Resilience Engineering in Healthcare

PI–RA: A Performance Indicator Resilience Assessment Framework for Emergency Departments

Master's Thesis

Alp Engin Kuvrag

Delft University of Technology

TU Delft

Resilience Engineering in Healthcare

PI–RA: A Performance Indicator Resilience Assessment Framework for Emergency Departments

by

Alp Engin Kuvrag

MSc Complex Systems Engineering and Management

Faculty of Technology, Policy & Management

Student Number:	5010535
Chair:	Perla Marang-van de Mheen
First Supervisor:	Ming Yang
Second Supervisor:	Yilin Huang
Date:	Sunday 7th December, 2025

TU Delft

Executive Summary

This thesis develops and tests a way to read emergency department (ED) performance indicators as direct evidence of resilient performance during disruptions, rather than as disconnected "better or worse" numbers. It focuses on how concrete work adaptations during COVID-19—like new isolation protocols, rapid assessment areas, and point-of-care testing—changed ED performance, and how those changes can be systematically translated into resilient performance profiles and quality trade-off narratives.

Background and problem

Emergency Departments operate under constant pressure to deliver fast, safe, and efficient care with finite resources. These pressures intensified during the COVID-19 pandemic, when EDs had to adjust their operations repeatedly while still maintaining core care functions. Performance indicators (PIs) such as length of stay, waiting times, and left-without-being-seen rates are widely used to monitor quality, but they are usually treated as isolated metrics or crude targets. Resilience Engineering and Safety-II emphasize how systems adapt under stress, yet existing tools typically produce qualitative capability profiles that are weakly linked to day-to-day operational performance.

The thesis identifies a central gap: there is no widely adopted method that uses routinely observed ED performance data to make resilience measurable, interpretable, and comparable, including its implications for the quality of care. As a result, resilience assessments often remain abstract, and they struggle to show concretely how disruptions and work adaptations affect real operations.

Research objective and questions

To address this gap, the thesis develops the Performance Indicator Resilience Assessment (PI–RA) framework, which links work adaptations to observable changes in performance indicators and to their associated quality trade-offs. PI–RA translates heterogeneous case evidence into a transparent read-out of *resilient performance*. In this thesis, resilient performance is interpreted using the resilience curve in Figure 1: a disruption pushes ED performance away from its usual level, after which the system may stabilize in a degraded state, recover back towards the baseline, or even improve beyond it. PI–RA can not measure the exact depth of the drop, but it uses before–after patterns in performance indicators to classify where the ED ends up on this curve—whether required operations remain degraded, move back onto a recovery trajectory, or improve with limited trade-offs—and what this implies for the quality of care.

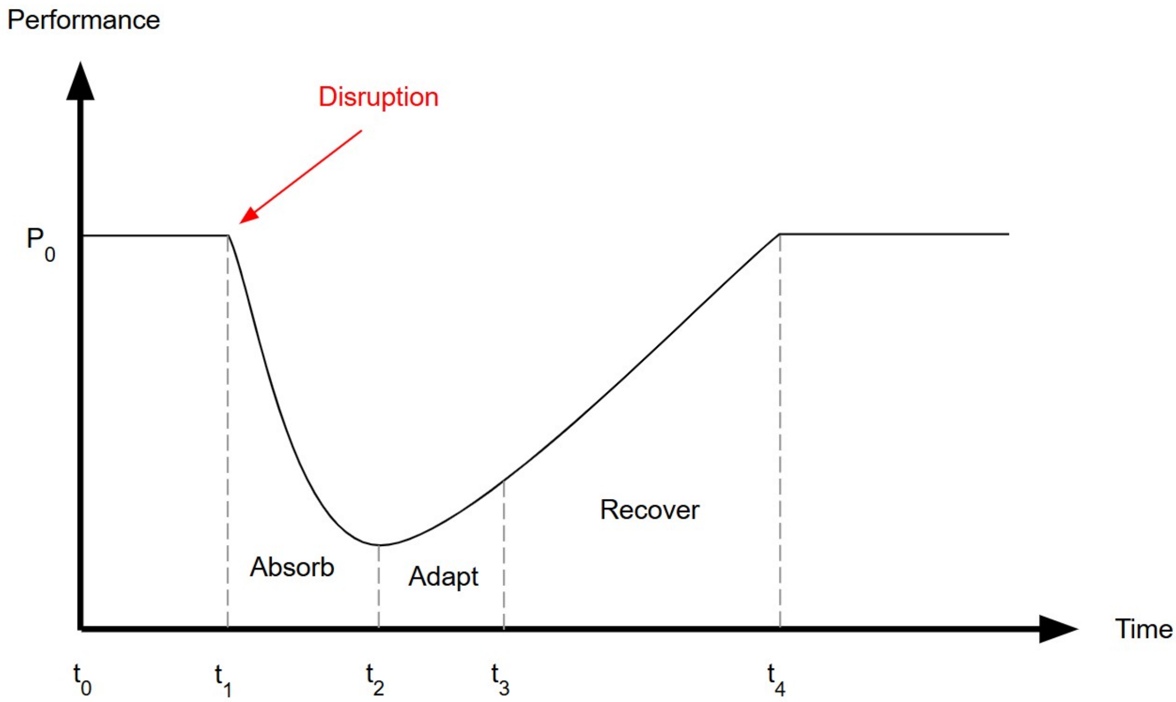

Figure 1: Resilience Curve: System performance trajectory across disruption, absorption, adaptation, and recovery

The objective is to use ED performance indicators to describe these performance changes and to make explicit which Institute of Medicine (IOM) quality dimensions—safety, timeliness, effectiveness, efficiency, equity, and patient-centredness—are improved, sacrificed, or left unchanged.

The main research question is:

How can observed changes in ED performance indicators be translated into a transparent assessment of resilient performance, and what do these assessments reveal about quality of ED care?

Three sub-questions structure the work:

1. ED workflow & critical function: What are the ED's critical function and workflow?
2. Performance & work adaptations: How can ED work adaptations in response to disruptions (e.g. COVID-19) be analyzed systematically to extract structured evidence and insights on performance?
3. Resilient performance: How can the structured evidence and insights on ED performance be used to assess the resilient performance of different work adaptations?

Conceptual foundations and Dutch ED context

The thesis first situates Dutch EDs within a strongly gatekept acute-care chain. General practitioners (GPs) and out-of-hours cooperatives filter non-urgent complaints, while Emergency Medical Services perform protocol-guided assessment and transport decisions in the field. Much triage, urgency assignment, and documentation happens before ED entry and is supported by standardised digital information exchange.

Given this structure, required operations are located mainly in the ED throughput and immediate post-ED phases. The thesis reconstructs a generic ED workflow and defines the ED's critical function as four tightly connected elements: rapid assessment, correct diagnosis, initiation of appropriate treatment, and management of patient flow such that these activities can be sustained under variable demand and disruption. These elements are mapped onto the workflow phases, clarifying where resilient performance must be maintained.

The literature review then covers healthcare quality measurement and the IOM quality dimensions, arguing that multiple indicators must be interpreted together and that single metrics can be politically misused or misread. This motivates the first phase of the PI–RA framework, which provides a structured way to read cases and interpret indicator changes in context, formulating consistent IF–THEN descriptions of how specific work adaptations influence performance and quality.

Methodology and PI-RA framework

Initially, the study intended to use Dutch COVID-era performance data to track how ED performance indicators changed across pandemic phases and to link these trends to locally implemented work adaptations. Practical barriers made this plan infeasible: suitable ED-level datasets were not accessible and timelines of Dutch ED adaptations were poorly documented. The thesis therefore pivots to an international case-based approach, analyzing detailed empirical studies that describe work adaptations and report pre–post changes in performance indicators.

To interpret these heterogeneous cases consistently, the thesis develops the three-phase Performance Indicator Resilience Assessment (PI–RA) framework (Figure 1.2). Phase 1, *structured performance change description*, turns each case into a standardized account of how performance and quality changed. The case is summarized with a performance focus; all reported indicators are classified using the Donabedian model and tagged with their relevant IOM quality dimensions; the direction and magnitude of change are coded using pragmatic percentage bins; contextual mechanisms and implications are interpreted; and the findings are condensed into a single IF–THEN statement that lists which indicators changed, by how much, and what this implies for different dimensions of care quality.

Phase 2, *system mapping and attribution*, uses the Dutch ED workflow and critical function model to locate where each work adaptation acts in the system and which parts of the critical function it targets. Performance indicators are assigned to workflow phases and critical-function elements, their IOM tags are retained, and each indicator is given an attribution weight w_i that reflects how strongly its observed change can credibly be linked to the adaptation rather than to other factors. The coded magnitude and direction of change are converted into signed indicator scores s_i, providing a common numeric scale for aggregation.

Phase 3, *resilient performance interpretation*, combines these weighted indicator scores into three composite measures: the performance change magnitude PC_M (a relevance-weighted mean summarizing the net direction and strength of change), the performance change breadth PC_B (the net number of weighted indicators improving versus worsening), and the trade-off coefficient TO_C (a normalized breadth measure between -1 and $+1$ that shows how dominant improvements or deteriorations are). Each work adaptation is positioned on the Performance Change Matrix using PC_M and TO_C, and this position is interpreted against the resilience curve and IOM quality tags to derive a resilient performance profile and its quality-of-care implications.

Application: three COVID-related work adaptations

The framework is applied to three detailed COVID-era work adaptations with traceable performance indicator changes:

- **Enhanced isolation protocol for fever or respiratory patients (FRPs)** – an extra isolation step at the front end to reduce contagion. Some end-to-end timeliness and efficiency indicators improve, but front-end waiting times increase, left-without-being-seen (LWBS) risks rise, and large equity gaps appear between patient groups. Once the intended safety benefit is taken into account, the overall resilient performance profile is almost neutral and entirely trade-off driven ($PC_M = -0.0375$, $PC_B = 0$, $TO_C = 0$): the ED's required operations remain in a slightly degraded state rather than clearly recovering.
- **Rapid Assessment Zone (RAZ)** – an alternative front-end assessment path that brings initial provider evaluation forward for many non-critical patients. RAZ strengthens rapid assessment and flow management, with clear improvements in timeliness and safety and no observed deterioration in other scored dimensions. The resulting scores ($PC_M = +0.60$, $PC_B = 3$, $TO_C = +1.0$) indicate that all relevant indicators move in the desired direction and that the adaptation helps the ED move onto a recovery trajectory.
- **Point-of-Care Testing (POCT)** – a bedside SARS-CoV-2 testing strategy in the middle of the throughput phase that directly informs admission and discharge decisions. POCT improves timeliness (shorter LOS and time to result), effectiveness (better diagnostic coverage and more targeted admissions), and safety (fewer intrahospital transfers), alongside modest efficiency gains, without visible trade-offs. It achieves the highest scores of the three cases ($PC_M \approx +0.67$, $PC_B = 5.5$, $TO_C = +1.0$), showing that required operations are strengthened across multiple quality dimensions while the system enters a clear recovery phase.

Taken together, the cases illustrate how PI–RA distinguishes a trade-off-heavy adaptation that fails to restore performance (FRPs) from two adaptations (RAZ and POCT) that support recovery.

Main answer, contributions, and implications

Overall, the thesis shows that the PI–RA framework can translate observed changes in ED performance indicators into transparent resilient-performance profiles. By linking indicators to a concrete ED workflow and critical function, reading their changes in terms of IOM quality dimensions, and combining them into the scores PC_M, PC_B and TO_C, the framework makes it visible whether a work adaptation leaves the ED in a degraded state, supports recovery, or improves required operations, and which aspects of care quality are gained or sacrificed in the process.

The main contributions are:

- **Conceptual:** making the ED's "required operations" explicit in the Dutch gatekept acute-care chain and anchoring them in a generic ED workflow linked to quality dimensions, so that routine indicators can be read as statements about how well this core function is maintained under disruption.
- **Methodological:** proposing PI–RA as a structured way to interpret heterogeneous indicators, from case reading and IF–THEN descriptions to system mapping, attribution weights, and composite performance-change scores that summarize net change and trade-offs.

- **Practical:** demonstrating the use of PI–RA on three COVID-19 work adaptations, distinguishing a trade-off-heavy response that fails to restore performance (FRPs) from two adaptations (RAZ and POCT) that support recovery, and—especially for POCT—improve the ED's critical function across multiple quality dimensions.

Finally, the thesis sketches directions for future work, including extending PI–RA towards empirical resilience curves and more data-driven weighting of indicators.

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
ACC	Acute Care Chain (Dutch acute care system phases)
ATP	Arrival-to-Provider time (time from ED arrival until first provider contact)
BIG	Wet op de Beroepen in de Individuele Gezondheidszorg (Individual Healthcare Professions Act)
BOR	Emergency Department Bed Occupancy Rate
CBS	Centraal Bureau voor de Statistiek (Statistics Netherlands)
ED	Emergency Department
EDLOS	Emergency Department Length of Stay
EMS	Emergency Medical Services
FGM	Franco-German Model (prehospital emergency care model)
FMS	Federatie Medisch Specialisten (Federation of Medical Specialists)
FRP	Enhanced isolation protocol for fever or respiratory patients
GP	General Practitioner
GPC	GP cooperative (out-of-hours GP clinic located at the ED)
HAP	Huisartsenpost (out-of-hours GP cooperative)
ICT	Informatie- en Communicatietechnologie (Information and Communication Technology)
ICU	Intensive Care Unit
IOM	Institute of Medicine
LCPS	Landelijk Coördinatiecentrum Patiënten Spreiding (National Coordination Centre for Patient Distribution)
LIS	Landelijk Letsel Informatie Systeem (national Injury Information System)
LOS	Length of Stay

Abbreviation	Definition
LPA	Landelijk Protocol Ambulancezorg (National Ambulance Care Protocol)
LWBS	Left Without Being Seen (patients leaving ED before physician assessment)
NHG	Nederlands Huisartsen Genootschap (Dutch College of General Practitioners)
OCEBM	Oxford Centre for Evidence-Based Medicine
OECD	Organisation for Economic Co-operation and Development
PI	Performance Indicator
PI-RA	Performance Indicator Resilience Assessment framework
POCT	Point-of-Care Testing
RAG	Resilience Assessment Grid
RAZ	Rapid Assessment Zone
RE	Resilience Engineering
RIC	Reduced Infection Contagion (assumed safety benefit category)
RIVM	Rijksinstituut voor Volksgezondheid en Milieu (National Institute for Public Health and the Environment)
ROAZ	Regionaal Overleg Acute Zorgketen (Regional Acute Care Consultation Body)
SEH	Spoedeisende Hulp (Dutch term for Emergency Department)
SOEP	Subjectief, Objectief, Evaluatie, Plan (structured clinical note format)
VWS	Ministerie van Volksgezondheid, Welzijn en Sport (Ministry of Health, Welfare and Sport)
VZVZ	Vereniging van Zorgaanbieders voor Zorgcommunicatie (operator of AORTA-LSP)
WA	Work Adaptation
WHO	World Health Organization
WT	Waiting Time (typically time from arrival to being called by provider)

Symbols

Symbol	Definition	Unit
s_i	Signed performance change score for indicator \$i\$ (from strong deterioration to strong improvement)	[-]
w_i	Relevance weight of performance indicator \$i\$	[-]
PC_M	Performance Change Magnitude (relevance-weighted mean of signed indicator scores)	[-]
PC_B	Performance Change Breadth (net number of weighted indicators improving vs deteriorating)	[-]
TO_C	Trade-off coefficient (normalized breadth index, from -1 to +1)	[-]
$P(t)$	Actual performance trajectory over time	[-]
R_{attained}	Attained resilience (area under the actual performance curve)	[performance·time]
R_{lost}	Lost resilience (area between desired and actual performance curves)	[performance·time]

PART I

Foundations

1

Introduction

Healthcare systems operate in complex and dynamic environments where events ranging from global pandemics to routine operational failures can severely impact patient care and safety. Furthermore, hospital departments such as the emergency department include multitasking, overcrowding, and interruptions, and these factors also harm patient safety (Källberg et al., 2015). Despite technological advancements and process improvements, healthcare organizations still struggle to maintain operational resilience in disruptive scenarios, which can hinder care delivery under stress. Resilience engineering provides a promising framework to address this challenge. Unlike traditional risk-based approaches that focus on preventing failure, resilience engineering emphasizes the ability of systems to anticipate, absorb, adapt to, and recover from disruptions (Yang et al., 2023).

1.1. Concepts Introduction

To understand the role of resilience engineering in healthcare, it is first necessary to define what resilience means. A widely cited definition describes it as “the intrinsic ability of a system (a clinic, ward, a hospital, a country) to adjust its functioning before, during, or following events (changes, disturbances, and opportunities) and thereby sustain required operations under both expected and unexpected conditions” (Chuang et al., 2020).

Patriarca et al. (2017) emphasize that resilience is something a system does, not something it has. They define resilience engineering (RE) as the study of what resilient performance looks like, how it can be measured or assessed, and how it can be improved. This perspective highlights resilience as a continuous process of learning and adaptation rather than a fixed system attribute.

In healthcare, unpredictable disturbances in timing, magnitude, duration, and character are common (Arcuri et al., 2022). Practitioners must often adjust their actions to manage unexpected events as they arise (Chuang et al., 2020). As a result, healthcare systems must continually strengthen their capacity to respond and adapt.

Traditional approaches to patient safety, known as Safety-I, focus on minimizing adverse events by identifying and eliminating errors. In contrast, the Safety-II perspective, proposed by Hollnagel, emphasizes understanding how things go right despite uncertainty and constraints. It shifts attention from detecting failures to exploring how everyday work succeeds through adaptation and resilience (Nemeth, 2019). While this already reframes safety as an adaptive capacity, Watt et al. (2019) further emphasize that Safety-II involves the ability to make dynamic trade-offs and adjust performance in response to changing demands and disturbances. For instance, in high-stakes situations like blood transfusion, this can mean weighing immediate clinical risks (e.g., infection) against the urgency of treatment, which shows that adaptation under pressure is not a deviation, but a routine and necessary part of safe care.

Building on these perspectives, this thesis conceptualises resilience as *changes in system performance over time*. Figure 1.1 illustrates this with a simplified performance curve. Before a disruption, between t_0 and t_1, the system operates around a desired performance level P_0; this flat segment represents routine functioning under normal conditions. When the disruption occurs at t_1, the system's ability to *absorb* the shock is reflected in the size and speed of the immediate performance drop during the interval $t_1 \rightarrow t_2$. The subsequent interval $t_2 \rightarrow t_3$ represents *adaptation*: during this phase the system reconfigures its functioning to halt further decline and create the conditions for improvement. Finally, in the interval $t_3 \rightarrow t_4$ the system *recovers*, and performance increases back towards the pre-disruption level P_0. The resulting displacement of performance over time forms what is commonly referred to as the resilience curve.

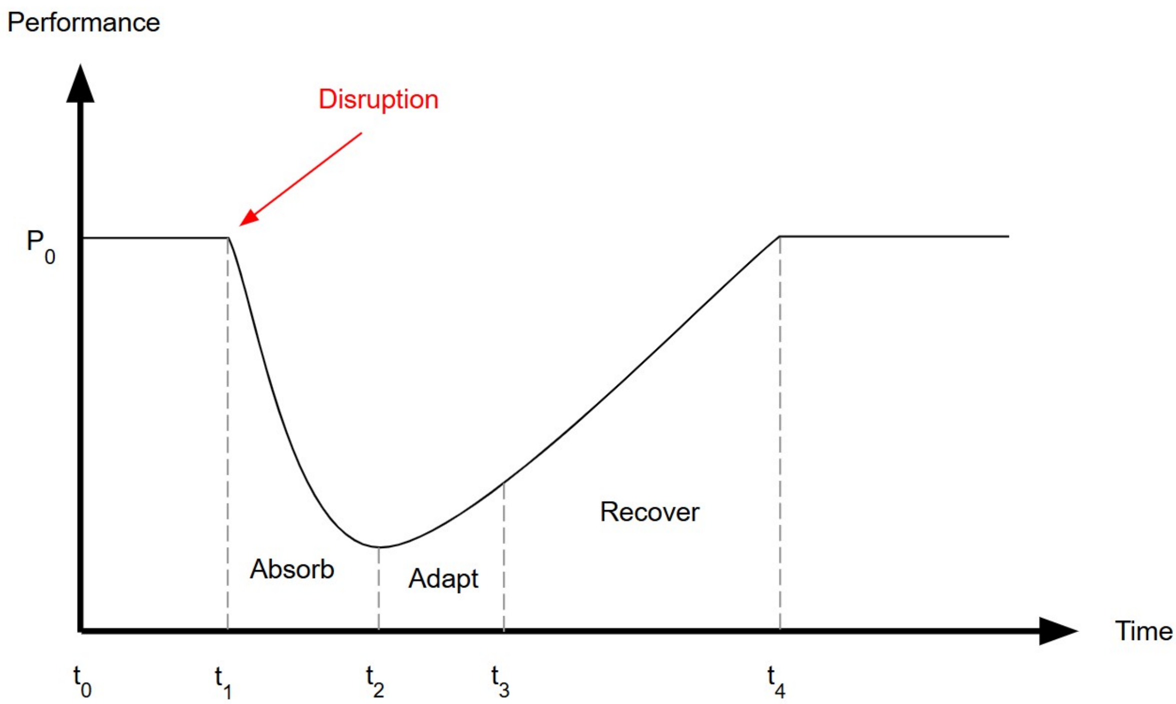

Figure 1.1: Resilience Curve: System performance trajectory across disruption, absorption, adaptation, and recovery

Designing resilient healthcare systems requires more than the ability to respond to disruptions; it also demands a clear understanding of what constitutes high-quality care. To lay this foundation, it is important to define what is meant by quality in the healthcare context. In this thesis, quality is defined based on frameworks from international health organizations. A joint report by the World Health Organization (WHO), the Organisation for Economic Co-operation and Development (OECD), and the World Bank defines healthcare quality as:

"the degree to which health services for individuals and populations increase the likelihood of desired health outcomes and are consistent with current professional knowledge" (World

Health Organization et al., 2018).

This definition emphasizes that quality is not simply a matter of intent or individual expertise, but the ability to consistently deliver care that leads to evidence-based, predictable outcomes. Building on this, the Agency for Healthcare Research and Quality (2015), drawing from the Institute of Medicine (IOM), identifies six core domains of healthcare quality:

- **Safety:** avoiding harm to patients from the care intended to help them.
- **Timeliness:** reducing unnecessary delays for both patients and providers.
- **Effectiveness:** providing scientifically supported services to those who will benefit, while avoiding underuse or misuse.
- **Efficiency:** avoiding waste of resources.
- **Equality:** ensuring care quality does not vary based on personal characteristics.
- **Patient-Centered:** aligning care with patient preferences and values.

Each of the six quality dimensions outlined by the Institute of Medicine can be assessed using performance indicators, which serve as measurable proxies for system performance. As defined by Jones et al. (2014), a quality indicator is "a measurable element of practice performance, for which there is evidence or consensus that it can be used to assess the quality, and hence the change in quality, of care provided." For example, indicators such as ED length of stay and time-to-initial-physician-assessment are commonly used to evaluate timeliness in emergency care (Feral-Pierssens et al., 2024).

However, reliance on such indicators is not without problems. Winslow (2020) shows how metrification and protocolized documentation requirements, introduced to improve accountability, can conflict with clinical judgment and direct patient care. In her study of sepsis quality metrics, hospitals that successfully reduced mortality were still penalized for failing to comply with abstracted reporting protocols — highlighting how regulatory efficiency was prioritized over clinical effectiveness. This reveals a misalignment between what gets measured and what constitutes true clinical quality, especially in dynamic and high-pressure environments like EDs. Cho et al. (2020) highlight that emergency departments are often overcrowded and exposed to unexpected factors, which complicates efforts to manage both care quality and operational efficiency. In such settings, attempts to streamline patient flow may place added pressure on clinical decision-making and staff resources. These trade-offs are central to understanding the adaptive capacity of healthcare systems — a core concern in resilience engineering.

1.2. Research Gap

The initial literature review revealed recurring limitations in how resilience engineering is applied in healthcare settings. The Resilience Assessment Grid (RAG) proposed by Hollnagel (2010) has emerged as a framework for assessing a system's capacity to manage varying conditions via the four abilities (respond, monitor, learn, anticipate). In Hollnagel's view, resilience is not directly measurable; instead, RAG profiles a system's "resilience potentials" over time through tailored self-assessments and facilitated discussions (Safi et al., 2022). Because RAG relies on self-assessment, findings largely reflect perceived capabilities rather than directly observed system behaviour.

