.” By making principle-to-outcome logic explicit, owned, and measurable, it provides a portable ethical spine that policy makers, auditors, and civil society can examine—supporting accountability in the very settings (time-pressed, tech-intensive, cross-jurisdictional) where it is most needed.

5.5.7 Where This Contribution Stops (and Why That’s a Strength)

Scope by design. EDMC is intentionally narrow and first-tier. Its job is to make ethical salience visible, tie principles to executable safeguards with named ownership, and record outcome thresholds—fast—so downstream specialists can do deeper work. That boundary keeps the system practical and prevents category mistakes.

Not a final arbiter of values. EDMC does not settle deep ethical disputes. It structures how disagreements are surfaced, reasoned, and documented (which principle

applies, which alternatives were considered, who owns the next step, what threshold gates a “proceed with conditions”). This makes contention tractable without pretending to erase it.

Not a substitute for specialists. EDMC does not replace legal counsel, risk engineers, data scientists, or domain auditors. It improves the hand-off to those functions by delivering a glass-box record: principles invoked, safeguards chosen, owners named, KPIs/thresholds set. In practice, that shortens cycles and raises the quality of downstream review.

Not a claim to statistical generalization (yet). At this validation phase, the evidence shows usability under time pressure, coherent PASO coverage, cautious stance-tightening, modest confidence gains, and stability under reasonable priority shifts. It does not claim population estimates or causal effects. Those belong to the next steps: multi-rater deployments, field pilots, and longitudinal outcome tracking.

Constraining EDMC to a first-tier, time-aware filter is precisely its comparative advantage. By avoiding overreach, it remains portable across sectors and jurisdictions, interoperable with existing governance (DPIA/AIIA, risk, assurance), and easy to adopt. In terms of the research questions, this focus preserves the strengths demonstrated for RQ5–RQ6 (translation into tools, fitness-for-use under constraints) without over-claiming what only future field studies can confirm.

Takeaway. EDMC is a disciplined starter layer, not an all-purpose oracle. That is a feature, not a bug: it prepares and strengthens downstream governance by turning values into verifiable action—clearly, quickly, and audibly.

5.5.8 Novelty and Real Merit of This Study

The study's core novelty is an engineered response to ethical blindness under time pressure. Rather than assuming reflective conditions, EDMC operationalizes noticing through concise attention prompts, explicit go/stop thresholds, and named stewardship—so ethical stakes are recognized and evidenced when minutes are scarce. The observed pattern (selective stance-tightening with only modest confidence uplift) is exactly what you would expect when salience and thresholds are designed in, not left to chance.

A second contribution is PASO as an integrative architecture. By binding Principles → Actions → Skills → Outcomes into one portable system, EDMC stays neutral on which principles a context adopts but is insistent about how they are enacted, owned, and measured. This “PASO DNA” turns ethical intent into traceable moves and testable effects, closing the common gap between aspiration and execution.

Methodologically, the project introduces a dual-validation approach suited to a practice-ready DBA phase: expert content validity (I-CVI/S-CVI) combined with cross-sector case robustness, including pre-specified reweighting presets. Together these show that EDMC is both recognizable to experts (coverage/coherence) and stable under realistic priority shifts (stakeholder vs. time-pressure emphasis).

Practically, EDMC integrates cleanly with existing governance. The glass-box decision record—with RACI cues and KPI/threshold hooks—slots into DPIAs/AIIAs, risk and controls, assurance, and board reporting. That makes adoption feasible without inventing a parallel bureaucracy.

The work also bridges theory to capability. EDMC turns episodic reflection into an organizational practice: decisions leave auditable footprints, outcomes feed learning loops, and the same scaffold scales to AI-intensive contexts (e.g., bias audits, HITL roles, rollback criteria tied to principles).

Finally, the comparative advantage is evidenced. *Table 30*, and Tables 36–38 show how EDMC addresses the field-level gaps identified in the review—fragmentation, limited adaptability, thin tooling, weak skills/ownership, and the near absence of time-pressure countermeasures. In short: EDMC is novel where it matters (engineered anti-blindness; PASO DNA), validated in ways that fit real use, and designed to plug into the systems organizations already run. With contributions in place, we now focus on limitations and delimitations, and how they influence scope and future steps.

5.6 Limitations and Boundary Conditions (and Why Claims are Modest)

As an emergent, practice-oriented DBA study, this work prioritized auditability and robustness over population inference (Edmondson and McManus, 2007; Lynham, 2002; Yin, 2018). The resulting claims are intentionally careful and bounded, and the limits below point directly to the next wave of evidence.

Panel composition and sampling. The expert panel (N = 30) was senior and practice-oriented, but skewed male (73.3%) and private-sector (70.0%). It was not a probability sample, so findings should not be generalized to a population. Future panels should widen geography, sector, and gender balance to stress-test portability and explore subgroup differences (Grant and Davis, 1997; Lynn, 1986).

Mode and scenario constraints. Public-facts vignettes (for the expert panel) and hypothetical researcher-led applications are well suited to content validity and usability testing, yet they do not substitute for longitudinal organizational studies where incentives, politics, and drift are at play (Yin, 2018). External validity requires field pilots—A/B roll-outs, before/after designs, and tracked outcomes.

Cultural breadth. Portability indicators were encouraging (global means \approx 4.0–4.3), but deeper non-Western validation remains a priority. Cross-cultural work should

explicitly examine hypernorms vs. local norms and document localization heuristics rather than assume transferability (Donaldson and Dunfee, 1999).

Change over time (AI/regulation). AI practice and data-protection regimes evolve quickly. EDMC will need periodic updates so attentional cues, thresholds, and hooks remain aligned with contemporary risk and regulation (Floridi, 2020).

Attribution limits. The observed effects, mean confidence +0.31 (0–10) and 12.2% stance-tightening, are consistent with EDMC’s mechanisms (salience cues, explicit go/stop gates), but a survey setting cannot isolate all confounds (e.g., framing effects, prior familiarity). Stronger attribution will require multi-rater scoring, inter-rater checks, and field outcomes (Hsieh and Shannon, 2005; Schwartz, 2016).

Threats to validity and alternative explanations. Results could be upwardly biased by format novelty (a clear record can create halo effects), demand characteristics (participants inferring “intended” answers), or researcher expectancy in single-rater case scoring. Mitigations included version-locked instruments, pre-specified anchors and weight presets, and a transparent audit trail (Appendices A–B, M–N). Even so, independent raters and in-situ deployment are needed to separate EDMC-specific effects from presentation or context effects (Shadish et al., 2002).

Reflexivity and positionality. As a practitioner-researcher designing for use under constraint, the author’s priors favor clarity and speed. To bound (not erase) this bias, we used directed content-analysis codebooks, pre-specified scoring anchors, version control for instruments and figures, reflexive memoing, and peer debriefs - Appendices F, H, M–N; (Alvesson and Sköldberg, 2017; Finlay, 2002).

Why the emphasis is descriptive (and appropriate). This phase deliberately emphasizes descriptive indicators and sensitivity checks, not hypothesis tests. That stance fits an emergent, build-and-validate cycle aimed at practice-ready utility (Edmondson and

McManus, 2007; Lynham, 2002). The next step is cumulative: field pilots, longitudinal outcomes (e.g., time-to-decision, incident rates, audit findings), and cross-cultural replications.

So what—credibility at this stage. Triangulation (survey + cases + sensitivity checks) and auditability (disclosed anchors, weights, formulas in Appendices A–B; version lock in M–N) provide the right kind of credibility for a DBA validation phase: clear enough to justify organizational pilots, careful enough to invite stronger tests next. In RQ terms, the limits above bound RQ6 (validation under constraints) to analytic generalization and near-term portability signals; Phase-II field work should extend those signals to causal and longitudinal claims.

5.7 Practice Implications (Preview; Elaborated in Chapter 6)

EDMC is ready for use as a first-tier, time-aware filter that strengthens—rather than replaces—downstream governance. Practically, make the EDMC Compass the front door for any material decision: capture a concise record that states the governing principle(s), the safeguard(s) to be implemented, the accountable owner with the right skills, and the outcome check with a review cadence. That single move embeds PASO DNA into day-to-day work and leaves a transparent trace boards, audit, and regulators can read.

To speed the first mile, ship a light v1.1 delivery pack: a one-look summary strip, a quick-reference cue card for time-pressure prompts, clearer RACI cues, a small library of KPI/mitigation exemplars, and gentle jurisdiction hooks. Standardize templates across the organization and digitize lightly by plugging the Compass into existing workflows (risk registers, DPIA/AIIA forms, incident run-books), aiming for traceability and repeatability, not a black-box app.

Treat every Compass entry as learning data. The loop is simple and agnostic to doctrine: PASO → brief outcome review → safeguards updated → refreshed PASO. Over time, that rhythm turns awareness into capability. Guard against ethics-washing by insisting on named owners with remit, explicit thresholds tied to evidence, routine spot-checks, and short board summaries of coverage and follow-through. Know the limits: EDMC adds less where outcome data are unavailable, owners lack authority, extended deliberation is appropriate, or local legality conflicts with hypernorms and escalation is refused. Even then, the Compass clarifies trade-offs and stewardship.

A short illustration. A team plans to launch an AI résumé screener ahead of campus season. They name fairness/non-discrimination as the principle; commit to a pre-launch bias audit, human-in-the-loop for ambiguous scores, and a rollback plan; assign a skilled model owner, a fairness reviewer, and the HR process owner; and set a simple gate (e.g., selection-rate parity must meet the agreed threshold before launch). Result: the timeline holds, but launch waits until fairness is demonstrated, ethical readiness becomes a criterion alongside technical readiness.

For implementation, target “systematic controls” on the study’s blindness-mitigation scale (structured prompts, thresholds and escalation, RACI ownership, KPIs tied to principles, and periodic review). The rubric, anchors, and presets are consolidated in the appendices for straightforward reuse. These moves operationalize the RQ6 verdict: usable under time pressure, auditable in practice, portable across contexts—and simple enough to become a habit.

5.8 Directions for Further Research

Building on this validation phase—pattern synthesis, an expert-panel stream, and researcher-led case applications—the next step is field deployment with longitudinal

evidence (Lynham, 2002; Yin, 2018). In practice, run EDMC as a first-tier filter in live workflows (e.g., AI launches, procurement, HR) and track decision quality, time-to-decision, incident rates, and audit findings before/after and against matched comparisons. Where feasible, use stepped-wedge or clustered roll-outs so implementation proceeds without disrupting operations (Shadish et al., 2002).

A second strand should test portability and localization explicitly: examine how hypernorms hold across jurisdictions and sectors, specify what can be adapted without diluting those hypernorms, and document the practical heuristics teams use to localize EDMC (Donaldson and Dunfee, 1999; Hofstede, 2001; Jobin et al., 2019). In AI-intensive domains, assess bias diagnostics, human-in-the-loop roles, rollback conditions, and model-risk hand-offs using fairness and quality KPIs in situ (Floridi, 2020; Zhang and Wade, 2022).

From a measurement perspective, report ICC(2,k) for rubric agreement and Kendall's W for rank concordance; extend construct validity by relating rubric dimensions to clarity, confidence, stance, and outcomes. Robustness checks can move beyond weight presets to bootstrapped CIs, leave-two-out scans, and anchor-perturbation tests (Aguinis and Solarino, 2019; Schwartz, 2016; Yin, 2018). Run a side-by-side test: teams using EDMC versus teams using the current policy/checklist approach; compare how often ethical issues are spotted early, how often safeguards are tightened or work is paused when standards aren't met, the time to reach a defensible decision, and the rate/severity of incidents and audit findings. Finally, co-design sector annexes (e.g., healthcare, public sector, safety-critical) with regulators and practitioners, pre-specify success criteria (adoption, decision quality, audit findings), and pilot before scaling.

5.9 Chapter Takeaway

Unlike many frameworks, EDMC converts fragmented theory into a glass-box, first-tier decision system organized by PASO—Principles → Actions → Skills → Outcomes. In one compact record it captures named principles, executable safeguards, accountable owners, and outcome thresholds, engineering noticing under time pressure. Across two independent evidence streams, the expert-panel read EDMC as clear and portable (global means 3.9–4.3), with good PASO coverage (S-CVI/Ave 0.82–0.87), cautious stance-tightening (12.2%) and modest confidence gains (+0.31/10). In the researcher-led stream, case composites were strong and stable under realistic priority shifts (≤ 2.50 percentage points), indicating robustness to stakeholder-emphasis or time-pressure reweighting.

What remains is practical and fixable: packaging for speed (summary strip, quick checklist, RACI cues, KPI exemplars, jurisdiction hooks) and sharper cultural adaptability—not a redesign of the core. In short: principled, actionable, auditable—at the speed of modern work. Chapter 6 turns these findings into a concrete implementation roadmap and a research agenda for scaling, localization, and digital deployment, carrying forward the RQ6 verdict that EDMC is usable under pressure, portable across contexts, and ready to embed in everyday governance.

CHAPTER VI: SUMMARY, IMPLICATIONS, AND RECOMMENDATIONS

6.1 Summary

This thesis tackled a durable leadership problem: how to make principled, auditable choices when time is short, information is noisy, and contexts differ. Many familiar frameworks illuminated parts of the journey yet proved too siloed, slow, or abstract for fast, high-stakes decisions; under pressure, ethical salience was often missed, rationales fragmented, and learning evaporated (Bazerman and Tenbrunsel, 2011).

The Ethical Decision-Making Compass (EDMC) offered a practical, simple, agnostic solution: a first-tier, time-aware decision architecture with PASO DNA—Principles (why), Actions (how), Skills (who), Outcomes (what). Not a checklist but a compass that orients attention at the moment of choice, EDMC links principles to executable safeguards, requires named and capable stewardship (the “S” is non-negotiable: ownership implies competence and, where needed, targeted upskilling or a qualified co-owner), and binds every choice to accountable outcomes, thresholds, KPIs, and a review cadence. By repeating this PASO rhythm at the point of decision, EDMC turns ethics from episodic reflection into organizational habit and culture: salience is noticed, action is competent, and lessons from outcomes update practice. Doctrine-agnostic on what to value, EDMC is uncompromising that whatever is named must be enacted, owned, and evidenced, strengthening (not replacing) downstream governance (Edmondson, 2018; Schein, 2010; Treviño and Nelson, 2016).

Two complementary evidence streams supported fitness-for-use. An expert-panel stream (N = 30) assessed usability and content coverage (I-CVI / S-CVI/Ave), while a researcher-led stream applied a transparent 0–3 rubric to eight public-facts cases and tested preset reweightings for robustness, all with version-locked materials for auditability (Aguinis and Solarino, 2019; Lynn, 1986; Morgan, 2014; Polit and Beck,

2006; Yin, 2018). Together, the signals converged: EDMC's PASO DNA was coherent, usable under pressure, portable across contexts, and supportive of clear, reviewable decisions—the intended signature of a first-tier ethical filter. Claims are credible but modest and appropriate to an applied validation: the panel was purposive, the cases used public facts, and scoring was single-analyst; nonetheless, version-locking, denominator transparency, and sensitivity checks support the conclusion that headline signals were not artifacts of handling choices (Aguinis and Solarino, 2019; Yin, 2018).

6.2 Implications

6.2.1 Theoretical Implications

Conceptually, EDMC connected theory and practice without diluting either side. By integrating Principles, Actions, Skills, and Outcomes into a single, functional architecture, it addressed the field's longstanding fragmentation: values were no longer ideals detached from execution; capability was named rather than assumed; outcomes were specified and reviewed instead of implied (Rest, 1986; Treviño, 1986). Behavioral insights about ethical blindness under time pressure were translated into design features—attentional cues, stakeholder sweeps, thresholds, and owned escalation—so noticing was engineered rather than left to chance (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012).

Framed this way, the thesis clarified a testable construct: ethical decision-making as a multidimensional, PASO-DNA process that enables timely, transparent, and accountable choices across cultures. That definition functions as mid-range theory: it is neutral about which principles a context adopts, but insistent that whatever is claimed ethically must be enacted (A), owned (S), and evidenced (O). Because the architecture is explicit and auditable, it supports analytic generalization (rather than statistical inference)

and yields hypotheses that future work can test—for example, that strengthening the Skills or Outcomes links will reduce blind-spot errors or narrow the gap between stated principles and realized practice (Edmondson and McManus, 2007; Yin, 2018).

6.2.2 Practical and Leadership Implications

For leaders, EDMC functioned as a first-tier ethical filter that could be used in urgent and deliberative settings alike. The PASO DNA made explicit why a decision was justified, how safeguards translated intent into action, who was responsible to act and escalate, and what evidence would demonstrate success. Because Outcomes were tied back to Principles and owned by named stewards, each decision created a learning loop: results informed adjustments and capability improved through practice.

In socio-technical environments, EDMC naturally incorporated human-in-the-loop roles, bias-audit checkpoints, and rollback conditions—guardrails that operationalized ethical intent without treating ethics as an afterthought (Floridi, 2020). The expert-panel results supported everyday deployability—high agreement on usability and strong global means for portability, integratability, time-pressure usability, and coherence—while the observed pattern of selective stance-tightening (12.2%) with modest confidence gains (+0.31/10) indicated that leaders could raise the bar on safeguards without inducing whiplash or overconfidence.

Practically, this pointed to a clear next step for organizations: embed the one-page decision record at the front door of material choices; pair it with a quick-reference cue card for time-pressure entry; and maintain a light library of RACI prompts and KPI/mitigation exemplars. The result is a repeatable, auditable rhythm—Principle → Action → Owner → Outcome → Review—that strengthens downstream governance rather than replacing it.

6.2.3 Policy and Governance Implications

EDMC raises the standard beyond compliance. It links “*Is it legal?*” to “*Is it right, who owns it, and how will we know?*” by producing a compact, glass-box decision record that boards, audit, and regulators can scrutinize. Because principles are tied to owned actions and outcome thresholds, claims can be evidenced rather than asserted, which helps deter ethics-washing and supports proportional disclosure (what can be shared publicly without breaching confidentiality). In higher-risk domains, EDMC records can serve as credible ex-ante due-diligence artifacts—a traceable rationale before launch—while sectoral standards (e.g., DPIAs/AI impact assessments, safety cases, labor and environmental regimes) retain their essential technical depth. In short, EDMC prepares and disciplines downstream governance; it does not replace it.

6.2.4 Societal Implications

Making trade-offs explicit and testable strengthens trust with employees, communities, and regulators. Because the PASO grammar is teachable and portable, frontline teams gain a shared language for hard choices, and multi-party collaborations (public–private–nonprofit) can compare options on the same footing: principle invoked → safeguard chosen → owner named → outcome to be demonstrated. That common format enables learning across organizations, improves accountability in data-heavy and AI-enabled settings, and makes stewardship visible—a precondition for public legitimacy. This claim is supported by the fact that the expert panel in Chapter 4 was geographically distributed by residence across Europe, North America, Asia-Pacific, the Middle East, and Africa (N = 30), indicating that evaluations of clarity, defensibility, and usability did not come from a single jurisdiction or governance culture.

6.2.5 Capacity, Empowerment, Awareness, and Trust

EDMC builds capability as a habit. Each pass through PASO makes four things more likely: teams notice the ethical dimension early, translate principles into concrete safeguards, assign skilled ownership with clear escalation, and check outcomes on a cadence. The result is less tunnel vision and more distributed accountability: ethical work moves to the edges of the organization, while reasoning stays visible and auditable.

In effect, EDMC builds both the ‘muscle’ of practice and the ‘brain’ of principle: repetition develops fluent execution and escalation discipline, while explicit principles keep judgment anchored and teachable.

Table 31 EDMC Contributions at a Glance

Source: Author (2025). Original table created from study data and materials.

Contribution	How Supported (EDMC Features)	Implications for Theory and Practice
Capacity / Skills	PASO prompts; KPI hooks; review cadence	Competence is developable; track skill growth over time
Empowerment	RACI ownership; escalation; clear hand-offs	Accountability at all levels; participatory ethics
Awareness	Saliency cues; thresholds; stakeholder sweeps	Fewer missed ethical dimensions; A/B-test noticing
Trust	Glass-box records; audit hooks	Transparency → legitimacy; stronger assurance

Note. “How EDMC does it” refers to PASO-DNA elements and the concise decision record; the right column summarizes implications (see Chapter 4 for indicators).

6.3 Recommendations

6.3.1 Organization-Wide Implementation Roadmap

Make EDMC the first-tier step for all material decisions, then scale in four light phases. The aim is a repeatable habit: for each decision, teams complete the PASO DNA—Principles, Actions, Skills/ownership, Outcomes/thresholds—and a one-glance decision-summary strip. Start small, learn quickly, and embed without bureaucracy. How the four phases work (in practice, not theory):

Pre-flight (weeks 0–2). Name a Sponsor (accountable senior leader) and a Steward (day-to-day owner). Pick 1–2 pilot domains with real decision volume (e.g., AI launch, supplier onboarding). Finalize the EDMC record + summary strip. Baseline a few indicators (e.g., coverage, PASO completeness, follow-through rates).