Practically, the RAG provides generic guidance and example questions that must be adapted to each specific setting and are typically administered via questionnaires, interviews, or workshops (Safi et al., 2022; Hollnagel, 2010). RAG can be used for repeated measures within a single system to support local improvement, but it is not designed for cross-system comparison and is not considered suitable by its originator for benchmarking across organisations Safi et al. (2022). These observations align with broader conclusions in the literature: while Safety-II and Resilience Engineering offer a compelling shift toward understanding how everyday work succeeds through adaptation, their practical operationalisation in healthcare remains limited; tools like RAG are widely used to profile resilience yet have rarely translated into implemented quality improvements or sustained organisational change (Safi et al., 2022; Chuang et al., 2020); and despite a decade of conceptual development, resilience in healthcare still lacks clear constructs, making it difficult to measure, test, and apply in practice (Berg et al., 2018; Patriarca et al., 2017). *Consistent with this, our search did not identify a widely adopted method that ties resilience assessment to routinely observed operational performance in a way that shows how resilience actually unfolds during disruptive events and what measurable consequences and quality implications it has for care.*

Research Gap: There is no widely adopted method that uses ED performance data to make resilience measurable, interpretable, and comparable, including its implications for the quality of care; as a result, current assessments tend to remain capability profiles without operational read-through to performance consequences.

1.3. Research Objective and Main Research Question

Healthcare systems, and particularly Emergency Departments, operate under constant pressure to meet multiple performance goals—shorter waits, patient safety, and prudent resource use. These goals are tracked with performance indicators that, while essential for assurance, expose real trade-offs (e.g., faster flow at the cost of diagnostic thoroughness or increased staff workload). During disruptions such as pandemics, variability and strain intensify, and the ED must remain resilient—i.e., sustain its required operations despite these pressures.

Motivation: These realities mean that observed changes in indicators must be read not only as "better or worse" numbers, but as signals about how required operations are being sustained—and what quality trade-offs are occurring—under disruption.

Objective: To address the gap identified, this thesis develops and tests a pragmatic evaluation framework that directly links work adaptations to their observable changes in performance indicators, focusing on the adaptation and recovery phases of the resilience curve introduced in Figure 1.1. The objective is to turn heterogeneous case evidence into a transparent, evidence-linked read-out of change in performance, and to explore how these changes can be used to empirically characterize resilience, so that assessments move beyond abstract capability profiles towards measurable operational consequences. Because different parts of emergency department work map onto different IOM quality dimensions, changes in specific activities can improve some dimensions while worsening others. The objective therefore also includes reporting an IOM trade-off narrative alongside the numerical results, ensuring that decisions reflect the multi-dimensional nature of quality rather than a single aggregate number.

This leads to the following main research question:

How can observed changes in ED performance indicators be translated into a transparent assessment of resilient performance, and what do these assessments reveal about quality of ED care?

To answer this question, the thesis develops and applies the *Performance Indicator Resilience*

Assessment (PI–RA) framework. PI–RA is a three-phase framework that uses quantitative Emergency Department (ED) performance indicators (PIs) to provide a transparent read-out of how ED performance changes during disruptions and work adaptations, and what these changes imply for resilience and quality of care. Figure 1.2 shows the PI–RA framework and its three phases.

- **Phase 1 – Structured performance change description.** Observed changes in ED performance indicators during a disruption or following a work adaptation are organised into a common format. Indicators are classified by Donabedian type and IOM quality dimension, their direction and magnitude of change are coded, and the main effects are summarised in concise IF–THEN statements.
- **Phase 2 – System mapping and attribution.** The coded indicator changes are then linked to the ED workflow and its critical function. This phase clarifies which parts of ED work each indicator reflects and how strongly the observed changes can be attributed to a specific disruption or work adaptation, providing the basis for aggregating heterogeneous indicators.
- **Phase 3 – Resilient performance interpretation.** Finally, the weighted indicator changes are combined into performance change scores that express the magnitude, net breadth, and trade-off balance of change (PC_M, PC_B, and TO_C). PC_M, and TO_C scores position each work adaptation on a Performance Change Matrix and are interpreted together with an IOM-based narrative of what improved, what deteriorated, and where, providing an interpretable read-out of resilience in practice.

The following research sub-questions structure the development and application of this framework, while the Research Approach and Methodology (Chapter 3) explains in detail how each phase of PI–RA was operationalized.

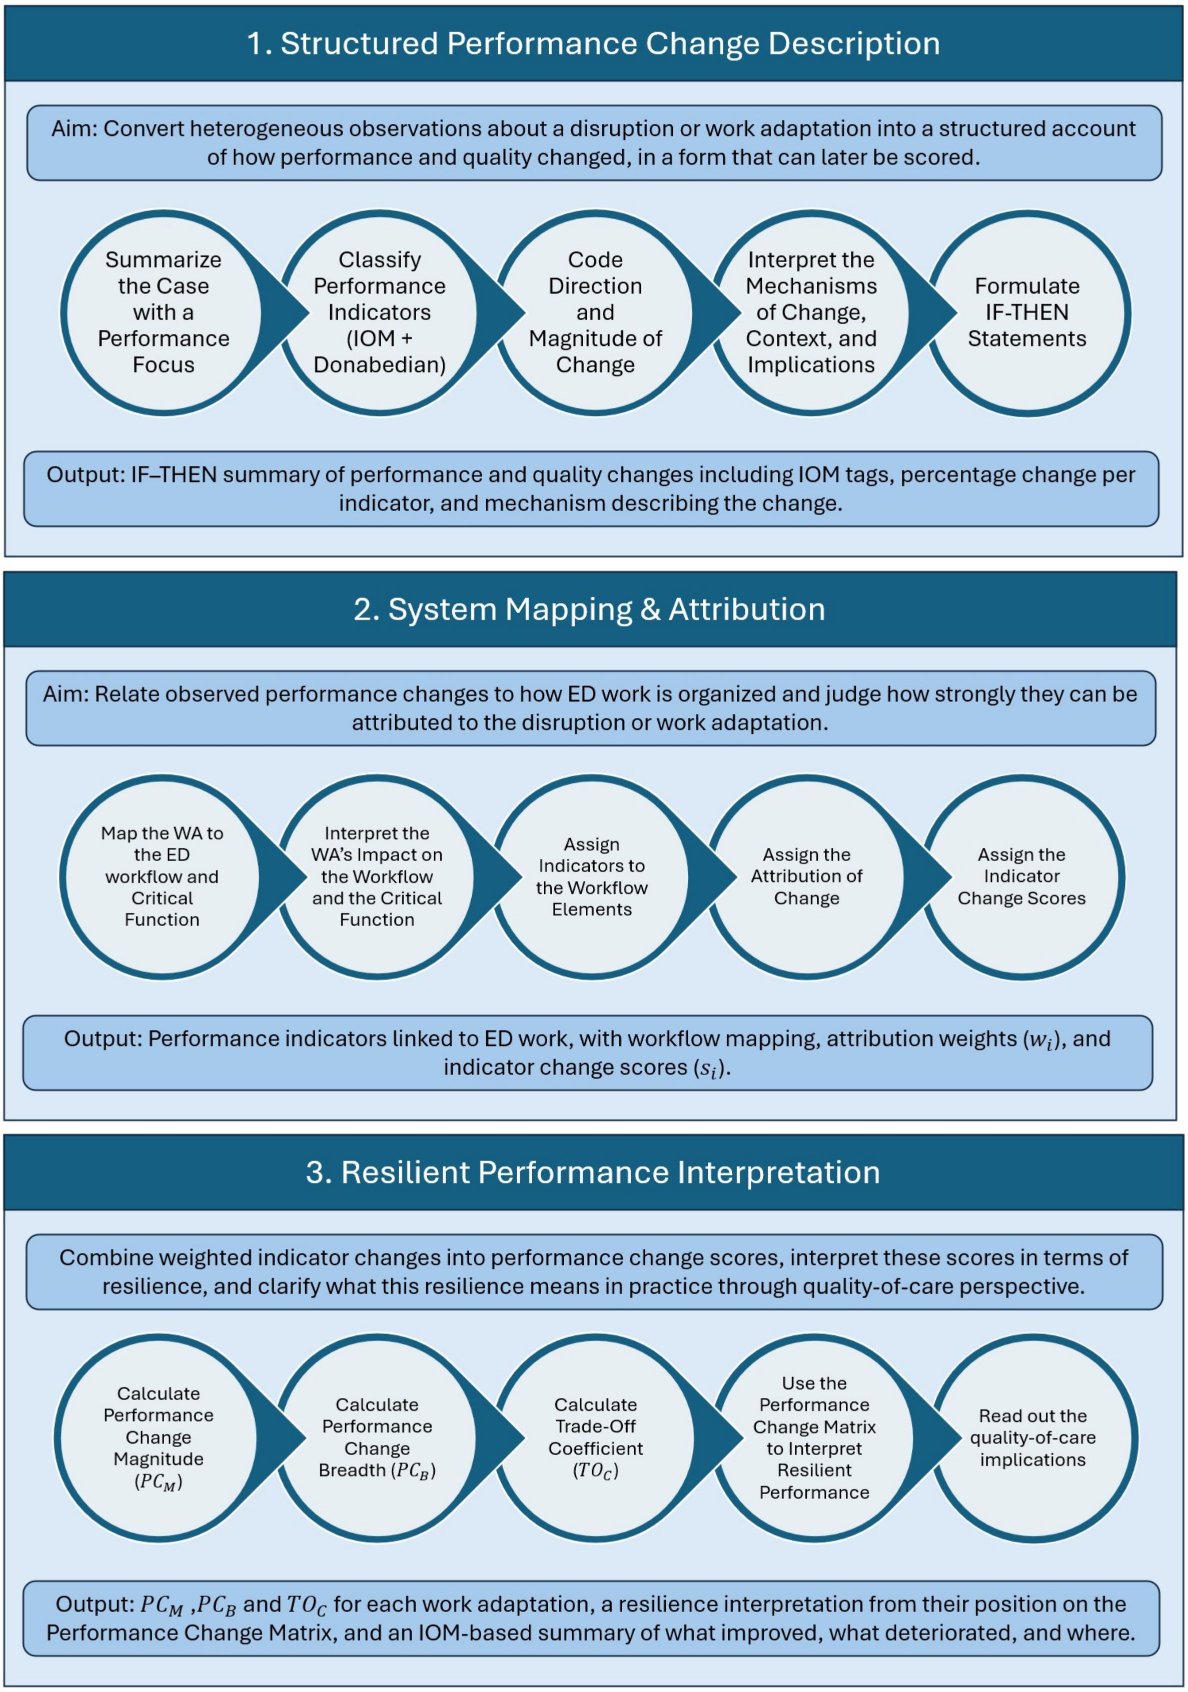

Figure 1.2: Performance Indicator - Resilience Assessment (PI-RA) Framework

1.4. Research Sub-Questions

To structure the development and application of the PI-RA framework, this thesis is guided by three interconnected sub-questions, summarised in Table 1.1.

Table 1.1: Research Sub-questions and Their Focus

Sub-question	Focus
SQ1: What are the ED's critical function and workflow?	This sub-question develops an operational description of the ED's critical function and workflow in the Dutch context. It identifies the main process phases and key steps, clarifies how patients move through the ED, and specifies what must be achieved for the critical function to be fulfilled. These definitions provide the reference system against which later performance changes and resilience assessments are interpreted.
SQ2: How can ED work adaptations in response to disruptions (e.g., COVID-19) be analyzed systematically to extract structured evidence and insights on performance?	This sub-question explains how empirical information on work adaptations and disruptions is converted into structured descriptions of performance change. It introduces a common coding scheme for performance indicators—using Donabedian and IOM classifications, direction and magnitude of change, and IF–THEN summaries—so that heterogeneous studies provide comparable input to the PI–RA framework.
SQ3: How can the structured evidence and insights on ED performance be used to assess the resilient performance of different work adaptations?	This sub-question uses the structured evidence from SQ2 together with the workflow model from SQ1 to construct and interpret performance change scores. It explains how indicators are linked to workflow elements and weighted, how the magnitude and net breadth scores (PC_M and PC_B) and the trade-off coefficient (TO_C) are calculated, and how their combination on the Performance Change Matrix, together with an IOM-based narrative of what improved and what deteriorated, is used to read resilient performance for each work adaptation.

1.5. Societal Relevance of the Study

The societal relevance of the study arises from the fact that it aims to increase public health resilience and support healthcare system efficiency. Pressures in healthcare management—such as misdirected patient assessments or shortages in ICU/ED beds and equipment—can have severe consequences on the quality of care, treatment outcomes, and even mortality rates (Drabecki et al., 2023). The fact that adverse events can result in harm and even death to many patients is already an important societal concern, but such instances also require financial and other resources to be allocated to mitigate harm and compensate victims of patient safety incidents (Anderson et al., 2016).

1.6. Relevance to CoSEM

This thesis contributes to the broader field of Complex Systems Engineering by examining how Emergency Departments function as adaptive systems that must maintain critical performance during disruptions. The focus on resilience, performance trade-offs, and systemic adaptation aligns with CoSEM's central objective of understanding and improving complex socio-technical systems under uncertainty. The research connects empirical performance data to a conceptual framework, reflecting the program's emphasis on bridging data-driven insights with systems thinking. The multi-actor dimension is also relevant: different stakeholders (e.g., clinicians, patients, administrators) interpret and prioritize healthcare quality dimensions differently, which directly relates to the challenges explored in the Managing Multi-actor Decision-making course. Additionally, while the thesis does not model networked structure in detail, it offers a simple, operational workflow model that suggests the ED as a networked system. This model identifies nodes (workflow elements and actors) and links (patient, information, and resource flows) that can be operationalised in future computer models. It builds on the resilience thinking introduced in the Design in Networked Systems course, which highlights resilience as a critical feature of complex systems. In this way, the thesis exemplifies the interdisciplinary, analytical, and systems-oriented mindset that defines the CoSEM program.

2

Theoretical Background and Literature Review

Chapter 2 builds the factual and conceptual scaffolding for the PI-RA framework. It first situates Dutch EDs within their emergency-care setting, contrasting the Franco-German and Anglo-American models to show how upstream filtering shapes the ED's role, then locates the Netherlands on that spectrum (strong GP gatekeeping, advanced EMS decision-making, frequent non-transport, and standardized data exchange) to clarify what arrives at the ED and what remains outside its operational control. Next, it profiles ED demand and arrival routes in the Netherlands and uses those inputs to derive an ED workflow tailored to the Dutch context by combining Asplin's input-throughput-output model with Claassen's acute-care chain phases—this becomes the backbone for mapping work adaptations and linking indicators to the critical function. The chapter then reviews how healthcare quality is measured with performance indicators, including IOM dimensions and the limits of single metrics, to motivate using multiple indicators and careful interpretation. Then it synthesizes COVID-era ED work adaptations in the Netherlands for potential use within the framework. That route was infeasible for analysis, but it gave insights on what kinds of adaptations to expect. Finally, the study pivots to other work adaptations with traceable PI changes (RAZ, POCT, FRP). These cases provided understanding of how a work adaptation perturbs the system with what implications, thereby helping to operationalize the framework and supplying the empirical input used in the thesis.

2.1. Emergency Care Systems

Understanding the function of an Emergency Department (ED) can be informed by analyzing the structure of the national healthcare and emergency medical services (EMS) system in which it operates. The organization of EMS—particularly how patients are triaged and transported—plays a significant role in shaping the ED's operational scope, workload, and strategic purpose. Two dominant structural models provide a comparative foundation for understanding these differences: the Franco-German model (FGM) and the Anglo-American model (AAM).

The Franco-German model emphasizes physician-led prehospital care. Medical staff are dispatched to the scene, equipped not only to stabilize but also to assess and often treat patients without requiring transport. This approach follows the principle of "bringing the hospital to the patient," and in many cases, the ED may be bypassed entirely if the attending field physician deems it appropriate (Garrone, 2011; Al-Shaqsi, 2010). For example, it is common in FGM systems for patients to be admitted directly to hospital wards, with the ED playing little to no role in their care pathway (Al-Shaqsi, 2010). Field physicians are authorized to exercise complex clinical judgment and provide treatment at the scene or in the patient's home, effectively filtering cases before they reach hospital infrastructure.

In contrast, the Anglo-American model structures EMS as a rapid transport system staffed predominantly by non-physician personnel, such as emergency medical technicians (EMTs) and paramedics. The overarching principle is "scoop and run"—transporting patients to the ED as quickly as possible, where diagnosis and treatment occur (Garrone, 2011). Because the Anglo-American EMS model prioritizes fast transport and offers only basic life-saving maneuvers, resources are instead concentrated in emergency departments to support patient admission and treatment (Garrone, 2011). Because patients in the Anglo-American model are typically transported directly to the ED with minimal prehospital filtering, functions such as triage, clinical assessment, and treatment initiation are generally expected to begin within the emergency department itself (Makrides et al., 2022; Garrone, 2011).

This divergence between models illustrates how upstream decision-making filters shape the ED's role. In FGM systems, much of the diagnostic filtering occurs prehospital, reducing the volume and urgency of cases reaching the ED. In contrast, AAM systems channel virtually all emergency cases into the ED, which is then expected to manage a broader array of conditions under high time pressure. Notably, hybrid reforms have emerged in response to growing concerns about ED overcrowding. For instance, the UK National Health Service, while largely following the Anglo-American model, introduced Emergency Care Practitioners (ECPs) to triage and treat patients at the community level or directly at the site of the incident, reducing unnecessary ambulance transports (Al-Shaqsi, 2010). Such adaptations demonstrate that even within a given structural model, health systems evolve mechanisms to recalibrate the burden placed on emergency departments.

Comparing these system-level structures offers a valuable lens through which to understand how the function of EDs is shaped across different national contexts.

2.1.1. Emergency Care System in Netherlands

The structure of the Dutch Emergency Medical Services (EMS) system closely aligns with the Franco-German model, as supported by a detailed perspective article by Backus et al. (2020) from the University of Groningen. Their study highlights how the Netherlands has developed a robust system of prehospital filtering, enabling emergency departments (EDs) to focus on more urgent, acute interventions. At the core of this system is the strong gatekeeping role of general practitioners (GPs), who serve as the first point of contact for non-life-threatening complaints. This GP-centered model reduces unnecessary ED visits and reflects a broader national effort to manage emergency care demand more efficiently. Additionally, over half of the country's primary care cooperatives are physically integrated with hospital EDs, further streamlining triage and routing decisions. Dutch ambulance nurses—who undergo an 18-month advanced training program—are licensed to provide autonomous treatment at the advanced life support level. Their competencies include clinical assessment, triage, and treatment initiation, which are guided by a standardized national protocol (Ambulancezorg Nederland, 2016). These capabilities mirror those typically expected from physicians in FGM systems and sharply contrast with the Anglo-American model, where EMS personnel generally provide only basic stabilization before transport. Notably, 23–25% of acute patients treated by Dutch EMS are not transported to the ED at all—indicating that meaningful care decisions are made in the field. Taken together, these elements demonstrate that the Dutch EMS system already performs an initial layer of filtering before hospital entry, functioning as a preliminary buffer that helps preserve ED capacity for the most critical cases. This model not only reduces crowding but also contributes to maintaining quality of care within emergency departments themselves. The broader European study by Rief et al. (2023) supports this characterization, classifying the Netherlands within the Franco-German model group. (This part is maybe for later usage, its just a connection I made between COVID regulations and ED system. What likely changed during COVID-19 was not the filtering logic for EDs, but the strictness and clarity of enforcement: While the COVID-19 pandemic prompted stricter screening of patients at the ED level, this practice was not a deviation from normal procedure in the Netherlands. Rather, it reflected a continuation—and intensification—of an already well-established gatekeeping system that discourages ED use for low-acuity complaints.)

2.2. Emergency Departments in Netherlands

To understand the critical function of Emergency Departments (EDs), it is important to first examine who enters the ED, how they arrive, and what this reveals about the broader structure and responsibilities of the emergency care system. In the Netherlands, patients typically access the ED via self-referral, general practitioner (GP) referral, or ambulance transport. These arrival types are closely linked to how the healthcare system filters and prioritizes acute cases, and thus offer valuable insight into where the ED fits within the wider acute care chain. This section reviews the main arrival pathways as a starting point for understanding how emergency care responsibilities are distributed and how those responsibilities begin to shift once the patient enters the ED.

2.2.1. ED Patients

To understand the structure and demand profile of Emergency Departments (EDs) in the Netherlands, it is useful to begin with a broad overview of who enters these departments and how. Recent figures from VeiligheidNL indicate that Dutch EDs received approximately 1.81 million visits in 2023, calculated using data from the national Injury Information System (LIS) (VZinfo.nl – Volksgezondheid en Zorg, 2025). This dataset captures both physical injury and acute illness presentations and offers a robust foundation for analyzing ED usage.

An analysis of time series data reveals a clear downward trend in ED visits since 2019. In that year, total visits stood at approximately 2,030,000. With the onset of the COVID-19 pandemic, this number dropped sharply by 15.3% to 1,720,000 in 2020, and further declined to 1,627,000 in 2021, representing a total drop of nearly 20% compared to pre-pandemic levels. This decline is attributed to pandemic-related factors such as mandatory remote work, closure of schools and sports facilities, and changes in daily activity patterns, including a greater hesitance to visit hospitals due to fear of infection or concern about burdening the system (VZinfo.nl – Volksgezondheid en Zorg, 2025). ED visits rebounded in 2022 by 13.1% to 1,840,000, then declined slightly again by 1.6% in 2023, settling at 1,810,000. These trends are illustrated in Figure 2.1 which visualizes shifts in overall ED usage and its breakdown by cause.

These patterns are not just statistical fluctuations; they reflect broader shifts in healthcare-seeking behavior, system demand, and operational pressure. In particular, the COVID-19 pandemic reshaped how and when people sought acute care, creating disruptions that may have affected how EDs managed resources, prioritized care, and maintained service continuity.

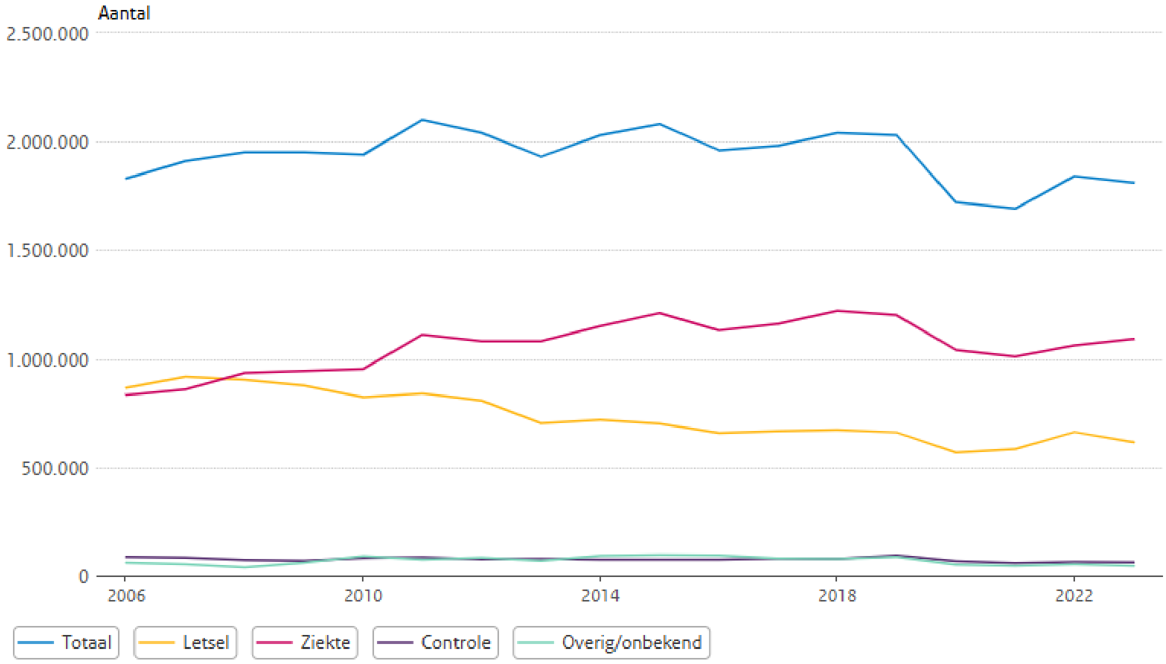

Figure 2.1: Trend in ED Visits, 2006–2023

While the primary focus of this thesis is on in-department functioning, these external factors provide important context for understanding variation in ED workload and performance.

Finally, ED visit data also reveal how patients access care through different entry routes. As illustrated in Figure 2.2, in 2022 approximately 29.9% of patients were self-referred, 39.4% were referred by a general practitioner (GP) without ambulance involvement, 14.8% arrived by ambulance after GP consultation, and 15.9% arrived via ambulance without prior GP contact. These figures highlight the structure of the Dutch emergency care chain, where gatekeeping and prehospital triage play an important role in determining who reaches the ED and how.

2.2.2. ED Entry Routes - Referral Types

Given this structured entry system, we now examine arrival types to explore whether they affect the clinical workflow and responsibilities of ED staff. By mapping out the characteristics of each arrival mode, we aim to understand whether these differences influence the department's core function. This analysis will help clarify the extent to which patient entry routes shape in-ED processes, and whether those processes vary meaningfully across arrival categories.

Self Referrals

Despite the accessibility of General Practitioners (GPs) in the Dutch healthcare system, a significant portion of ED visits still originate from self-referred patients. In 2022, approximately 29.9% of ED visits in the Netherlands were self-referred (VZinfo.nl – Volksgezondheid en Zorg, 2025). While this figure underscores that self-referral remains a common entry pathway, it also raises concerns about appropriateness and the efficient use of emergency care resources.

Empirical studies have shown that many self-referred patients seek care for conditions that could be managed by GPs at lower cost and with similar outcomes (Kraaijvanger et al., 2016). In a Dutch community teaching hospital, the percentage of self-referred ED visits considered appropriate ranged between 48.1% and 58.8%, depending on the assessment method used. These findings highlight a recurring mismatch between the clinical necessity of ED visits and

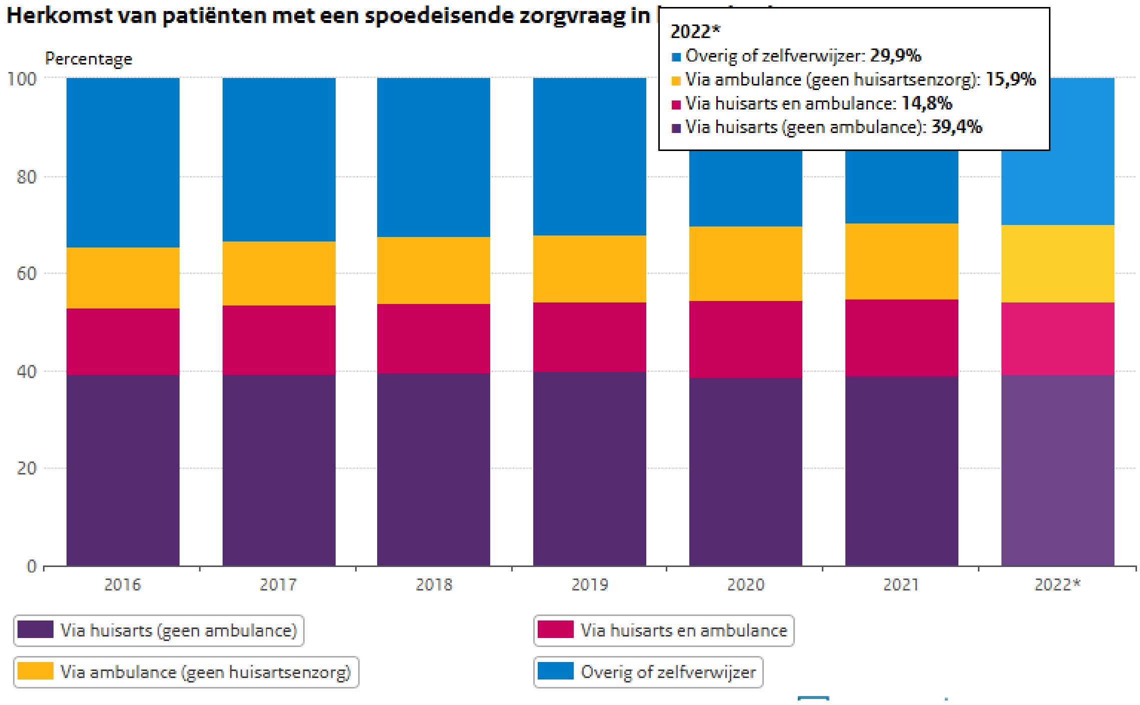

Figure 2.2: Distribution of Entry Routes to the Emergency Department

actual patient behavior. Notably, 76.7% of surveyed patients indicated they would return to the ED under similar circumstances, regardless of whether their visit had been medically appropriate.

Several factors appear to drive this pattern (Rooijen et al., 2013):

- Many patients bypass the GP gatekeeping system believing they will receive faster or better care at the ED.
- Convenience-related motives such as proximity and ease of access also play a role.
- Nearly half of patients were unsure where to seek care for specific complaints, all contributing to avoidable ED use.

These motivations suggest that self-referral is not merely a matter of necessity, but also one of perception, habit, and structural convenience—indicating that it is likely to remain a persistent feature of the Dutch emergency care landscape.

In addition, self-referrals in the Netherlands are typically young to middle-aged men presenting with low-urgency, trauma-related complaints (Rutten et al., 2017). This demographic is also the group most likely to engage in physically demanding or risk-prone activities which makes them more vulnerable to seasonal accident surges. Two Dutch studies illustrate how such events can suddenly strain ED operations. Brand et al. (2014) showed that during a 10-day period of snow and ice, the number of fracture cases more than doubled compared to a control period, with a sharp rise among patients aged 31–60. Likewise, Lieshout et al. (2010) reported a 5.5-fold increase in distal radius fractures during a rare cold spell that enabled nationwide natural ice skating, which overloaded EDs and related hospital services.

Although these studies did not explicitly differentiate by arrival type, the fact that self-referrals are predominantly low-acuity trauma patients makes it likely that many of the additional cases represented self-referrals rather than ambulance arrivals. In practice, such patients are often brought in directly by friends or relatives immediately after the accident. This highlights the link between environmental or recreational events, surges in trauma cases, and the dynamics of self-referrals, which can in turn disrupt the balance of ED workflows.

Ambulance Arrival

Ambulance transport constitutes a significant entry point into Emergency Departments (EDs) in the Netherlands, accounting for approximately 30.7% of all ED visits (VZinfo.nl – Volksgezondheid en Zorg, 2025). In contrast to self-referrals, patients arriving by ambulance typically receive structured, protocol-based prehospital care before reaching the ED. Dutch ambulance crews consist of a specialized nurse and a trained driver, with nurses undergoing an intensive 18-month program that authorizes them to administer care at the advanced life support level (Backus et al., 2020).

This prehospital model reflects the Dutch emergency system's alignment with the Franco-German tradition, which emphasizes early clinical decision-making and treatment in the field. Under the Landelijk Protocol Ambulancezorg (LPA), ambulance personnel assess patients, collect vital signs and histories, initiate treatments when necessary, and—crucially—decide whether hospital transport is even required. Notably, about 23% of ambulance deployments are resolved without ED transfer (Backus et al., 2020).

One of the most critical elements for ED operations is the structured handover process. Upon arrival, ambulance teams formally transfer both the patient and prehospital clinical data to ED staff. This typically includes observations, urgency assessments, treatments administered, and other relevant information (Guasconi et al., 2022). A more detailed specification of this process is provided in the Landelijk Protocol Ambulancezorg (LPA9), which requires that handovers and pre-arrival notifications follow the (I)SBAR(R) structure. This consists of Situation (current condition and vital signs), Background (medical history, allergies, medications, special circumstances), Assessment (findings, working diagnosis, treatments given), and Recommendation (expected or required care upon arrival), with Identification and Repeat as optional elements (Ambulancezorg Nederland, 2023). At a minimum, the Situation and Recommendation together with the estimated arrival time must be communicated before arrival (Ambulancezorg Nederland, 2023). This standardized structure ensures that ED staff are aware of the incoming patient's condition and expected needs, allowing them to mobilize appropriate resources for time-critical cases.

General Practitioner Referral

In the Netherlands, approximately 54.2% of Emergency Department (ED) patients arrive through referrals from General Practitioners (GPs), making this the most common arrival pathway (VZinfo.nl – Volksgezondheid en Zorg, 2025). GPs and their cooperatives (Huisartsenposten, or HAPs) serve as the primary point of contact for non-emergency care, assessing complaints and issuing ED referrals when hospital-level treatment is warranted. During regular hours, patients consult their own GP; outside these hours, care is managed by regional HAPs (Valk et al., 2014). In both cases, the GP determines the urgency of the situation and initiates access to secondary care.

A defining feature of the Dutch system is the efficient and structured exchange of patient data between primary care and hospitals. Referral information—such as patient identification (BSN and NAW), urgency level, reason for referral, previous emergency visit reports, SOEP notes (subjective complaints, objective findings, evaluation, and treatment plan), test results from the past four months, current medication and medication history, allergy information, treat-

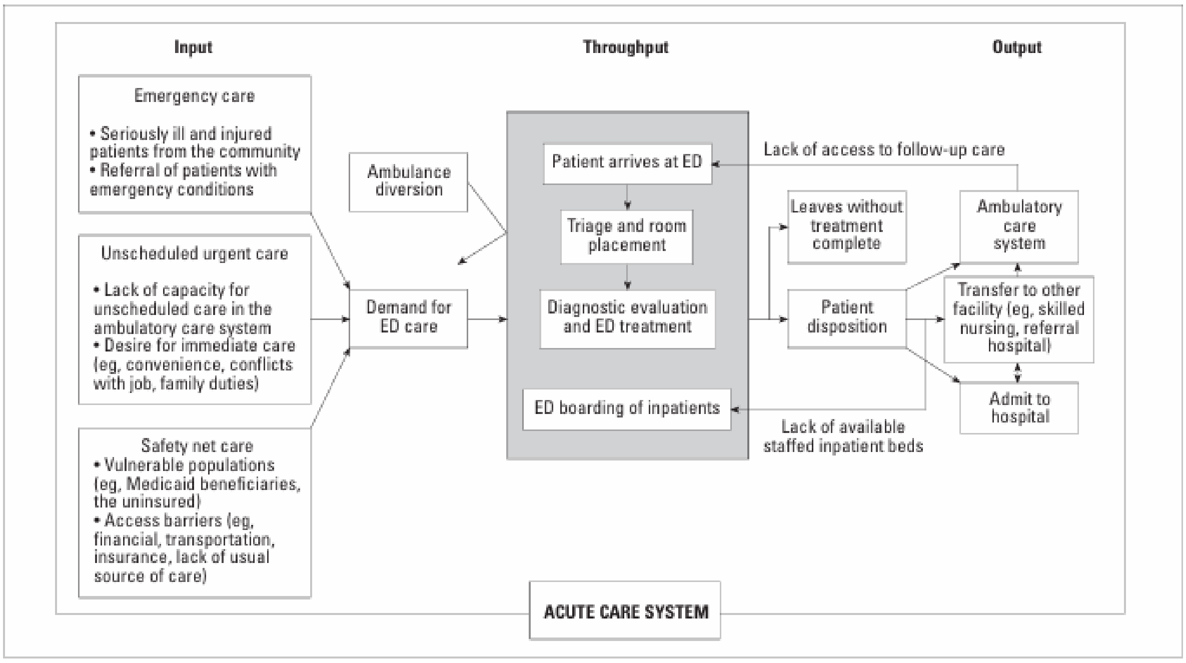

Figure 2.3: The input-throughput-output conceptual model of ED crowding (Asplin et al., 2003)

ment limitations (e.g. advance directives), and recent treatments—is transmitted directly to the ED's information system and can be imported into the patient file. This structured data exchange has been reported by both EDs and HAPs to improve efficiency, reduce duplication, and minimize errors (AORTA-LSP (VZVZ), n.d.; Nederlands Huisartsen Genootschap (NHG), 2022).

While this section has identified the three main ED arrival types, the actual pathways through which patients access emergency care are far more complex. The Directive on Data Exchange Acute Care (Nederlands Huisartsen Genootschap (NHG), 2022) describes 14 different acute care scenarios, each shaped by variations in urgency, provider contact, and referral mechanisms. However, as this study will focus specifically on the care and treatment responsibilities that occur within the ED itself to define the ED's core function, a more detailed breakdown of arrival types is not required here. Importantly, the structured transfer of patient information—whether through GP referrals or ambulance handovers—means that key administrative and documentation tasks are completed before the patient reaches the ED. These upstream processes help clarify which responsibilities lie outside the department's core function. The rationale for emphasizing in-ED responsibilities will be further explained in the next section.

2.2.3. Emergency Department Workflow

Defining the critical function of Emergency Departments (EDs) requires a detailed understanding of the processes that occur within the ED itself.

A widely cited model proposed by Asplin et al. (2003) divides ED operations into three core phases: input, throughput, and output which can be seen in Figure 2.3. This structure offers a pragmatic lens through which ED workflow can be analyzed and improved.

According to Asplin et al. (2003), the input phase can be defined as "any condition, event, or system characteristic that contributes to the demand for ED services". In the context of this thesis, these activities will encompass activities prior to ED arrival, including the referral process and patient transport (e.g., self-referral, GP referral, ambulance dispatch). The throughput phase includes triage, diagnostic evaluations, and care delivery that occur inside the ED. Key factors influencing throughput efficiency include the cohesiveness of care teams, ED layout, staff-to-patient ratios, efficiency of diagnostic testing (e.g., radiology, lab), accessibility of medical records, quality of documentation systems, and availability of specialty consultations (Asplin et al., 2003).

The output phase is defined as all activities related to discharging a patient, admitting them to the hospital, or transferring them to another facility (Asplin et al., 2003). Delays in this phase—particularly due to inpatient bed shortages—can lead to the ”boarding” of admitted patients in the ED. This practice consumes ED staff resources, reduces capacity for new patients, and is a major cause of ambulance diversion (Asplin et al., 2003).

A recent study by Claassen et al. (2025) provides a representative overview of adult patients and their journey through the Dutch Acute Care Chain (ACC) which is divided into four distinct phases: the pre-referral phase, the referral phase, the ED phase, and the post-ED phase. This structure has been used as a conceptual foundation for constructing a localized ED workflow framework in this study.

The pre-referral and referral phases describe the patient’s journey before arriving at the ED. The pre-referral phase includes the initial onset and duration of symptoms, any prior contact with a healthcare provider, and prescribed medication. The referral phase captures the circumstances of the referral, such as the time and source of the referral (e.g., GP, EMS, or self-referral), as well as the urgency level assigned by either the GPC or EMS.

The post-ED phase focuses on patient outcomes following discharge from the ED. These include the type of disposition (e.g., home discharge, hospital admission, or transfer), the occurrence of any adverse events (including 30-day all-cause mortality), and follow-up care arrangements.

Throughput Phase

The throughput phase of Emergency Department (ED) operations encompasses the core clinical processes that occur once a patient physically enters the department. These activities include triage, diagnostics, treatment, and coordination of care up to the point of discharge or hospital admission.

However, not every ED is equipped to manage all emergency conditions at all times. According to the Breedveld Working Group, there is considerable variation in material resources and professional capabilities across EDs in the Netherlands (Werkgroep Kwaliteitsindeling SEH, 2009). To account for this, EDs are classified into three profile levels based on their available specializations and infrastructure (Gezondheidsraad, 2012). As shown in Table 2.1, high-resource EDs must be capable of providing acute care across multiple specialties—such as cardiology, neurology, surgery, orthopedics, pediatrics, obstetrics, and psychiatry—while lower-profile EDs are expected to refer patients outside their scope to more specialized centers.

In practice, this means that throughput activities vary not only by patient condition but also by institutional capability. Nevertheless, triage, stabilization, resuscitation, and initiating treatment remains a universal responsibility within the ED throughput function.

Furthermore, ED nurses play a critical role in the throughput phase. Their responsibilities include performing triage, supporting or initiating stabilization interventions, conducting diag-

Table 2.1: Comparison of ED Types by Task Profile - Translated from Gezondheidsraad (2012)

ED Type	Basic	Profile	Complete
Patient Care	• Triage • Stabilization • Resuscitation, incl. airway management • Initiating treatment or referral	• Triage • Stabilization • Resuscitation, incl. airway management • Referral/initiating treatment • Full treatment for a selected group of patients	• Triage • Stabilization • Resuscitation, incl. airway management • Referral/initiating treatment • Full treatment for all patients
Coordination Tasks	No	No	Yes
(Evaluation) Research	Participation	• Participation • Carrying out tasks for own profile	• Initiating (academic center) • Executing • Coordinating (academic center)
Education	• Identifying learning needs	• Identifying learning needs • Providing education for own profile	• Generating learning requests • Providing education • Organizing education

nostics and treatments based on established protocols, and coordinating care with other disciplines under protocol or physician direction (Werkgroep Kwaliteitsindeling SEH, 2009).

The question of who should conduct triage in the ED remains under discussion (Gezondheidsraad, 2012). At the same time, multiple triage models are currently in use across the Dutch acute care system, which complicates standardization (Nederlandse Zorgautoriteit, 2023). Nevertheless, directing patients to the correct care pathway early in the process has been shown to improve outcomes (Gezondheidsraad, 2012).

2.3. Measuring Healthcare Quality with Performance Indicators

A foundational definition of quality in healthcare has been proposed by the U.S. Institute of Medicine (IOM), describing it as “the degree to which services for individuals and populations increase the likelihood of desired outcomes and are consistent with current professional knowledge.” This definition acknowledges that performance in healthcare can be assessed along a scale and that quality evaluation should consider both individual and population-level outcomes. It also emphasizes that improving healthcare must be grounded in scientific evidence. When the scientific evidence is not available, it should be grounded on expert consensus (Greaney, 2009). This emphasis on expert-backed judgment is further reinforced by Soldatenkova et al. (2023), who argue that selected performance indicators must be supported or recommended by the research community. Using widely accepted indicators not only increases their likelihood of adoption by ED managers and clinical professionals, but also facilitates benchmarking and comparability across different emergency departments. Moreover, they help ensure alignment with established performance measurement frameworks, making the evaluation more consistent and actionable at both local and system-wide levels (Soldatenkova et al., 2023)

The primary rationale for monitoring healthcare quality is to detect opportunities for improvement when current performance does not meet expected standards. The act of measuring and reporting on healthcare quality is widely recognized as a basis for driving improvements in service delivery (Greaney, 2009). According to Greaney, this monitoring process involves assessing current performance—including the perspectives of patients—against predefined expectations. It requires the definition of appropriate indicators, the development of supporting information systems, and the systematic evaluation and analysis of the data gathered (Greaney, 2009). The ultimate goal is to identify gaps between actual and desired performance and to use these insights to both address weaknesses and enhance existing strengths (Greaney, 2009).

In the context of emergency care, performance indicators play a critical role in helping ED managers pinpoint areas of the operation that require attention and define strategies to respond to rapid environmental changes (Mehrolhassani et al., 2025). Measuring ED performance can also facilitate the identification and elimination of non-value-added procedures. For this reason, establishing transparent, reliable, achievable, and clinically appropriate key performance indicators (KPIs) is a fundamental step in assessing ED performance (Mehrolhassani et al., 2025). With this in mind, this section first turns to the foundations of healthcare performance measurement, before moving toward the selection of appropriate indicators for ED evaluation.

2.3.1. Foundations of Healthcare Quality Measurement

The IOM's conceptualization of healthcare quality has also been operationalized into six measurable dimensions, which are frequently used in modern indicator frameworks. The definitions for these six dimensions as provided by Agency for Healthcare Research and Quality (2015) are:

- **Safety:** avoiding harm to patients from the care intended to help them.
- **Timeliness:** reducing unnecessary delays for both patients and providers.
- **Effectiveness:** providing scientifically supported services to those who will benefit, while avoiding underuse or misuse.
- **Efficiency:** avoiding waste of resources.
- **Equality:** ensuring care quality does not vary based on personal characteristics.
- **Patient-Centered:** aligning care with patient preferences and values.

Each of these domains can be operationalized through performance indicators, which act as measurable proxies for the system's ability to deliver high-quality care across various settings and contexts. This kind of domain-based mapping has been applied in practice by Aaronson et al. (2015), who categorized a large set of emergency care indicators according to the IOM's six quality dimensions. Their approach will be explained in more detail in Section 2.4.3.

Another widely applied approach to assessing healthcare quality is the tripartite model developed by Donabedian (2005). This model evaluates quality through three interrelated categories:

- **Structure:** referring to the resources, equipment, and organizational infrastructure in place;
- **Process:** representing the actions taken during patient care delivery;
- **Outcomes:** reflecting the end results of care in terms of health improvement or deterioration.

This categorization is widely used in the literature as a way to group and interpret performance indicators (Greaney, 2009; Bos et al., 2015; Madsen et al., 2015; Moes et al., 2019; Aaronson et al., 2015). The structure-process-outcome classification helps define the scope of quality assessment by distinguishing among different approaches to assessment.

Performance indicators in healthcare can be further classified along several dimensions. As shown by Greaney (2009) and illustrated in Figure 2.4, indicators are first grouped into generic and disease-specific categories. Generic indicators apply broadly across all patient groups and conditions, measuring general aspects of care delivery. In contrast, disease-specific indicators are tied to particular illnesses and measure care processes relevant to those specific conditions. Beyond this core distinction, indicators can also be categorized by the type of care, such as preventive, acute, or chronic care, depending on the healthcare phase they are intended to evaluate. Additionally, they may be grouped by the function of care, which includes activities like screening, diagnosis, treatment, and follow-up.

2.3.2. Political Use and Misapplication of Indicators

Although performance indicators are intended to support quality improvement, their application is not always aligned with clinical priorities. A case study by Moes et al. (2019) describes

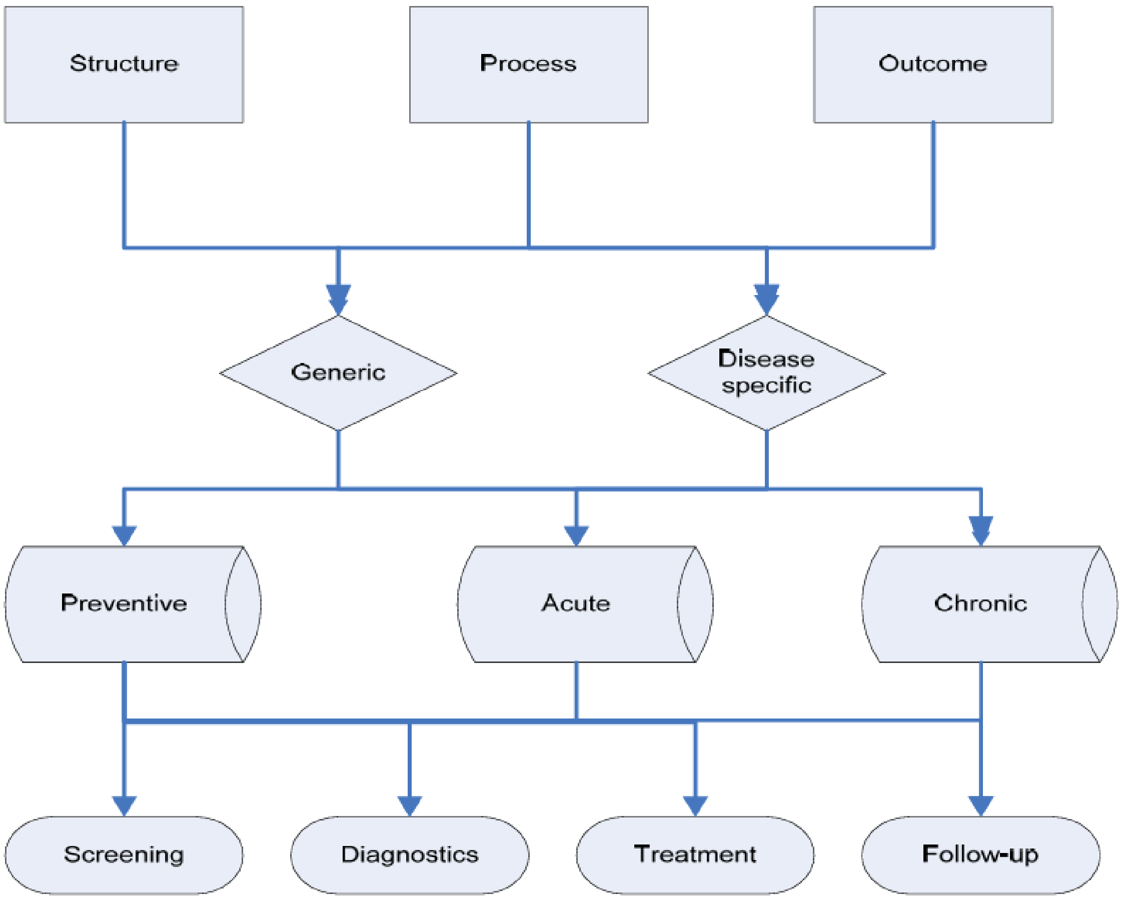

Figure 2.4: Types of Indicators (Greaney, 2009)

how private insurers in the Netherlands attempted to reform emergency care using rigid, top-down quality metrics. These indicators were drawn from clinical guidelines and research, but were applied with the goal of increasing efficiency and administrative control, rather than improving care quality. They were used primarily to justify the centralization of services and the enforcement of volume norms (Moes et al., 2019). This argument is further supported by Mehrolhassani et al. (2025), who observe that ED performance evaluation has historically prioritized financial metrics. JohnGreaney (2009) similarly warns against the misuse of indicators for economic or regulatory goals, emphasizing that performance measurement should remain focused on improving patient outcomes.