Pilot (weeks 3–12). Run 10–20 live decisions per domain. Offer short coaching “office hours.” Mid-pilot and end-pilot, share a brief status snapshot (coverage, completeness, early learning). Exit when completion is high and at least one decision shows tightened safeguards or a pause based on thresholds.

Scale (months 4–9). Extend to 4–6 domains. Add light digitization (templates, pick-lists, minimal workflow). Begin board-level status snapshots (counts; fit-with-workflow; % records with thresholds and named owners; % outcome reviews on time). Shift coaching toward peer exemplars.

Embed (months 10–12). Codify in policy and onboarding. Add vendor and incident pathways. Schedule a light external spot-review. Publish a v1.2 pack with updated exemplars and lessons learned. *Table 32* summarizes these phases at a glance—timeframe, key actions/deliverables, and exit criteria—so sponsors and teams can track progress and know exactly when to move to the next phase.

Table 32 EDMC Implementation Phases

Source: Author (2025). Original table created from study data and materials.

Phase	Timeframe	Key Actions and Deliverables	Exit Criteria
Pre-Flight	Weeks 0–2	Assign Sponsor and Steward; select pilot domains; finalize templates; define KPIs	Roles named; templates frozen; pilots identified; KPIs baselined
Pilot	Weeks 3–12	Run 10–20 decisions in pilot domains; train users; review outcomes; office hours	≥80% EDMC completion; fit with workflow demonstrated; high PASO
Scale	Months 4–9	Expand to 4–6 domains; digitize lightly; board-level reporting; standardize KPIs	≥75% coverage; board/CSR/AI report; hand-off to risk/assurance
Embed	Months 10–12	Policy integration; onboarding/training; vendor/incident flows; external review	Policy codified; recurring training; external spot-review; v1.2 release

Note. “Completion” = all PASO fields + decision-summary strip; “status snapshot” = one-page update on coverage, completeness, follow-through, and lessons learned.

Anticipated obstacles (and how to keep momentum): The usual pitfalls—performative paperwork, bottlenecks, local resistance, over-standardization, authority gaps, and thin data—are manageable when thresholds and ownership are required, the record stays concise, and principled local adaptation is allowed. *Table 33* lists common barriers and practical mitigations mapped to PASO DNA so teams can act immediately.

Table 33 Barriers and Mitigation

Source: Author (2025). Original table created from study data and materials.

Barrier	Mitigation Strategy
Ethics-washing	Require thresholds, named owners, and evidence links; audit a sample monthly
Bottlenecks	Keep the record concise; use the decision-summary strip; share exemplars
Local resistance	Start with one unit; showcase wins; use exemplars to reduce effort
Over-standardization	Allow local adaptation while preserving EDMC’s PASO-DNA logic
Authority gaps	Provide an escalation path to the Sponsor where owners lack remit
Data poverty	Use proxy metrics now; document a plan to improve data

Note. Mitigations map to PASO DNA: salience/thresholds (awareness), RACI roles (ownership), proxies + review cadence (outcomes). Audit a small random sample monthly and publish a brief lessons-learned note quarterly. See Section 6.4 for the evaluation cadence.

6.3.2 For AI, Data, and Socio-Technical Programs

Make the EDMC decision record the one-page launch gate and handoff artifact for DPIAs/AI impact assessments, model-risk governance, safety cases, and ESG reporting (NIST, 2023). In practical terms, every build/release should be able to pass a simple, visible test: does the PASO-DNA chain close? That means (i) the principle(s) invoked are named (e.g., fairness, proportionality, dignity), (ii) actions are concrete (pre-launch bias audit, red-team test, shadow mode, rollback plan), (iii) skills/ownership are explicit (model owner, risk steward, escalation contact), and (iv) outcomes are tied to principle-linked KPIs with a monitoring cadence (e.g., parity thresholds, privacy leakage bounds, harm triage SLA). Example to picture: before a résumé-screening model goes live, the EDMC record states Fairness as the governing principle; Actions include pre-

launch bias audit and human-in-the-loop review for ambiguous scores; Skills names the ML lead, the fairness reviewer, and the HR process owner; Outcomes set a go/stop line (e.g., selection-rate parity ≥ 0.80 and bias delta ≤ 2 pp across gender/ethnicity). If monitoring later drops below the line, the record already names who pauses the system and how rollback works. The benefit isn't paperwork—it's glass-box interoperability: the same page feeds the DPIA/AI-impact template, model-risk register, and board report with consistent, auditable fields.

6.3.3 For Boards, Audit, and Policy/Governance Bodies

Use EDMC to raise the ceiling beyond compliance. For any high-risk initiative, require a concise EDMC record as part of the pack. In a 90-second readout, directors should be able to see four things: the governing principle(s), the go/stop threshold, the named steward, and the date the outcome will be reviewed. Emphasize thresholds and stewardship over narrative alone; then sample a risk-based subset each quarter to verify that outcomes were reviewed on time and that “proceed with conditions” led to the conditions being met. Publish a short committee-level summary, coverage, notable signals, follow-through, to deter ethics-washing and normalize evidence-based accountability. For regulators and policy bodies, treat EDMC-style artifacts as credible ex-ante due diligence (the why/how/who/what), while leaving technical specifics to the sector frameworks they already oversee.

6.3.4 For Researchers, Educators, and Partners

To build cumulative evidence, focus on field deployment with longitudinal tracking—use EDMC as the first-tier filter in live workflows and assess decision quality, time-to-decision, incident rates, and audit findings through stepped-wedge or clustered

rollouts. Combine this with lightweight A/B tests of engineered noticing (such as salience prompts in the summary strip) to measure improvements in “proceed with conditions/pause” decisions and decreases in missed risks. Pursue cross-cultural validation and localization, documenting what remains constant (hypernorms) and what varies by context. Enhance measurement science by reporting inter-rater reliability (for example, ICC(2,k)), rank concordance (Kendall’s W), and linking construct validity between rubric dimensions and outcomes like clarity, confidence, and stance, along with sensitivity analyses beyond preset weights. Treat digital embedding as research: prototype simple integrations into DPIA/AI-impact templates or risk registers, and evaluate their completeness, friction, and auditability without dulling their importance. Finally, conduct policy pilots with boards and regulators where the single-page EDMC record functions as an ethical front-end to existing regimes, and facilitate open teaching and data sharing (such as anonymized records, coding guides, micro-sims) to support learning of the PASO grammar through practice.

The thread running through all three audiences is the same: close the PASO-DNA loop on one page, make ownership and thresholds explicit, and tie choices to evidence and review. That’s how EDMC moves ethics from aspiration to accountable practice, at the speed of modern work.

6.4 Implementation and Evaluation: How We’ll Know EDMC + PASO Is Working

EDMC’s PASO DNA architecture—Principles → Actions → Skills → Outcomes, acts as an embedded Compass. It sits at the moment of choice, not after it, making ethical salience explicit, translating principles into safeguards, naming capable ownership, and defining how success will be checked before a material decision advances. The aim is to develop a repeatable decision-making habit that counters ethical blindness under pressure

(Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012). Over time, the Compass develops both the muscle of reliable execution and the brain of principled judgment—automatic where it helps, reflective where it matters.

What “working” looks like. Every material decision first passes the Compass. The governing Principle is clearly named; the Actions are concrete (safeguards, escalation); the Skills required are detailed and assigned to an accountable owner with defined remit; and the Outcomes have measurable criteria and review schedules. When standards aren’t yet met, the decision proceeds only with enhanced safeguards—or it is paused. A short, living record makes the reasoning and responsibilities clear. Brief scheduled reviews turn experience into improved guidance, helping the habit strengthen over time (Edmondson and McManus, 2007).

Light signals, strong discipline. Track several operational signals to steer: coverage (how many material decisions used the Compass), pace (fit with workflow), completeness (P→A→S→O linked, no missing owners/thresholds), follow-through (agreed safeguards implemented on time), auditability (spot-checks pass), and learning (PASO materials updated from reviews). Where AI/data risk is in scope, include fairness/quality thresholds with explicit no-go gates (NIST, 2023). The goal is course-correction, not bureaucracy.

Roles and cadence. A Sponsor sets the standard and removes obstacles; a Steward monitors these signals and coaches teams; Domain leads ensure local PASO guidance remains relevant and conduct brief outcome reviews. Quarterly, leaders receive a one-page update on coverage, key signals, exceptions, and the changes resulting from these updates.

Apply proportionate rigor. Use pragmatic designs—such as before/after in early domains, matched-unit comparisons, or a stepped rollout—along with small A/B checks

for prompts intended to enhance noticing and appropriate tightening without delaying teams (Shadish et al., 2002; Yin, 2018). Keep templates version-locked and make summaries clear about what was counted and how (Lynham, 2002).

Adaptable, spine-true. Localize examples, KPI sources, and escalation paths; keep the PASO spine intact so capability and accountability remain clear across regions and teams.

Early signs you're on the right track include more confirmed approvals or brief pauses when standards aren't yet met; clearer ownership with designated skills and responsibilities; fewer blind spots during quick decisions; smoother transitions to legal, compliance, and assurance teams; and a visible learning loop as PASO guidance is refined.

Why a Compass? Because orientation is dynamic. Contexts shift; the PASO DNA does not. It keeps attention, action, capability, and evidence aligned—so principled choices stay operational, auditable, and adaptive at the speed of real work.

6.5 Limitations and Boundary Conditions

Claims here are deliberately modest and appropriate to an early, practice-focused validation. The expert panel (N = 30) was senior but not statistically representative; public-facts scenarios cannot fully capture incentives, politics, or data frictions; and single-analyst case scoring motivates inter-rater checks in future work. Broader use beyond Western settings will require deeper localization, and both AI and regulatory landscapes will evolve—EDMC should evolve with them. These limits do not diminish the contribution; they define scope and set a clear agenda for field tests (Lynham, 2002; Yin, 2018).

Impact also depends on human and cultural preconditions. EDMC's PASO-DNA works best where leaders sponsor the effort, psychological safety enables escalation, and coaching supports brief reviews. In organizations with compliance-only habits or low trust, adoption can stall or devolve into paperwork. The transition therefore requires cultural alignment—visible sponsorship, role-modeling, and short learning loops—not only procedural updates (Edmondson, 2018; Kotter, 1996; Schein, 2010). Sections 6.3–6.4 outline practical roles (Sponsor/Steward) and a light review cadence to support that shift.

6.6 Conclusions

When time is short and stakes are high, aspiration alone is not enough. EDMC's PASO-DNA—Principles (Why), Actions (How), Skills (Who), Outcomes (What)—functions as an organizational Compass rather than a checklist. It is doctrine-agnostic about which principles a community adopts, yet uncompromising that whatever is named must be enacted, owned, and evidenced. Used as a first-tier filter for material choices, the Compass orients attention when it matters most, links intent to safeguards and stewardship, and leaves a transparent trace that can be examined, questioned, and improved.

EDMC does not replace legal regimes or domain expertise; it prepares and disciplines them. Nor does it erase disagreement; it structures how disagreement is surfaced, reasoned, decided, and recorded. Its distinctive features—engineered salience cues, go/stop gates, named stewardship, and outcome thresholds—address ethical blindness under speed and cognitive load while leaving an auditable trail (Bazerman and Tenbrunsel, 2011; Palazzo et al., 2012). Organizations can adopt EDMC now as the first pass for important decisions using clear roles and measurable standards; its modular form

allows straightforward digital embedding, and the resulting records serve as credible evidence of ethical diligence for boards, audit, and regulators.

Cultural architecture: repetition builds capacity. Because the Compass works through repetition, the same moves recur: state the principle, specify the action, assign a skilled owner, define the check, it does more than guide single decisions; it builds ethical muscle memory. Each pass makes salience easier to notice, strengthens the expectation that ownership implies capability, and normalizes brief reviews that connect outcomes back to values. Over time, the architecture becomes living DNA—an organizational habit of seeing clearly, acting responsibly, and learning in public. Consistent use builds shared habits and psychological safety, turning resistance into learn-by-doing momentum across teams (Edmondson, 2018; Kotter, 1996; Schein, 2010). The habit is also personal: leaders and teams develop the skills ethical judgment demands, listening across differences, reading context under pressure, weighing trade-offs without losing sight of first principles, and holding the tension between what’s expedient and what’s right.

Building on this culture-by-design, the EDMC framework offers distinct contributions:

- Doctrine-agnostic structure. EDMC does not prescribe what to value—it insists that whatever is valued be made explicit, actionable, and accountable. This design enables use in multinational firms, the public sector, cross-cultural collaborations, and AI governance settings operating under different regulatory expectations, without flattening distinct ethical commitments. Many frameworks offered these elements in fragments; EDMC binds them—simply, visibly, and at the moment of choice.

- Transparency as infrastructure. By requiring that principles, actions, skills, and outcomes be documented and traceable, EDMC makes ethical decision-making auditable—not punitively, but for learning. When choices can be examined, patterns

emerge; when patterns are visible, improvement becomes possible (Palazzo et al., 2012). Transparency is not exposure; it is the infrastructure of trust—producing credible evidence for boards, audit, and regulators.

- Capability before accountability. Accountability without capability is theater. Before asking “Did you act ethically?” EDMC asks “Do you have the skills to recognize ethical complexity, the language to name it, and the organizational support to act on it?” By foregrounding skill development—especially listening, context-reading, and trade-off navigation—the framework treats ethical leadership as a learnable practice, not an innate virtue (Edmondson, 2018).

Looking ahead. Priorities include cumulative field evidence across sectors, deeper cross-cultural testing and localization, stronger reliability and measurement science, and lightweight digital supports that preserve principle-to-outcome traceability while reducing friction. Practitioners can move now: record a Compass entry for every material decision; ensure PASO is complete; use attentional cues and go/stop gates to counter blind spots; run brief outcome reviews that update practice; integrate with—rather than replace—legal, risk, and assurance functions; and scale with light delivery aids (summary strips, quick-reference cues, RACI prompts, KPI exemplars, jurisdiction hooks).

An invitation: a Compass for uncertain terrain. In an age where technologies move faster than wisdom, where AI makes decisions at scale before ethics can catch up, and where complexity outpaces our ability to predict consequences (Floridi and Cowls, 2019; Mittelstadt et al., 2016), we need compasses, not maps. Maps assume stable terrain; compasses work in fog. EDMC is a compass for ethical leadership in uncertainty—dynamic enough for shifting contexts, steady enough to hold true north, and simple enough to use every day by anyone willing to name what matters and act accordingly. It does not promise certainty; it promises direction and orientation. It does not eliminate

dilemmas; it makes them visible and navigable. It does not prescribe values; it demands that whatever values we claim, we enact, own, and evidence.

Repeated, practiced, embedded—this is how architecture becomes DNA, and how DNA becomes culture. In the end, ethical leadership is not about perfection; it is about direction and orientation—the willingness to check our heading, correct our course, learn, build skills and capability, and remain accountable to the principles we claim to serve. EDMC helps build cultures worth inhabiting by keeping orientation dynamic and the PASO-DNA agnostic and steady, turning principled intent into traceable action and shared learning at the speed of modern work.

APPENDIX A: RESEARCHER-LED CASE PACKETS (PUBLIC CASES) — ONE-LINE SCORING RATIONALES (C1–C8)

Scope. This appendix records the assessor’s ≤15-word rationales that accompany the 0–3 rubric scores for the eight publicly documented cases (C1–C8) used in the researcher-led applications.

Not the panel vignettes. Expert-panel hypothetical vignettes appear separately in Appendix D.

How to use. For each case and each metric, enter the 0–3 score and a one-line rationale. Computation rules, anchors, and weight presets are in Appendix B; case-level results are summarized in Table 27–28 in the main text. Public source notes for C1–C8 are listed in Appendix O (Public-Facts Case Source Notes).

Scale anchors (brief). 0 = not evident · 1 = minimal · 2 = mostly · 3 = fully (see Appendix B, Table B1).

Terminology. PASO uses Skills; any “Capabilities” wording refers to the same construct.

Formatting. Begin each rationale with the score in square brackets, e.g., [2] Principles link to actions; skills named but thin.

A.1 Case C1 — Shein hazardous chemicals (EU REACH)

A.1.1 One-line rationale per metric (0–3)

PASO linkage: [] ...

Blindness mitigation: [] ...

Cultural adaptability: [] ...

Time-pressure suitability: [] ...

Implementation guidance: [] ...

Outcomes evaluability: [] ...

Composite (0–3 / 0–100): ____ / ____ Weight preset: Equal Stakeholder Time-pressure

A.1.2 Conditions to proceed (if any)

[≤3 bullets, thresholded safeguards tied to the case]

A.1.3 Assessor notes (optional)

[1–2 bullets] Initials/date: ____ / ____

A.2 Case C2 — Shell Jackdaw gas-field protests

A.2.1 One-line rationale per metric

PASO linkage: [] ...

Blindness mitigation: [] ...

Cultural adaptability: [] ...

Time-pressure suitability: [] ...

Implementation guidance: [] ...

Outcomes evaluability: [] ...

Composite: ____ / ____ Weight preset: Equal Stakeholder Time-pressure

A.2.2 Conditions to proceed

[≤3 bullets]

A.2.3 Assessor notes

[] Initials/date: ____ / ____

A.3 Case C3 — Foxconn labor violations

A.3.1 One-line rationale per metric

PASO linkage: [] ...

Blindness mitigation: [] ...

Cultural adaptability: [] ...

Time-pressure suitability: [] ...

Implementation guidance: [] ...

Outcomes evaluability: [] ...

Composite: ____ / ____ Weight preset: Equal Stakeholder Time-pressure

A.3.2 Conditions to proceed

[≤3 bullets]

A.3.3 Assessor notes

[] Initials/date: ____ / ____

A.4 Case C4 — Starbucks union violations

A.4.1 One-line rationale per metric

PASO linkage: [] ...

Blindness mitigation: [] ...

Cultural adaptability: [] ...

Time-pressure suitability: [] ...

Implementation guidance: [] ...

Outcomes evaluability: [] ...

Composite: ____ / ____ Weight preset: Equal Stakeholder Time-pressure

A.4.2 Conditions to proceed

[≤3 bullets]

A.4.3 Assessor notes

[] Initials/date: ____ / ____

A.5 Case C5 — Google gender discrimination settlement

A.5.1 One-line rationale per metric

PASO linkage: [] ...

Blindness mitigation: [] ...

Cultural adaptability: [] ...

Time-pressure suitability: [] ...

Implementation guidance: [] ...

Outcomes evaluability: [] ...

Composite: ____ / ____ Weight preset: Equal Stakeholder Time-pressure

A.5.2 Conditions to proceed

[≤3 bullets]

A.5.3 Assessor notes

[] Initials/date: ____ / ____

A.6 Case C6 — Tech sector age-bias trend

A.6.1 One-line rationale per metric

PASO linkage: [] ...

Blindness mitigation: [] ...

Cultural adaptability: [] ...

Time-pressure suitability: [] ...

Implementation guidance: [] ...

Outcomes evaluability: [] ...

Composite: ____ / ____ Weight preset: Equal Stakeholder Time-pressure

A.6.2 Conditions to proceed

[≤3 bullets]

A.6.3 Assessor notes

[] Initials/date: ____ / ____

A.7 Case C7 — Amazon AI hiring tool bias

A.7.1 One-line rationale per metric

PASO linkage: [] ...

Blindness mitigation: [] ...

Cultural adaptability: [] ...

Time-pressure suitability: [] ...

Implementation guidance: [] ...

Outcomes evaluability: [] ...

Composite: ____ / ____ Weight preset: Equal Stakeholder Time-pressure

A.7.2 Conditions to proceed

[≤3 bullets, e.g., Proceed only if selection-rate parity ≥ 0.80, appeal path live]

A.7.3 Assessor notes

[] Initials/date: ____ / ____

A.8 Case C8 — OpenAI leadership and governance gap

A.8.1 One-line rationale per metric

PASO linkage: [] ...

Blindness mitigation: [] ...

Cultural adaptability: [] ...

Time-pressure suitability: [] ...

Implementation guidance: [] ...

Outcomes evaluability: [] ...

Composite: ____ / ____ Weight preset: Equal Stakeholder Time-pressure

A.8.2 Conditions to proceed

[≤3 bullets]

A.8.3 Assessor notes

[] Initials/date: ____ / ____

APPENDIX B: SCORING RUBRIC, WEIGHTS, AND COMPUTATION (0–3 anchors)

This appendix contains everything needed to score the researcher-led case applications: the EDMC rubric anchors (0–3), the composite/weight formulas actually used in Chapter 4, a blank worked-example sheet (optional), and a brief reliability note. Expert-panel outputs (Likert/CVI, heatmaps, global items) are reported separately in Appendices G (tables) and H (figures).

Contents and navigation

- B.0 How to score (at a glance)
- B.1 Rubric anchors (0–3)
- B.2 Where to find the researcher-led case results (pointer only; no duplication)
- B.3 Composite formulas and weight presets used in Section 4.3
- B.4 Reliability/validity snapshot (not computed in this phase)
- B.5 How the rubric was used in the eight-case study

Terminology note. PASO uses “Skills.” If any earlier material says “Capabilities,” read it as “Skills.”