A big issue was the use of volume norms that equated higher patient volumes with better quality, despite limited supporting evidence. This approach risked penalizing smaller EDs that delivered competent care but failed to meet arbitrary thresholds (Moes et al., 2019).

Another concern relates to how indicators are interpreted and applied. In some cases, simplified scoring systems overlook local context and treatment capacity, leaving little room for explanation or adaptation. This can undermine the reflective purpose of performance indicators. Moes et al. (2019) argue that such simplification turns indicators into blunt instruments for control, rather than tools for improvement.

2.3.3. Need for Multiple Indicators

Relying on one or two performance indicators is insufficient for evaluating the quality of emergency care. As Sørup et al. (2013) argue, although individual metrics may capture isolated aspects of performance, their narrow focus can lead to unintended and counterproductive outcomes. For instance, in efforts to stay within accepted upper thresholds for ED length of stay, patients are sometimes transferred to other wards before receiving adequate treatment. While this may satisfy target metrics, it compromises care quality, increases healthcare costs, and places additional strain on staff. In such cases, reported performance appears to improve, even as actual outcomes worsen (Sørup et al., 2013).

To avoid such distortions, performance measurement systems should rely on a set of indicators that together provide a more accurate picture of how an ED functions. Singular indicators can obscure weaknesses in other dimensions of care, while multiple complementary metrics help capture a broader and more realistic view of performance under everyday conditions (Sørup et al., 2013).

Sørup et al. (2013) also highlight that despite the widespread use of ED indicators, there is still no consensus on which ones are most reliable, evidence-based, or clearly defined. This lack of standardization continues to limit the field.

2.4. COVID-19–Related Work Adaptations in the Netherlands

The COVID-19 pandemic confronted the Dutch healthcare system with significant disruptions that required rapid adaptation. Emergency Departments, situated at the intersection of acute care provision, had to adjust to these pressures under conditions of high uncertainty. To understand these developments, this section examines two complementary dimensions: first, the broader governance measures—including legal frameworks, regulatory adjustments, and policy responses—that shaped the environment of acute care delivery; and second, the actual operational adaptations implemented within EDs themselves.

2.4.1. Governance Measures: Legal, Regulatory, and Policy Responses

This subsection reviews the formal governance context that framed emergency care during the pandemic. It considers the national legal measures, temporary regulations, and policy initiatives introduced to sustain acute care delivery under crisis conditions. Since no regulations were specifically directed at Emergency Departments, these measures primarily affected the acute care chain as a whole rather than EDs in isolation. For this reason, the subsection briefly outlines these governance measures to give the reader an impression of how the Dutch response to COVID-19 was organized. While not designed specifically for EDs, some of these measures still had indirect consequences for ED operations, making it relevant to discuss them here.

Recognition of COVID-19 under the Public Health Act

The coronavirus was officially recognized under the Dutch Public Health Act as a Group A infectious disease by the Minister van Volksgezondheid, Welzijn en Sport (VWS) (2020b) through a regulation dated 28 January 2020, published in the Staatscourant on 31 January 2020. The Public Health Act (Wet Publieke Gezondheid) is the main legislative framework for combating infectious diseases in the Netherlands. Diseases classified in Group A are considered the most severe, and this designation activates the strongest legal instruments available to the state, including mandatory reporting, enforcement, and centralized coordination. By placing COVID-19 in this category, the Dutch government ensured that the full set of legal measures could be applied to manage the emerging outbreak.

The Minister of Medical Care provided an accompanying explanation outlining three key implications of this designation. First, physicians were obliged to report any suspected or confirmed cases immediately to the municipal health service, which in turn reported to the RIVM. Second, the classification allowed for direct implementation of control measures such as isolation of (suspected) infected individuals, coordinated at the municipal level under the authority of the regional safety chair. Third, responsibility for directing the national response was formally assigned to the Minister, enabling a uniform and centrally managed approach to combating the virus. While these provisions were not specific to Emergency Departments, they shaped the overall acute care environment in which EDs operated during the pandemic.

Temporary Workforce Regulations (BIG Register Extensions)

Healthcare professionals in the Netherlands are regulated under the BIG Act, which requires nurses, doctors, and other Article-3 professions to periodically re-register to demonstrate that they remain qualified to practice. Failure to re-register within the required period normally results in loss of registration and the right to independently perform reserved medical acts.

In response to the COVID-19 pandemic, the Minister of Health temporarily relaxed these requirements. First, on 18 March 2020 (BIG-register, 2020b), re-registration deadlines were suspended, meaning that healthcare professionals could continue practicing even if their registration would otherwise have expired. This prevented the administrative burden of re-registration from removing active staff from the workforce at a time when every available professional was needed.

Later, on 27 October 2020 (BIG-register, 2020a), the Ministry extended these measures to allow former nurses and doctors whose BIG registrations had lapsed after 2016 (for nurses) or 2018 (for doctors) to return to work without re-registering. While they could not perform all tasks independently, they could contribute under defined conditions, thereby broadening the pool of available staff. This arrangement remained in place until 1 August 2022, after which former registrants once again lost the right to independently perform reserved medical acts.

These temporary measures were not specific to Emergency Departments, but they influenced hospital operations more broadly by helping maintain staffing capacity during the pandemic.

Policy Framework for Ensuring Acute Care Delivery

During the second wave of the COVID-19 pandemic in autumn 2020, hospital admissions rose sharply and threatened to overwhelm the system. Regular care could no longer be delivered in full, staff were exhausted, and choices about postponing non-urgent care were already being made. In this context, the Minister of Health introduced a temporary policy framework (Minister voor Medische Zorg, 2020) in 23 Oktober 2020 to safeguard the continuity of acute care. The aim was to guarantee that time-critical and life-saving treatments remained available within a six-week horizon and that access was distributed fairly across the country.

The framework was structured around five main pillars. First, the Landelijk Coördinatiecentrum Patiënten Spreiding (LCPS), the national coordination center for patient distribution, provided regional forecasts of expected COVID-19 admissions for the each coming week. Based on these forecasts, hospitals were instructed to scale up COVID-19 capacity and reduce other care accordingly. Second, the Federatie Medisch Specialisten (FMS) set guidance for prioritizing non-COVID care, requiring clinicians to determine which planned treatments had to continue within six weeks, while leaving final judgment to physicians. Third, the framework emphasized regional collaboration through the Regionaal Overleg Acute Zorgketen (ROAZ), which are networks coordinating acute care providers in each region. ROAZ regions were tasked with aligning capacity and patient flows across hospitals, independent clinics, general practitioners, nursing homes, and home care.

The fourth pillar was transparency in accessibility. The Dutch Healthcare Authority (NZa) was tasked with monitoring and sharing information on the availability of critical care, using hospital and LCPS data to help ROAZ regions steer capacity. Finally, the fifth pillar involved financial safeguards. Existing agreements ensured that providers would be compensated for lost income and additional costs, allowing them to concentrate on delivering care without the barrier of financial uncertainty. In addition to these pillars, the Ministry of Health also pursued concrete capacity expansions.

By early 2021, approximately 350 new ICU beds and 700 clinical beds had been created, alongside efforts to recruit and train additional staff. The Ministry further explored concentrating COVID-19 patients in fewer hospitals and arranged mutual agreements with Germany for cross-border patient distribution.

In summary, this framework did not provide regulations directed specifically at Emergency Departments, but some of its measures may have indirectly influenced ED workloads. Increasing hospital and ICU capacity, for example, was not aimed at EDs themselves, but creating downstream space could help reduce crowding and pressure in EDs. However, as these adjustments primarily concerned hospital-wide capacity management, their detailed analysis falls outside the scope of this thesis.

Corona Opt-In: Data Access for Emergency and Primary Care

Under normal circumstances, Dutch privacy regulations require patients to give explicit consent before their professional summary (PS) can be accessed by healthcare providers outside the general practitioner's office. This system, often referred to as "opt-in," meant that out-of-hours GP services (HAPs) and Emergency Departments (SEHs) could not automatically view essential patient information unless the patient had already registered consent in advance.

At the onset of the COVID-19 pandemic, this requirement was considered a barrier to timely and effective acute care. To address this, the government introduced a temporary measure in April 2020 (Minister van Volksgezondheid, Welzijn en Sport (VWS), 2020a), commonly known as the Corona Opt-in. This allowed HAPs and EDs to access patient summaries even for individuals who had not previously recorded a choice. Access was still conditional: if the patient was capable, verbal consent had to be requested at the point of care.

For emergency care providers, this change significantly improved the speed and safety of triage and treatment. Immediate access to professional summaries meant that clinicians no longer had to contact the patient's GP before acting, which would have been impractical under pandemic conditions. Instead, they could rely directly on up-to-date medical information, reducing delays and supporting more informed decisions in acute settings.

The measure was subsequently defended in Parliament (Minister voor Medische Zorg, 2021) in 30 April 2021, reflecting its continued relevance throughout the pandemic. Ultimately, the Corona Opt-in expired on 4 April 2023, at which point professional summaries again became unavailable for patients who had not given explicit consent.

Although always intended as a temporary exception, the Corona Opt-in illustrates one of the few national regulatory adaptations during COVID-19 that directly shaped Emergency Department operations. By enabling real-time access to essential patient data, it supported faster triage and more efficient emergency management under conditions of critical system pressure.

Absence of ED-Specific Regulations

The legal and policy review did not identify any regulations during the COVID-19 pandemic that were specifically directed at Emergency Departments. Most measures introduced by the Dutch government applied to the acute care chain as a whole, or targeted hospitals with a particular focus on intensive care capacity. For example, the temporary policy framework for ensuring acute care was later supplemented with provisions for ICU-dependent critical planned care in 19 November 2021 (Minister van Volksgezondheid, Welzijn en Sport (VWS), 2021).

These findings align with the analysis of O'Connor et al. (2021), who examined how Dutch EDs adapted their work practices during the pandemic. Their study noted that, while ICU capacity for COVID-19 patients was centrally monitored and coordinated at the national level by the Landelijk Coördinatiecentrum Patiënten Spreiding (LCPS), there was no comparable national guidance on surge capacity management for EDs. As a result, hospitals had to reorganize their ED operations independently, leading to institution-specific solutions and local variation (O'Connor et al., 2021)

This absence of national ED-specific regulation provides the rationale for turning to empirical studies of ED work adaptations. These studies offer insights into the kinds of organizational adjustments that EDs adopted in response to pandemic pressures, and are therefore discussed in the following subsection.

2.4.2. Emergency Department Operational Adaptations in the Netherlands

O'Connor et al. (2021) conducted a nation-wide, cross-sectional questionnaire study to assess how Dutch EDs prepared for and responded to the first wave of COVID-19 in the absence of ED-specific national regulatory guidance. The study was carried out between the first and second COVID-19 surges (July–September 2020), reflecting adaptations made during the peak period of February–April 2020. Responses were collected from 66 EDs, covering 80% of all Dutch EDs, across 58 hospital organizations. The study therefore offers a comprehensive overview of ED operations during the crisis. Importantly, the authors note that there was no consensus in the Netherlands on ED surge capacity, infrastructure, or work processes, and that regional differences in infection rates—ranging from 501 per 100,000 in the south to 60 per 100,000 in the north—meant that preparations and adaptations varied by local context.

Within this framework, the main work adaptations reported are:

1. **Expansion of treatment capacity**: Around 70% of EDs expanded their treatment capacity to accommodate the surge of patients. The median number of treatment spaces per ED increased by 8 (IQR 4–10), on top of a pre-pandemic median of 17 spaces. This amounted to a 49% median increase (IQR 33–73%).
2. **Segregation of suspect-COVID-19 patients from non-COVID-19 patients**: Physically separating suspected COVID-19 patients from non-COVID-19 patients, either within the ED or in alternative spaces.
3. **Redirection of less urgent patients**: Reported in 63% of EDs. Patients with non-urgent conditions, such as minor traumatic injuries, were redirected to outpatient departments. In addition, 18% of EDs streamlined hospital admissions and ICU transfers, which reduced ED length of stay.
4. **Changes in triage protocol**:To facilitate segregation, many EDs introduced formal pre-entry screening for COVID-19 before patients entered the main department. The most common method was a symptom-based screening list, used alone in 65% of EDs. In 20% of EDs, symptom-based screening was combined with radiological imaging (such as a chest X-ray or CT scan). A very small minority (2%) relied solely on CT imaging. Emergency Physicians (EPs) played a central role in these adapted triage processes: in the majority of hospitals employing them, EPs were directly involved in the triage and segregation of suspected COVID-19 patients, and often coordinated this activity within the ED.
    a. **Use of radiological imaging for COVID-19 screening**: Chest X-rays or CT scans were incorporated into screening in 20% of EDs in combination with symptom checklists. In 2% of EDs, CT imaging was used as the sole method of screening.
    b. **Screening before ED entry**: Patients were screened before entering the ED using symptom checklists, chest X-rays, CT scans, or combinations of these methods.
5. **Expansion of workforce**: Reported in 82% of EDs. This primarily involved recruiting additional ED nurses (53%) and redeploying nurses from other hospital departments (61%). Physicians from multiple specialties were also directly involved in COVID-19 care at the ED, including emergency medicine (86%), internal medicine (85%), pulmonology (82%), anesthesiology (40%), geriatrics (36%), and surgery (35%), among others.
    a. **Improved interdisciplinary collaboration**: Pandemic conditions encouraged closer collaboration between specialties, which respondents indicated was valuable to maintain.
6. **Faster admission process:** In 18% of the Emergency Departments (EDs), a faster process for admitting patients to hospital wards and intensive care units was implemented. This was a key logistical change that resulted in a shortened length of stay in the ED.
7. **Expanded roles for Emergency Physicians (EPs)**: Among the responding hospitals, 85% reported employing EPs. In those hospitals, EPs were directly involved in treating COVID-19 patients. Furthermore, 88% reported EPs coordinating the ED, 82% had them

overseeing triage and segregation, and 94% confirmed that EPs held a formal role in the hospital crisis management team, either at a strategic or operational level.

A notable conclusion of the study is that several of these adaptations transitioned from temporary crisis responses into permanent structural changes. O'Connor et al. (2021) emphasize that e-health applications became more firmly embedded, infection prevention practices were strengthened, interdisciplinary collaboration was improved, infrastructural changes were made to segregate infectious patients, and less urgent patients were permanently redirected to GPs or outpatient departments. This shows that the pandemic not only triggered short-term adjustments but also accelerated long-term institutional learning and resilience in Dutch EDs, embedding new routines and practices that continue beyond the acute crisis.

Work Adaptations in Haaglanden Medical Centre ED

To illustrate how these nationwide work adaptations were operationalized in practice, this section draws on the study by Linden et al. (2023), which explained how the Haaglanden Medical Centre (HMC) Emergency Department in The Hague adapted its operations during the first wave of the COVID-19 pandemic; and investigated the association between the COVID-19 surge and ED patient flow during the pandemic's first peak in 2020 compared to similar periods in 2018 and 2019. The work adaptations described by the study are as follows:

Triage protocols were altered: whereas before February 2020 an ED nurse directed patients to either the GP cooperative (GPC) or the ED, from February onwards this initial triage was performed by a GPC assistant using predefined criteria. This adjustment shifted part of the first screening step outside the ED itself.

Protective measures were also expanded. In 2020, ED staff worked in protective gear when approaching the majority of patients, treating all those with suspicious complaints as potentially COVID-19 positive.

The hospital itself was reconfigured to increase capacity. Treatment rooms at the policlinics were used to expand ED space, and admission units for suspected COVID-19 patients were added. Outpatient visits were cancelled, and surgical procedures were postponed or delayed, particularly elective ones. These measures freed resources and staff capacity for acute care.

To further strengthen capacity, staffing was expanded and supplemented. Medical specialists worked side by side with emergency physicians in the ED, nurse assistants were hired to support ED, ICU, and inpatient wards, and health practitioners with critical care skills, including retired nurses, were recruited and trained to assist.

In addition, the **GPC was used as an alternative site for stable patients with suspected COVID-19.** After triage, those patients deemed suitable for GP-level care were redirected to the GPC, while others were treated at the ED.

Finally, **an extra CT scanner was installed** at the ED to improve diagnostic capacity.

Even though these operational adjustments were described in the context of a single level-one hospital in The Hague, many of them reflect broader patterns observed across Dutch emergency departments. As highlighted by O'Connor et al. (2021), Dutch EDs implemented similar measures nationwide during the first COVID-19 wave, though the extent and form of adaptations varied by hospital and region. As stated in the previous section, this is because there was no uniform national framework for ED surge capacity, meaning that hospitals had to tailor their operational responses to their own circumstances(O'Connor et al., 2021).

Preparations Prevented Overcrowding

The search to identify operational adaptations in Dutch EDs showed that preparations prevented overcrowding during the first COVID-19 peak. Linden et al. (2023) observed that although the number of patients with respiratory complaints increased significantly, the overall number of ED visits remained unchanged compared to 2018 and 2019. This is due to the change in the case-mix, where the presentations for other conditions, such as chest pain and syncope (fainting) decreased. Despite a longer length of stay for patients with respiratory complaints who required hospital admission, significantly less crowding occurred compared to earlier years.

Even thoudh total demand did not change, the authors linked this reduction in crowding to a combination of hospital preparations and public health messaging such as the stay-at-home policy, postponement of elective surgery, and patients' fear of infection. (Linden et al., 2023). As they conclude, "advanced warning and its associated preparations in the Netherlands prevented significant delays in ED **throughput** during the first Covid-19 peak" (Linden et al., 2023).

Even though these findings come from a single level-one hospital in The Hague, the source notes that they are generalizable for the Netherlands as they align with broader trends reported in the Netherlands. This direction also appears consistent with national survey data.(O'Connor et al., 2021) reported that the majority of hospital organizations (52%) experienced no crowding during the first surge, and 41% reported only occasional crowding. This suggests that the local findings in The Hague reflect a broader pattern across the country.

Bundled Adaptations and Limits for Performance Assessment

While both Linden et al. (2023) and O'Connor et al. (2021) provide valuable insights into the operational responses of Dutch EDs during the COVID-19 pandemic, neither study offers a feasible basis for detailed performance assessment. Their limitations stem either from the way work adaptations are presented or from gaps in available performance indicator (PI) data, making it difficult to establish clear links between specific measures and outcomes.

The study by Linden et al. (2023) reports crowding indicators and ED length of stay, however, these performance metrics are measured as the outcome of the entire bundle of adaptations implemented at the hospital, rather than as the effect of individual measures. This bundling makes it impossible to disentangle which specific adaptation influenced a given PI. Furthermore, the fact that the study only assessed two performance indicators—crowding and length of stay—renders any analysis extremely limited, as it leaves out the broader range of indicators that would be necessary for a meaningful resilience assessment.

This weakness is already represented in the research gaps defined. As the impact of an adaptation on changes in PIs is not traceable due to this bundling, it resonates with the first gap "Because "managing variability" in an ED is enacted through work adaptations that produce observable changes in performance indicators, we need tools relate to, and ideally help navigate, those indicator trade-offs in practice". Furthermore, this study also lacks sufficiently broad PI set to trace how concrete ED work changes affect performance for real (Yeah ED LOS maybe decreased lets say, but it decreased because we increased the waiting times or something that lead to a reduction of people seeking care (LWBS), ED LOS reduction should not be interpreted as improvement without understanding the full story) which relates to the second gap identified.

This weakness maps directly onto the gaps identified. Because multiple ED work adaptations are reported as a single bundle, the link between any specific adaptation and the observed PI changes cannot be traced—this is Gap 1: if "managing variability" is enacted through concrete adaptations that yield observable PI shifts, our assessment must relate to those adaptation-specific trade-offs in practice. Moreover, reading the outcome through only two PIs is too narrow to reveal what is actually happening operationally which related to Gap 2. For instance, ED LOS might drop not because flow improved, but because waiting times rose and more patients left without being seen (LWBS). In such cases, an LOS reduction should not be interpreted as improvement without the fuller story—hence the need for a broader, evidence-informed PI set tied to different parts of the ED workflow and function.

The nationwide survey conducted byO'Connor et al. (2021) presents a comprehensive list of operational adaptations implemented across Dutch EDs, but it does not report on corresponding changes in performance indicators. In addition, the study does not provide precise timing for when the different adaptations were implemented. This means that even if PI data were available from other sources, it would not be possible to align adaptation dates with PI trends in order to assess causal relationships between specific work adaptations and changes in performance indicators.

Beyond the methodological limits of these two studies, there is also a structural problem in the availability of data. As will be discussed in greater detail in the results section, national databases such as RIVM, INFOVZ, and CBS do not record ED performance indicator data in a way that would enable systematic resilience or performance analysis. This absence of reliable PI data makes it fundamentally impossible to link operational adaptations to outcomes on a national scale.

For these reasons, this thesis shifts its focus in the next section toward studies that have examined more narrowly defined work adaptations and their measurable effects on performance indicators. By drawing from such focused analyses, it becomes possible to enhance the understanding of how particular operational changes influence ED performance and to build a more concrete basis for resilience assessment.

2.5. Focused Work Adaptations with Performance Assessments

Building on the limitations identified in the previous subsection, this part turns to the few studies that have examined specific work adaptations in greater detail and linked them to measurable performance outcomes. Unlike the broader bundled changes, these focused analyses provide concrete examples of work adaptations with travable perforamnce changes.

2.5.1. Enhanced Isolation Protocol for Fever or Respiratory Patients (FRPs)

Kim et al. (2022) studied the impact of an enhanced isolation protocol in the emergency department of a large tertiary teaching hospital in Seoul, South Korea, during the COVID-19 pandemic. Before COVID-19, the hospital's isolation measures were applied only to patients with known airborne or droplet infections such as tuberculosis or measles. With the arrival of COVID-19, the protocol was significantly expanded. All patients presenting with fever (≥37.5°C), respiratory symptoms such as cough or dyspnea, or relevant exposure histories (travel or confirmed case contact) were automatically directed to isolation. This shift aimed to reduce nosocomial transmission but created a substantial change in ED workflows, as a much larger proportion of patients required placement in limited isolation spaces.

The study analysed ED visits during March–July 2019 and the same months in 2020 to assess the effects of this adaptation. Several performance indicators were measured; Admitted to Inpatient Care, ED Bed Occupancy Rate (BOR), ED Lenght of Stay (LOS), Waiting Time (WT), Left Without Being Seen (LWBS), Mortality Rate, Discharged out of ED, Transfer to Another ED. The findings showed that while overall ED length of stay and occupancy decreased, patients meeting the new isolation criteria experienced significant delays and were more likely to leave before evaluation. Waiting times increased most sharply for fever and respiratory patients, and LWBS rates were disproportionately higher for this group compared to non-febrile patients. The authors interpreted this as a bottleneck effect caused by the expanded isolation criteria: although general crowding was alleviated, the limited capacity of isolation zones led to longer waits and greater attrition among the very patients at highest risk.

The study further highlighted three broader implications. First, the bottleneck at isolation zones explained why crowding indicators for the ED as a whole improved, while outcomes for FRPs worsened. Second, the higher LWBS rates among this group raised concerns about patient equity. Finally, the authors identified a public health risk: symptomatic patients leaving the ED untreated posed a danger of continued infection spread within the community. These findings positioned the adaptation as a trade-off—effective for infection prevention inside the ED but introducing new vulnerabilities in terms of patient care and broader public safety.