B.0 How to score (at a glance)

Score each metric 0–3 using Table B1. Allow N/A only when a metric truly does not apply; exclude N/A from denominators when averaging. Report means to 2 decimals and disclose the effective N (N_{eff} = after excluding N/A). Default to Equal weights unless Section B.3 specifies a preset.

Reporting footer (whenever you summarize): version = v1.0 (Appendix N); N_{eff} disclosed; weights = Equal unless stated.

Use the blank scoring template in Appendix B (Table 35) and copy/paste it into any case packet in Appendix A when needed.

B.1 Rubric anchors (0–3)

Table 34 Appendix B1: EDMC Effectiveness Metrics—0–3 Anchors (Overview)

Source: Author (2025). Original table created from study data and materials.

Metric	0 = Not evident	1 = Minimally evidenced	2 = Adequately evidenced	3 = Clearly and consistently evidenced
PASO linkage (Principles→Actions→Skills→Outcomes)	Elements missing or disconnected; no traceability	Some links; gaps remain; partial traceability	All four present; mostly coherent and traceable	Fully coherent chain; explicit traceability and rationale
Blindness mitigation (under pressure)	No prompts/guards for salience; risks unflagged	Generic prompts; limited salience cues	Case-specific prompts; common blind spots surfaced	Targeted prompts with red flags; salience reliably triggered
Cultural adaptability (hypernorm + local fit)	Universalism/relativism only; no adaptation path	Mentions culture but no mechanism	Hypernorms stated; basic adaptation steps	Clear hypernorms plus concrete local adaptation mechanism
Time-pressure suitability (usable fast)	Assumes long deliberation; not usable quickly	Some simplification; unclear under time limits	Abbreviated flow possible; essentials covered	Time-aware cues; rapid path preserves ethical checks
Implementation guidance (how to do it)	Vague principles; no steps/roles/artifacts	Steps or roles present, not both	Steps, roles, artifacts named; feasible	Clear steps, owners, artifacts; practice-ready
Outcomes evaluability (KPIs, conditions, feedback)	No measures/conditions	Measures or conditions, not both	KPIs and conditions present; basic loop	KPIs, go/stop conditions, and feedback loop integrated

Computation note (applies throughout B): score each metric 0–3; allow N/A only when truly inapplicable; exclude N/A from means; round to 2 d.p.; default Equal weights.

Alignment to PASO and gap dimensions (orientation):

1. PASO linkage; 2) Blindness mitigation; 3) Cultural adaptability; 4) Time-pressure suitability; 5) Implementation guidance; 6) Outcomes evaluability.

These six metrics align to PASO and informed the survey items in Appendix C; the expert panel, however, used Likert/CVI (not the 0–3 rubric).

B.2 Where to find the researcher-led case results (pointer only; no duplication)

Case composites (Equal, Stakeholder, Time-pressure) for C1–C8 are in the main text: Table 27 (scores/composites) and Table 28 (deltas vs. Equal). One-line rationales per metric are in Appendix A. To avoid duplication, no numbers are repeated here.

B.3 Composite formulas and weight presets used in Section 4.3

Reporting rules

- Per-metric mean (0–3): average of available rater scores for that metric; exclude N/A; round 2 d.p.
- Per-rater composite (per case): mean of that rater’s non-N/A metric scores.
- Case composite (0–3): mean of the six per-metric means (or a weighted mean if weights are specified).
- Rescale to 0–100: $(\text{composite}_{0-3} \div 3) \times 100$.
- Always disclose N_{eff} (effective N after excluding N/A).

Explicit formulas

- Per-metric mean (0–3) = average of available rater scores for that metric (drop N/A).
- Case composite (0–3, Equal) = average of the six per-metric means across included metrics.
- Weighted composite (0–3) = $\Sigma (w_m \times \text{per-metric mean}_m)$, summing only over included metrics; if any metric is N/A, renormalize remaining weights to sum to 1.
- Rescale (0–100) = $(\text{composite}_{0-3} \div 3) \times 100$.

B.3.1 Weight presets (used in Section 4.3)

Order of metrics = PASO, Blindness, Cultural, Time, Implementation, Outcomes.

- Equal (neutral baseline): 1/6 each.
- Stakeholder emphasis (↑ Cultural, Outcomes): 0.20, 0.10, 0.20, 0.10, 0.15, 0.25.
- Time-pressure emphasis (↑ Time, Blindness, Implementation): 0.10, 0.25, 0.10, 0.30, 0.15, 0.10.

Computation note: If any metric is N/A, drop it and renormalize the remaining weights to sum to 1 before computing the weighted mean.

Table 35 Appendix B2: Worked Example–Per-Metric Scoring Sheet (Blank Template)
Source: Author (2025). Original table created from study data and materials.

Metric	Anchor 0 — Not evident	Anchor 1 — Emerging	Anchor 2 — Adequate	Anchor 3 — Strong	R1 (0–3/N/A)	R2 (0–3/N/A)	R3 (0–3/N/A)	Mean (0–3)*	Notes (≤15 words)
PASO linkage	Unlinked; not traceable	Links unclear	Clear chain; minor gaps	Explicit, auditable chain					
Blindness mitigation	No prompts	Generic statements	Specific prompts; some cues	Red flags; pre-mortem; works under pressure					
Cultural adaptability	No mechanism	Mentions culture	Hypernorms + basic steps	Clear guide; deviations documented					
Time-pressure suitability	Slow-only	Time limits named	Abbreviated steps + heuristics	Built-in prompts; rapid path					
Implementation guidance	No steps/owners	High-level steps	Steps, owners, artifacts	Full workflow + templates					
Outcomes evaluability	No KPIs	Outcomes named only	KPIs + process checks	KPIs, thresholds, audit trail					

*Mean excludes N/A; round to 2 d.p.

B.4 Reliability/validity snapshot (not computed in this phase)

Researcher-led tallies used a single assessor; therefore, inter-rater statistics (e.g., ICC, Kendall's W) were not computed. Content validity indices (I-CVI / S-CVI/Ave) pertain to the expert-panel instrument and are reported in Appendix G2 (table) and Appendix H2 (figure); they are not recomputed here.

Future research (brief):

- Multi-rater replication with ICC(2,k) and/or Kendall's W on the 0–3 rubric.
- Short test–retest to gauge scoring stability.
- Field settings to extend validity beyond public-facts cases.

B.5 How the rubric was used in the 8-case study

The 0–3 rubric in Table B1, the composite rules in Section B.3, and the disclosed weight presets in Section B.3.1 are the exact procedures used to score the eight publicly documented cases reported in Chapter 4. For each case, the assessor recorded a 0–3 score per metric with a ≤ 15 -word rationale (Appendix A), computed per-metric means (excluding any N/A), produced a case composite (Equal weights by default), and, where shown, recomputed composites under the Stakeholder and Time-pressure presets. For plotting, composites were optionally rescaled to 0–100. Expert-panel topics overlap with three domains, but the panel used Likert/CVI (see Appendices G–H), not the 0–3 rubric.

APPENDIX C: EXPERT-PANEL PROTOCOL AND SURVEY INSTRUMENT (INDEX AND VERBATIM)

[Expert Panel] This appendix provides the panel instrument structure (index) and the verbatim, version-locked PDF used in the study. (Sections C.1–C.5; C.5 is appended as an archival PDF.)

Contents and navigation

C.1 Instrument overview (compact)

C.2 Item index (construct map)

C.3 PASO micro-block wording (concise)

C.4 Survey flow (screens/sections) + versioning/governance

C.5 Verbatim instrument (appended PDF)

C.1 Instrument overview (compact)

The expert-panel instrument evaluates EDMC’s clarity, usability under time pressure, portability, and content validity (PASO). It collects: baseline stance/confidence per case; post-EDMC ratings (Q1–Q8), a PASO micro-block (P/A/S/O), post-stance/confidence, a single efficiency item, and a global wrap-up. Items map to the six metrics in Section 3.2; instrumentation details are in Section 3.7. Verbatim wording is provided in Appendix C.5 (appended PDF).

Case selection note. To minimize prior-knowledge bias and keep the evaluation focused on EDMC itself, the panel used hypothetical, broadly structured vignettes (see Section 3.8.1 for the methodological rationale and citations). Researcher-led applications using public cases are separate and reported in Chapter 4.

C.2 Item index (construct map)

This index lists the instrument structure; verbatim text appears in C.5.

Intro and consent (see C.5)

Participant code (NDA-XX; verification only; removed at anonymization)

Demographics (privacy-safe bins)

4–9. Case 1

Baseline: stance (+ conditions if any), confidence (0–10), brief rationale

After EDMC: Q1–Q8 (1–5 Likert), PASO P/A/S/O (1–5), post-stance, post-confidence (0–10), brief comment (optional), efficiency item

10–15. Case 2 — same fields

16–21. Case 3 — same fields

22–28. Global wrap-up — overall effectiveness; sufficiency count; learning/efficiency (Case 1→3); portability; integratability; time-pressure usability; conceptual coherence; optional tech/AI; limits; one improvement.

Construct alignments (examples).

Time-pressure suitability → Q4; Defensibility/Integratability → Q5, Q7; Portability → Q6; Actionability (PASO linkage + implementation) → Q1 + PASO block; Intention to use → Q8; Outcomes evaluability → PASO “O” + “conditions to proceed”; Judgment effects → stance change, confidence uplift, efficiency.

C.3 PASO micro-block wording (concise)

Terminology note. PASO uses Skills; any “Capabilities” wording should be read as the same construct.

P — Principles: Core ethical principles are explicit and, where relevant, norms/policies/standards are referenced.

A — Actions: Concrete actions with responsible roles and needed artifacts (docs/systems) are stated and feasible.

S — Skills (Capabilities): Required skills/resources/ownership to execute the actions are identified (RACI hooks).

O — Outcomes: Intended outcomes, KPIs, thresholds/mitigations, and review cadence are specified for auditability.

C.4 Survey flow (screens/sections)

Fixed order under realistic time pressure (~15–20 minutes total):

Baseline (per case) → 2) View EDMC case note → 3) Post-EDMC items (Q1–Q8 + PASO + post stance/confidence + efficiency) × three cases → 4) Global wrap-up.

Versioning and governance. Instrument and cases were version-locked; the fieldwork window and minor link fix are documented in Appendix N. Consent, anonymity, and data handling follow Appendices K–L.

C.5 Verbatim instrument (appended PDF)

Instrument tag: EDMC Expert Evaluation – Baseline EDMC + PASO CVI (v1.0)

Version-lock: 23 Sep 2025, 12:30 CEST (Europe/Zurich)

Cases: AI Hiring Bias; Supplier Labor Compliance; Data Privacy Breach (Case-3 link fixed 23 Sep 2025, 12:15 CEST; no wording changes).

Legend (printed above the first PASO micro-block in the PDF):

PASO = Principles (why) · Actions (how) · Skills (who) · Outcomes (what).

Appendix C.5 — Verbatim Survey Instrument (PDF, appended at end of thesis)

To preserve original layout and pagination, the full verbatim survey is included as a separate PDF appended as the final section of the compiled thesis PDF. Instrument tag: “EDMC Expert Evaluation – Baseline EDMC + PASO CVI (v1.0)”. This pointer page is included so Appendix C.5 appears in the Table of Contents; the verbatim pages follow after the last appendix.

APPENDIX D: EXPERT-PANEL CASE NOTES (HYPOTHETICAL VIGNETTES)

Terminology note. In this thesis, the concise EDMC decision record is implemented as a case note for study purposes. The terms are used interchangeably.

Scope. These are the three one-page, anonymized vignettes shown to the expert panel for rating EDMC's clarity, usability under time pressure, and PASO coverage. They are hypothetical by design to avoid bias from prior knowledge. Researcher-led public cases are reported separately (Appendix A).

Format used with panel. Brief context; domains; recommended actions; outcomes to monitor (with example thresholds); traceability & go/stop conditions; time-aware note. No real organizations are named.

D.1 AI Hiring Bias — ML resume screening (hypothetical)

Context (2–3 lines). A resume screener is slated for campus season. Early pilots suggest adverse impact across gender in two roles. Product wants to launch; HR is concerned about parity and transparency.

Domains. Fairness • Non-discrimination • Transparency • Due process • Accountability

Recommended actions (5–8).

Run pre-launch bias audit on training and recent inference outputs; document adverse-impact ratio.

Add human-in-the-loop review for borderline scores; publish candidate-facing appeal path.

Narrow features to job-related signals; remove proxies with known bias leakage.

Create rollback plan and kill-switch criteria if parity thresholds fail in production.

Brief hiring managers; issue guidance on appropriate use/limits.

Outcomes to monitor (examples).

Selection-rate parity ≥ 0.80 across gender and major groups; trend weekly.

Appeal volume and reversal rate; time-to-decision vs. baseline.

Drift alerts on feature distributions; retrain cadence met.

Traceability & conditions.

Proceed only if pre-launch parity ≥ 0.80 and appeal path is live.

Owners: ML Lead (model), HR/Talent Ops (process), Compliance/Ethics (oversight).

Review gates: T+14d, T+60d parity check; escalate if breach persists.

Time-aware note. Summary strip and quick checklist support a go / proceed-with-conditions / pause decision in ≤ 30 minutes.

D.2 Supplier Labor Compliance — multi-tier risk (hypothetical)

Context (2–3 lines). A tier-2 supplier shows red flags on hours and documentation during rapid expansion. Orders are time-critical; remediation options exist but require joint action.

Domains. Human rights • Non-exploitation • Duty of care • Transparency • Corrective action

Recommended actions (5–8).

Immediate on-site verification via independent auditor; protect worker anonymity.

Issue corrective action plan (CAP) with timelines and incentives/penalties.

Map upstream/downstream tiers; freeze expansion to flagged line items.

Establish worker grievance channel and hotline access.

Engage peer buyers for an aligned CAP to avoid audit churn.

Outcomes to monitor (examples).

CAP closure on time; verified reduction in excessive hours; proper pay/records.

Grievance volume/resolution time; recurrence rate post-closure.

Re-audit pass rate within 90 days.

Traceability & conditions.

Proceed-with-conditions only if CAP is signed and independent verification is scheduled.

Owners: Procurement Lead, ESG/Compliance, Supplier GM.

Review gates: 30/60/90-day checkpoints; escalate to suspend if misses.

Time-aware note. Decision framed as continue with safeguards vs. pause purchases pending verification.

D.3 Data Privacy Breach — SaaS cross-tenant exposure (hypothetical)

Context (2–3 lines). A config flaw allowed cross-tenant viewing of limited profile fields (~36h window). Hotfix applied; scope under investigation; EU/US jurisdictions in play.

Domains. Privacy • Harm minimization • Accountability • Transparency • Due process

Recommended actions (5–8).

Contain: disable/patch endpoint; revoke tokens; rotate credentials; add temporary WAF rule.

Validate depth: reproduce exploit safely; test for adjacent misconfigurations; add CI guardrails.

Preserve evidence; run scope analysis; maintain chain-of-custody.

Regulatory path: assess GDPR Art. 33; notify lead DPA within 72h if criteria met; track US triggers.

Prepare clear user notices and support channel; phishing-risk guidance.

Define rollback/kill-switch if re-exposure indicators appear.

Outcomes to monitor (examples).

Time-to-containment $\leq 2h$; zero residual exposure after fix.

On-time regulator notices; 100% affected users notified ≤ 7 days.

Zero repeat cross-tenant events over 30 days; MTTA/MTTR improving.

Traceability & conditions.

Stay-live-with-conditions only if monitoring is active and rollback criteria are set; else regional pause.

Owners: CISO, DPO, Legal, Product.

Review gates: T+24h, T+72h, T+14d status with updated scope/impact.

Time-aware note. Prompts surface essential checks regardless of time budget; posture is stay-live-with-conditions vs. pause.

APPENDIX E: LITERATURE SYNTHESIS TABLES

This appendix closes the loop between the Chapter 2 review (Table 37) and the Chapter 4 results by situating EDMC’s empirically evidenced features alongside classical, structured, and contemporary applied approaches.

Table 36 PASO Mapping Across Reviewed Frameworks (P/A/S/O)
Source: Author (2025). Original table created from study data and materials.

Framework	Principles (P)	Actions (A)	Skills (S)	Outcomes (O)
Rest’s Four-Component Model (1986)	Moral sensitivity, judgment, motivation, character (Rest, 1986)	Promote moral awareness and stepwise ethical decision-making in education/organizations (Rest, 1986)	Moral reasoning; ethical judgment (Rest, 1986)	Ethical awareness; improved ethical reasoning (Rest, 1986)
Treviño’s Interactionist Model (1986)	Person–situation interaction; ethical climate and context (Treviño, 1986)	Integrate ethical awareness into corporate decision processes (Kaptein, 2008)	Ethical climate assessment; decision-making under pressure (Kaptein, 2008)	Stronger ethical culture; embedded decision practices (Kaptein, 2008)
Ferrell and Gresham’s Contingency Model (1985)	Personal norms, corporate ethics, external contingencies (Ferrell and Gresham, 1985)	Balance ethical considerations in marketing/business contexts (Jones, 1991)	Marketing ethics; stakeholder engagement (Ferrell and Gresham, 1985)	More balanced decisions in marketing/business (Jones, 1991)
Velasquez’s Business Ethics Model (1982)	Stakeholder accountability; fairness; responsibility (Velasquez, 1982)	Align business strategy with stakeholder obligations (Donaldson and Dunfee, 1999)	Stakeholder analysis; ethical strategy design (Velasquez, 1982)	Corporate responsibility; improved stakeholder relations (Velasquez, 1982)
AI and Digital Ethics (Floridi, 2013; 2020)	Transparency; fairness; accountability in digital ecosystems (Floridi, 2013; Mittelstadt et al., 2016)	Advocate ethical AI development and governance practices (Floridi, 2013)	AI ethics assessment; compliance literacy (Floridi, 2020)	Ethical AI governance; reduced bias in AI decisions (Floridi, 2020)
Virtue ethics (Hursthouse, 1999; MacIntyre, 1984)	Integrity; honesty; character (Hursthouse, 1999)	Encourage character formation in leadership/professional ethics (Kristjánsson, 2020)	Emotional intelligence; character development (Hursthouse, 1999)	Character-driven decision-making; moral leadership (Kristjánsson, 2020)
Deontological ethics (Kant, 1785)	Duty-based obligations; universal principles (Kant, 1785; Alexander and Moore, 2016)	Establish clear duties in law/healthcare applications (Alexander and Moore, 2016)	Duty-based ethical reasoning; critical analysis (Alexander and Moore, 2016)	Rights protection; enhanced policy/legal compliance (Alexander and Moore, 2016)
Utilitarian ethics (Bentham, 1789; Mill, 1863)	Maximization of overall benefit (Bentham, 1789; Mill, 1863)	Apply cost–benefit analysis in policy-making (Driver, 2012)	Forecasting; evaluative balancing (Bentham, 1789; Mill, 1863)	More efficient resource allocation; decision rationality (Driver, 2012)
Stakeholder theory (Freeman, 1984)	Balancing corporate responsibility with stakeholder interests (Freeman, 1984)	Encourage CSR and sustainable business practices (Freeman, 1984)	Conflict resolution; corporate governance competence (Freeman, 1984)	Ethical corporate strategies; increased stakeholder trust (Freeman, 1984)
Ethical leadership in public administration (Denhardt and Denhardt, 2003)	Transparency; accountability; stewardship (Denhardt and Denhardt, 2003)	Implement ethical governance in public institutions (Denhardt and Denhardt, 2003)	Cultural competence; policy ethics (Denhardt and Denhardt, 2003)	Improved public trust and accountability (Denhardt and Denhardt, 2003)

Transformational leadership (Bass, 1985)	Values-led, vision-driven ethical change (Bass, 1985)	Lead ethical transformation via vision and example (Bass, 1985)	Transformational influence; ethical leadership (Bass, 1985)	Stronger ethical leadership; organization-wide change (Bass, 1985)
Servant leadership (Greenleaf, 1977)	Service to others; humility; stakeholder well-being (Greenleaf, 1977)	Encourage ethical leadership and employee well-being (Van Dierendonck, 2011)	Servant leadership competencies; ethical influence (Greenleaf, 1977)	Employee well-being; enhanced organizational ethics (Van Dierendonck, 2011)
Triple Bottom Line (Elkington, 1997)	Sustainability pillars: people, planet, profit (Elkington, 1997)	Drive sustainability strategy and reporting (Elkington, 1997)	Sustainability strategy; metrics literacy (Elkington, 1997)	Sustainable growth; improved environmental/social impact (Elkington, 1997)
General Data Protection Regulation (GDPR, 2016)	Rights-based data protection; privacy by design (GDPR, 2016)	Ensure compliance via documentation, DPIAs, controls (Floridi, 2013)	Data protection competence; regulatory compliance (GDPR, 2016)	Demonstrable compliance; auditable accountability (GDPR, 2016)
Big Data and Smart Cities ethics (Kitchin, 2014)	Ethics of datafication, governance, and public space (Kitchin, 2014)	Develop governance for AI/smart-city projects (Zuboff, 2019)	Data governance; bias mitigation practices (Kitchin, 2014)	Ethical AI/smart-city implementation (Kitchin, 2014)
Hunt and Vitell's Theory of Marketing Ethics (1986)	Personal moral codes; cultural/situational factors (Hunt and Vitell, 1986)	Explain marketing decision processes and evaluation (Hunt and Vitell, 1986)	Marketing ethics judgement; cross-cultural reasoning (Hunt and Vitell, 1986)	Greater consumer trust; more ethical marketing strategies (Hunt and Vitell, 1986)

Note: PASO anchors = Principles (P), Actions (A), Skills (S), Outcomes (O).

Table 37 Comparative Analysis of Ethical Decision-Making Frameworks (Across Six Dimensions)

Source: Author (2025). Original table created from study data and materials.

Framework	Philosophical foundation	Implementation guidance	Cultural inclusivity	Adaptability under time constraints	Competency requirements	Evaluation mechanisms
Deontological ethics (Kant, 1785)	Strong universal principles	Limited practical application	Western-centric	Low adaptability under pressure	Assumes sophisticated reasoning	Limited outcome assessment
Utilitarian ethics (Bentham, 1789; Mill, 1863)	Strong consequentialist basis	Calculation-focused approach	Culturally limited perspective	Moderate adaptability	Quantitative assessment skills required	Outcome-focused evaluation
Virtue ethics (Hursthouse, 1999; MacIntyre, 1984)	Character-centered approach	Limited decision processes	Variable cultural application	Low time-pressure adaptation	Emphasis on character development	Limited structured assessment
Stakeholder theory (Freeman, 1984)	Pluralistic responsibilities	Stakeholder-analysis process	Moderate cultural adaptation	Moderate time adaptation	Stakeholder engagement skills	Multi-perspective evaluation
Rest's Four-Component Model (1986)	Cognitive moral development	Structured decision stages	Limited cultural adaptation	Low time-pressure adaptation	Cognitive/ethical judgement skills	Process-oriented assessment
Treviño's Interactionist Model (1986)	Person-situation integration	Contextual decision process	Moderate cultural sensitivity	Moderate time adaptation	Individual + situational awareness	Limited outcome measures
Ferrell and Gresham's Contingency Model (1985)	Contingency-based foundation	Moderate implementation structure	Limited cultural adaptation	Moderate time adaptation	Org + individual factor analysis	Multi-level evaluation

Velasquez's Business Ethics Model (1982)	Integrative business ethics	Balanced decision process	Moderate cultural sensitivity	Moderate time adaptation	Ethical reasoning in strategy	Comprehensive assessment approach
Hunt and Vitell's Theory of Marketing Ethics (1986)	Deontological + teleological synthesis	Theory-guided process	Limited cultural adaptation	Low time-pressure adaptation	Normative evaluation skills	Dual-evaluation methodology
Transformational leadership (Bass, 1985)	Values/vision orientation	Influence-based guidance	Variable cultural application	Moderate crisis adaptation	Inspirational leadership skills	Vision-alignment assessment
Servant leadership (Greenleaf, 1977)	Service-centered foundation	Community-focused guidance	Moderate cultural adaptability	Moderate time adaptation	Service-orientation competencies	Stakeholder-benefit assessment
Ethical leadership in public administration (Denhardt and Denhardt, 2003)	Public-interest orientation	Institutional process focus	Contextual adaptability	Moderate time adaptation	Public-service competencies	Accountability-based evaluation
Triple Bottom Line (Elkington, 1997)	Multi-value integration	Sustainability-focused process	Moderate global applicability	Moderate time adaptation	Multi-dimensional analysis skills	Comprehensive metrics approach
AI and Digital Ethics (Floridi, 2013; 2020)	Information-ethics foundation	Technology-focused guidelines	Limited cultural adaptation	Variable time adaptation	Technical + ethical expertise	Technology-impact assessment
General Data Protection Regulation (GDPR, 2016)	Rights-based foundation	Detailed compliance processes	European-centered approach	Complex implementation timelines	Legal + technical expertise	Compliance-focused assessment
Big Data and Smart Cities ethics (Kitchin, 2014)	Urban/tech-ethics foundation	Context-specific guidelines	Variable urban inclusivity	Moderate time adaptation	Urban planning + tech-ethics skills	Urban-impact assessment

Table Note: Column guide: "Philosophical foundation" = underlying lens; "Implementation guidance" = explicit process/tools; "Cultural inclusivity" = portability across contexts; "Adaptability under time constraints" = performance in high-tempo decisions; "Competency requirements" = skills needed; "Evaluation mechanisms" = how outcomes are assessed.

Table 38 Comprehensive Review Across 16 Definitions, Components, Strengths, Gaps, and EDMC Relevance

Source: Author (2025). Original table created from study data and materials.

Framework	Definition	Key components	Application domains	Key strengths	Limitations and gaps	Typical challenges	Practical examples / metrics	Interdisciplinary applications	Relevance to EDMC
Rest's Four-Component Model (1986)	Ethical decision-making involves moral sensitivity, judgment, motivation, and character (Rest, 1986)	Sensitivity; judgment; motivation; character (Rest, 1986)	Psychology; education; organizational ethics (Rest, 1986)	Structured cognitive stages; clear individual focus (Sonenshein, 2007)	Under-emphasizes organizational influences (Sonenshein, 2007)	Limited scalability to complex organizational contexts (Sonenshein, 2007)	Moral sensitivity surveys; moral judgment tests (Rest, 1986)	Psychology; ethics education (Rest, 1986)	PASO: A,S. EDMC: adds context and trade-off syntheses to individual steps.

Treviño's Interactionist Model (1986)	Ethical decisions arise from interaction of individual traits and situational context (Treviño, 1986)	Person; situation; ethical climate (Treviño, 1986)	Organizational ethics; governance (Kaptein, 2008)	Integrates multiple influences (Kaptein, 2008)	Lack of standardized measurement (Kaptein, 2008)	Measuring situational impacts consistently (Kaptein, 2008)	Ethical climate surveys; ethics programs (Kaptein, 2008)	Business ethics; organizational governance (Kaptein, 2008)	PASO: S,A. EDMC: reconciles person-situation with principled constraints and outcomes.
Ferrell and Gresham's Contingency Model (1985)	Ethical decisions are contingent on individual and organizational factors (Ferrell and Gresham, 1985)	Personal norms; corporate ethics; external influences	Marketing ethics; business ethics; CSR (Jones, 1991)	Context analysis; environmental contingencies (Jones, 1991)	Limited depth of moral reasoning (Jones, 1991)	Balancing ethics under external pressure	CSR/marketing ethics audits	Marketing; corporate responsibility	PASO: A,O. EDMC: layers Principles and Skills onto contingency steps.
Velasquez's Business Ethics Model (1982)	Ethical business actions based on stakeholder accountability and principles (Velasquez, 1982)	Stakeholder accountability; fairness; responsibility	Corporate governance; CSR (Donaldson and Dunfee, 1999)	Stakeholder-centric orientation	Neglects individual moral cognition (Donaldson and Dunfee, 1999)	Aligning principles with practice at scale	Stakeholder impact analysis; CSR reporting	Business ethics; stakeholder management	PASO: P,A. EDMC: arbitrates stakeholder conflicts via multi-perspective synthesis.
AI and Digital Ethics (Floridi, 2013; 2020)	Addresses ethical implications of digital systems (fairness, transparency, accountability)	Transparency; fairness; accountability (Floridi, 2013; Mittelstadt et al., 2016)	AI governance; digital ethics	Contemporary guidance for AI ethics	Lack of standardized ethical metrics (Mittelstadt et al., 2016)	Quantifying/prioritizing principles in systems	GDPR alignment; AI ethics audits	Tech governance; info ethics	PASO: P,A. EDMC: connects principles to cross-stakeholder outcomes.
Virtue ethics (Hursthouse, 1999; MacIntyre, 1984)	Emphasizes character in ethical leadership and practice	Character development; phronesis	Education; healthcare; leadership development (Kristjánsson, 2020)	Motivation and moral tenor for action	Less concrete decision guidance (Kristjánsson, 2020)	Assessing character; scaling across cultures (Annas, 2011)	Ethics training; character surveys	Education; health; leadership	PASO: S,P. EDMC: adds stepwise Actions and auditability Outcomes.

Deontological ethics (Kant, 1785)	Duty-based rules and universal principles	Rights; duties; universalizability (Kant, 1785; Alexander and Moore, 2016)	Law; healthcare; human rights	Clear constraints and rights protection	Rigidity; conflicts of duty (McNaughton and Rawling, 2018; Ross, 1930)	Applying duties in complex scenarios	Compliance audits; policy alignment	Law; health; policy	PASO: P. EDMC: balances rights with consequences and context.
Utilitarianism (Bentham, 1789; Mill, 1863)	Chooses actions that maximize overall benefit	Utility calculus; consequences	Public health; policy; AI ethics (Driver, 2012)	Pragmatic, outcome-oriented evaluation	Risk of problematic trade-offs; quantification issues (Driver, 2012)	Valuing minority impacts; long-term effects	DALYs/cost-benefit analyzes	Health; economics; tech policy	PASO: O,A. EDMC: checks utility against rights and virtues.
Stakeholder theory (Freeman, 1984)	Balances responsibilities across stakeholder groups	Stakeholder mapping; salience	Governance; environmental policy; social enterprise	Broadens moral scope	Prioritization under conflict; limited metrics	Weighing competing claims across contexts	Stakeholder surveys; sustainability reports	Governance; CSR; entrepreneurship	PASO: O,A. EDMC: provides prioritization logic and documentation.
Ethical leadership in public administration (Denhardt and Denhardt, 2003)	Promotes transparency, accountability, stewardship	Public-interest orientation; stewardship	Anti-corruption; government transparency (Huberts, 2018)	Accountability norms; public trust	Bureaucratic resistance; digital governance gaps (Van Wart, 2013)	Sustaining reforms in complex institutions	Transparency indices; citizen surveys	Public management; policy	PASO: S,O. EDMC: adds procedural Actions and trade-off analysis.
Transformational leadership (Bass, 1985)	Inspires ethical change via vision and example	Vision; values alignment; influence	Business leadership; organizational psychology; education (Van Dierendonck, 2011)	Culture-shaping leverage	Equity/justice blind spots; sustainability under pressure (Eagly, 2007; Judge and Piccolo, n.d.)	Maintaining momentum; leader dependence	Engagement surveys; change metrics	Management; education	PASO: S,O. EDMC: distributes ethical agency and codifies steps.
Servant leadership (Greenleaf, 1977)	Prioritizes follower well-being and community	Service; humility; stewardship	Non-profits; community leadership; education (Van Dierendonck, 2011)	Strong ethical motivation; trust	Fewer governance/accountability mechanisms (Hunter et al., 2013)	Implementing in competitive settings	Employee well-being; community impact	Social sector; education; HR	PASO: S,O. EDMC: adds governance Actions and documentation.
Triple Bottom Line	Balances people,	Social; environmental;	Corporate sustainability	Multi-value framing	Standardization gaps; greenwashing	Prioritizing across	GRI metrics; B-Corp	Finance; operations; ESG	PASO: O,A. EDMC:

(Elkington, 1997)	planet, profit	economic pillars	bility reporting (Dyllick and Muff, 2016)		risk (Slaper and Hall, 2011)	pillars; verification	certification		clarifies trade-offs and evidence needs.
General Data Protection Regulation (GDPR, 2016)	Legal and ethical framework for data protection	Rights; lawful bases; DPIAs; accountability	Data privacy regulation (Floridi, 2013)	Enforceable safeguards; documentation	Legalistic focus over proactive ethics; resource intensity (Tikkinen-Piri et al., 2018)	Cross-jurisdiction implementation	Compliance rates; breach statistics	Tech law; compliance; governance	PASO: P.A.O. EDMC: balances rights with utility; records rationale.
Big Data and Smart Cities ethics (Kitchin, 2014)	Addresses AI/data ethics in urban governance	Data governance; bias/privacy; consent	Smart-city projects (Zuboff, 2019)	Context-specific guidance	Regulatory fragmentation; bias/privacy risks (Mökander et al., 2023)	Keeping pace with tech change	PIA; AI bias audits	Urban planning; public policy	PASO: A.O. EDMC: integrates stakeholder trade-offs with safeguards.
Hunt and Vitell's Theory of Marketing Ethics (1986)	Ethical decisions reflect personal codes and cultural/situational factors	Personal morals; cultural; situational inputs; evaluation	Marketing; consumer ethics; cross-cultural analysis	Recognizes cultural/societal influence; adaptable lens	Limited generalizability beyond marketing	Applying outside marketing; limited validation	Consumer behavior analyzes; ad ethics reviews	Business ethics; consumer psychology	PASO: P.O. EDMC: extends beyond discipline; structures steps and evidence.

Table Note. Abbreviations: GDPR = General Data Protection Regulation; DPIA = Data Protection Impact Assessment; PIA = Privacy Impact Assessment; DALY = Disability-Adjusted Life Year; GRI = Global Reporting Initiative; CSR = Corporate Social Responsibility; ESG = Environmental, Social, and Governance; B Corp = Benefit Corporation certification.

Table 39 Comparative Analysis — EDMC Versus Prior Ethical Decision-Making Frameworks (Closing the Loop)

Source: Author (2025). Original table created from study data and materials.

Feature / Dimension	Classical Theories (Deontology, Utilitarianism, Virtue, Stakeholder)	Structured Models (Rest; Treviño; Ferrell and Gresham, etc.)	Contemporary Applied (GDPR, AI Ethics, TBL, etc.)	EDMC (This Study)
Normative principles	✓ Strong, but abstract	● Present, often implicit	● Domain-specific, sometimes fragmented	✓ Integrated, portable, context-adaptable
Actionable steps / tools	✗ Lacking	✓ Stepwise, often slow	✓ Checklists, siloed	✓ Time-aware toolkit; auditable artifacts

Skill / ownership development	X Assumed, not developed	● Named, not operationalized	X Rarely addressed	✓ Explicit roles (RACI), ownership, escalation
Outcome measurement / KPIs	X Rarely specified	● Sometimes present	✓ Compliance-oriented, narrow	✓ KPIs, thresholds, learning loops
Cultural / contextual adaptability	● Universalist, limited adaptation	● Some context	● Jurisdictional, not global	✓ Portable: hypernorms + local hooks
Time-pressure functionality	X Not addressed	X Not addressed	X Not addressed	✓ Tiered, time-aware protocols
Blindness mitigation	X Not addressed	X Not addressed	X Not addressed	✓ Attentional cues; go/stop gates
Auditability / traceability	X Not auditable	● Sometimes traceable	✓ For compliance, not ethics	✓ Glass-box traceability end-to-end
Integration with governance	X Not designed for integration	● Sometimes fits	✓ Compliance fit	✓ Drop-in for risk, compliance, AI, ESG
Empirical validation	X Theoretical	● Some studies	● Policy-driven	✓ Dual validation: panel + case robustness

Legend: ✓ = explicitly addressed and strong; ● = partially addressed; X = not addressed.

APPENDIX F: ANALYSIS PLAN AND DERIVED METRICS

Scope. Defines the analyses, exclusion/missingness handling, and derived metrics for both evidence streams: the expert-panel survey (Likert/CVI, stance/confidence/efficiency) and the researcher-led case composites (0–3 rubric). Computation/conventions here match Chapter 4 and the version-locked artifacts in Appendices M–N.

F.1 Purpose

Specify the quantitative and qualitative analyses, denominator rules, and reproducibility conventions used in reporting.

F.2 Quantitative indicators (expert-panel strand): Report whole-percent ≥ 4 ; means to 2 d.p.; always show N (effective N where applicable).

≥ 4 agreement (Likert 1–5): share of responses rated 4 or 5. Report the effective N used; exclude N/A from denominators.

Stance-change rate: distribution of Proceed / Proceed with conditions / Pause; report the share with Post \neq Baseline.

Confidence uplift: mean(PostConfidence – BaselineConfidence) on a 0–10 scale.

Efficiency shift: proportions across Significantly less / Moderately less / About the same / More / Uncertain.

PASO CVI: I-CVI per PASO item (percent ≥ 4); S-CVI/Ave = mean of the four I-CVIs (P/A/S/O).

N/A handling: exclude N/A values from denominators; disclose N for every statistic.

F.3 Rubric composites (researcher-led strand; see Appendix B)

Per-metric mean (0–3): average of assessor scores for that metric; exclude N/A; round to 2 d.p.

Per-rater composite (0–3): mean of that rater's non-N/A metric scores for a case.

Case composite (0–3, equal): mean of the six per-metric means (equal weights by default).

Rescale to 0–100: $(\text{composite}_{0-3} \div 3) \times 100$.

Weights/presets: Equal / Stakeholder / Time-pressure presets as defined in Appendix B; if any metric is N/A, renormalize remaining weights to sum to 1.

F.3a Composite computation (detail; applies to C1–C8)

- Equal-weight (all six metrics present).

Composite (0–100) = $((m_1 + m_2 + m_3 + m_4 + m_5 + m_6) \div 18) \times 100$.

(Each metric m_i is 0–3; $6 \times 3 = 18$.)

- Equal-weight with N/A (K included metrics).

Composite (0–100) = $((\Sigma \text{ included metrics}) \div (3 \times K)) \times 100$.

(Exclude truly inapplicable metrics from the denominator.)

- Weighted presets (order = PASO, Blindness, Cultural, Time, Implementation, Outcomes).

- 1) Compute the weighted 0–3 mean over the included metrics:

Composite (0–3, weighted) = $\Sigma (w_i \times m_i)$.

If any metric is N/A, renormalize the remaining weights so they sum to 1.

- 2) Rescale to 0–100: Composite (0–100, weighted) = (Composite (0–3, weighted) ÷ 3) × 100.

Presets are defined in Appendix B.3 (Equal; Stakeholder-emphasis; Time-pressure-emphasis).

Reporting rules. Round means to 2 d.p.; disclose the effective N. Where used, show which preset is applied by default (Equal unless stated). For weight-sensitivity as a robustness probe (Aguinis and Solarino, 2019).

F.4 Optional inter-rater reliability (future phases)

ICC(2,k): average-measures, absolute-agreement, two-way random; report 95% CI and interpretation.

Kendall's W (optional): ordinal concordance check across raters.

F.5 Qualitative analysis

Directed content analysis with PASO-guided initial codes and allowance for inductive additions; focused/thematic synthesis for usability, portability, time-pressure, and “conditions to proceed.” Short anonymized exemplar quotes (≤ 20 words) reported in summary tables.

F.6 Exclusions, N/A, and missing data

No imputation. N/A values remain in the dataset but are excluded from denominators by design.

Exclusions: mis-coded IDs excluded; corrected resubmissions included (see Appendix M).

N/A: excluded from denominators; always show effective N.

Missing: no imputation; report item-level N.

Version integrity: any link fixes logged; item wording unchanged post version-lock (Appendix N).

F.7 Reporting conventions

Means to 2 d.p.; whole percentages for $\% \geq 4$; always display N (effective N where applicable).

Optional 95% CIs for key means.

Visuals are in Appendix H (e.g., heatmap for Q1–Q8; PASO coverage chart).

Reproducibility: formula-only workflow; summary tables in Appendix G reference the same version-locked sources.

F.8 Sensitivity analyses

Re-run headline indicators excluding all EXCLUDE/SUPERSEDED rows (Appendix M). Flag if any threshold is met:

(i) S-CVI/Ave shift ≥ 0.05 ; (ii) median confidence shift ≥ 1 point (0–10); (iii) material change in qualitative themes.

If none trigger → state “immaterial impact.” If any trigger → present both sets side-by-side with a brief comment.

APPENDIX G: EXPERT-PANEL RESULTS — SUMMARY TABLES AND
TIMELINE

[Expert-Panel] One-page summary tables (G1–G6) and fieldwork timeline supporting Chapter 4.

G1. Key indicators by case

Table 40 Appendix G1: Expert-Panel Key Indicators by Case

Source: Author (2025). Original table created from study data and materials.

Metric	Case 1	Case 2	Case 3
N (responses)	30	30	30
%≥4 Actionable pathway	90.0%	83.3%	93.3%
%≥4 Ethical dimensions	86.7%	86.7%	93.3%
%≥4 Defensible to stakeholders	86.7%	86.7%	90.0%
Confidence uplift (post – baseline, 0–10)	+0.37	+0.03	+0.53
Stance-change % (baseline ≠ post)	33.3%	36.7%	33.3%
Efficiency — less time	63.3%	66.7%	73.3%
Efficiency — same	0%	0%	0%
Efficiency — more time	23.3%	20.0%	13.3%
Efficiency — uncertain	3.3%	0%	3.3%

Computation: %≥4 = count(4–5) ÷ N valid (N/A excluded). Confidence uplift = mean(Post – Baseline) on 0–10. Stance-change % = share whose post stance differs from baseline (per case). N per case = 30.