2.5.2. Rapid Assessment Zone (RAZ)

Faber et al. (2023) describe the implementation of a Rapid Assessment Zone (RAZ) in urban community hospital emergency department in the United States during the COVID-19 pandemic. The intervention was introduced as a response to worsening overcrowding, long wait times, and high rates of patients leaving without being seen (LWBS) that were intensified by pandemic pressures. The RAZ was created by redesigning the existing triage area into eight dedicated rapid assessment bays and reallocating staff to the front end of the ED. This vertical care model allowed lower-acuity patients to be evaluated and treated while seated, without occupying a traditional ED bed, conserving resources for higher-acuity cases. The redesign integrated multiple steps of the patient journey into a single location: initial nursing and provider assessment, order entry, phlebotomy, medication administration, and registration were all carried out in the RAZ before patients returned to the waiting area. The explicit process goal was to have patients assessed and their workup initiated within 20 minutes of arrival. A contingency plan was included whereby RAZ nurses could initiate protocol orders if no provider was available within this time frame.

During implementation, staffing was adapted to match patient volume. Initially, the plan involved one or two nurses assigned to the RAZ. However, higher-than-expected numbers of patients suitable for vertical care led to the redeployment of a third nurse from the main ED during peak times. This was made possible by closing a section of the main ED that had been relieved by the RAZ's capacity to manage lower-acuity patients.

Providers and nurses received education on their specific roles in the new model, the types of patients to be managed in the RAZ, and the practice of concurrent provider–nurse assessment. In the early implementation phase, a clinical expert was present on-site to provide support. Sustainment of the new model was ensured through daily performance reviews, structured debrief meetings, and issue log books completed by ED nurses in the first weeks. These strategies ensured consistency and supported the long-term integration of the RAZ model.

The effects of the intervention were assessed by comparing six months of data before and after implementation. Reported performance indicators included LWBS, Time to Seen by a

Doctor, length of stay for discharged patients, and length of stay for admitted patients. All four showed statistically significant improvement, with marked reductions in LWBS, shorter Time to Seen by a Doctor, and reduced ED LOS for both discharged and admitted patients.

2.5.3. Point-of-Care Testing (POCT)

A recurring challenge during the COVID-19 pandemic was the delay associated with centralized laboratory testing, where turnaround times could extend from several hours to more than a day. Several studies evaluated the impact of introducing point-of-care testing (POCT) as a workflow adaptation to address this issue. While the exact technologies and baseline methods differed—ranging from standard RT-PCR to various laboratory-based NAATs—the underlying shift was the same: replacing longer laboratory-dependent processes with faster bedside diagnostics. Despite these differences, the studies converge on the finding that faster testing facilitated earlier clinical decisions, leading to improvements in emergency department (ED) performance.

Mortazavi et al. (2022) conducted a retrospective observational study in Sweden, assessing the sequential introduction of rapid antigen testing and the VitaPCR molecular POCT to replace central laboratory RT-PCR, which had turnaround times of 12–24 hours. The rationale for this transition was to reduce delays in patient flow and optimize use of hospital resources during the pandemic. The study reported several performance-related outcomes. For patients testing positive at the ED, rapid testing was associated with a shorter ED length of stay, while test-negative patients experienced a reduction in hospital length of stay, from 6.6 to 5.1 days on average. The authors also observed that negative patients were more often admitted directly to the appropriate specialized wards, reducing misplacement in COVID-19 wards. This led to a substantial decline in intrahospital transfers, which fell from 33.2% to 15.9% for test-negative patients. The study concluded that faster diagnostics not only improved patient throughput but also enhanced safety by reducing unnecessary transfers and enabling immediate initiation of appropriate therapy.

Baron et al. (2022) investigated a French ED through a before–after study, where laboratory-based molecular assays with effective turnaround times of around four hours were replaced by the ID NOW POCT, which delivered results in 5–13 minutes at the bedside. The primary motivation for this shift was to mitigate ED overcrowding and improve patient flow. The study found that the median ED length of stay decreased from 276 minutes in the pre-POCT period to 208 minutes with POCT, and the proportion of patients discharged within four hours rose significantly from 38.3% to 61.3%. Time to result was also markedly reduced, dropping from 261 minutes to 112 minutes. The authors emphasized that the ability to discharge patients within four hours is a recognized quality indicator in emergency care, and that this improvement remained significant even after adjustment for other variables. Although hospitalization rates increased during the POCT period, the authors attributed this to differences in patient characteristics between the study periods rather than a direct effect of POCT. Overall, the findings highlighted that faster testing directly translated into more timely ED management and reduced overcrowding.

Taken together, these cases address the research gap by tracing reported changes from specific ED work adaptations to their causal impacts on performance indicators. They give the base material and the practical intuition needed to build and apply PI-RA framework: what to look for, where in the workflow/critical function the change sits, and which PIs move (and why). Combined with the insights from the Introduction and the Literature Review, we now move to the Methodology to explain how we (i) construct the framework, (ii) extract and structure the empirical material needed to use it, and (iii) apply it to answer the three sub-questions.

PART II

Framework Development

3

Research Approach and Methodology

This chapter explains how the Performance Indicator Resilience Assessment (PI-RA) framework was developed and applied. Figure 3.1 provides an overview of the three main phases: (1) structured performance change description, (2) system mapping and attribution, and (3) resilient performance interpretation. Together, these phases show how observed changes in performance indicators can be organised, linked to the ED system, and summarised into a transparent read-out of resilience and its quality-of-care implications.

Section 3.1 explains how to establish the ED workflow and critical function for the Dutch context, which serves as the operational backbone for the framework. Section 3.2 then explains how to carry out the first phase of PI–RA by turning heterogeneous reports of disruptions and work adaptations into structured descriptions of performance change, including IF–THEN summaries. Section 3.3 explains how to implement the second and third phases: it uses the workflow model to map and weight indicator changes, calculates the performance change scores PC_M and PC_B and the trade-off coefficient TO_C, and interprets their combination on the Performance Change Matrix together with the associated quality-of-care trade-offs.

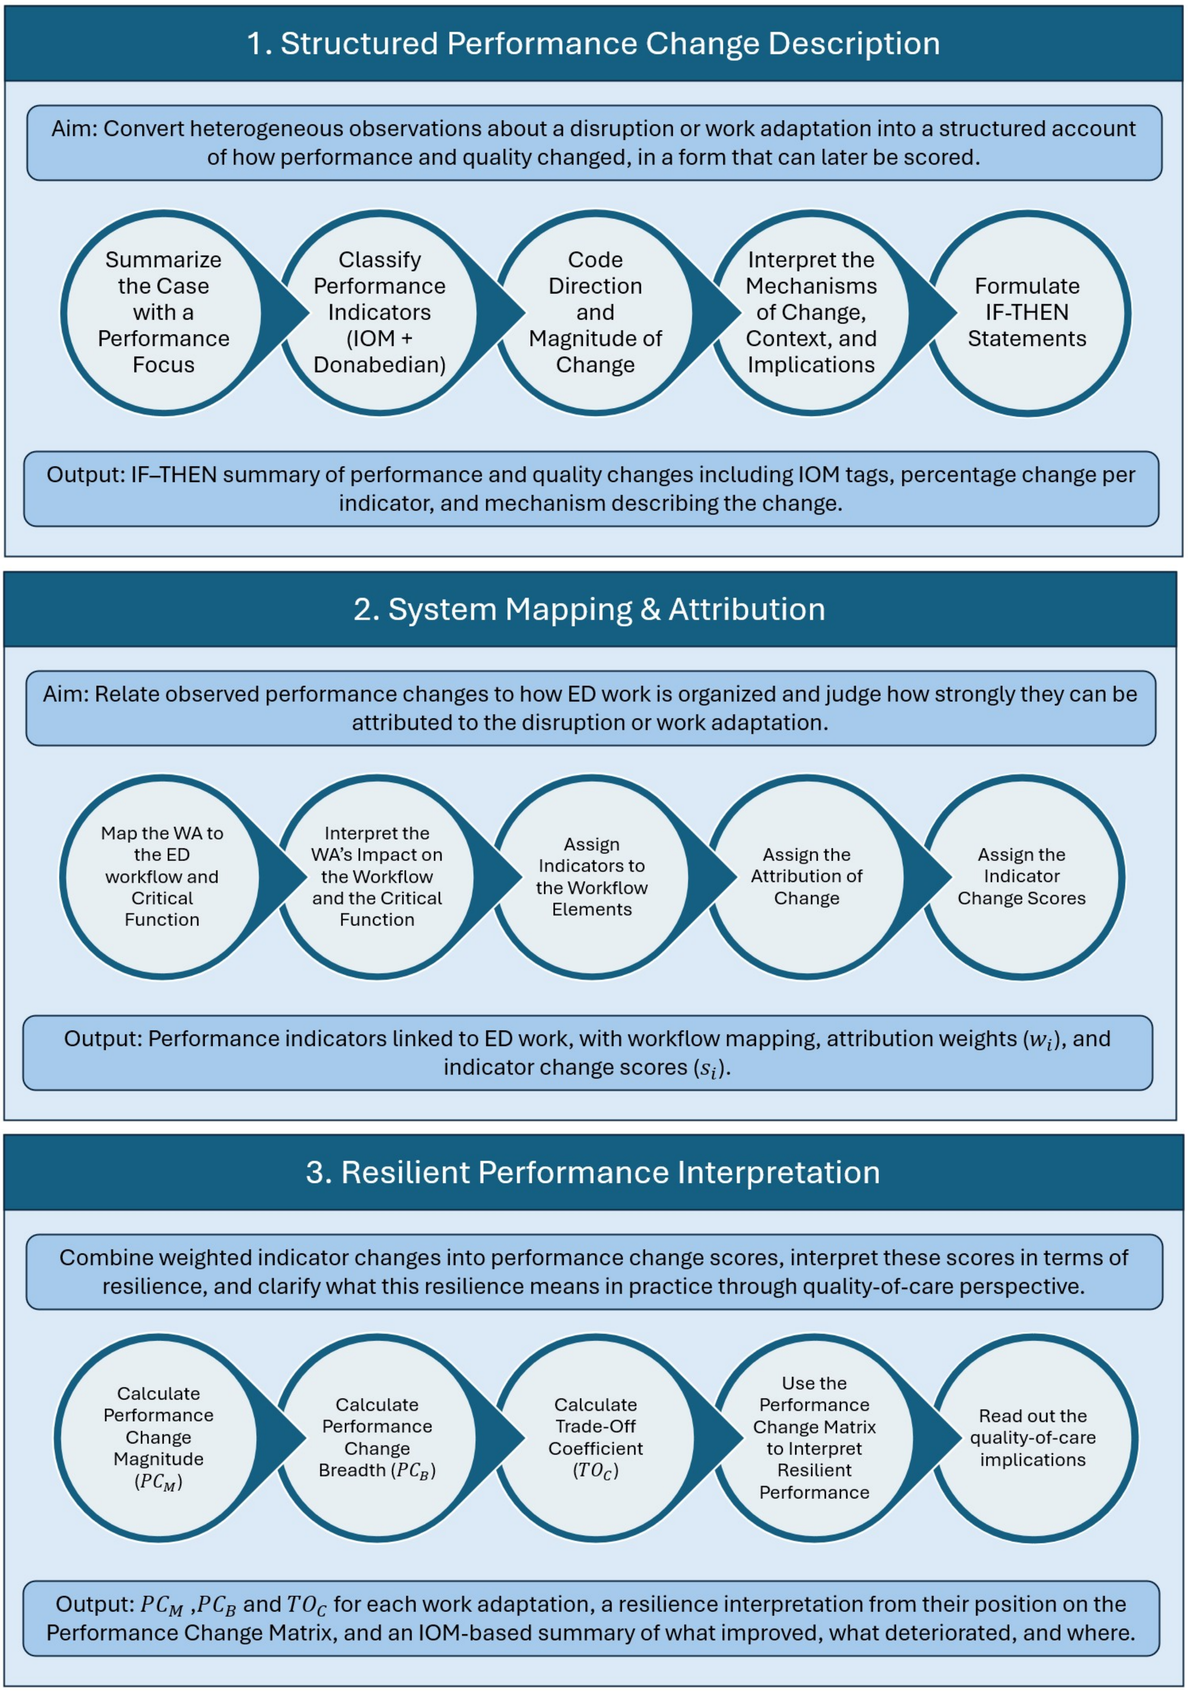

Figure 3.1: Performance Indicator - Resilience Assessment (PI-RA) Framework

3.1. Sub-Question 1 - ED Workflow and Critical Function

This section explains the method used to answer the first sub-question: "What are the ED's critical function and workflow?" It describes how we map the ED workflow and define its critical function within the Dutch emergency care context. Defining the workflow and the critical function sets the reference point for the performance analysis. With these definitions, we can identify which parts of performance may be affected. Without them, we could only speak in vague terms (for example, saying that a work adaptation affects timeliness of treatment). A clear ED workflow lets us show exactly where an adaptation acts and which later steps in the workflow are also affected. This makes the effects visible and supports future modelling. Linking the workflow to the critical function then allows us to trace how changes to specific steps in the workflow influence the ED's ability to deliver its critical function.

First, a contextual analysis of Dutch Emergency Departments was conducted to understand their role within the healthcare system. This included: (i) situating the Netherlands on the Franco-German vs. Anglo-American emergency care spectrum to clarify upstream filtering and gatekeeping; (ii) characterizing the ED demand context (who presents and under what circumstances); (iii) mapping arrival routes and referral mechanisms; and (iv) delineating ED workflow elements by distinguishing activities handled upstream versus those within ED operations. These elements are presented in Chapter 2 (Theoretical Background and Literature Review) in this sequence, reflecting the analytical progression applied in this thesis. Taken together, this produced a system-level understanding of what Dutch EDs are expected to do and where their operational control ends.

To tailor the ED workflow to the Dutch context, we then outline the common elements and phases of an ED. To do this, we draw on two complementary conceptual frameworks:

- **Asplin et al. (2003) ED Model:** Divides ED operations into three phases — input, throughput, and output — providing a high-level structure for understanding ED functioning.
- **Claassen et al. (2025) Framework:** Offers a detailed representation of patient journeys within the Dutch Acute Care Chain (ACC), distinguishing four phases: pre-referral, referral, ED, and post-ED.

Both frameworks are introduced and discussed in section 2.2.3.

The Netherlands-specific ED workflow was developed by combining the structure from these frameworks with findings from the contextual analysis explained above. The model positions the ED within the broader acute care chain. More importantly, it sets out the operational phases in a clear, structured way and helps identify the ED's critical function. A clear delineation of responsibilities and system context is needed for valid assessment of performance and resilience.

This stepwise methodology ensured that the definition of the ED's critical function was grounded in both empirical context and conceptual modeling, while also laying the foundation for performance and resilience assessments in later sub-questions.

3.2. Sub-Question 2 - Gathering Structured Performance Change and Relevant Insights

This section addresses the second sub-question: *"How can ED work adaptations in response to disruptions (e.g., COVID-19) be analyzed systematically to extract structured evidence and insights on performance?"*. It explains how we collect the inputs for the PI-RA framework and how we code and analyze them in a consistent way so the evaluation uses a common method and language. The study so far focused on the Dutch context, and the intended data collection and subsequent analysis were planned to focus on ED performance in the Netherlands. However, several complications arose that changed the trajectory of the thesis.

3.2.1. Initial Plan - PI Data analysis over COVID-19 Phases

The original design for Sub-question 2 was to track how Emergency Department (ED) performance changed during the COVID-19 pandemic in the Netherlands. The approach combined two elements: (i) mapping changes in performance indicators (PIs) over time, and (ii) linking those changes to the work adaptations introduced by EDs in response to the disruption. Together, these steps were intended to build a time-based picture of resilience by tracing how performance indicators changed across the disruption: from the initial performance drop after COVID-19 onset (absorption, t_1–t_2) to the subsequent adaptation and recovery phases when work adaptations were introduced (t_2–t_4), as conceptualized in the resilience curve in Figure 3.2. To carry this out, we first defined a set of PIs to measure ED performance. The method for selecting them and the final list are provided in Appendix A.

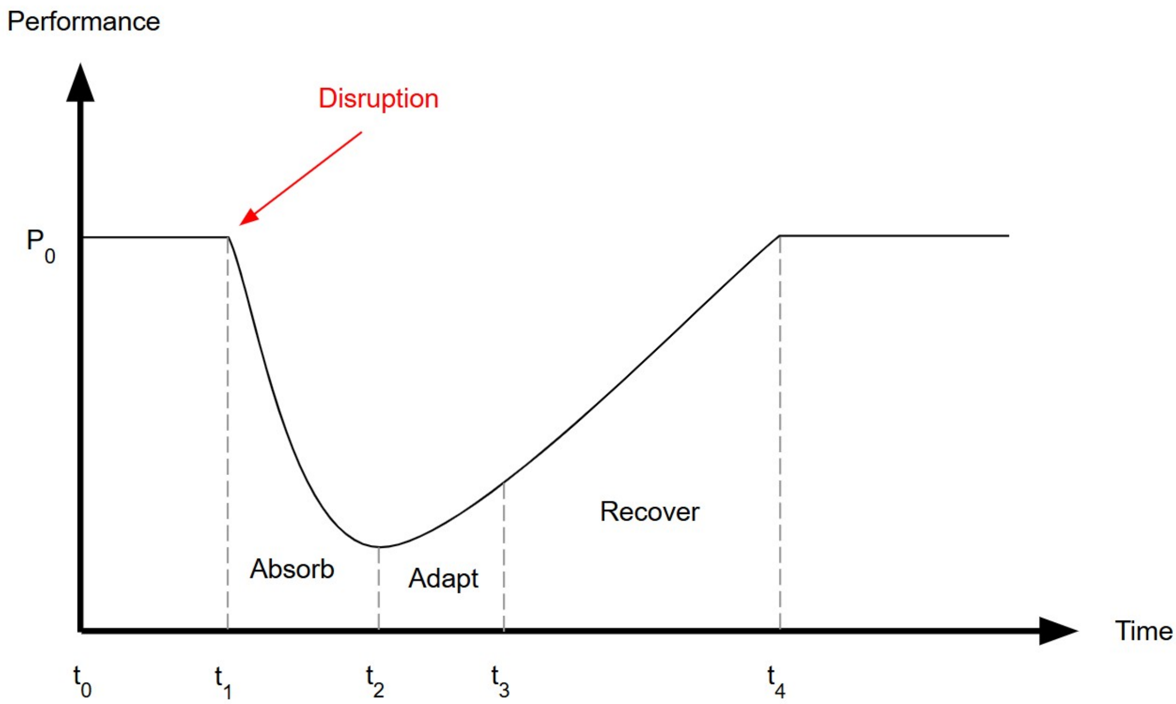

Figure 3.2: Resilience Curve: System performance trajectory across disruption, absorption, adaptation, and recovery

1) Gather Performance Indicators Across COVID Phases

The first step was to collect ED performance indicators (PIs) across the COVID-19 phases in the Netherlands. As noted above, temporal variation in PIs was expected to reflect the effects of the pandemic pressures followed by the work adaptations.

To operationalize this, national PI sources were explored. Databases from the Dutch National Institute for Public Health and the Environment (RIVM), VZINFO, and the Central Bureau of Statistics (CBS) were screened for the selected indicators in Appendix A. When this search did not yield results, the parties were contacted directly. RIVM advised that the request be directed to VZINFO – Acute Zorg. VZINFO then referred the request to CBS, specifically the department responsible for microdata (Catalogus microdata | CBS).

CBS informed us that access would require an upfront fee of €3,000 plus €200 per month, with no clear guarantee that the requested ED-level PI data would be available. Given these barriers, this option was deemed infeasible for the scope of the thesis.

The data request form can be seen in Appendix B.

2) Work Adaptations Across Phases

The second step was to link shifts in PIs to organizational or process changes made during COVID-19 (e.g., revised triage, surge staffing, altered patient pathways). As highlighted in the literature review, O'Connor et al. (2021) documented a variety of such work adaptations implemented in Dutch EDs during COVID-19.

The plan was to map these adaptations onto the COVID-19 phases and examine how they coincided with or influenced PI trends. First, regulations or policy changes aimed at EDs were analyzed, as they would provide exact implementation dates. However, the literature review revealed that no such measures were applied in the Netherlands (see Section 2.4.1), and EDs were expected to adjust their own operations (see Section 2.4.2). Furthermore, studies describing the work adaptations were limited, with most reported adaptations concentrated in the first wave and without a detailed account of what was implemented when.

Transition Away from the Initial Plan

The intended logic of the initial plan was clear: performance indicators would be tracked across pandemic, and these temporal shifts would then be interpreted through the impact of work adaptations.

However, the lack of PI datasets made this approach infeasible. At the same time, systematic information on ED-focused regulations, policy changes, or process-level adaptations was absent, with the few identified examples concentrated in the first wave. This dual absence in both PI data and clear timelines of which adaptations were implemented when meant the mapping could not be executed reliably.

Recognizing these limitations, the study adapted methodologically to analyzing case studies of ED work adaptations during COVID-19. These sources provided more detailed accounts of when adaptations were implemented, observed changes in some performance indicators, and context to interpret their impact on performance. A key limitation remained: there were no Netherlands-based studies that tracked how a specific adaptation led to changes in performance indicators over time. The few studies that made such links were conducted outside the Netherlands.

3.2.2. New Plan - International Case-Based Approach

For this revised plan, the analysis uses international empirical studies that describe organizational and process-level work adaptations and report associated changes in quantitative performance indicators (PIs) before and after implementation. These studies give detailed accounts of how specific adaptations were introduced and how selected PIs responded over time. This lets us analyze the impact of work adaptations on performance using more grounded evidence than the initial plan of mapping adaptations to PI trends. The trade-off is that each case uses a limited set of PIs, which narrows the breadth of performance insights.

Because these cases come from outside the Netherlands, what we can conclude about performance and resilience within the Netherlands is limited. We therefore interpret the adaptations functionally rather than in a Dutch-specific context. We assess the effects of the adaptations using established quality dimensions, which allows the lessons to generalize beyond the immediate case settings.

In order to systematically evaluate each work adaptation and its impact on performance, the case studies are processed using the steps shown in Figure 3.3, which represent the first phase of the PI-RA framework. The aim is to turn heterogeneous study reports into a structured description of how each adaptation changed measurable performance and quality, in a form that can later be used for scoring. Central to this is the construction of an IF–THEN statement for each work adaptation, which condenses the key performance effects into one consistent statement. The individual steps of this phase are detailed in the list below.

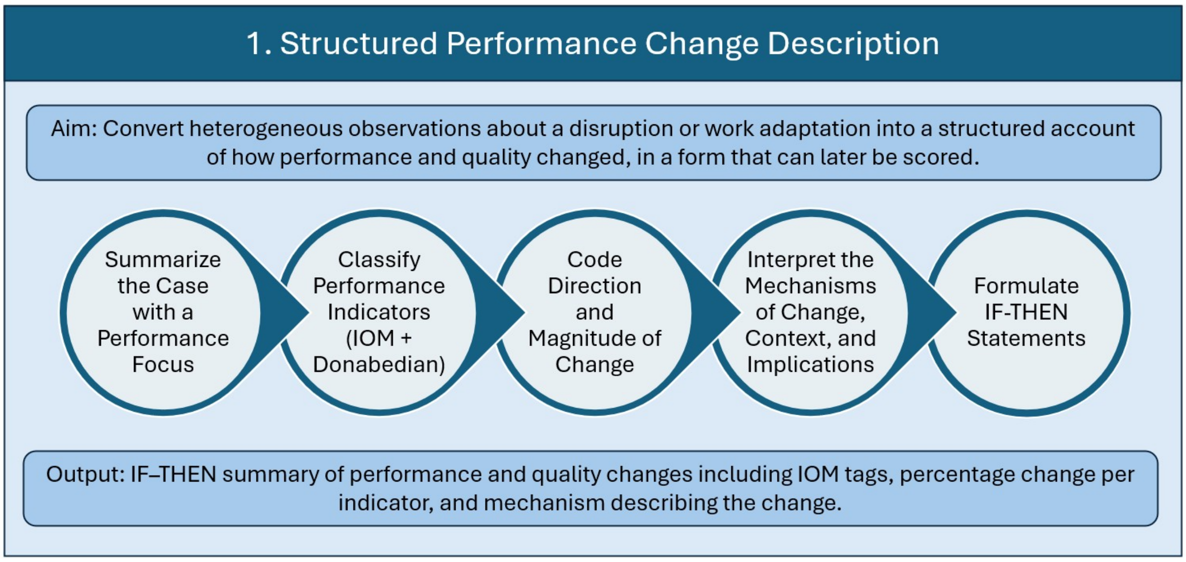

Figure 3.3: Phase 1 of PI-RA: Structured Performance Change Description

1. **Summarize the case with a performance focus.**
   The study is briefly summarized in terms of the emergency department setting, the work adaptation that was implemented, and the main performance issues it was intended to address (e.g. crowding, delays, infection risk).
2. **Classify Performance Indicators.**
   All performance indicators reported in the study are listed and classified under Donabedian. Each indicator is then linked to the relevant IOM quality dimension(s) it measures.
   - **Donabedian classification:** This element sets the angle from which performance

is assessed and defines the scope of quality assessment. Practically, it requires coding each performance indicator into one of three categories: structure (how the ED or hospital is prepared to perform—resources, capabilities, readiness), process (how care is actually delivered—the actions taken during care), and outcome (the overall results or success of the system). Using this classification lets us systematically organize indicators and consistently interpret whether we are measuring preparedness, execution, or results.

- **IOM quality dimension classification:** Each indicator is then linked to the relevant IOM quality dimension (i.e., Safety, Effectiveness, Patient-Centeredness, Timeliness, Efficiency, Equity). In this thesis, these IOM tags provide the value lens through which performance is interpreted: because each indicator is associated with a quality dimension, any change in that indicator can be read as a statement about which aspects of care quality improve or deteriorate.

3. **Code direction and magnitude of change.**
For each indicator, the reported change between the pre- and post-adaptation periods is translated into (i) a direction (increase or decrease, interpreted as improvement or deterioration depending on what is desirable for that indicator) and (ii) a magnitude category using the percentage bins (low, medium, high, very high). This provides a comparable description of how strongly each indicator moved and in which quality dimension; the details of this magnitude classification are described in the following subsection.

4. **Interpret the mechanisms of change, context, and implications.**
The indicator changes are interpreted in light of the case description: how the adaptation is expected to work, and what contextual factors may influence the results (e.g. case mix, boarding, concurrent policies). This step extracts detailed insights into what the observed changes in indicators mean for the functioning of the ED in that case, providing the basis for later connecting these findings to the ED workflow and critical function. Without this interpretation, the subsequent IF–THEN statements would only capture a partial picture of the observed changes in performance.

5. **Formulate an IF–THEN statements.**
Finally, the findings are condensed into a standardised IF–THEN statement of the form:

   If [work adaptation A] occurred (for each performance indicator);

   - performance indicator X [increased or decreased] by [magnitude], indicating [improvement or deterioration] in [IOM quality dimension];
   - .....
   - performance indicator Y [increased or decreased] by [magnitude], indicating [improvement or deterioration] in [IOM quality dimension];

   The statement lists all relevant indicators. In this way, one compact IF–THEN statement captures which indicators changed, by how much, and what these changes imply for different dimensions of care quality.

Magnitude Classification and Directionality

Evaluating “how much” a performance indicator (PI) changed is not straightforward in the ED context. Pure percentages can mislead because the same relative change can mean very different things operationally. For example, reducing ED Length of Stay (ED LOS) by 2 hours when the baseline is 10 hours is a major improvement for overall flow, yet a small relative change on paper. By contrast, cutting the waiting time to initial physician from 8 to 4 minutes looks like a 50% reduction, but the absolute gain is only 4 minutes—clearly not comparable to saving two hours of ED LOS. The same problem appears with mean differences: a 30-minute decrease in ED LOS from 10 hours and a 30-minute decrease in waiting time from 40 to 10 minutes have the same absolute delta, but the latter is often more consequential (e.g., ”leaving without being seen” risk or patient safety risk due to late provider encounter). These examples show that neither percentages nor raw means capture the operational importance tied to where in the patient journey the change occurs.

Using percentage points (pp) also fails to resolve this. Consider outcomes of very different salience: increasing mortality from 1% to 3% is “only” +2 pp, yet it is a 300% relative increase and carries severe clinical implications. Meanwhile, bed occupancy rate (BOR) increasing from 40% to 60% is +20 pp but only a 50% relative increase, and an ED might still function adequately at 60% occupancy. In short, pp can overweight benign shifts in high-base, low-risk metrics and underweight critical shifts in low-base, high-risk metrics. Standardized effect sizes (e.g., Cohen’s h for proportions, Hedges’ g/Cohen’s d for means) help with statistical comparability but still do not align magnitude with operational/clinical value; a change can be “large” statistically while modest in practical impact—or vice versa.

Given these issues, this thesis adopts a pragmatic magnitude classification for interpreting PI changes, grounded in common sense, the literature reviewed on PI behavior, and how changes tend to affect ED performance and risk. The schema below is used consistently throughout the results and interpretation:

- No change: 0–1%
- Low: 1–5%
- Medium: 5–20%
- High: 20–75%
- Very high: >75%

This classification is intentionally simple and transparent. It supports cross-metric readability without claiming to encode clinical importance perfectly. Throughout the thesis, whenever relevant, results are accompanied by contextual notes about known implications about changes in PIs to avoid over-interpreting the percentage label alone.

Directionality: Magnitude labels are paired with directionality, since an increase is desirable for some PIs (↑ desired) while a decrease is desirable for others (↓ desired). For each PI, the desired direction is applied consistently. Reported changes are interpreted as improvements if they move in the desired direction and as deterioration if they move opposite to it.

Future discussion and research: The limitations of percentage-based bins will be revisited in the Discussion. A more value-aware magnitude bin selection—co-developed with clinical and operational experts—could refine thresholds by PI type so that magnitude categories better reflect real-world ED performance and safety impact. For the purposes of this thesis, however, the above classification will be used consistently to interpret changes across PIs.

This allows us to summarize the impact of the work adaptation into concise, consistent statements that serve as inputs to the PI-RA framework.

3.3. Sub-Question 3 - Interpreting Performance Change as Resilient Behavior

This section addresses the third sub-question: *How can structured evidence and insights on ED performance be translated into quantitative and qualitative measures of performance change that support a transparent assessment of resilient performance?*

Resilience was defined as "the intrinsic ability of a system (a clinic, ward, a hospital, a country) to adjust its functioning before, during, or following events (changes, disturbances, and opportunities) and thereby sustain required operations under both expected and unexpected conditions." In this thesis, required operations are understood as the ED workflow and its critical function. The act of "adjusting the functioning" corresponds to the work adaptations (WAs) analysed here, which have observable consequences for performance indicators (PIs). Tracking how these indicators change is therefore the main way this thesis assesses resilient performance.

In the resilience curve in Figure 3.4, these work adaptations occur in the adaptation and recovery segments (between t_2 and t_4), where the system reorganises and performance begins to improve after the initial disruption. The case studies on which this thesis is based report PI changes by comparing periods before and after a work adaptation, and do not provide data on the performance drop following the disruption itself (absorption phase between t_1 and t_2). As a result, the analysis in this sub-question primarily characterizes how performance changes during adaptation and recovery, or how the ED climbs back up the resilience curve. Conceptually, however, the same scoring logic could also be applied to the absorption phase if performance data before disruption were to be available.

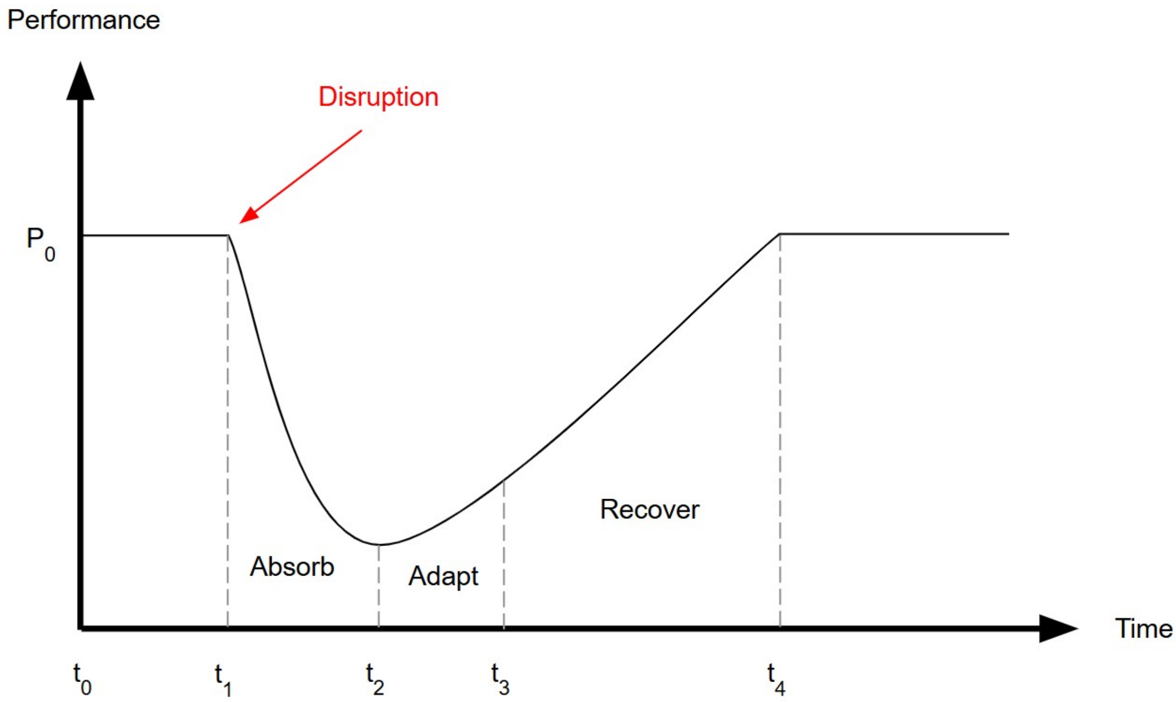

Figure 3.4: Resilience Curve: System performance trajectory across disruption, absorption, adaptation, and recovery

Below, we describe the practical steps required to apply this framework in the ED context. After outlining these application steps, sections 3.3.1 and 3.3.2 then detail the numerical procedures: how attribution weights and indicator change scores are constructed, how the performance change scores PC_M and PC_B and the trade-off coefficient TO_C are calculated, and how these scores are interpreted using the Performance Change Matrix.

System Mapping and Attribution

Figure 3.5 highlights the *system mapping and attribution* phase of the PI-RA framework. In this phase, the ED workflow and critical function model from Sub-question 1 is used to see where a disruption or work adaptation actually affects the system. It clarifies which parts of ED work each indicator reflects and how plausibly the observed changes can be linked to the change in work, providing the basis for later aggregation. The steps that make up this phase are described in the following list.

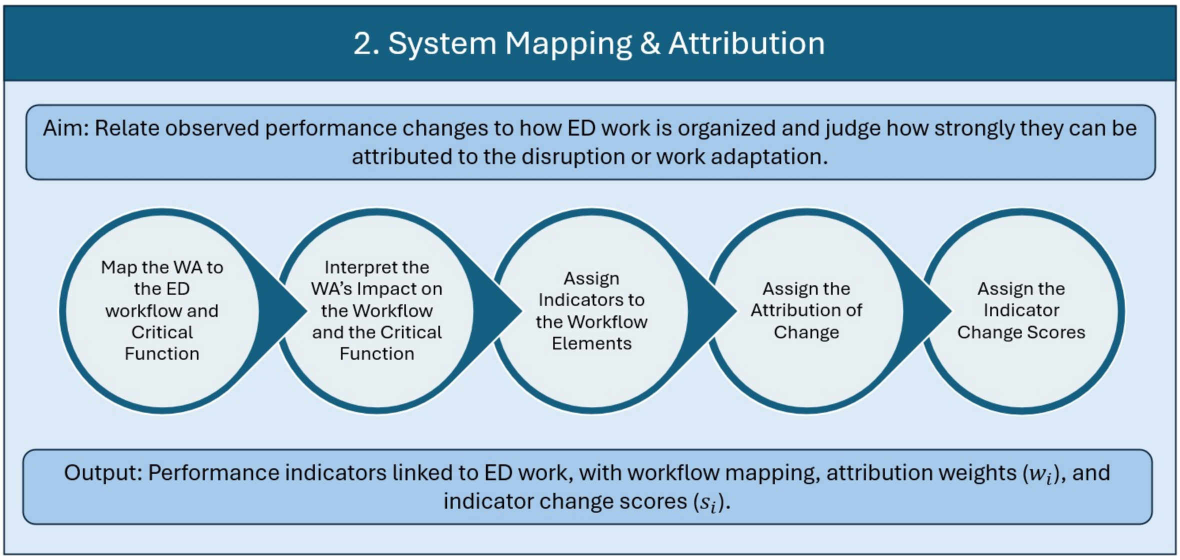

Figure 3.5: Phase 2 of PI-RA: System Mapping and Attribution

1. **Map the WA to the ED workflow and the critical function.**
   Identify the phase(s) and node(s) in the ED workflow where the WA acts and specify which elements of the critical function it addresses. This establishes *what* in the system the WA targets and anchors interpretation in the operational backbone from SQ1.
2. **Interpret the WA's impact on the workflow and the critical function.**
   Describe the expected mechanism of change across the workflow (propagation, bottlenecks relieved or created) and its implications for the critical function. This makes explicit *how* the WA influenced required operations.
3. **Assign indicators to workflow elements and the critical function.**
   For each reported performance indicator, specify where in the ED workflow it is measured (e.g. triage, diagnostics, discharge) and which part of the critical function it reflects (rapid assessment, diagnosis, treatment, flow). Retain the IOM quality tags and Donabedian classification from SQ2 so that each indicator remains linked to its quality dimension and type.
4. **Assign the Attribution of Change.**
   For each indicator, judge how much of the observed change can credibly be attributed to the work adaptation rather than to other factors. On this basis, assign an attribution

weight w_i (e.g. 1.0 for direct attribution, 0.5 for mixed attribution, 0 for off-path effects). These weights later determine how strongly each indicator contributes to the aggregated scores.

5. **Assign the Indicator Change Scores.**
Using the magnitude categories and directions coded in SQ2, convert each indicator's change into a signed indicator score s_i. The absolute value of s_i reflects the magnitude category (no, low, medium, high, very high), while the sign reflects whether the change is desirable or undesirable given the indicator's desired direction. These scores provide a common numeric scale for aggregating heterogeneous indicators in the next phase.

Steps 4 and 5 will be further detailed in Section 3.3.1.

Resilient Performance Interpretation

Figure 3.6 shows the final *resilient performance interpretation* phase of the PI–RA framework. Once indicator changes have been located in the system and judged for attribution, this phase compresses that information into a concise description of overall performance change. It explains how Performance Change Magnitude (PC_M), Performance Change Breadth (PC_B), and the trade-off coefficient (TO_C) are calculated and how their combination is read on the Performance Change Matrix and along the resilience curve. The steps in this phase are detailed in the list below.

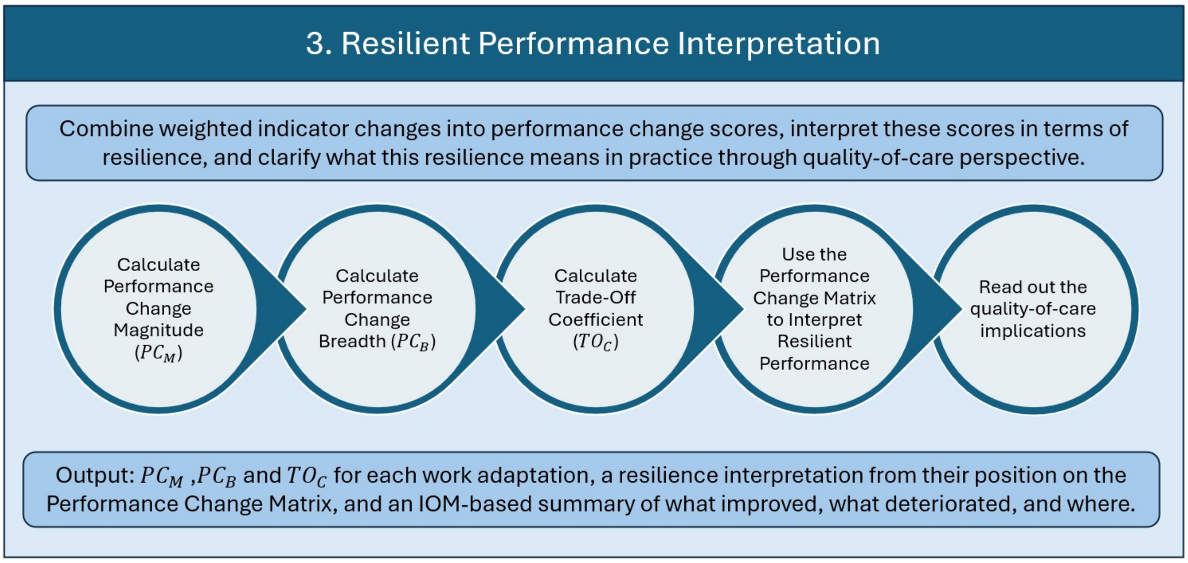

Figure 3.6: Phase 3 of PI-RA: Resilient Performance Interpretation

1. **Calculate Performance Change Magnitude (PC_M).**
Using the attribution weights w_i and the signed indicator change scores s_i, compute the performance change magnitude PC_M as a relevance-weighted mean (Eq. 3.1). This score summarizes the *net* direction and magnitude of change across all relevant indicators: positive values mean that, on average, required operations perform better than before; negative values mean they perform worse; and the absolute size of PC_M indicates how strong this overall improvement or deterioration is.

2. **Calculate Performance Change Breadth (PC_B).**
Next, compute the performance change breadth PC_B using the same weights w_i and

the sign of each indicator score, $\text{sign}(s_i)$ (Eq. 3.2). PC_B expresses the *net breadth* of change in units of weighted indicators: positive values mean that, after accounting for the relevance weights, more indicators improve than deteriorate on net; negative values mean that deteriorations dominate on net; and zero indicates that improvements and deteriorations balance each other.

3. **Calculate the trade-off coefficient (TO_C).**
   To characterize how strongly gains are accompanied by losses, normalize PC_B by the total relevance weight to obtain the trade-off coefficient TO_C (Eq. 3.3). This coefficient is bounded between -1 and $+1$ and indicates how dominant improvements or deteriorations are across the relevant indicators. Values of TO_C close to $+1$ mean that almost all weighted indicators move in the desired direction and trade-offs are limited; values near 0 indicate that improvements and deteriorations largely offset each other, so performance change is driven by trade-offs; and values close to -1 mean that almost all weighted indicators move in the undesired direction.

4. **Use the Performance Change Matrix to Interpret Resilient Performance.**
   The two scores PC_M and TO_C are then combined by plotting the work adaptation as a point on the Performance Change Matrix, with PC_M on the horizontal axis and TO_C on the vertical axis. The four regions of the matrix distinguish broad net improvement, trade-off-heavy improvement, broad net deterioration, and trade-off-heavy deterioration. In this way, the matrix provides a compact description of resilient performance: the combination of PC_M and TO_C indicates whether the observed pattern resembles an upward, stable, or downward movement on the resilience curve, and whether that change is broadly shared or driven by trade-offs between different parts of ED performance. The absolute net breadth PC_B is used alongside the matrix, in the narrative interpretation, to indicate how widely the observed pattern of change extends across the measured aspects of ED functioning.

5. **Read out the quality-of-care implications.**
   Summarize the material trade-offs by IOM dimensions and tie them back to the workflow location and critical-function elements affected (what improved, what deteriorated, and where). This makes explicit how the pattern of changes behind a given (PC_M, TO_C) position on the Performance Change Matrix reflects resilient performance in practice.

Detailed descriptions of how to calculate PC_M, PC_B, and TO_C are provided in Section 3.3.1, and the Performance Change Matrix and how to interpret it is provided in Section 3.3.2.

3.3.1. Performance Change Magnitude and Breadth & Trade-Off Coefficient

Building on the resilience definition unpacked above, this subsection sets out how we translate case evidence into a comparable read-out of resilient performance. The idea is straightforward: for a given WA, we read the observed PI changes, weight them by how closely they speak to where the WA acts in the workflow, and combine them into a normalized score that reflects the net balance of gains and losses for the ED's critical function, alongside a complementary breadth read-out that summarizes the weighted net balance of indicators improving versus worsening.

This part aims to quantify the change in performance associated with each work adaptation (WA) by aggregating indicator movements, while accounting for (i) where each indicator sits in the ED workflow and critical function, and (ii) how strongly the observed change in that indicator can be attributed to the WA, as captured by the relevance weights introduced below. These performance-change scores are then later interpreted as expressions of resilient performance.

Indicator Relevance (by Workflow Phase and Mechanism)

To keep the score faithful to what the work adaptation (WA) is plausibly causing, each reported PI is interpreted through its *causal linkage* to the WA's mechanism and its placement on the ED workflow. The question is: *why should this metric move if the WA is operating, and what evidence links the change to the WA rather than to other factors?* Workflow phase can inform that judgment (e.g., timing and sequence), but phase alone does not determine the weight. Indicators without a credible causal link in this context are set aside.

Concretely, for each WA, classify every reported indicator i as:

- **Directly attributable ($w_i = 1.0$).**
  The indicator has a clear, specific causal link to the WA (mechanism and timing). This can occur at the intervention node *or downstream* when the mechanism explains the connection (e.g., defined handoffs that predictably affect the metric).
- **Mixed attribution ($w_i = 0.5$).**
  The indicator plausibly reflects the WA *and* other concurrent drivers (e.g., volume/acuity shifts, concurrent policies, secular trends). The link remains credible but diluted; the metric is informative yet down-weighted.
- **Off-path / no credible link ($w_i = 0$).**
  The indicator lacks a defensible causal connection to the WA's mechanism in this context and is excluded from scoring.

This attribution-first weighting keeps the score focused on what the adaptation is set up to achieve, while still acknowledging meaningful spillovers and filtering out noise.

Magnitude Bins

In Section 3.2.2, Magnitude Classification and Directionality, each indicator's observed percentage change was classified to a bin score that encodes its magnitude of change. Here, we map these bins to indicator change scores (absolute values $|s_i|$) that quantify how large the change is. The scoring values are chosen on an ordinal scale that increasingly up-weights larger changes. Each step gets a larger jump than the previous one (0.1 → 0.2 → 0.3 → 0.4), which reflects the assumption that progressively larger percentage shifts in ED indicators are increasingly difficult to achieve and therefore should contribute more strongly to the aggregate score. The sign of s_i records the desirability for that indicator: assign $s_i > 0$ when the change in PI is in the desired direction (improvement), and assign $s_i < 0$ when the change is in the undesired direction (detrimental).

A score of 0 means the indicator did not change and therefore contributes no directional push to the aggregate. This is not a negative judgment: on this scale, 0 denotes stability, and stability is itself a valid resilient outcome when the aim is to sustain required operations. Positive values indicate improvement, negative values indicate deterioration, and 0 indicates stability.

The magnitude-to-score mapping can be seen in Table 3.1

Figure 3.7 visualizes this mapping by plotting the percentage-change magnitude bins on the horizontal axis and the corresponding indicator change scores on the vertical axis. The step-

Table 3.1: Magnitude bins and corresponding absolute indicator change scores.

Magnitude Bins (Percentage Change in PI)	Indicator Change Score (s_i Absolute Value)
Very High	1.0
High	0.6
Medium	0.3
Low	0.1
No change	0

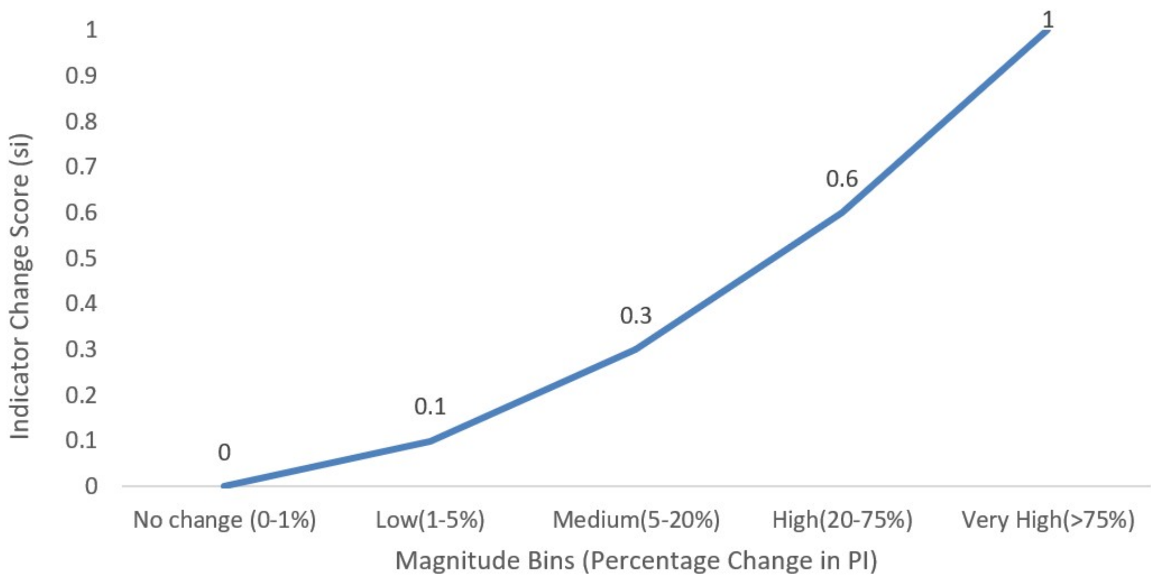

Figure 3.7: Mapping of PI Percentage-Change Magnitude Bins to Indicator Change Scores

wise, non-linear increase in scores makes explicit that larger percentage shifts in an indicator are weighted disproportionately more heavily in the aggregate.