Note: Efficiency percentages are calculated as a share of all respondents (N=30) per case. A small number left the efficiency item blank (≈10%), so rows may sum to ~90%.

“Same = 0%” means no respondent selected “About the same.”

G2. PASO coverage and scale averages (I-CVI / S-CVI)

Table 41 Appendix G2: PASO Content Validity Coverage by Case

Source: Author (2025). Original table created from study data and materials.

PASO micro-item (I-CVI)	Case 1	Case 2	Case 3
Principles	86.7%	86.7%	83.3%
Actions	86.7%	83.3%	90.0%
Skills	80.0%	73.3%	83.3%
Outcomes	86.7%	83.3%	90.0%
S-CVI/Ave (PASO)	85.0%	81.7%	86.7%

Computation: I-CVI(P/A/S/O) = %≥4 per PASO item; S-CVI/Ave = mean of the four I-

CVIs. N/A excluded. (Table 22).

G3. Global wrap-up items (means, 1–5)

Table 42 Appendix G3: Global Wrap-Up Item Means

Source: Author (2025). Original table created from study data and materials.

Global item	Mean (1–5)
Portability	4.09
Integratability	3.96
Time-pressure usability	4.20
Conceptual coherence	4.23
Tech/AI controls	4.19

Computation: Each item is asked once after Case 3; mean across respondents (N≈30).

G4a. “Sufficient for a defensible decision” (0–3 cases)

Table 43 Appendix G4a: Sufficiency Distribution Across the Three Cases

Source: Author (2025). Original table created from study data and materials.

Cases sufficient	Count	Share
0 cases	0	0%
1 case	4	13%
2 cases	8	27%
3 cases	18	60%

Computation: Share per category = count ÷ N valid (N≈30); N/A excluded.

G4b. PASO consistency across the three cases

Table 44 Appendix G4b: PASO Consistency Distribution (All 3 / 2 / Not Present)

Source: Author (2025). Original table created from study data and materials.

Category	Share
All 3 cases	73.3%
2 cases	23.3%
Not present	3.3%

Computation: Share per category = count ÷ N valid (N≈30); N/A excluded.

G5. Post-confidence (means, 0–10)

Table 45 Appendix G5: Post-Framework Confidence by Case

Source: Author (2025). Original table created from study data and materials.

Case	Mean
Case 1	8.20
Case 2	8.07
Case 3	8.47

Computation: Mean of 0–10 confidence after EDMC per case; N valid; no imputation.

G6. Q1–Q8 item means (1–5), pooled across cases

Table 46 Appendix G6: Q1–Q8 Means Pooled Across Cases 1–3

Source: Author (2025). Original table created from study data and materials.

Item	Mean
Q1 Actionable pathway	4.33
Q2 Ethical dimensions	4.23
Q3 Prompted considerations	3.83
Q4 Time-pressure usability	4.20
Q5 Defensible to stakeholders	4.25
Q6 Portability	4.10
Q7 Integratability	3.96
Q8 Intention to use	4.09

Computation: Mean on 1–5 for each item across all three case blocks; N/A excluded.

APPENDIX H: EXPERT-PANEL FIGURES AND METHODS NOTES

[Expert-Panel] Visuals (H1–H7) with concise computation notes corresponding to the expert-panel results.

Computation note (applies to H1–H7). %≥4 = count(4–5) ÷ N valid; N/A excluded.

Means are shown to 2 d.p. N per case = 30 unless noted. No imputation. For the “Efficiency” item, categories are Less / Same / More / Uncertain; see Appendix G1 footnote regarding 0% “Same.”

H1. Heatmap of agreement (≥4/5) by item × case

H1 takeaway. EDMC scores high on actionability, ethical salience, and defensibility across cases; Case 3 trends strongest overall.

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8
Case 1	90	86.7	73.3	86.7	86.7	84.4	82.2	81.1
Case 2	83.3	86.7	63.3	83.3	86.7	83.3	80	83.3
Case 3	93.3	93.3	83.3	86.7	90	83.3	86.7	83.3

Figure 17 Appendix H1: Heatmap of Agreement (≥4/5) by Item × Case
 Source: Author (2025). Original figure created from study data and materials.

What it shows. For each case (rows) and survey item Q1–Q8 (columns), the share of respondents giving 4–5 (darker = stronger agreement).

Computation. %≥4 as defined above.

Source. Computed from instrument responses; see Appendix G6 (pooled Q1–Q8 means) and Appendix G1 (headline indicators).

H2. PASO coverage by case (I-CVI = %≥4)

H2 takeaway. Actions and Outcomes are consistently strong; Skills is the relative bottleneck across cases.

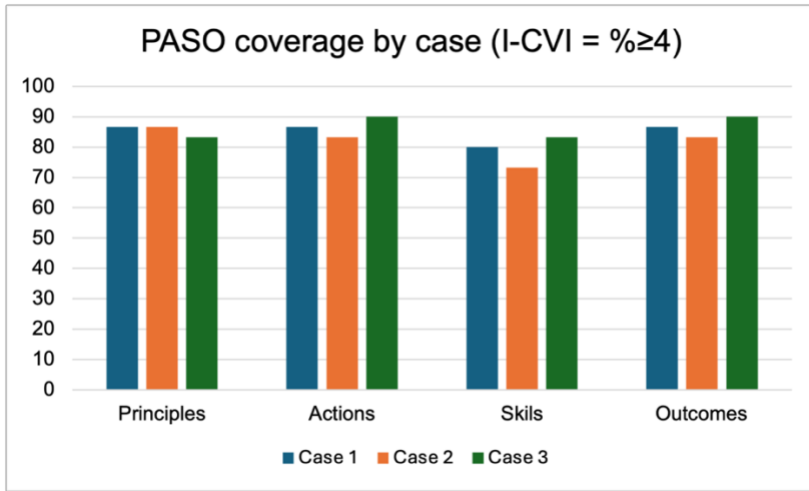


Figure 18 Appendix H2: PASO Coverage by Case (I-CVI = %≥4)
 Source: Author (2025). Original figure created from study data and materials.

What it shows. Item-level content validity (Principles, Actions, Skills, Outcomes) per case; S-CVI/Ave = mean of the four.

Computation. I-CVI(P/A/S/O) = %≥4; S-CVI/Ave = mean of I-CVIs; N = 30; N/A excluded.

Source. Appendix G2.

H3. Efficiency impact by case

H3 takeaway. Most respondents report less time to a decision with EDMC; “same time” is negligible; “more time” is a minority.

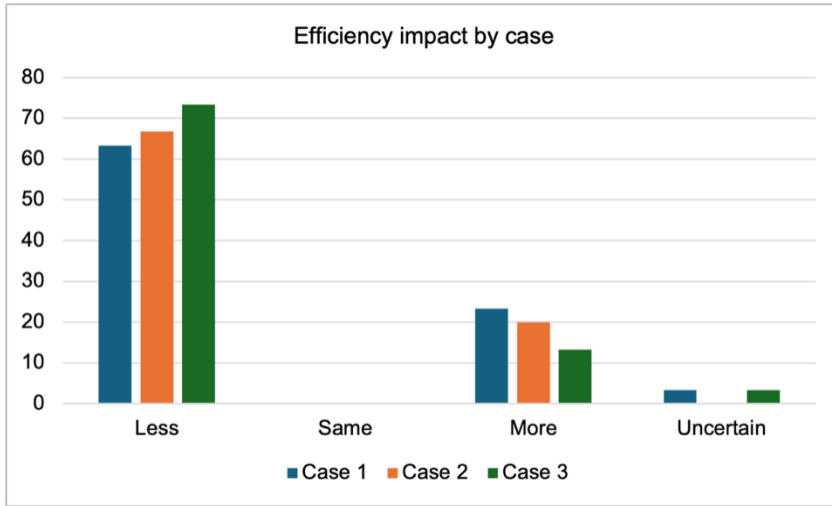


Figure 19 Appendix H3: Efficiency Impact by Case
 Source: Author (2025). Original figure created from study data and materials.

What it shows. Shares reporting Less / Same / More / Uncertain time to decision after using EDMC, by case.

Computation. Share = count(category) ÷ N valid; N = 30; N/A excluded.

Values. C1: 63.3 / 0.0 / 23.3 / 3.3; C2: 66.7 / 0.0 / 20.0 / 0.0; C3: 73.3 / 0.0 / 13.3 / 3.3.

Source. Appendix G1 (“Efficiency” rows). Note. 0% “Same” is consistent with Appendix G1 footnote.

H4. Confidence uplift (Post – Baseline)

H4: Confidence increases most in Case 3 (+0.53), modestly in Case 1 (+0.37), and barely in Case 2 (+0.03).



Figure 20 Appendix H4: Confidence Uplift by Case
Source: Author (2025). Original figure created from study data and materials.

What it shows. Mean change in confidence (0–10) from Baseline to Post, per case.

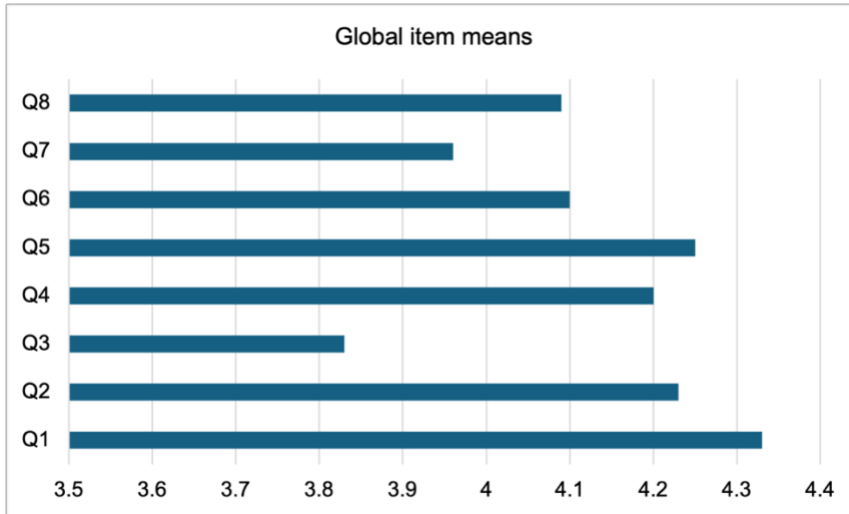
Computation. Mean(Post – Baseline), N = 30; N/A excluded.

Values. C1 = +0.37; C2 = +0.03; C3 = +0.53.

Source. Appendix G1.

H5. Global item means (Q1–Q8 across all cases)

H5: Time-pressure usability and conceptual coherence lead; integratability is solid but leaves room for workflow alignment.



*Figure 21 Appendix H5: Global Item Means (Q1–Q8 Pooled Across Cases)
Source: Author (2025). Original figure created from study data and materials.*

What it shows. Average 1–5 rating for each Q1–Q8 across all three cases.

Computation. Mean on 1–5 per item across cases; $N \approx 30$; N/A excluded.

Values. Q1 4.33; Q2 4.23; Q3 3.83; Q4 4.20; Q5 4.25; Q6 4.10; Q7 3.96; Q8 4.09.

Source. Appendix G6.

H6. PASO consistency and sufficiency (global distributions)

H6 takeaway. ~60% found EDMC sufficient in all three cases; ~73% saw PASO present in all three—evidence of consistency and coverage.

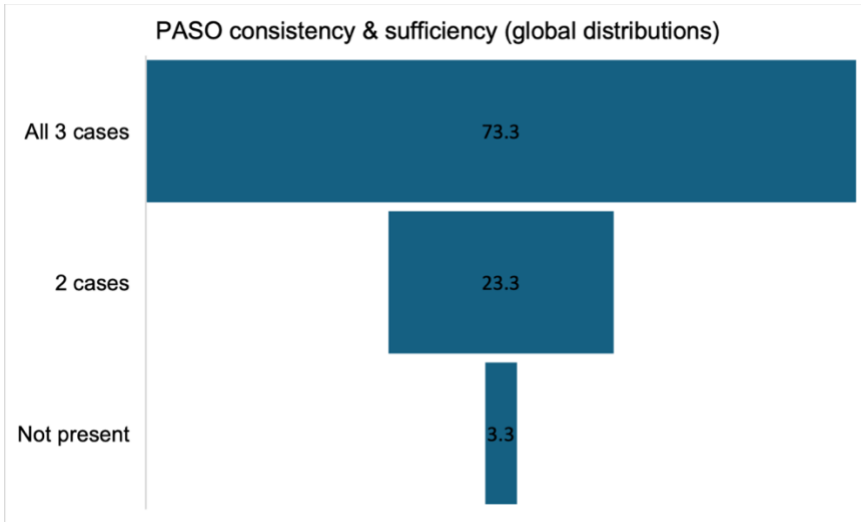


Figure 22 Appendix H6: (a) PASO Consistency and Sufficiency (All 3 / 2 / Not Present)
 Source: Author (2025). Original figure created from study data and materials.

What it shows. How often respondents reported all four PASO elements present: All 3 cases (73.3%), 2 cases (23.3%), Not present (3.3%).

Computation. Share per category = count ÷ N valid; N/A excluded.

Source. Appendix G4b.




Figure 23 Appendix H6: (b) PASO Sufficient and Defensible Decision (0-3 Cases)
 Source: Author (2025). Original figure created from study data and materials.

What it shows. How many cases respondents felt EDMC was sufficient: 0 (0%), 1 (13%), 2 (27%), 3 (60%).

Computation. Share per category = count ÷ N valid; N/A excluded.

Source. Appendix G4a.

H7. Global items (means, 1–5)

H7 takeaway. Time-pressure usability and conceptual coherence are highest; portability and tech/AI relevance are strong; integratability is solid.

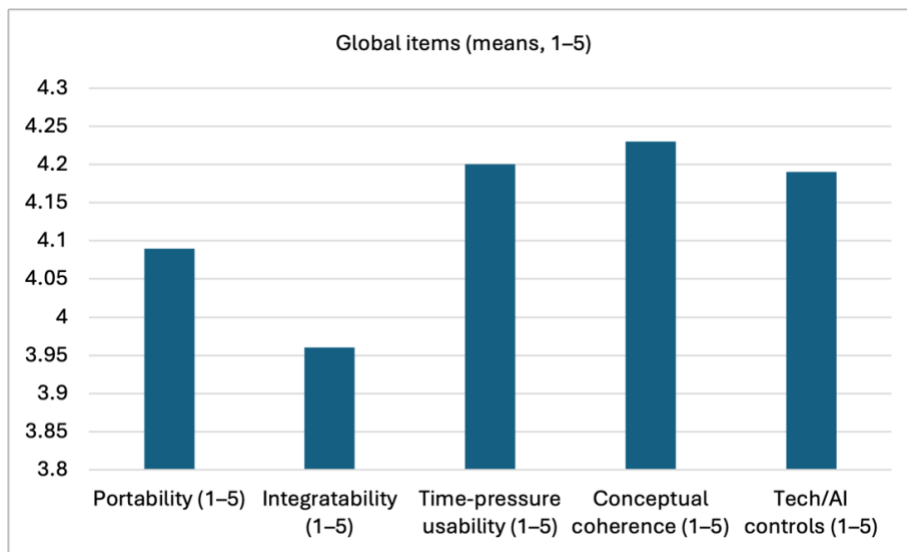


Figure 24 Appendix H7: Global Wrap-Up Items (Means, 1–5)

Source: Author (2025). Original figure created from study data and materials.

What it shows. Bar chart of the five global items: Portability 4.09; Integratability 3.96;

Time-pressure usability 4.20; Conceptual coherence 4.23; Tech/AI controls 4.19.

Computation. Mean on 1–5; N = 30; N/A excluded.

Source. Appendix G3.

APPENDIX I: DATA AVAILABILITY AND RETENTION

[Both / Administrative] Storage, access, retention period, and de-identification approach.

- Storage and scope. De-identified survey data and analysis outputs are stored in SSBM-managed Google Drive.
- Access. Principal Investigator (PI) for thesis analysis; examiners upon request. No public release of raw rows.
- Sharing on request. Aggregated tables/figures are shareable; de-identified extracts may be shared with examiners for verification.
- Formats. Google Sheets / Excel (analysis), PDF (appendices/figures). No custom code; all computations are formula-only (see Appendix F for methods; Appendix G–H for outputs).
- Retention. Up to five (5) years from degree award or per SSBM policy (whichever is shorter). [Insert confirmed period here before final submission.]
- Deletion. Direct identifiers were never collected. The NDA code ↔ person mapping (maintained outside the dataset) is destroyed at anonymization (see Appendix K).
- Provenance. Datasets comprise (i) expert-panel survey responses to hypothetical vignettes and (ii) researcher-led scoring notes for eight public cases; derived indicators/tables are reproducible via documented spreadsheet formulas (Appendix F).
- Contact. PI (see Appendix K) for inspection or examiner access.

APPENDIX J: GLOSSARY AND ABBREVIATIONS

EDMC — Ethical Decision-Making Compass (the decision tool evaluated).

PASO — Principles, Actions, Skills, Outcomes (EDMC's organizing model).

CVI / I-CVI / S-CVI/Ave — Content Validity Index: item-level (I-CVI = % \geq 4 on 1–5) and scale-average (S-CVI/Ave = mean of item I-CVIs).

KPI — Key Performance Indicator.

DPIA — Data Protection Impact Assessment.

NDA — Non-Disclosure Agreement.

Hypernorms — Trans-cultural baseline norms used to guide local adaptation.

RACI — Responsibility, Accountability, Consulted, Informed (role clarity).

Likert % \geq 4 — Share of ratings “4 or 5” on a 1–5 scale (N/A excluded).

Stance change — Baseline vs. Post decision category differs (Proceed / Proceed with conditions / Pause).

Confidence uplift — Post – Baseline confidence (0–10).

N / N_{eff} — Number of valid responses; effective N after excluding N/A.

N/A — Not applicable; excluded from denominators by design.

ICC(2,k) — Intraclass correlation (two-way random, absolute-agreement, average-measures).

Kendall's W — Rank-order concordance (0–1).

GDPR / Swiss FADP — EU General Data Protection Regulation / Swiss Federal Act on Data Protection.

De-identified / Pseudonymous — De-identified: no direct identifiers in the dataset.

Pseudonymous: NDA code used only to verify consent before anonymization.

Version-lock (v1.0) — Final instrument and case notes locked on 23 Sep 2025 (Appendix N).

Integratability — The ease of integrating EDMC into current decision/approval processes with minor adjustments. Operationalized via survey Q7 (“I could apply this within our existing process ...”). Distinct from Portability (cross-culture/jurisdiction fit).

Terminology policy (PASO “S”). In EDMC, S = Skills. When prior literature or strategy terms use capability/capabilities, we (i) keep the original wording for accuracy and (ii) treat it as part of PASO Skills unless it explicitly refers to organizational capability (e.g., capability maturity models), which we retain as such. Throughout the thesis, person- or team-level “capabilities” are standardized to skills.

Terminology note. In this thesis, the concise EDMC decision record is implemented as a case note for study purposes. The terms are used interchangeably.

APPENDIX K: PRIVACY AND DATA HANDLING (GDPR / SWISS FADP)

K.1 Participant-facing summary

Lawful basis. Consent for participation in an academic study; withdrawal possible any time before anonymization.

Data collected. Survey responses (baseline/post stance and confidence, Q1–Q8, PASO micro-items, efficiency), brief free-text comments, and privacy-safe demographics. No names/emails collected.

Participant code. An NDA-XX code verifies consent/access; it is not stored with identity and is removed at anonymization.

Storage/security. SSBM Google Drive; PI-restricted access; analysis via formula-only spreadsheets.

Sharing. Aggregated results may appear in the thesis/academic outputs; no public release of raw rows.

Retention. De-identified data retained for up to five (5) years or per SSBM policy (whichever is shorter).

Rights. GDPR/FADP rights (access, rectification, erasure, restriction, portability, objection, as applicable). Rights may be limited after anonymization because records can no longer be linked back.

Contact. PI — Indira Bunic, carrieny1@hotmail.com. Institution: Swiss School of Business and Management (SSBM), Geneva.

K.2 Detailed notice

Controller and contact. Indira Bunic, SSBM, Geneva (email above). DPO (if applicable): [insert name/email].

Purpose. Evaluate EDMC via expert review of three case notes and associated survey responses.

Data categories.

Direct identifiers: none in the survey.

Pseudonymous identifier: NDA code (used only to verify consent and enable withdrawal before anonymization).

Content data: stance choices; confidence (0–10); Q1–Q8 items; PASO micro-items; efficiency item; global wrap-up; short comments.

Demographics (coarse/optional): gender (optional), age band, education, sector, region, years of experience, role, domain.

Special-category data: not sought; participants are asked not to include sensitive details in free text.

Lawful basis and safeguards.

GDPR: Art. 6(1)(a) (consent); research safeguards per Art. 89.