Aggregate with relevance weighting and normalize

We compute a relevance-weighted mean, normalized by the sum of non-zero weights. Normalizing by the total number of indicators rather than by the relevant ones would let non important indicators dilute the score.

To do that, we build directly on the two elements defined above, (i) the relevance weights w_i and (ii) the indicator change scores s_i to compute **Performance Change Magnitude (PC_M)** for a given work adaptation:

$$PC_M = \frac{\sum_i (w_i \cdot s_i)}{\sum_i w_i} \tag{3.1}$$

The resulting PC_M score is unitless and normalized by the sum of non-zero weights, so it reflects the average weighted direction and magnitude of change across the indicators that are actually expected to move. Because the indicator change scores $|s_i|$ range from 0 to 1, the aggregated PC_M is bounded between -1 and $+1$, where negative values indicate net deterioration and positive values indicate net improvement, with $|s_i| = 1$ representing an extreme case in which all relevant indicators deteriorate or improve very strongly. In this thesis, PC_M is used to summarise the net change in performance for each work adaptation within its own case. Although the normalization places all work adaptations on a common numerical scale, the scores are not used as strict benchmarks across EDs, because the underlying cases differ in their baseline (pre-disruption) performance levels and in where along the disruption–recovery trajectory the measurement window begins. In principle, if comparable baseline data and aligned disruption timing were available across EDs, PC_M could support more direct comparison; here, it is interpreted primarily as a relative read-out of performance change within each case. To aid interpretation, we can classify the PC_Ms into the following tiers:

- $PC_M > 0$: performance improves;
- $PC_M = 0$: there is no change in performance;
- $PC_M < 0$: performance deteriorates.

A limitation of PC_M is that it reflects only the *magnitude* of performance change (i.e. how the system's required operations sustained, improved, or worsened on average), but it does not capture *how many* indicators moved in the desired direction (how broadly ED workflow elements and the ED critical function were affected). To address this, we also report a **Performance Change Breadth (PC_B)** score, defined as

$$PC_B = \sum_i w_i \operatorname{sign}(s_i), \tag{3.2}$$

which expresses the *net breadth* of change in units of weighted indicators. Positive values of PC_B mean that, after accounting for the relevance weights, more indicators improve than deteriorate on net; negative values mean that deteriorations dominate on net; and a value of $PC_B = 0$ indicates that improvements and deteriorations balance exactly (or that there was no change in the weighted set of indicators). For example, $PC_B = 3$ means that, on net, the equivalent of three weighted indicators improve. However, this absolute net breadth does not reveal how strongly gains are accompanied by losses: a value of $PC_B = 3$ could arise from five improving and two deteriorating indicators or from ten improving and seven deteriorating indicators. In both situations the net breadth is the same, even though the second configuration involves a heavier pattern of trade-offs across the indicator set.

To make this balance between improvements and deteriorations more transparent, and to place breadth information on the same $[-1, 1]$ scale as PC_M, we normalise PC_B by the total relevance weight and define a **trade-off coefficient (TO_C)**:

$$TO_C = \frac{PC_B}{\sum_i w_i} = \frac{\sum_i w_i \operatorname{sign}(s_i)}{\sum_i w_i}. \tag{3.3}$$

Because $\operatorname{sign}(s_i)$ can only take the values $-1, 0,$ or $+1$, the trade-off coefficient TO_C is bounded between -1 and $+1$. For interpretation, this thesis uses the following qualitative reading:

- $TO_C > 0$: on net, more (weighted) indicators improve than deteriorate. Values close to $+1$ indicate that almost all relevant indicators move in the desired direction and trade-offs are minimal, whereas values closer to 0 indicate that the gains are increasingly accompanied by deteriorations.
- $TO_C = 0$: improvements and deteriorations balance out in terms of spread; performance change is dominated by trade-offs.

- $TO_C < 0$: on net, more (weighted) indicators deteriorate than improve. Values close to -1 indicate that almost all relevant indicators move in the undesired direction, while values closer to 0 indicate that some improvements remain but are outnumbered by losses.

Taken together, PC_M and TO_C therefore provide two complementary views on performance change. PC_M summarizes *how strong* the net change in performance is across the relevant indicators, while TO_C summarizes *how aligned* those changes are in terms of improvements versus determinations, that is, how strongly the overall pattern is shaped by trade-offs. This makes the pair (PC_M, TO_C) well suited for joint visualisation in the Performance Change Matrix in the next section, where PC_M forms the horizontal axis (net magnitude of change) and TO_C forms the vertical axis (dominance of improvements versus deteriorations). Because TO_C is normalized, it does not by itself show how many weighted indicators are involved in the change; for that reason, the absolute net breadth PC_B is retained alongside the matrix and used in the narrative interpretation of resilient performance to indicate how wide the impact of a work adaptation extends across the measured aspects of ED functioning.

3.3.2. Performance Change Matrix

The two scores introduced above — performance change magnitude (PC_M) and the trade-off coefficient (TO_C) — capture complementary aspects of how a work adaptation affects ED performance. On their own, each score is a one-dimensional summary: PC_M reflects the net magnitude and direction of change in the required operations, while TO_C reflects how strongly improvements dominate over deteriorations, that is, how trade-off-heavy the overall pattern of change is. To interpret resilient behaviour better, it is useful to see how these two dimensions combine. The *Performance Change Matrix* (Figure 3.8) therefore places each work adaptation as a point in a two-dimensional space, showing at a glance whether performance is improving or deteriorating and whether those changes occur with limited trade-offs or are largely driven by trade-offs.

In Figure 3.8, the horizontal axis represents PC_M and the vertical axis represents TO_C, both ranging from -1 to $+1$. The origin $(PC_M = 0,\ TO_C = 0)$ marks a situation in which there is no net change in performance and improvements and deteriorations balance each other out in the weighted indicator set. For the purposes of the matrix, the interpretation focuses primarily on the *sign* of each score: the space is divided into four regions based on whether PC_M and TO_C are positive or negative.

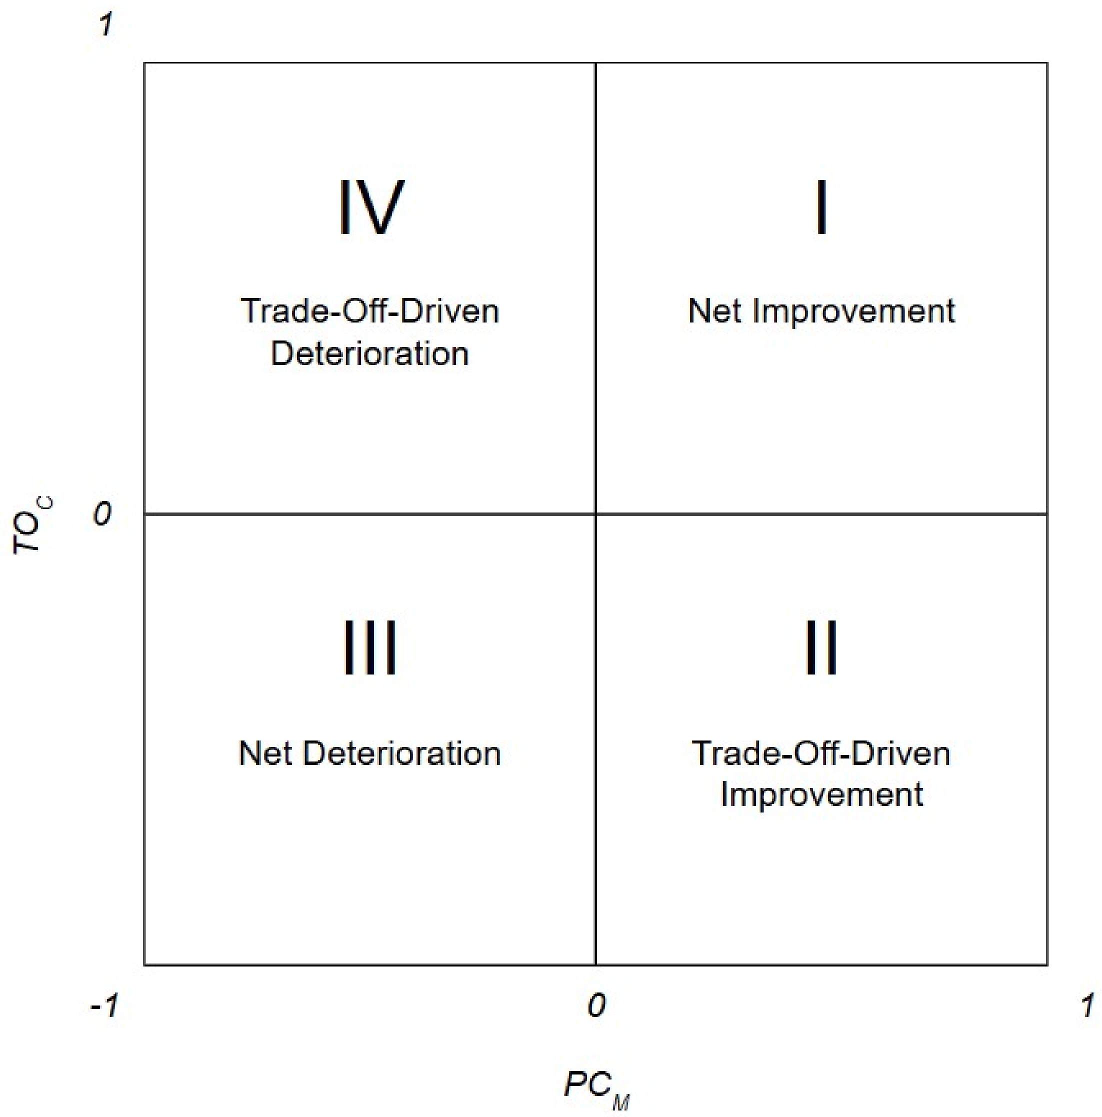

Figure 3.8: Performance Change Matrix

The four regions can then be read as follows:

Region I: $PC_M > 0, TO_C > 0$ --- net improvement.

Performance improves overall, and more (weighted) indicators improve than worsen. On the resilience curve (see Figure 3.4), this corresponds to a clear upward recovery segment in which care moves in the desired direction. Within this region of Figure 3.8, movement from left to right reflects increasing net performance gains as PC_M grows more positive, while movement from bottom to top reflects a pattern in which improvements increasingly dominate over deteriorations as TO_C approaches $+1$. As PC_M and TO_C move closer to $+1$, the work adaptation's position shifts toward the upper-right corner of Figure 3.8, indicating that improvements become both stronger and less dependent on trade-offs. *Region I therefore represents the desired outcome of a work adaptation*: ideally, changes to ED work should place the system in this region, indicating that resilience is expressed by putting ED performance into the recovery segment, through a net improvement in performance in which gains clearly outweigh any losses, rather than through trade-off-heavy changes.

Region II: $PC_M > 0, TO_C < 0$ --- trade-off-driven improvement.

The net magnitude is positive, but more (weighted) indicators deteriorate than improve. This pattern arises when a limited set of indicators improve strongly, while a larger set shows smaller deterioration. In resilience terms, the curve still moves upward, but the recovery is selective and achieved by accepting losses elsewhere in the system. Improvements are therefore "bought" through trade-offs between different aspects of ED performance.

Within this region, horizontal movement from left to right reflects increasing net performance gains: higher values of PC_M indicate that, despite the deteriorations, the overall balance of change becomes more favourable. Vertical movement from top to bottom reflects increasingly heavy trade-offs: as TO_C moves further below zero, a growing share of the weighted indicators deteriorates relative to those that improve. Near the top boundary of Region II (values of TO_C just below zero), improvements and deteriorations are closer in number or weight, so losses only slightly outweigh gains. Towards the lower boundary (more negative TO_C), the improvement becomes more narrowly concentrated in a small subset of indicators, while a larger share of the indicator set worsens.

Ideally, when a work adaptation falls in Region II rather than Region I, it should be located as far towards the upper-right corner of the region as possible, corresponding to substantial net improvement (PC_M large and positive) with as few and as limited trade-offs as possible (values of TO_C close to zero). By contrast, positions closer to the lower-left corner represent low net gains that rely on many or severe trade-offs.

Although Region II still mainly would corresponds to a positive recovery segment on the resilience curve (between t_2 and t_4 in Figure 3.4), the expression of resilience is constrained. Performance rises compared to the disrupted level, but this recovery is uneven: some parts of the ED benefit, while others deteriorate. Adaptations located in Region II are therefore interpreted as instances where resilience is expressed through trade-off-driven improvement rather than through net improvement with limited trade-offs.

Region III: $PC_M < 0, TO_C < 0$ --- net deterioration.

Performance worsens overall, and most indicators move in the undesired direction. In terms of the resilience curve, this region primarily would correspond to the absorption phase between t_1 and t_2 in Figure 3.4, where the disruption pushes the ED away from its desired performance level P_0. The further a point lies toward the lower-left corner of Figure 3.8, the deeper the deterioration: larger negative values of PC_M indicate a stronger drop in performance, while more negative values of TO_C indicate that this drop is driven by a larger share of the weighted indicators moving in the undesired direction.

Closer to the origin ($PC_M \approx 0$, $TO_C \approx 0$), the ED can be seen as more resilient in the sense that the overall deterioration remains limited. A small negative value of PC_M means that performance is drop is limited. Values of TO_C closer to zero indicate that, while more indicators move in the undesired direction, not all parts of the system deteriorate in the same way: some indicators may remain relatively stable or even show reduced pressure because the disruption shifts load or activity elsewhere. By contrast, moving further down and left within Region III means that performance drops more and more aspects of ED functioning are pulled into this downward movement.

In addition to characterizing the absorption phase, Region III can also represent a failed or counterproductive work adaptation when assessed after implementation: despite the intervention, performance remains worse than the pre-implementation level. In that situation, a point in Region III corresponds to an adaptation and recovery phase that continues to trend downward—or at best flattens at a reduced performance level—rather than reversing towards recovery on the resilience curve.

Region IV: $PC_M < 0, TO_C > 0$ --- trade-off-driven deterioration.

Performance worsens overall, but more (weighted) indicators improve than worsen. This pattern arises when a relatively small set of indicators deteriorates strongly, while a larger set shows only modest improvements. In other words, gains are outnumbered by a few substantial losses, so the net performance change remains negative.

Within this region of Figure 3.8, horizontal movement from right to left reflects increasing severity of net deterioration: more negative values of PC_M indicate that the deteriorations pull the average further down. Vertical movement from bottom to top reflects that deterioration becomes more concentrated and severe rather than more widespread: as TO_C increases (while remaining positive), more indicators show some improvement, so the remaining deteriorating indicators must worsen more strongly, or in more heavily weighted parts of the system, to keep PC_M below zero. Positions near the centre of the matrix represent mild net decline with several small improvements and several small deteriorations. By contrast, points closer to the upper-left corner (strongly negative PC_M and strongly positive TO_C) would imply that many indicators improve while a very small subset deteriorates extremely strongly. Such configurations are expected to be rare in practice, because the large number of improvements would normally push PC_M upward towards zero rather than allowing it to move further into the negative range.

If such a pattern does appear, it will most often correspond to a failed work adaptation or failed recovery attempt. The adaptation generates improvements in several aspects of performance, but the magnitude of deterioration in a small set of dimensions is strong enough to outweigh these gains. In other words, the trade-offs are unfavourable: small or moderate improvements in some areas are purchased at the cost of substantial deterioration elsewhere, so overall performance still declines.

Region IV: $PC_M < 0, TO_C > 0$ --- trade-off-driven deterioration.

Performance worsens overall, but more (weighted) indicators improve than worsen. This pattern can arise when a relatively small set of indicators deteriorates strongly, while a larger set shows only modest improvements. In other words, small gains outnumber few substantial losses, but the net performance change remains negative.

Within this region of Figure 3.8, horizontal movement from right to left reflects increasing severity of net deterioration: more negative values of PC_M indicate that the deteriorations pull the average further down. Vertical movement from bottom to top reflects that improvements become more dominant in terms of the number or weight of indicators: as TO_C increases (while remaining positive), a growing share of the weighted indicators shows some improvement, so the remaining deteriorating indicators must worsen more strongly, or in more heavily weighted parts of the system, to keep PC_M below zero. Positions near the centre of the matrix (moderately negative PC_M and TO_C just above zero) represent mild net decline with several small improvements and several small deteriorations. By contrast, points closer to the upper-left corner (strongly negative PC_M and strongly positive TO_C) would imply that many indicators improve slightly while a very small subset deteriorates extremely strongly. Such configurations are expected to be rare in practice, because the large number of improvements would normally push PC_M upward towards zero rather than allowing it to move further into the negative range.

If such a pattern does appear, it will most often correspond to a failed work adaptation or failed recovery attempt. The adaptation generates improvements in several aspects of performance, but the magnitude of deterioration in a small set of dimensions is strong enough to outweigh these gains. In other words, the trade-offs are unfavourable: small or moderate improvements in some areas are purchased at the cost of substantial deterioration elsewhere, so overall performance still declines.

Interpreting adaptations in this matrix helps to link the quantitative scores back to resilient performance. Movements horizontally across Figure 3.8 primarily represent changes in the net strength and direction of performance change (PC_M), while movements vertically modify this pattern by indicating how dominant improvements or deteriorations are across the weighted indicators (TO_C) — that is, how strongly the overall pattern is shaped by trade-offs. In principle, if PC_M and TO_C were computed repeatedly over time, the trajectory of a system across the matrix would trace how it moves along the resilience curve, from initial deterioration towards recovery and whether that recovery is achieved with limited or heavy trade-offs. Methodological implications of such time-varying trajectories are further discussed in the Discussion, as the current cases can only compute one PC_M and one TO_C value (with PC_B reported alongside for absolute net breadth).

3.3.3. Donabedian Classification

As noted earlier in Section 2.3.1, the Donabedian structure–process–outcome model adds stage-specific context for each PI: structure (referring to the resources, equipment, and organizational infrastructure in place), process (representing the actions taken during patient care delivery), and outcomes (reflecting the end results of care in terms of health improvement or deterioration).

For analysing change in performance, process and outcome indicators would be more informative because they are the indicators that move when day-to-day operations or care results change (e.g. waiting times, LOS, mortality). Structural indicators (such as bed numbers, or fixed staffing establishment) usually remain stable over the time. As a result, we expect fewer structural indicators to show up in a performance-change analysis. However, if a disruption or work adaptation clearly alters structural features, those structural indicators should then also be included, because they do reflect a change in the system's performance capacity.

4

ED Workflow and Critical Function

4.1. ED Workflow

Using the conceptual models of Claassen et al. (2025) and Asplin et al. (2003), along with the findings from prior sections, Figure 4.1 presents a customized ED workflow framework adapted to the Dutch acute care context. The aim of this adaptation is to build on internationally recognized models while tailoring the workflow to Dutch referral structures and data exchange practices.

In this adaptation, the input phase from Asplin corresponds to the referral phase in Claassen's model. It is represented here by the main arrival pathways into Dutch EDs: self-referral, GP referral, and ambulance arrival. These routes reflect how demand for ED services is generated, but the operational responsibility for referral decisions and patient transport lies primarily with GPs and EMS. For this reason, the input/referral domain is included in the framework for completeness but not a central focus of ED performance assessment.

Between input and throughput, this study introduces a Pre-ED phase. While not explicitly present in either source model, this phase reflects the practical steps of patient and data transfer prior to ED entry—including registration, handover of medical records, and digital data transfer from GPs or EMS. Acknowledging this phase makes the model more representative of actual patient flow.

The throughput phase is defined by Asplin et al. (2003) and is divided into two sub-phases. Phase 1 includes triage, room placement, and the initial provider evaluation, which determine the urgency and trajectory of care. Phase 2 involves diagnostic testing and ED-based treatment. Asplin et al. (2003) note that the second phase of throughput is typically the longest and most resource-intensive. Both phases are included in the adapted workflow model as they are described in the original framework. Additionally, the inclusion of inpatient bed transfer reflects a key bottleneck in ED operations: the inability to promptly move admitted patients to hospital wards. This delay ties up ED treatment spaces and consumes nursing and physician resources, reducing the department's overall capacity to manage new patients (Asplin et al., 2003). However, this bottleneck is not included in either Phase 1 or Phase 2 of the original model. Given its heavy relevance for ED throughput, it is added in this study's workflow model as a separate box outside Phase 1 and Phase 2, but still within the throughput domain.

Finally, the output phase of Asplin corresponds with Claassen's post-ED phase. This includes

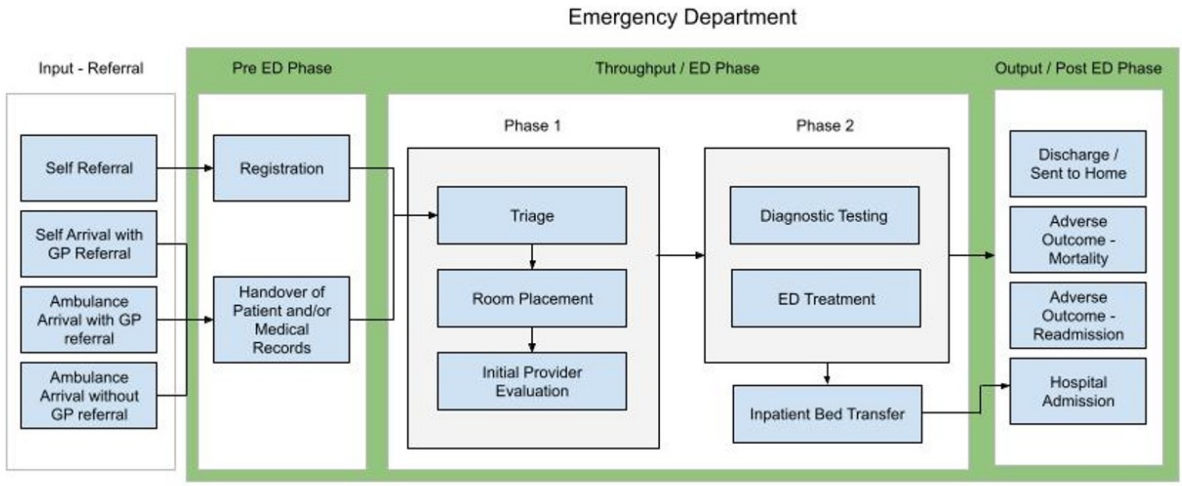

Figure 4.1: Acute Care Chain workflow focusing Emergency Departments Role

discharge home, hospital admission, or transfer to another facility, along with adverse post-discharge outcomes such as readmission or 30-day mortality. While simpler in structure, this phase captures essential patient outcomes and connects ED processes to the broader care chain.

By combining Asplin's generic input–throughput–output model with Claassen's Dutch ACC perspective, this framework situates the ED within its national context while maintaining international comparability. The adapted model highlights throughput and output as the core ED domains of responsibility, while also accounting for referral pathways and data transfer processes that shape ED demand without being direct ED functions.

4.1.1. Focus of ED in Dutch Context

In the Dutch healthcare system, the input phase is not primarily managed by the ED. Patients with non-urgent complaints are expected to first consult their General Practitioner (GP), who acts as a gatekeeper and decides whether referral to hospital care is required. In acute cases, Emergency Medical Services (EMS) facilitate patient transport and support the referral process through registration, physical handover, and digital transfer of clinical data prior to ED entry. These activities are strongly supported by the Netherlands' highly developed health IT infrastructure, which enables standardized and mostly automated data exchange between GPs, EMS, and hospitals (Nederlands Huisartsen Genootschap (NHG), 2022). As a result, much of the patient registration and administrative workload associated with ED entry is effectively streamlined outside the ED itself.