Swiss FADP: consent; research safeguards aligned to FADP principles (proportionality, purpose limitation, data minimization).

Processors and systems. Google Forms/Drive (data capture/storage); NDA management (e.g., DocuSign/Adobe/university system); university/PI email.

International transfers. Possible under Google's Standard Contractual Clauses (SCCs).

No public sharing of raw data; only de-identified aggregates appear in thesis/outputs.

Security. Role-based access; encryption at rest/in transit; logical separation of (a) survey dataset, (b) NDA repository, and (c) mapping key. The mapping key is destroyed at anonymization.

Retention and deletion. Workflow: active (pseudonymous) → cleaning/analysis → anonymization (remove mapping; scrub incidental identifiers) → archive (de-identified) for the retention period → secure deletion. See Appendix I for retention details.

Rights and withdrawal. GDPR/FADP rights apply while data are still linkable. To withdraw, email the PI with your NDA code. After anonymization, records cannot be located or erased.

Automated decision-making. None; no profiling.

Supervisory authorities. Switzerland: FDPIC. EU/EEA: your national data protection authority (and FDPIC, as applicable).

Data minimization. Do not include names/employers/confidential details in free text; demographics are banded.

Provenance and transparency. The dataset includes (i) expert-panel survey responses to hypothetical vignettes (Appendix D) and (ii) researcher-led scoring notes for eight public cases (Appendix A). Methods and derived metrics are documented in Appendix F; outputs in Appendices G–H.

(This appendix is informational and does not constitute legal advice.)

APPENDIX L: CONSENT / NDA WORDING AND PARTICIPANT-CODE PROCESS

Cross-reference. For the full privacy and data-handling notice, see Appendix K.

L.1 Consent summary (survey preamble text)

By proceeding, you confirm you have read the information about this study and consent to participate. Participation is voluntary; there are no right/wrong answers. Responses are analyzed in aggregate; no names/emails are collected. You may withdraw before anonymization by emailing the PI with your participant code.

Contact: Indira Bunic — carrieny1@hotmail.com; Institution: Swiss School of Business and Management (SSBM), Geneva.

L.2 NDA usage

The EDMC case notes are shared under NDA to protect intellectual property. Redistribution is prohibited. The NDA governs access to the stimulus materials only; survey responses remain de-identified.

L.3 Participant codes (NDA-XX)

Each invitee receives a unique NDA-XX code by email. Entering the code at the start of the survey verifies consent and access. The code↔person mapping is stored separately (admin-only) and is destroyed at anonymization.

APPENDIX M: EXCLUSION RULES AND SENSITIVITY CHECKS

Cross-reference. For the version-lock and collection window, see Appendix N.

Context. Two admin issues occurred: (a) three different experts mistakenly entered participant code “07”; (b) one expert (ID-02) initially had a Case-3 link issue and later re-submitted correctly. The invitee actually assigned to 07 withdrew and did not participate.

M.1 Rules (applied consistently)

Mis-coded / duplicate code. Exclude any row using an incorrect or duplicate participant code. If the expert re-submits with the correct code, analyze only the corrected row; mark the mis-coded row EXCLUDE/SUPERSEDED.

Corrections / re-submissions (same person). Earlier row = SUPERSEDED/EXCLUDE once a corrected row is received.

N/A handling. Exclude N/A from denominators; report effective N.

M.2 Actions taken

Code 07 mistakes (three separate people). Each mis-coded “07” row was removed; each person re-submitted under their assigned code and those corrected rows were included.

The genuine ID-07 invitee withdrew (no valid submission).

ID-02 link issue. Corrected re-submission received 29 Sep 2025 and included; the original row marked SUPERSEDED/EXCLUDE.

Post-close policy. Form closed 29 Sep 2025, 14:40 CEST. Any later time-stamped submissions are LATE/EXCLUDE (form not re-opened).

M.3 Sensitivity design (SENS-1)

Re-run key summaries excluding all EXCLUDE/SUPERSEDED rows and note if any trigger:

(i) S-CVI/Ave shift ≥ 0.05 ; (ii) median confidence/clarity shift ≥ 1 point; or (iii) qualitative themes change materially.

Decision rule. If none trigger \rightarrow retain inclusive set and note “immaterial impact.” If any trigger \rightarrow report both sets and comment.

Outcome (this study). No triggers; immaterial impact. Inclusive set retained. See Appendix G (footnote in G1) for the stance-change note and consistency with locked figures.

M.4 Timestamps for transparency

Case-3 link fix: 23 Sep 2025, 12:15 CEST (wording unchanged).

Collection closed: 29 Sep 2025, 14:40 CEST (see Appendix N).

APPENDIX N: VERSION LOCK AND FIELDWORK WINDOW

Instrument tag. EDMC Expert Evaluation – Baseline EDMC + PASO CVI (v1.0)

Version-lock. 23 Sep 2025, 12:30 CEST (Europe/Zurich).

Cases included. AI Hiring Bias; Supplier Labor Compliance; Data Privacy Breach —

Case-3 link fixed 23 Sep 2025, 12:15 CEST (wording unchanged).

Data-collection window

Opened: 22 Sep 2025 (invitation issued; Form set to Accepting responses = On).

Closed: 29 Sep 2025, 14:40 CEST (Form set to Accepting responses = Off).

Event log (for transparency)

Case-3 link fix: 23 Sep 2025, 12:15 CEST — no wording changes.

Instrument version-lock: 23 Sep 2025, 12:30 CEST — v1.0.

Fieldwork window: 22–29 Sep 2025 — closed 29 Sep 2025, 14:40 CEST.

Event	Timestamp	Note
Case-3 link fix	23 Sep 2025, 12:15 CEST	No wording changes
Instrument version-lock	23 Sep 2025, 12:30 CEST	v1.0
Fieldwork window	22–29 Sep 2025	Closed 29 Sep 2025, 14:40 CEST

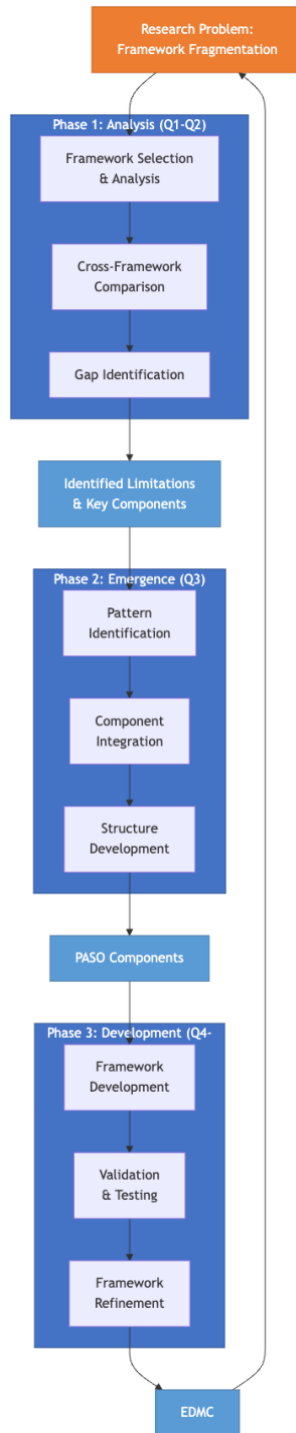
Cross-references. Data cleaning/sensitivity → Appendix M. Dashboard/reproducibility

→ Appendix G. Examiner access to case notes → Appendix D.


APPENDIX O: RESEARCH JOURNEY FIGURES

Scope. Visuals that show how the study moved from literature to PASO synthesis and into EDMC validation across two evidence streams.

Figure 13: The three-phase research approach (Initial → Bridge → Emergent) enables continuous refinement through overlapping inquiry waves (Edmondson & McManus, 2007). Phase 1 (Analysis) identifies strengths and limitations in existing frameworks. Phase 2 (Bridge) integrates these insights into the PASO components. Phase 3 (Emergent) develops and validates the EDMC architecture. This cyclical approach supports the emergence of patterns through recursive analysis and addresses framework fragmentation.



*Figure 13: EDMC Framework Development: Three-Phase Cyclical Research Process
Source: Author (2025). Original figure created from study data and materials.*



*Figure 25 Literature Review Journey—Seven Steps (Chapter II)
Source: Author (2025). Original figure created from study data and materials.*

Note: Numeric prefixes inside the boxes (e.g., “4.2.4”) reflect the Research Proposal’s Chapter IV numbering and do not map to this thesis’s chapter numbers.

APPENDIX P: EXTENDED METHODS SCAFFOLDS

Scope. Quick index of reusable scaffolds and where they live, so others can replicate or adapt the approach.

PASO plan scaffold → Appendix B (Table B1 notes).

EDMC one-page decision record (structure) → Appendix D (case-note layout).

RACI and escalation prompts → Appendix B (metric anchors).

KPI/threshold exemplars → Appendix D (Outcomes sections).

Weight presets and composite formulas → Appendix B.3.

Coding memo / audit-trail pattern → Appendix F (analysis plan).

Note. These scaffolds are aligned to EDMC's PASO DNA (Principles, Actions, Skills, Outcomes) to keep values, execution, ownership, and evidence in the same line of sight during design, evaluation, and audit.

REFERENCES

- Aguinis, H. and Bradley, K.J. (2014) 'Best practice recommendations for designing and implementing experimental vignette methodology studies', *Organizational Research Methods*, 17(4), pp. 351–371.
<https://doi.org/10.1177/1094428113514858>
- Aguinis, H. and Solarino, A.M. (2019) 'Transparency and replicability in qualitative research: The case of interviews with elite informants', *Organizational Research Methods*, 22(4), pp. 1291–1315. <https://doi.org/10.1177/1094428119840967>
- Alexander, L. and Moore, M. (2016) 'Deontological ethics', in Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2016). Available at:
<https://plato.stanford.edu/archives/win2016/entries/ethics-deontological/>
(Accessed: 25 December 2025).
- Alvesson, M. and Sandberg, J. (2013) *Constructing Research Questions: Doing Interesting Research*. London: SAGE.
- Alvesson, M. and Sköldbberg, K. (2017) *Reflexive Methodology: New Vistas for Qualitative Research*. 3rd edn. London: SAGE.
- Ambrose, J. (2022) 'Shell's Jackdaw gasfield approved despite climate concerns', *The Guardian*, 1 June. Available at:
<https://www.theguardian.com/business/2022/jun/01/shells-jackdaw-gasfield-approved-despite-climate-concerns> (Accessed: 25 December 2025).
- Babalola, M.T., Stouten, J., Euwema, M.C. and Ovadje, F. (2023) 'A social identity perspective on ethical leadership and employee psychological well-being', *Journal of Organizational Behavior*, 44(2), pp. 134–149.
<https://doi.org/10.1002/job.2623>

- Bass, B.M. (1985) *Leadership and Performance Beyond Expectations*. New York: Free Press.
- Bazerman, M.H. and Tenbrunsel, A.E. (2011) *Blind Spots: Why We Fail to Do What's Right and What to Do About It*. Princeton, NJ: Princeton University Press.
- Bentham, J. (1789) *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.
- Bhattarai, A. (2022) 'Google to pay \$118 million to settle gender discrimination lawsuit', *The Washington Post*, 12 June. Available at: <https://www.washingtonpost.com/business/2022/06/12/google-gender-discrimination-settlement/> (Accessed: 25 December 2025).
- Bloomberg News (2023) 'Apple supplier Foxconn violated China labor rules, report says', *Bloomberg*, 15 December. Available at: <https://www.bloomberg.com/news/articles/2023-12-15/apple-supplier-foxconn-violated-china-labor-rules-report-says> (Accessed: 25 December 2025).
- Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', *Qualitative Research in Psychology*, 3(2), pp. 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Brown, M.E. and Treviño, L.K. (2006) 'Ethical leadership: A review and future directions', *The Leadership Quarterly*, 17(6), pp. 595–616. <https://doi.org/10.1016/j.leaqua.2006.10.004>
- Brown, T. (2008) 'Design thinking', *Harvard Business Review*, 86(6), pp. 84–92.
- Carroll, A.B. (1991) 'The pyramid of corporate social responsibility: Toward the moral management of organizational stakeholders', *Business Horizons*, 34(4), pp. 39–48. [https://doi.org/10.1016/0007-6813\(91\)90005-G](https://doi.org/10.1016/0007-6813(91)90005-G)

- Charmaz, K. (2006) *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis*. London: SAGE.
- Craft, J.L. (2013) 'A review of the empirical ethical decision-making literature: 2004–2011', *Journal of Business Ethics*, 117(2), pp. 221–259.
<https://doi.org/10.1007/s10551-012-1518-9>
- Creswell, J.W. and Plano Clark, V.L. (2018) *Designing and Conducting Mixed Methods Research*. 3rd edn. Thousand Oaks, CA: SAGE.
- Dastin, J. (2018) 'Amazon scraps secret AI recruiting tool that showed bias against women', *Reuters*, 9 October. Available at:
<https://www.reuters.com/world/us/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-2018-10-10/> (Accessed: 25 December 2025).
- Davis, M.S. (1971) 'That's interesting! Towards a phenomenology of sociology and a sociology of phenomenology', *Philosophy of the Social Sciences*, 1(2), pp. 309–344. <https://doi.org/10.1177/004839317100100211>
- Del Popolo Cristaldi, F., Palmiotti, G.P., Cellini, N. and Sarlo, M. (2021) 'Pulling the lever in a hurry: The influence of impulsivity and sensitivity to reward on moral decision-making under time pressure', *Psychology*, 12(1), pp. 1–9.
<https://doi.org/10.4236/psych.2021.1211062>
- Deloitte (2023) *The State of Ethics and Trust in Technology*. Available at:
<https://www.deloitte.com/us/en/about/governance/technology-trust-ethics-annual-report.html> (Accessed: 25 December 2025).
- Denhardt, R.B. and Denhardt, J.V. (2003) *The New Public Service: Serving, Not Steering*. Armonk, NY: M.E. Sharpe.
- Denzin, N.K. (1978) *The Research Act: A Theoretical Introduction to Sociological Methods*. 2nd edn. New York: McGraw-Hill.

- DeVellis, R.F. (2017) *Scale Development: Theory and Applications*. 4th edn. Thousand Oaks, CA: SAGE.
- Dillman, D.A., Smyth, J.D. and Christian, L.M. (2014) *Internet, Phone, Mail, and Mixed-Mode Surveys: The Tailored Design Method*. 4th edn. Hoboken, NJ: John Wiley & Sons.
- Dixon-Woods, M., Agarwal, S., Jones, D., Young, B. and Sutton, A. (2005) 'Synthesizing qualitative and quantitative evidence: A review of possible methods', *Journal of Health Services Research & Policy*, 10(1), pp. 45–52.
<https://doi.org/10.1258/1355819054308576>
- Donaldson, T. and Dunfee, T.W. (1999) *Ties That Bind: A Social Contracts Approach to Business Ethics*. Boston, MA: Harvard Business School Press.
- Dorst, K. (2015) *Frame Innovation: Create New Thinking by Design*. Cambridge, MA: MIT Press.
- Driver, J. (2012) *Consequentialism*. London: Routledge.
- Eagly, A.H. (2007) 'Female leadership advantage and disadvantage: Resolving the contradictions', *Psychology of Women Quarterly*, 31(1), pp. 1–12.
<https://doi.org/10.1111/j.1471-6402.2007.00326.x>
- Edmondson, A.C. (2018) *The Fearless Organization: Creating Psychological Safety in the Workplace for Learning, Innovation, and Growth*. Hoboken, NJ: John Wiley & Sons.
- Edmondson, A.C. and McManus, S.E. (2007) 'Methodological fit in management field research', *Academy of Management Review*, 32(4), pp. 1246–1264.
<https://doi.org/10.5465/amr.2007.26586086>

- Eisenbeiss, S.A. (2012) 'Re-thinking ethical leadership: An interdisciplinary integrative approach', *The Leadership Quarterly*, 23(5), pp. 791–808.
<https://doi.org/10.1016/j.leaqua.2012.03.001>
- Eisenhardt, K.M. (1989) 'Making fast strategic decisions in high-velocity environments', *Academy of Management Journal*, 32(3), pp. 543–576.
<https://doi.org/10.2307/256434>
- Eisenhardt, K.M., Graebner, M.E. and Sonenshein, S. (2016) 'Grand challenges and inductive methods: Rigor without rigor mortis', *Academy of Management Journal*, 59(4), pp. 1113–1123. <https://doi.org/10.5465/amj.2016.4004>
- Elkington, J. (1997) *Cannibals with Forks: The Triple Bottom Line of 21st Century Business*. Oxford: Capstone.
- Emsler, A.K., Krzikalla, C. and Abele-Brehm, A. (2021) 'Higher, faster, further: Occupational ethical decision-making under time pressure in type A versus type B personalities', *Psychology*, 12(10), pp. 1662–1679.
<https://doi.org/10.4236/psych.2021.1211062>
- Ete, Z., Demirtas, O., Akdogan, A.A. and Arslan, A. (2023) 'Ethical leadership in organizations: Current theoretical trends and future research directions', *Journal of Business Ethics*, 184(1), pp. 213–241. <https://doi.org/10.1007/s10551-022-05216-6>
- European Parliament and Council (2016) 'Regulation (EU) 2016/679 ... (General Data Protection Regulation)', *Official Journal of the European Union*, L 119, pp. 1–88. Available at: <https://eur-lex.europa.eu/eli/reg/2016/679/oj> (Accessed: 25 December 2025).

- Eva, N., Robin, M., Sendjaya, S., van Dierendonck, D. and Liden, R.C. (2019) 'Servant leadership: A systematic review and call for future research', *The Leadership Quarterly*, 30(1), pp. 111–132. <https://doi.org/10.1016/j.leaqua.2018.07.004>
- Evans, S.C., Roberts, M.C., Keeley, J.W., Blossom, J.B., Amaro, C.M., Garcia, A.M., Stough, C.O., Canter, K.S., Robles, R., Reed, G.M. and Medina-Mora, M.E. (2015) 'Vignette methodologies for studying clinicians' decision-making: Validity, utility, and application in ICD-11 field studies', *International Journal of Clinical and Health Psychology*, 15(2), pp. 160–170. <https://doi.org/10.1016/j.ijchp.2015.02.002>
- Ferrell, O.C., Fraedrich, J. and Ferrell, L. (2013) *Business Ethics: Ethical Decision Making and Cases*. 9th edn. Stamford, CT: Cengage Learning.
- Ferrell, O.C. and Gresham, L.G. (1985) 'A contingency framework for understanding ethical decision making in marketing', *Journal of Marketing*, 49(3), pp. 87–96. <https://doi.org/10.1177/002224298504900308>
- Finlay, L. (2002) 'Outing the researcher: The provenance, process, and practice of reflexivity', *Qualitative Health Research*, 12(4), pp. 531–545. <https://doi.org/10.1177/104973202129120052>
- Flammer, C. (2013) 'Corporate social responsibility and shareholder reaction: The environmental awareness of investors', *Academy of Management Journal*, 56(3), pp. 758–781. <https://doi.org/10.5465/amj.2011.0746>
- Floridi, L. (2013) *The Ethics of Information*. Oxford: Oxford University Press.
- Floridi, L. (2020) 'Translating principles into practices of digital ethics: Five risks of being unethical', *Philosophy & Technology*, 33(1), pp. 155–164. <https://doi.org/10.1007/s13347-019-00354-6>

- Floridi, L. (2020) 'Artificial intelligence as a public service: Learning from Amsterdam and Helsinki', *Philosophy & Technology*, 33(3), pp. 541–546.
<https://doi.org/10.1007/s13347-020-00413-8>
- Floridi, L., Cows, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P. and Vayena, E. (2018) 'AI ethics: Mapping the debate', *Science and Engineering Ethics*, 24(5), pp. 1537–1540. <https://doi.org/10.1007/s11948-018-00665-7>
- Floridi, L., Cows, J., King, T.C. and Taddeo, M. (2020) 'AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations', *Minds and Machines*, 28(4), pp. 689–707.
<https://doi.org/10.1007/s11023-018-9482-5>
- Forbes Business Council (2023) 'The role of ethical leadership in long-term organizational success', *Forbes*, 8 September. Available at:
<https://www.forbes.com/councils/forbesbusinesscouncil/2023/09/08/the-role-of-ethical-leadership-in-long-term-organizational-success/> (Accessed: 25 December 2025).
- Freeman, R.E. (1984) *Strategic Management: A Stakeholder Approach*. Boston, MA: Pitman.
- Gigerenzer, G. and Gaissmaier, W. (2011) 'Heuristic decision making', *Annual Review of Psychology*, 62, pp. 451–482. <https://doi.org/10.1146/annurev-psych-120709-145346>
- Gilligan, C. (1982) *In a Different Voice: Psychological Theory and Women's Development*. Cambridge, MA: Harvard University 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. <https://doi.org/10.1177/1094428112452151>
- Grant, J.S. and Davis, L.L. (1997) 'Selection and use of content experts for instrument development', *Research in Nursing & Health*, 20(3), pp. 269–274. [https://doi.org/10.1002/\(SICI\)1098-240X\(199706\)20:3<269::AID-NUR9>3.0.CO;2-G](https://doi.org/10.1002/(SICI)1098-240X(199706)20:3<269::AID-NUR9>3.0.CO;2-G)
- Greenleaf, R.K. (1977) *Servant Leadership: A Journey into the Nature of Legitimate Power and Greatness*. New York: Paulist Press.
- Greenwood, M., Bernardi, C. and Bernardi, D. (2022) 'Measuring ethical performance in organizations: Challenges and new directions', *Business Ethics Quarterly*, 32(2), pp. 269–295. <https://doi.org/10.1017/beq.2021.38>
- Harrison, J. (2020) 'Accountability and ethics in decision-making', *Journal of Business Ethics*, 162(4), pp. 715–729. <https://doi.org/10.1007/s10551-018-3995-2>
- Harrison, J.S., Freeman, R.E. and Abreu, M.C.S. (2019) 'Stakeholder theory as an ethical approach to effective management: Applying the theory to multiple contexts', *Revista Brasileira de Gestão de Negócios*, 21(1), pp. 1–16. <https://doi.org/10.7819/rbgn.v21i1.4001>
- Harvard Professional Development (2023) 'What is ethical leadership and why is it important?', *Harvard Division of Continuing Education*. Available at: <https://professional.dce.harvard.edu/blog/what-is-ethical-leadership-and-why-is-it-important/> (Accessed: 25 December 2025).
- Heaven, W.D. (2023) 'The OpenAI drama shows why AI governance is a mess', *MIT Technology Review*, 22 November. Available at:

- <https://www.technologyreview.com/2023/11/22/1084277/openai-drama-ai-governance-mess/> (Accessed: 25 December 2025).
- Held, V. (2006) *The Ethics of Care: Personal, Political, and Global*. New York, NY: Oxford University Press.
- Hevner, A.R., March, S.T., Park, J. and Ram, S. (2004) 'Design science in information systems research', *MIS Quarterly*, 28(1), pp. 75–105.
<https://doi.org/10.2307/25148625>
- Hofstede, G. (2001) *Culture's Consequences: Comparing Values, Behaviors, Institutions, and Organizations across Nations*. 2nd edn. Thousand Oaks, CA: SAGE.
- Hsieh, H.F. and Shannon, S.E. (2005) 'Three approaches to qualitative content analysis', *Qualitative Health Research*, 15(9), pp. 1277–1288. <https://doi.org/10.1177/1049732305276687>
- Hunt, S.D. and Vitell, S. (1986) 'A general theory of marketing ethics', *Journal of Macromarketing*, 6(1), pp. 5–16. <https://doi.org/10.1177/027614678600600103>
- Hunter, E.M., Neubert, M.J., Perry, S.J., Witt, L.A., Penney, L.M. and Weinberger, E. (2013) 'Servant leaders inspire servant followers: Antecedents and outcomes for employees and the organization', *The Leadership Quarterly*, 24(2), pp. 316–331. <https://doi.org/10.1016/j.leaqua.2012.12.001>
- Hursthouse, R. (1999) *On Virtue Ethics*. Oxford: Oxford University Press.
- IBM Security (2024) *Cost of a Data Breach Report 2024*. Armonk, NY: IBM Security/Ponemon Institute. Available at: <https://www.ibm.com/reports/data-breach> (Accessed: 25 December 2025).
- Israel, M. (2015) *Research Ethics and Integrity for Social Scientists: Beyond Regulatory Compliance*. 2nd edn. London: SAGE.

- Jacobson, I., Booch, G. and Rumbaugh, J. (1999) *The Unified Software Development Process*. Boston, MA: Addison-Wesley.
- Jick, T.D. (1979) 'Mixing qualitative and quantitative methods: Triangulation in action', *Administrative Science Quarterly*, 24(4), pp. 602–611. <https://doi.org/10.2307/2392366>
- Jobin, A., Ienca, M. and Vayena, E. (2019) 'The global landscape of AI ethics guidelines', *Nature Machine Intelligence*, 1(9), pp. 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Johnson, M. (2021) 'Bridging theory and practice in ethical leadership', *Journal of Business Ethics*, 174(1), pp. 1–15. <https://doi.org/10.1007/s10551-020-04553-7>
- Jones, T.M. (1991) 'Ethical decision making by individuals in organizations: An issue-contingent model', *Academy of Management Review*, 16(2), pp. 366–395. <https://doi.org/10.5465/amr.1991.4278958>
- Kant, I. (1785) *Groundwork of the Metaphysics of Morals*. New York: Harper & Row.
- Kaptein, M. (2008) 'Developing and testing a measure for the ethical culture of organizations: The Corporate Ethical Virtues model', *Journal of Organizational Behavior*, 29(7), pp. 923–947. <https://doi.org/10.1002/job.520>
- Kaptein, M. (2017) 'The battle for business ethics: A struggle theory', *Journal of Business Ethics*, 144(2), pp. 343–361. <https://doi.org/10.1007/s10551-015-2826-1>
- Kaptein, M. (2020) 'The effectiveness of ethics programs: The role of scope, composition, and implementation', *Journal of Business Ethics*, 165(1), pp. 1–18. <https://doi.org/10.1007/s10551-018-4084-5>
- Kaptein, M. (2023) 'The moral organization: A comprehensive framework for understanding and improving organizational ethics', *Journal of Business Ethics*, 183(4), pp. 979–1003. <https://doi.org/10.1007/s10551-022-05216-6>

- Kim, K. and Patel, P.C. (2018) 'The impact of cultural intelligence on global ethical decision-making', *Journal of Business Ethics*, 148(1), pp. 1–16. <https://doi.org/10.1007/s10551-015-2993-7>
- Kitchin, R. (2014) *The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences*. London: SAGE.
- Kolb, D.A. (1984) *Experiential Learning: Experience as the Source of Learning and Development*. Englewood Cliffs, NJ: Prentice-Hall.
- Kotter, J.P. (1996) *Leading Change*. Boston, MA: Harvard Business School Press.
- Kristjánsson, K. (2020) 'Virtue ethics in organizations: A critical assessment', *Journal of Business Ethics*, 167(1), pp. 1–14. <https://doi.org/10.1007/s10551-019-04181-7>
- Krosnick, J.A. and Presser, S. (2010) 'Question and questionnaire design', in Marsden, P.V. and Wright, J.D. (eds.) *Handbook of Survey Research*. Bingley: Emerald, pp. 263–313.
- Krüger, P. (2015) 'Corporate goodness and shareholder wealth', *Journal of Financial Economics*, 115(2), pp. 304–329. <https://doi.org/10.1016/j.jfineco.2014.09.008>
- Lehnert, K., Craft, J., Singh, N. and Park, Y. (2015) 'Ethical decision-making theory: An integrated approach', *Journal of Business Ethics*, 134(3), pp. 429–443. <https://doi.org/10.1007/s10551-014-2430-4>
- Lincoln, Y.S. and Guba, E.G. (1985) *Naturalistic Inquiry*. Beverly Hills, CA: SAGE.
- Lynham, S.A. (2002) 'The general method of theory-building research in applied disciplines', *Advances in Developing Human Resources*, 4(3), pp. 221–241. <https://doi.org/10.1177/1523422302043002>
- Lynn, M.R. (1986) 'Determination and quantification of content validity', *Nursing Research*, 35(6), pp. 382–385. <https://doi.org/10.1097/00006199-198611000-00017>

- MacDougall, A.E., Bagdasarov, Z., Johnson, J.F. and Mumford, M.D. (2015) 'Managing workplace ethics: An extended conceptualization of ethical sensemaking and the facilitative role of human resources', *Research in Personnel and Human Resources Management*, 33, pp. 121–189. <https://doi.org/10.1108/S0742-730120150000033004>
- MacIntyre, A. (1984) *After Virtue: A Study in Moral Theory*. 2nd edn. Notre Dame, IN: University of Notre Dame Press.
- MacIntyre, A. (2007) *After Virtue: A Study in Moral Theory*. 3rd edn. Notre Dame, IN: University of Notre Dame Press.
- Martin, K.D., Borah, A. and Palmatier, R.W. (2021) 'The dark side of AI: Understanding ethical concerns in marketing analytics', *Journal of Marketing*, 85(3), pp. 85–103. <https://doi.org/10.1177/0022242920984096>
- McClaren, N. (2015) 'A literature review on the Hunt–Vitell model of ethics in marketing', *Journal of Business Ethics*, 128(3), pp. 735–752. <https://doi.org/10.1007/s10551-014-2081-9>
- Miles, M.B., Huberman, A.M. and Saldaña, J. (2014) *Qualitative Data Analysis: A Methods Sourcebook*. 3rd edn. Thousand Oaks, CA: SAGE.
- Mill, J.S. (1863) *Utilitarianism*. London: Parker, Son, and Bourn.
- Mittelstadt, B.D., Allo, P., Taddeo, M., Wachter, S. and Floridi, L. (2016) 'The ethics of algorithms: Mapping the debate', *Big Data & Society*, 3(2), pp. 1–21. <https://doi.org/10.1177/2053951716679679>
- Mo, S., Ling, C.D. and Jin, X. (2023) 'How and when ethical leadership promotes employee creativity: A moderated mediation model', *International Journal of Environmental Research and Public Health*, 20(1), p. 75. <https://doi.org/10.3390/ijerph20010075>

- Mökander, J., Axente, M. and Floridi, L. (2023) 'AI ethics and digital governance: A critical analysis', *AI & Society*, 38(1), pp. 341–356. <https://doi.org/10.1007/s00146-022-01418-7>
- Moore, C. and Gino, F. (2013) 'Ethically adrift: How others pull our moral compass from true north, and how we can fix it', *Research in Organizational Behavior*, 33, pp. 1–20. <https://doi.org/10.1016/j.riob.2013.08.001>
- Morgan, D.L. (2014) 'Pragmatism as a paradigm for social research', *Qualitative Inquiry*, 20(8), pp. 1045–1053. <https://doi.org/10.1177/1077800413513733>
- Morse, J.M. (1991) 'Approaches to qualitative-quantitative methodological triangulation', *Nursing Research*, 40(2), pp. 120–123. <https://doi.org/10.1097/00006199-199103000-00014>
- National Labor Relations Board (2023) 'NLRB Region 3 Buffalo wins administrative law judge decision requiring Starbucks to remedy 500 labor violations'. Press release, 1 March. Available at: <https://www.nlr.gov/news-outreach/news-story/nlr-region-3-buffalo-wins-administrative-law-judge-decision-requiring-starbucks-to-remedy-500-labor-violations> (Accessed: 25 December 2025).
- Newman, A., Round, H., Bhattacharya, S. and Roy, A. (2017) 'Ethical climates in organizations: A review and research agenda', *Business Ethics Quarterly*, 27(4), pp. 475–512. <https://doi.org/10.1017/beq.2017.30>
- Newman, A., Eva, N., Bavik, Y.L. and Shao, B. (2023a) 'Ethical leadership: A systematic review and future research agenda', *The Leadership Quarterly*, 34, 101673. <https://doi.org/10.1016/j.leaqua.2022.101673>
- Newman, A., Round, H., Wang, S. and Mount, M. (2023b) 'The moderating role of ethical leadership in unethical pro-organizational behavior', *Journal of Business Ethics*, 178(1), pp. 23–41. <https://doi.org/10.1007/s10551-021-04965-2>

- NIST (2023) *AI Risk Management Framework (AI RMF 1.0)*. Gaithersburg, MD: National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- OECD (2019) ‘Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449)’. Paris: OECD. Available at: <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449> (Accessed: 25 December 2025).
- O’Fallon, M.J. and Butterfield, K.D. (2005) ‘A review of the empirical ethical decision-making literature: 1996–2003’, *Journal of Business Ethics*, 59(4), pp. 375–413. <https://doi.org/10.1007/s10551-005-2929-7>
- Palazzo, G., Krings, F. and Hoffrage, U. (2012) ‘Ethical blindness’, *Journal of Business Ethics*, 109(1), pp. 1–17. <https://doi.org/10.1007/s10551-012-1253-4>
- Peppers, K., Tuunanen, T., Rothenberger, M.A. and Chatterjee, S. (2007) ‘A design science research methodology for information systems research’, *Journal of Management Information Systems*, 24(3), pp. 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
- Podsakoff, P.M., MacKenzie, S.B., Lee, J.-Y. and Podsakoff, N.P. (2003) ‘Common method biases in behavioral research: A critical review of the literature and recommended remedies’, *Journal of Applied Psychology*, 88(5), pp. 879–903. <https://doi.org/10.1037/0021-9010.88.5.879>
- Polit, D.F. and Beck, C.T. (2006) ‘The content validity index: Are you sure you know what’s being reported? Critique and recommendations’, *Research in Nursing & Health*, 29(5), pp. 489–497. <https://doi.org/10.1002/nur.20147>

- Polit, D.F., Beck, C.T. and Owen, S.V. (2007) 'Is the CVI an acceptable indicator of content validity? Appraisal and recommendations', *Research in Nursing & Health*, 30(4), pp. 459–467. <https://doi.org/10.1002/nur.20199>
- Rest, J.R. (1986) *Moral Development: Advances in Research and Theory*. New York: Praeger.
- Rest, J., Narvaez, D., Bebeau, M.J. and Thoma, S.J. (1999) *Postconventional Moral Thinking: A Neo-Kohlbergian Approach*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Reuters (2023) 'Shein supplier audits find fewer serious violations'. *Reuters Graphics*. Available at: <https://www.reuters.com/graphics/SHEIN-IPO/ESG/xmvjndgrjvr/> (Accessed: 25 December 2025).
- Saldaña, J. (2021) *The Coding Manual for Qualitative Researchers*. 4th edn. London: SAGE.
- Schein, E.H. (2010) *Organizational Culture and Leadership*. 4th edn. San Francisco, CA: Jossey-Bass.
- Scherer, A.G. and Palazzo, G. (2011) 'The new political role of business in a globalized world: A review of a new perspective on CSR and its implications for the firm, governance, and democracy', *Journal of Management Studies*, 48(4), pp. 899–931. <https://doi.org/10.1111/j.1467-6486.2010.00950.x>
- Schön, D.A. (1983) *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic Books.
- Schwartz, M.S. (2016) 'Ethical decision-making theory: An integrated approach', *Journal of Business Ethics*, 139(4), pp. 755–776. <https://doi.org/10.1007/s10551-015-2886-8>

- Schwartz, M.S. (2023) *Business Ethics: An Ethical Decision-Making Approach*. Hoboken, NJ: Wiley.
- Shadish, W.R., Cook, T.D. and Campbell, D.T. (2002) *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Boston, MA: Houghton Mifflin.
- Slaper, T.F. and Hall, T.J. (2011) ‘The triple bottom line: What is it and how does it work?’, *Indiana Business Review*, 86(1), pp. 4–8.
- Sonenshein, S. (2007) ‘The role of construction, intuition, and justification in responding to ethical issues at work: The sensemaking-intuition model’, *Academy of Management Review*, 32(4), pp. 1022–1040. <https://doi.org/10.5465/amr.2007.26585677>
- Tenbrunsel, A.E. and Smith-Crowe, K. (2008) ‘Ethical decision making: Where we’ve been and where we’re going’, *Academy of Management Annals*, 2(1), pp. 545–607. <https://doi.org/10.1080/19416520802211677>
- Timmermans, S. and Tavory, I. (2012) ‘Theory construction in qualitative research: From grounded theory to abductive analysis’, *Sociological Theory*, 30(3), pp. 167–186. <https://doi.org/10.1177/0735275112457914>
- Tremblay, D., Turcotte, A., Touati, N. et al. (2022) ‘Research vignettes to collect qualitative data: A scoping review’, *BMJ Open*, 12(7), e057095. <https://doi.org/10.1136/bmjopen-2021-057095>
- Treviño, L.K. (1986) ‘Ethical decision making in organizations: A person-situation interactionist model’, *Academy of Management Review*, 11(3), pp. 601–617. <https://doi.org/10.5465/amr.1986.4306235>
- Treviño, L.K. and Nelson, K.A. (2016) *Managing Business Ethics: Straight Talk about How to Do It Right*. 7th edn. Hoboken, NJ: Wiley.

- Treviño, L.K. and Nelson, K.A. (2021) *Managing Business Ethics: Straight Talk about How to Do It Right*. 8th edn. Hoboken, NJ: John Wiley & Sons.
- Treviño, L.K., Nieuwenboer, N.A. and Kish-Gephart, J.J. (2014a) '(Un)ethical behavior in organizations', *Annual Review of Psychology*, 65, pp. 635–660. <https://doi.org/10.1146/annurev-psych-113011-143745>
- Treviño, L.K., Weaver, G.R. and Reynolds, S.J. (2014b) 'Behavioral ethics in organizations: A review', *Journal of Management*, 40(1), pp. 1–20. <https://doi.org/10.1177/0149206313495411>
- U.S. Equal Employment Opportunity Commission (2024) *Fiscal Year 2024 Annual Performance Report*. Washington, DC: EEOC. Available at: <https://www.eeoc.gov/eeoc-fy-2024-annual-performance-report> (Accessed: 25 December 2025).
- Van de Ven, A.H. (2007) *Engaged Scholarship: A Guide for Organizational and Social Research*. Oxford: Oxford University Press.
- Van Wart, M. (2013) 'Lessons from leadership theory and the contemporary challenges of leaders', *Public Administration Review*, 73(4), pp. 553–565. <https://doi.org/10.1111/puar.12040>
- Velasquez, M. (1982) *Business Ethics: Concepts and Cases*. Englewood Cliffs, NJ: Prentice Hall.
- Wedell-Wedellsborg, T. (2020) *What's Your Problem? To Solve Your Toughest Problems, Change the Problems You Solve*. Boston, MA: Harvard Business Review Press.
- Werhane, P.H. (2008) 'Mental models, moral imagination and system thinking in the age of globalization', *Journal of Business Ethics*, 78(3), pp. 463–474. <https://doi.org/10.1007/s10551-007-9367-6>

- Whetstone, J.T. (2001) 'How virtue fits within business ethics', *Journal of Business Ethics*, 33(2), pp. 101–114. <https://doi.org/10.1023/A:1017588429287>
- Whittier, T.R., Van Horn, D.J. et al. (2006) 'A process for creating multimetric indices for large-scale aquatic surveys', *Journal of the North American Benthological Society*, 25(4), pp. 822–841. [https://doi.org/10.1899/0887-3593\(2006\)025\[0822:APFCMI\]2.0.CO;2](https://doi.org/10.1899/0887-3593(2006)025[0822:APFCMI]2.0.CO;2)
- Woiceshyn, J. (2011) 'A model for ethical decision making in business: Reasoning, intuition, and rational moral principles', *Journal of Business Ethics*, 104(3), pp. 311–323. <https://doi.org/10.1007/s10551-011-0910-1>
- Yin, R.K. (2018) *Case Study Research and Applications: Design and Methods*. 6th edn. Thousand Oaks, CA: SAGE.
- Zhang, D. and Wade, M. (2022) 'The governance of digital ethics in organizations', *MIS Quarterly*, 46(3), pp. 1189–1218. <https://doi.org/10.25300/MISQ/2022/16213>

EDMC Expert Panel — Case Note Evaluation (3 Cases)

👏 **Welcome, and thank you for participating!**

This doctoral research develops and validates the Ethical Decision-Making Compass (EDMC), a first-tier ethical filter for real-world leadership decisions. As part of the multi-method validation, this expert review assesses how the EDMC case note supports timely, transparent, and auditable decisions across contexts.

What you'll do (≈15–20 minutes). For three short, public-source cases you will: (1) record a quick baseline view, (2) review an EDMC case note, and (3) rate clarity, sufficiency, defensibility, portability, and usability—noting any change in stance or confidence.

Anonymity. We don't collect names or emails. You will enter a participant code (from your NDA email) to verify consent; the code is stored for admin and removed before analysis. Results are used only for academic research and reported in aggregate.

Answer candidly. There are no right or wrong answers. Lower confidence or a changed stance is useful evidence.

Scales. Agreement items use 1–5 (Strongly disagree → Strongly agree). Confidence uses 0–10.

Time. 15–20 minutes (one sitting). If you're not signed in to Google, progress may not save.

Materials. Case notes are provided for study use only under NDA; please don't redistribute.

On the next page, please enter your participant code (from your NDA email). We welcome diverse professional and cultural perspectives—they're essential to this validation.

* Indicates required question

Privacy notice

Privacy & Informed Consent:

Controller/PI: Indira Bunic — carrieny1@hotmail.com.

Purpose: Academic research to evaluate the EDMC framework.

What you'll do: One online survey (≈15–20 minutes) reviewing three public-source cases and an EDMC case note.

Risks & discomforts: Minimal. You may skip any question you prefer not to answer.

Benefits: No direct personal benefit; your input advances research on ethical decision-making.

Voluntary participation: Participation is voluntary. You may stop at any time without penalty.

Eligibility: Adults (18+) only.

Data collected: Survey answers and an **NDA participant code** (no names/emails).

Confidentiality: Results are reported only in aggregate; no individual will be identified.

Data handling: Survey data are stored securely (Google Forms/Drive) with limited access. NDAs are collected **outside** the survey and stored separately; the **code-to-person mapping key is never stored with survey responses**.

Lawful basis: Consent to participate; scientific research safeguards (GDPR Art. 6(1) (a), Art. 89).

International transfers: Google may process data outside the EU using Standard Contractual Clauses.

Retention: Data are **anonymized after analysis** and retained in de-identified form for **[X years]** per university policy; the mapping key is **destroyed at anonymization**.

Withdrawal: You may withdraw **before anonymization** by emailing your NDA code to carrieny1@hotmail.com.

Ethics approval: [Insert Ethics/IRB approval reference if available].

(By selecting "I consent," you confirm you've read this summary.)

Consent question (MCQ, Required ON):

I consent to participate in this study.

Data handling: Survey data are stored securely (Google Forms/Drive) with limited access.

NDA's are managed outside the survey and stored separately. The NDA code→person mapping key is kept separate from survey data, restricted to the PI, and destroyed at anonymization.

This study adheres to SSBM Geneva's research ethics policy; if the school requires a formal approval ID, this notice will be updated and communicated to participants.

1. **Participant code (from your NDA email) ***

Please enter your code exactly as shown (e.g., NDA-07), don't enter only numbers, but NDA-X!

Consent

2. **I consent to participate in this study ***

Mark only one oval.

Yes

No → Decline / End section

Demographic Information

To help interpret results across contexts, please answer the three short questions below.

Responses are reported only in aggregate.

3. **Demographic Information Gender:**

Mark only one oval.

- Male
- Female
- Non-binary
- Prefer not to say
- Prefer to self-describe

4. **Demographic Information Age:**

Mark only one oval.

- 18-29
- 30-39
- 40-49
- 50-59
- 60 and above

5. **Highest Level of Education:**

Mark only one oval.

- Bachelor's degree
- Master's degree
- Doctorate (PhD, DBA, etc.)
- Other: _____

6. **Field of Expertise: ***

Mark only one oval.

- Leadership
- Technology
- Organizational Behavior
- Ethics
- Other: _____

7. **Years of Experience in the Field: ***

Mark only one oval.

- 0-5 years
- 6-10 years
- 11-15 years
- 16-20 years
- 21 years and above

8. **Current Role/Position: ***

Mark only one oval.

- Academic (Professor, Researcher, etc.)
- Industry Professional (Manager, Executive, etc.)
- Executive/Board
- Senior management
- Specialist/Advisor
- Other

9. **Primary domain:** *

Mark only one oval.

- Risk/Audit
- Operations
- HR/People
- Technology/Data/AI
- Governance/Legal/Compliance
- Sustainability/ESG
- Healthcare/Public sector
- Education
- Finance/Banking
- Other

10. **Sector:** *

Mark only one oval.

- Public
- Private
- Nonprofit/NGO
- Academia
- Multilateral/International Org
- Other

11. **Country/Region of Residence: ***

Mark only one oval.

- Africa
- Asia-Pacific
- Europa
- Latin America & Caribbean
- Middle East
- North America

Case 1 Vignette (read)

AI Hiring Bias (Technology / HR)

Context. A multinational software company pilots an AI résumé-screening tool for junior engineers. After ~6 weeks (~2,400 applications), women advance to interview at 61% the rate of men (risk ratio 0.61). Among those interviewed, test scores show no gender difference. Training data includes the firm’s historical hires plus public tech résumés. Media scrutiny on AI bias is high. US/EU contexts apply; EU AI Act conformity assessments expected in 12–18 months. Vendor offers bias-mitigation options (+6–8 weeks). Manual screening is feasible but slows hiring ~3–4 weeks per req.

Decision question. After reviewing the situation, should the organization proceed as proposed, proceed with conditions (e.g., bias audit + HITL + thresholds), or pause/do not proceed?

Instruction. Answer the baseline questions before viewing the EDMC case note.

(Confidential – EDMC study/NDA.)

12. **Baseline decision for Case 1 – choose one ***

Mark only one oval.

- Proceed
- Do not proceed
- Proceed with conditions (describe below)

13. **If you chose “Proceed with conditions”, list the key condition(s)**

14. **Confidence in your baseline decision (0–10) ***

Labels: 0 = Not confident ... 10 = Completely confident

Mark only one oval.

0 1 2 3 4 5 6 7 8 9 10

15. **Brief rationale for your baseline decision (2–3 sentences)**

Case 1 — EDMC case note (read now)

Please open and review the one-page EDMC case note for Case

[Case 1](#)

Study materials are for your personal review only – please don't download, copy, or redistribute.

Mobile tip: Press and hold the link → "Open in New Tab," then return here to answer.

Case 1 — After EDMC (ratings & comments)

16. **The framework as applied in this case synthesized the case's complexity into an actionable pathway—linking values, concrete actions (with responsible roles and artefacts), needed capabilities, and measurable outcomes—so that I could formulate and substantiate a defensible course of action with clearly stated conditions.** *

Actionable pathway = values (Why?) → actions with roles/artefacts (How?) → capabilities (Who?) → outcomes/KPIs & mitigations (What?).

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

17. **The framework as applied in this case made the ethically salient dimensions explicit and case-specific (e.g., rights/harms, fairness/equity, duties/obligations, conflicts of interest/power, affected stakeholders).** *

"Ethical dimensions" = clearly named issues relevant to this case (not generic labels), with enough specificity to guide analysis and trade-offs.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

18. **The framework prompted additional ethically relevant considerations beyond my initial view.** *

Rate higher only if it added new, material considerations; if your initial view already covered them, rate lower.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

19. **Within a short, practical time window, the framework enabled an ethics-first preliminary judgement that can then be aligned with compliance requirements.** *

Ethics-first = values, stakeholder impacts, risks/mitigations, and conditions to proceed—compatible with, but not dependent on, detailed policies/laws.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

20. **The justification produced using this framework would withstand reasonable scrutiny from key stakeholders and oversight bodies.** *

Think transparency of assumptions, traceability from values → actions → outcomes, and auditable artefacts.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

21. **As applied here, the framework is applicable across cultures, sectors, and jurisdictions with minimal adaptation.** *

Hypernorms/values are explicit; local requirements are captured as conditions to proceed; language avoids context-specific bias.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

22. **I could apply this framework within our existing decision-making and approval processes with only minor adjustments.** *

Minor adjustments = small tweaks (e.g., add a field or checklist item). No new processes.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

23. **I would use this framework in my organization.** *

For decisions of similar complexity.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

24. **The EDMC case note made the core ethical principles explicit and, where relevant, referenced applicable norms/policies/standards.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

25. **The EDMC case note stated concrete actions with responsible roles and needed artefacts (docs/systems), feasible in practice.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

26. **The EDMC case note identified required capabilities/skills/resources to execute the actions.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

27. **The EDMC case note specified intended outcomes, KPIs, and key risk mitigations/conditions so progress and safeguards are auditable.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

Before the next question, please read the following option definitions.

Proceed as proposed – go ahead now as planned.

- Meaning: ethically acceptable as-is; no extra safeguards needed.
- When: values respected; actions/roles/artefacts clear; capabilities exist; KPIs & follow-up in place; residual risk low.

Proceed with conditions – go ahead only if specific safeguards/mitigations are added.

- Meaning: acceptable if gaps are closed (risk, transparency, fairness, governance).
- Typical conditions: stakeholder consultation; DPIA/ethics review; bias/robustness test; escalation path; add KPI; training/capability; independent oversight.

Pause / do not proceed – stop or defer until substantial issues are addressed.

- Meaning: not acceptable in current form.
- When: high risk of harm/rights infringement; conflicts of interest; missing capabilities; legal/regulatory red flags; unacceptable residual risk.

If none fits exactly, choose “Other”, and add a brief note (e.g., pilot only; redesign; seek independent opinion).

28. What is your current stance after applying the framework to this case? *

After EDMC (ratings & comments)

Mark only one oval.

- Proceed as proposed
- Proceed with conditions
- Pause / do not proceed
- Other: _____

29. If you selected ‘Proceed with conditions’, list the 1–3 key conditions required to proceed.

Answer only if you selected ‘Proceed with conditions’ above (list 1–3 brief conditions).

30. **How confident are you in your course of action after applying the framework to this case?** *

Mark only one oval.

0 1 2 3 4 5 6 7 8 9 10

Not Completely confident

31. **Brief comment (e.g., what shifted, or anything missing/over-emphasized?)**
1–3 sentences are enough.

32. **To what extent did applying the EDMC framework affect the efficiency of your decision-making for this case, compared with your initial approach before using any structured framework?** *

Mark only one oval.

- To a significant extent – less time required ($\approx 50\%$ + reduction)
- To a moderate extent – less time required ($\approx 10\text{--}49\%$ reduction)
- No meaningful effect ($\pm 10\%$)
- To a moderate extent – more time required ($\approx 10\text{--}49\%$ increase)
- To a significant extent – more time required ($\approx 50\%$ + increase)
- Uncertain / unable to assess

36. **Brief rationale for your baseline decision (2–3 sentences)**

Case 2 – EDMC case note (read now)

Please open and review the one-page EDMC case note for Case

[Case 2](#)

Study materials are for your personal review only – please don't download, copy, or redistribute.

Mobile tip: Press and hold the link → “Open in New Tab,” then return here to answer.

Case 2 – After EDMC (ratings & comments)

37. **The framework as applied in this case synthesized the case's complexity into an actionable pathway—linking values, concrete actions (with responsible roles and artefacts), needed capabilities, and measurable outcomes—so that I could formulate and substantiate a defensible course of action with clearly stated conditions.** *

Actionable pathway = values (Why?) → actions with roles/artefacts (How?) → capabilities (Who?) → outcomes/KPIs & mitigations (What?).

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

38. **The framework as applied in this case made the ethically salient dimensions explicit and case-specific (e.g., rights/harms, fairness/equity, duties/obligations, conflicts of interest/power, affected stakeholders).** *

"Ethical dimensions" = clearly named issues relevant to this case (not generic labels), with enough specificity to guide analysis and trade-offs.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

39. **The framework prompted additional ethically relevant considerations beyond my initial view.** *

Rate higher only if it added new, material considerations; if your initial view already covered them, rate lower.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

40. **Within a short, practical time window, the framework enabled an ethics-first preliminary judgement that can then be aligned with compliance requirements.** *

Ethics-first = values, stakeholder impacts, risks/mitigations, and conditions to proceed—compatible with, but not dependent on, detailed policies/laws.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

41. **The justification produced using this framework would withstand reasonable scrutiny from key stakeholders and oversight bodies.** *

Think transparency of assumptions, traceability from values → actions → outcomes, and auditable artefacts.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

42. **As applied here, the framework is applicable across cultures, sectors, and jurisdictions with minimal adaptation.** *

Hypernorms/values are explicit; local requirements are captured as conditions to proceed; language avoids context-specific bias.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

43. **I could apply this framework within our existing decision-making and approval processes with only minor adjustments.** *

Minor adjustments = small tweaks (e.g., add a field or checklist item). No new processes.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

44. **I would use this framework in my organization.** *

For decisions of similar complexity.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

45. **The EDMC case note made the core ethical principles explicit and, where relevant, referenced applicable norms/policies/standards.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

46. **The EDMC case note stated concrete actions with responsible roles and needed artefacts (docs/systems), feasible in practice.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

47. **The EDMC case note identified required capabilities/skills/resources to execute the actions.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

48. **The EDMC case note specified intended outcomes, KPIs, and key risk mitigations/conditions so progress and safeguards are auditable.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

Before the next question, please read the following option definitions.

Proceed as proposed – go ahead now as planned.

- Meaning: ethically acceptable as-is; no extra safeguards needed.
- When: values respected; actions/roles/artefacts clear; capabilities exist; KPIs & follow-up in place; residual risk low.

Proceed with conditions – go ahead only if specific safeguards/mitigations are added.

- Meaning: acceptable if gaps are closed (risk, transparency, fairness, governance).
- Typical conditions: stakeholder consultation; DPIA/ethics review; bias/robustness test; escalation path; add KPI; training/capability; independent oversight.

Pause / do not proceed – stop or defer until substantial issues are addressed.

- Meaning: not acceptable in current form.
- When: high risk of harm/rights infringement; conflicts of interest; missing capabilities; legal/regulatory red flags; unacceptable residual risk.

If none fits exactly, choose “Other”, and add a brief note (e.g., pilot only; redesign; seek independent opinion).

49. What is your current stance after applying the framework to this case? *

After EDMC (ratings & comments)

Mark only one oval.

- Proceed as proposed
- Proceed with conditions
- Pause / do not proceed
- Other: _____

50. If you selected ‘Proceed with conditions’, list the 1–3 key conditions required to proceed.

Answer only if you selected ‘Proceed with conditions’ above (list 1–3 brief conditions).

51. **How confident are you in your course of action after applying the framework to this case?** *

Mark only one oval.

1 2 3 4 5 6 7 8 9 10

Not Completely confident

52. **Brief comment (e.g., what shifted, or anything missing/over-emphasized?)**
1–3 sentences are enough.

53. **To what extent did applying the EDMC framework affect the efficiency of your decision-making for this case, compared with your initial approach before using any structured framework?** *

Mark only one oval.

- To a significant extent – less time required ($\approx 50\%$ + reduction)
- To a moderate extent – less time required ($\approx 10\text{--}49\%$ reduction)
- No meaningful effect ($\pm 10\%$)
- To a moderate extent – more time required ($\approx 10\text{--}49\%$ increase)
- To a significant extent – more time required ($\approx 50\%$ + increase)
- Uncertain / unable to assess

57. **Brief rationale for your baseline decision (2–3 sentences)**

Case 3 – EDMC case note (read now)

Please open and review the one-page EDMC case note for Case

[Case 3](#)

Study materials are for your personal review only – please don't download, copy, or redistribute.

Mobile tip: Press and hold the link → “Open in New Tab,” then return here to answer.

Case 3 – After EDMC (ratings & comments)

58. **The framework as applied in this case synthesized the case's complexity * into an actionable pathway—linking values, concrete actions (with responsible roles and artefacts), needed capabilities, and measurable outcomes—so that I could formulate and substantiate a defensible course of action with clearly stated conditions.**

Actionable pathway = values (Why?) → actions with roles/artefacts (How?) → capabilities (Who?) → outcomes/KPIs & mitigations (What?).

Mark only one oval.

1 2 3 4 5

Strongly Strongly agree

59. **The framework as applied in this case made the ethically salient dimensions explicit and case-specific (e.g., rights/harms, fairness/equity, duties/obligations, conflicts of interest/power, affected stakeholders).** *

"Ethical dimensions" = clearly named issues relevant to this case (not generic labels), with enough specificity to guide analysis and trade-offs.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

60. **The framework prompted additional ethically relevant considerations beyond my initial view.** *

Rate higher only if it added new, material considerations; if your initial view already covered them, rate lower.

Mark only one oval.

1 2 3 4 5

Stro Strong agree

61. **Within a short, practical time window, the framework enabled an ethics-first preliminary judgement that can then be aligned with compliance requirements.** *

Ethics-first = values, stakeholder impacts, risks/mitigations, and conditions to proceed—compatible with, but not dependent on, detailed policies/laws.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

62. **The justification produced using this framework would withstand reasonable scrutiny from key stakeholders and oversight bodies.** *

Think transparency of assumptions, traceability from values → actions → outcomes, and auditable artefacts.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

63. **As applied here, the framework is applicable across cultures, sectors, and jurisdictions with minimal adaptation.** *

Hypernorms/values are explicit; local requirements are captured as conditions to proceed; language avoids context-specific bias.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

64. **I could apply this framework within our existing decision-making and approval processes with only minor adjustments.** *

Minor adjustments = small tweaks (e.g., add a field or checklist item). No new processes.

Mark only one oval.

1 2 3 4 5

Stro Strong agree

65. **I would use this framework in my organization. ***

For decisions of similar complexity.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

66. **The EDMC case note made the core ethical principles explicit and, where relevant, referenced applicable norms/policies/standards. ***

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

67. **The EDMC case note stated concrete actions with responsible roles and needed artefacts (docs/systems), feasible in practice. ***

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

68. **The EDMC case note identified required capabilities/skills/resources to execute the actions.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

69. **The EDMC case note specified intended outcomes, KPIs, and key risk mitigations/conditions so progress and safeguards are auditable.** *

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

Before the next question, please read the following option definitions.

Proceed as proposed – go ahead now as planned.

- Meaning: ethically acceptable as-is; no extra safeguards needed.
- When: values respected; actions/roles/artefacts clear; capabilities exist; KPIs & follow-up in place; residual risk low.

Proceed with conditions – go ahead only if specific safeguards/mitigations are added.

- Meaning: acceptable if gaps are closed (risk, transparency, fairness, governance).
- Typical conditions: stakeholder consultation; DPIA/ethics review; bias/robustness test; escalation path; add KPI; training/capability; independent oversight.

Pause / do not proceed – stop or defer until substantial issues are addressed.

- Meaning: not acceptable in current form.
- When: high risk of harm/rights infringement; conflicts of interest; missing capabilities; legal/regulatory red flags; unacceptable residual risk.

If none fits exactly, choose "Other", and add a brief note (e.g., pilot only; redesign; seek independent opinion).

70. **What is your current stance after applying the framework to this case? ***

After EDMC (ratings & comments)

Mark only one oval.

- Proceed as proposed
- Proceed with conditions
- Pause / do not proceed
- Other: _____

71. If you selected 'Proceed with conditions', list the 1–3 key conditions required to proceed.

Answer only if you selected 'Proceed with conditions' above (list 1–3 brief conditions).

72. **How confident are you in your course of action after applying the framework to this case? ***

Mark only one oval.

0 1 2 3 4 5 6 7 8 9 10

Not Completely confident

73. **Brief comment (e.g., what shifted, or anything missing/over-emphasized?)**

1–3 sentences are enough.

74. **To what extent did applying the EDMC framework affect the efficiency of your decision-making for this case, compared with your initial approach before using any structured framework?** *

Mark only one oval.

- To a significant extent – less time required ($\approx 50\%$ + reduction)
- To a moderate extent – less time required ($\approx 10\text{--}49\%$ reduction)
- No meaningful effect ($\pm 10\%$)
- To a moderate extent – more time required ($\approx 10\text{--}49\%$ increase)
- To a significant extent – more time required ($\approx 50\%$ + increase)
- Uncertain / unable to assess

Global wrap-up

You've completed three baseline decisions and reviewed three EDMC case notes. Please give your overall view of clarity, practicality, cross-cultural portability, and tech relevance.

75. **Across the three cases, the framework enabled an actionable, defensible decision pathway.** *

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

76. **In how many of the three cases did the framework feel sufficient for a defensible course of action?** *

Mark only one oval.

- 0
- 1
- 2
- 3

77. **My speed and ease of using the framework improved from Case 1 to Case 3.** *

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

78. **Across contexts and jurisdictions, the EDMC case note felt portable (cross-cultural / cross-jurisdiction).** *

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

79. **I could use the EDMC case note in our existing decision/approval process with only minor tweaks.** *

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

80. **I could apply the EDMC case note under time pressure (about 15 minutes) and still make a defensible decision.** *

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

81. **The EDMC case note made the logic from principles → actions (with roles/artefacts) → capabilities → outcomes clear (conceptual coherence).** *

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

82. **For technology/AI decisions, the EDMC case note can capture needed data/model/privacy controls when relevant.**

(If not applicable to your work, you may skip this item.)

Mark only one oval.

1 2 3 4 5

Stro Strongly agree

83. **Across the three cases, how consistently did each PASO element appear?"** *

(Options: All 3 cases / 2 cases / 1 case / Not present – keep as is.)

Principles = why; Actions/Roles/Artefacts = how/by whom/with what; Capabilities = who/skills/resources; Outcomes = results + mitigations.

Mark only one oval per row.

	All 3 cases	2 cases	1 case	Not present
Principles explicit & traceable	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Actions/roles/artefacts specified	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Capabilities identified	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Outcomes/KPIs & mitigations stated	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

84. **If anything would limit the practical use of the EDMC case note in your context (e.g., approval flow, legal wording, tooling), describe it briefly (1–2 sentences).**


Please avoid naming specific people or organizations.

85. **What one change would most improve the usefulness of the EDMC case note in your setting?**

1–2 concrete suggestions are ideal.

Closing Statement

Thank you for completing the questionnaire. Your feedback directly strengthens the EDMC framework and its practical usefulness across contexts.

 **We appreciate your expertise and time. Thank you for supporting ethical leadership and decision-making.**

“Leadership is not about making the easy choice, but the right one. Ethics is the compass that keeps us on course.”

Ethical Decision-Making Compass (EDMC)

Privacy note. Responses are anonymous (no names/emails collected). Your NDA code is used only to verify consent and is removed before analysis. Data are stored securely, reported only in aggregate, and used solely for this research.