For this reason, this study excludes Pre-ED activities from its performance assessment. Including them would risk conflating ED performance with processes controlled by GPs, EMS, or IT systems, rather than by the ED. Consequently, indicators measuring performance in this domain—such as ambulance arrival times—are not considered in the analysis. The focus instead is placed on the throughput and output phases, where EDs have direct control over processes and can demonstrate adaptive capacity under disruption.

This adapted model in Figure 4.1 will guide the selection of core tasks within the ED and help define its critical function. By focusing on the Throughput and Output phases, the model will support a structured, context-aware approach to evaluating ED performance under disruption and mapping that performance to resilience capacities.

4.2. Emergency Department Critical Function

While all the phases are essential to the success and overall functioning of Emergency Departments, the ED's critical function lies in the throughput phase—where unique, high-pressure care is delivered that cannot be postponed, rerouted, or replicated elsewhere in the health-care system. It is here that the ED fulfills its distinct role: **rapidly assessing, correctly diagnosing, and initiating treatment for acute conditions under conditions of uncertainty and urgency.** This function involves concentrated diagnostic work, real-time clinical decision-making, and multidisciplinary coordination—all under the constraints of time, resource availability, and even incomplete information. These features define the ED as a site of irreplaceable medical intervention, and make the throughput phase the core of its performance.

However, the ED's ability to carry out this function is directly shaped by its **ability to manage capacity and flow through the department.** Although flow management is not a clinical task in itself, it is integral to the ED's critical function because it enables the other three elements—rapid assessment, correct diagnosis, and initiation of treatment—to operate and connect reliably. If admitted patients cannot be moved promptly to hospital wards, treatment spaces and staff resources remain blocked. In such cases, the ED's ability to rapid assessment, correct diagnosis, and ultimately initiating treatment for new acute patients is impaired.

Furthermore, this definition is enacted step-by-step within the throughput phase of our workflow in Figure 4.1. **Rapidly assessing** is operationalized in Phase 1 by *triage* and *initial provider evaluation* (quick risk sorting → focused clinical assessment), and because some interventions (e.g., resuscitation and stabilization) must occur immediately—sometimes even before formal triage—this function also encompasses zero-wait actions at initial contact. **Correct diagnosis** and **initiating treatment** occur in Phase 2. In this phase, *diagnostic testing* provides the actionable evidence to refine or establish a working diagnosis and guide care. *ED treatment* then delivers the corresponding ED-based interventions. **Managing flow** is the enabling thread throughout the whole ED Workflow: Before Phase 1 (allocating space and sequencing arrivals), within Phase 1 (room placement), between phases (handoffs and task progression), and after Phase 2 (inpatient bed transfer to relieve bottlenecks and sustain capacity).

This explicit mapping shows the critical function squarely captures the ED's distinctive role within throughput, as identified in our framework and the literature. While the output boxes (discharge, admission, readmission, mortality) are not themselves elements of the critical function, they remain the primary reflections of ED performance because they are tightly driven by how effectively these throughput steps and flow management are executed.

Reflection on IOM Dimensions

Building on this definition and the mapping, different parts of the ED's critical function align with different IOM dimensions of quality. **Rapidly assessing** speaks to **timeliness** under conditions of uncertainty and urgency. It is tightly linked to **efficiency**: if ED resources (including personnel) are not used efficiently, timeliness suffers. Practically, this translates to **managing capacity and flow** through the department, where **timeliness and efficiency** are emphasized.

Correctly diagnosing underscores **effectiveness**, both of the diagnostic work and of the subsequent ED treatment that is initiated for acute conditions. The impact of **correct diagnosis**

and effective care becomes visible in the output/post-ED phase, where patients are **safely** discharged home or admitted to the hospital for the care they need.

When this intended flow is not achieved, adverse outcomes such as re-admissions or mortality can occur. These outcomes do not arise solely from failures within the ED; patient factors and disease progression also play roles. However, when adverse outcomes are attributable to the ED's inability to rapidly assess, correctly diagnose, and initiate treatment—or to manage flow—they reflect ineffective care. In IOM terms, this is also a failure of **safety**—avoiding harm to patients from the care intended to help them—because **ineffective ED care increases the risk that the care process itself does not prevent, or even contributes to, harm.**

In short: Timely, efficient ED (rapidly assessing) → Effective ED (correctly diagnosing and initiating adequate treatment) → Safe ED (minimizing adverse outcomes).

Underrepresented IOM Dimensions

Overcrowding is the central challenge that disrupts ED performance, and it can be most directly assessed through measures of timeliness and efficiency and its outcomes are observed in effectiveness and safety. The critical issue under disruptive conditions therefore lies in sustaining patient flow and preventing bottlenecks. The literature reinforces this point: Asplin et al. (2003) highlight that "key factors influencing throughput efficiency include the cohesiveness of care teams, ED layout, staff-to-patient ratios, efficiency of diagnostic testing (e.g., radiology, lab), accessibility of medical records, quality of documentation systems, and availability of specialty consultations." Asplin's framework provides a strong illustration of how overcrowding impacts ED operations, though other studies also emphasize that managing overcrowding is a critical determinant of resilience.

Equality: ensuring care quality does not vary based on personal characteristics

The Dutch healthcare system already achieves a high degree of equity in access and quality, which makes the inclusion of equity less relevant for ED performance. Strong evidence demonstrates that access to care in the Netherlands does not vary systematically by personal characteristics. The State of Health in the EU: The Netherlands Country Health Profile 2023 (OECD/European Observatory on Health Systems and Policies, 2023) highlights that unmet medical needs in the Netherlands are exceptionally low—only 0.2% of the population reported unmet needs in 2022, the lowest rate in the European Union and identical to pre-pandemic levels. Crucially, there is virtually no difference in unmet need between the top and bottom 20% income groups, indicating equitable access across socioeconomic strata. This outcome is supported by the system's below-average reliance on out-of-pocket payments, universal and mandatory insurance coverage, and legal requirements obliging insurers to accept all applicants regardless of pre-existing conditions. With 99.9% of the population insured and 85% of health expenditure covered by government and compulsory insurance schemes, structural inequities in access are minimized. Given this context, and no representation on the critical function, equity is not a pressing determinant of ED performance in the Netherlands.

Patient-Centered: aligning care with patient preferences and values

Patient-centered care primarily rely on patient surveys and are inherently subjective. Evidence also suggests that patient satisfaction is not a reliable proxy for the quality of emergency care. Fenton et al. (2012) found that higher patient satisfaction was associated with increased healthcare utilization, higher expenditures (including drug prescriptions), and even higher mortality rates. Similarly, Jerant et al. (2019) confirmed that patients reporting higher satisfaction with clinicians exhibited an elevated risk of short-term mortality, even after case-mix adjustment.

These findings highlight that patient satisfaction measures do not necessarily translate into better healthcare outcomes. In the context of emergency medicine, patients ultimately value competent, effective, and lifesaving treatment over positive interpersonal experiences, making patient-centeredness a weak basis for evaluating the ED's performance and again does not have any direct ties to the critical function.

Chapter 4 clarified the ED workflow and articulated the critical function that anchors the rest of the thesis. By specifying the operational backbone of the framework—the sequence of steps and the core capability required to **rapidly assess, correctly diagnose, and initiate treatment**, together with the enabling role of **managing flow**—we defined how this part of the system should perform. This gives us a concrete reference against which to judge work adaptations. With this foundation in place, the next chapter turns to case analyses. These cases provide the empirical evidence to test how each adaptation aligns with the expected performance of the critical function and its supporting workflow.

PART III

Application & Results

5

Work Adaptations on ED Performance: What Changed and How

This chapter operationalises Phase 1 – Structured performance change description of the PI-RA framework (Figure 5.1). It assembles the empirical input for the later phases by analysing international empirical studies that describe organisational or process-level work adaptations in emergency departments and report associated changes in quantitative performance indicators (PIs) before and after implementation. These cases provide detailed accounts of how specific adaptations were introduced and how selected indicators responded over time, allowing a grounded analysis of change in ED performance. The trade-off is that each case reports only a limited set of PIs, which narrows the breadth of performance insights.

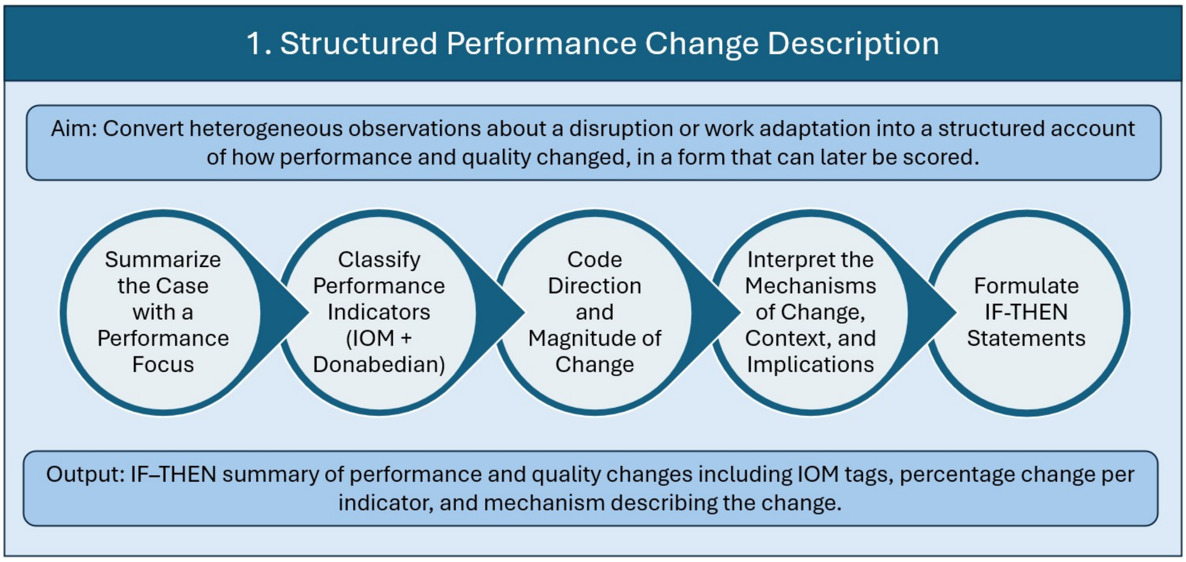

Figure 5.1: Phase 1 of PI-RA: Structured Performance Change Description

For each work adaptation, the Phase-1 steps shown in the framework are applied in sequence. The case is summarised with a performance focus, the reported indicators are classified using the Donabedian and IOM schemes from Chapter 3, the direction and magnitude of each PI change are coded, and the observed pattern is interpreted in light of the adaptation's mechanism and context. These elements are then condensed into a single IF–THEN statement per work adaptation that lists which indicators changed, by how much (magnitude bin and direction), and what these changes imply for different IOM quality dimensions. The set of IF–THEN statements produced in this chapter forms the structured performance-change input for the system-level analysis in the subsequent chapters.

5.1. Enhanced Isolation Protocol for Fever or Respiratory Patients (FRPs)

This work adaptation and the performance indicator data are drawn from Kim et al. (2022), who studied changes in ED operations at a tertiary academic hospital in Seoul, South Korea during the early COVID-19 pandemic. The study analyzed emergency department visits before (March to July 2019) and after (March to July 2020) the COVID-19 pandemic. The hospital introduced an expanded isolation protocol requiring all patients with fever (≥37.5 °C) or respiratory symptoms (FRPs) to be streamed into an isolation zone before entering a medical care area. The intent was to protect staff and other patients from contagion. In this isolation zone only basic diagnostics (physical exam, plain CXR, labs, ECG) were available, while advanced imaging (CT, MRI, echo) was deferred until release. The effects of this work adaptation on performance indicators summarized in Table 5.1.

Table 5.1: Summary of Performance Indicators from Kim et al. (2022)

PI	Change / Evidence	Direction & Magnitude
ED BOR	66.2% → 46.5% (-29.8%); $p < 0.001$	Decreases (High)
ED LOS	7.06 hrs → 6.64 hrs (–5.9%); $p < 0.001$	Decreases (Medium)
WT for FRPs	24.0 min → 45.0 min (+87.5%); $p < 0.001$	Increases (Very High)
LWBS for FRPs	2.8% → 19.2% (+585.7%); $p < 0.001$	Increases (Very High)
WT for non-FRPs	23.0 min → 26.0 min (+13.0%); $p < 0.001$	Increases (Medium)
LWBS for non-FRPs	5.7% → 10.0% (+75.4%); $p < 0.001$	Increases (Very High)
LWBS (overall)	5.1% → 12.1% (+137%); $p < 0.001$	Increases (Very High)
Mortality Rate	0.4% → 0.4%	No Change
Discharged out of ED	69.7% → 61.6% (–11.6%); $p < 0.001$	Decreases (Medium)
Transfer to Another ED	3.8% → 2.7% (-28.9%); $p < 0.001$	Decreases (High)
Admitted to Inpatient Care	21.0% → 23.2% (+10.5%); $p < 0.001$	Increases (Medium)

Below is the classification of the performance indicators from Kim et al. (2022) using the definitions of IOM dimensions and Donabedian categories as provided in Section 2.3.1, Foundations of Healthcare Quality Measurement.

- **ED Bed Occupancy Rate (ED BOR) [Efficiency | Process]**: percentage of ED beds occupied at a given time; beds are resources and this reflects resource use (Efficiency);

reflects conditions created by the placement actions during care delivery (Process).

- **ED Length of Stay (ED LOS) [Timeliness | Process]**: time from arrival/triage to ED departure; captures the time in completing the visit (Timeliness); measures the duration of actions during care delivery across the ED pathway (Process).
- **Left Without Beeing Seen (LWBS)[Safety | Outcome]:**
- **Waiting Time (WT) [Timeliness | Process]**: time from triage to entry into a treatment area / initial provider evaluation; measures access delay before care proceeds / during room placement (Timeliness); is a queuing step within the actions during care delivery (Process).
- **Mortality Rate [Safety | Outcome]**: proportion of patients who die; reflects harm from/after care (Safety); represents the end result of care in terms of health deterioration (Outcome).
- **Discharged out of ED; Transfer to another ED; Admitted to Inpatient Care [Efficiency | Outcome]**: final disposition at ED exit; indicates throughput and bed-allocation consequences reflecting use of resources (Efficiency); records the end result of the care episode (Outcome).

Next, the case analysis turns to mechanisms and consequences.

Isolation workflow constraints and mechanism of delay

The isolation zone allowed only basic tests; advanced imaging was postponed, so isolated patients could not progress diagnostically until release. This design increases front-end friction and extends the time to clinical action for FRPs, independent of whole-ED load. Waiting times rose substantially for FRPs (24 min → 45 min, +87.5%), reflecting queuing delays at the isolation entry point where capacity and diagnostics were constrained.

Local bottleneck evidenced by divergent signals

While ED LOS (7.06 → 6.64 h(-0.42 h)) and BOR (66.2% → 46.5% (−29.8%)) fell which indicates a system relief, FRP LWBS rose by 586% (2.8% → 19.2%). This divergence is a signature of a bottleneck localized to the isolation queue, not generalized crowding. The authors explicitly interpret the finding as insufficient isolation capacity under broader eligibility (for isolation).

Proportion shifts and demand context

Although the share of FRPs among arrivals decreased post-COVID, LWBS among FRPs surged. That pattern further supports a selection/throughput problem at isolation rather than a volume surge.

Disproportionate impact and who is affected

The multivariable analysis shows that FRP status (aOR 1.76) and the post-COVID period (aOR 2.29) each independently increased the odds of patients LWBS. Here, aOR stands for adjusted odds ratio: an aOR of 1.76 means that, after adjusting for other factors, FRP patients have 76% higher odds to LWBS than non-FRPs; an aOR of 2.29 means that patients seen in the post-COVID period have about 129% higher odds to LWBS compared to the pre-COVID period.

Crucially, the analysis also found a significant interaction (FRP×post), meaning they enhanced each other's impact. In other words, FRPs in the post-COVID period were at a disproportionately higher risk of LWBS than would be expected from the effect of FRP status or the post-

COVID period alone. This makes clear that the new isolation protocol amplified an existing tendency of FRPs (often lower-acuity, mild-symptom patients) to leave early, turning it into a much more severe problem after the policy change.

The problems associated with higher LWBS are: Higher chance of re-visitation, since patients' problems are often left unresolved; LWBS who require isolation (i.e., those with a higher likelihood of infection) risks the spread of infection.

Case Mix and Disposition Patterns

Admissions increased (21.0% → 23.2%), home discharges decreased (69.7% → 61.6%), transfers decreased (3.8% → 2.7%); mortality was unchanged (0.4%). These changes reflect altered case mix and routing, instead of consequences of the ED processes.

System-boundary recommendations

To mitigate public-health risk from infectious LWBS, the authors suggest tracking/supervision systems for FRPs who leave and parallel, rapid-access fever services/24-h clinics — pointing to design options outside the ED boundary that complement isolation.

Safety Gains due to Enhanced Isolation Protocol

Kim et al. (2022) mainly analyzed the adverse effects of the enhanced isolation protocols but did not address the safety benefits that come from reduced contagion risks. To provide a more complete picture, it is important to acknowledge these safety benefits, even though they are not captured by any PI.

Unintended Improvements

Although the protocol was introduced to improve safety, LWBS increased after implementation. That change contributed to lower ED BOR and ED LOS, which had impact on efficiency and timeliness gains.

Building on the mapping and these insights, the findings are distilled into a set of IF–THEN statement that capture the operational logic of the adaptation and its measurable consequences for ED performance.

IF *patients with fever or respiratory symptoms were streamed to a capacity-limited isolation zone with restricted diagnostics,* ***THEN:***

- **ED Bed Occupancy Rate** decreased by 29.8% (High), indicating a improvement in **Efficiency**.
- **ED Length of Stay** decreased by 5.9% (Medium), indicating improvement in **Timeliness** (end-to-end).
- **Waiting Time** increased for FRPs by 87.5% (Very High) and for non-FRPs by 13% (Medium), indicating deterioration in **Timeliness** (front-end).
- **LWBS** increased by 137% (Very High), indicating deterioration in **Safety**.
- **Disparity in WT and LWBS between FRPs and non-FRPs** is 573.1% (Very High) for WT (FRPs +87.5% vs non-FRPs +13.0%) and is 676.8% (Very High) for LWBS (FRPs +585.7% vs non-FRPs +75.4%), indicating deterioration in **Equity**.
- **Infection contagion reduced**: It is the intended **safety** benefit however is not measured through any PI.
- **Admissions, Home Discharges, Transfers, Mortality:** Changes in these indicators reflect the altered case mix, not a performance consequence of the work adaptation.

Isolating all FRPs reduced overall ED occupancy and LOS (Efficiency↑) but created a capacity-limited isolation queue that delayed access (WT↑) and increased early departures (LWBS↑) especially among low-acuity FRPs (Timeliness↓, Safety/Equity↓). The evidence points to a local bottleneck induced by broader eligibility plus constrained isolation capacity and restricted diagnostics.

5.2. Rapid Assessment Zone (RAZ)

The study by Faber et al. (2023) was conducted at the Emergency Department (ED) of Mercy Health–Fairfield Hospital in Fairfield, Ohio. The work adaptation, which was the implementation of a Rapid Assessment Zone (RAZ), took place on 1 February 2021. To evaluate its impact, the study examined data from the six months prior to this date (August 2020 to January 2021) and compared them with the six months after implementation (February 2021 to July 2021). Although the study does not explicitly frame the RAZ as a COVID-19 intervention, but rather as a response to longstanding challenges "the underlying problems the RAZ was designed to solve—ED crowding and high rates of patients leaving without being seen (LWBS)—were noted by the hospital back in 2019, before the pandemic", it remains highly relevant to this analysis. Additionally, the authors anecdotally noted frequent ED boarding as a back-end constraint which is a recognized contributor to ED crowding. This is because the period of implementation coincides with the third COVID-19 wave in the USA. Moreover, the study does refer to the pandemic as an important contextual factor shaping ED operations. Finally, it represents a concrete ED work adaptation undertaken during the COVID-19 period and evaluated in terms of its impact on performance indicators.

The adaptation involved the creation of a Rapid Assessment Zone through the adoption of a vertical care model. This redesign targeted the ED's front-end flow. An existing eight-bay treatment and triage zone was repurposed into eight rapid assessment rooms. At the same time, the roles of existing ED staff—nurses, providers, and technicians—were shifted and redefined to form a front-end team, without the need for additional hires. This team focused on rapidly evaluating and treating patients, particularly those of lower acuity, who could be managed within the vertical care model. Vertical care was defined as "evaluating and treating patients without the use of a physical emergency department room" and assigning them to "virtual beds." In practice, this allowed the ED to preserve monitored beds for higher-acuity patients and reduce downstream congestion, without increasing staff numbers or extending provider hours. In addition to staff reallocation, new patient flow processes were created to expedite assessment, testing, and treatment, with the explicit aim of initiating a patient's workup within twenty minutes of arrival.

The changes on performance indicators after the work adaptation implementation can be seen in Table 5.2.

Differing from the FRP case, Faber et al. (2023) report one additional indicator, categorized and defined as:

- **Arrival-to-provider [Timeliness | Process]**: Time from ED arrival to initial provider evaluation; captures the immediate front-end access delay before care proceeds (Timeliness);is a queuing step within the actions during care delivery (Process) *(Conceptually similar to Waiting Time in the FRP case, but starts at arrival rather than post-triage.)*

Now, we turn to the mechanisms and consequences of the RAZ implementation.

Table 5.2: Summary of Performance Indicators from Faber et al. (2023)

PI	Change / Evidence	Direction & Magnitude
ED LOS (Admitted Patients)	395 min $\rightarrow$ 332 min (-15.9%); $p < .01$	Decreases (Medium)
ED LOS (Discharged Patients)	205 min $\rightarrow$ 163 min (-20.5%); $p < .01$	Decreases (High)
Arrival-to-Provider	28 min $\rightarrow$ 11 min (-60.7%); $p < .01$	Decreases (High)
LWBS	5.64% $\rightarrow$ 2.55% (-54.8%); $p < .01$	Decreases (High)
Admitted to Inpatient Case	26.4% $\rightarrow$ 25.13% (-4.8%); $p < .01$	Decreases (Low)
Discharged out of ED	64.47% $\rightarrow$ 69.06% (+7.1%); $p < .01$	Increases (Low)

Timeliness gains despite downstream constraints

The separation of front-end processes mitigated crowding effects even under conditions of inpatient boarding pressure. The study site acknowledged that inpatient boarding was common during the study period; nevertheless, both admitted and discharged length of stay decreased significantly after the implementation of the RAZ (from 395 min to 332 min & 205 min to 163 min respectively). This indicates that front-end interventions are capable of improving throughput and timeliness even when back-end constraints exist.

Timeliness   Safety: faster assessment halved LWBS

Timeliness also directly enhanced safety through a reduction in LWBS by 54.8%. The study explicitly attributes this change to the sharp reduction in physician initial assessment times. The relationship is clear: faster front-end access led to fewer patients leaving without treatment, thereby reducing safety risks.

Multiple Factors at Play

While the RAZ contributed to shorter ED LOS, there were other factors at play. In the post period, overall ED attendance was higher, which would normally push LOS upward. The case mix also shifted toward lower acuity, with fewer ESI 2 patients and more ESI 3–4 patients, which would be expected to shorten processing time. The observed reduction in LOS therefore reflects multiple forces: increased demand that tends to lengthen stays, a less severe patient mix that tends to shorten them, and the adaptation operating alongside these factors.

Sustainment indicates systemic embedding

The time-series plots presented in the study showed that improvements were immediate at the time of implementation and remained stable throughout the six-month follow-up period. This indicates that the RAZ functioned not as a temporary “shock absorber” during surge conditions but as a durable and embedded workflow change.

Process governance as a key success factor

The implementation was supported by daily and weekly debriefs, on-shift coaching, and scripted communication for staff and patients. These governance mechanisms contributed significantly to the maintenance of improvements and should be recognized as important drivers of the intervention's success.

Taken together, these observations point toward a set of operational patterns that can be formalized as IF–THEN statement, expressing how the Rapid Assessment Zone translated into concrete changes in emergency department performance:

IF the ED implemented a Rapid Assessment Zone (RAZ) with vertical flow and early provider evaluation, **THEN:**

- **Arrival-to-provider** decreased by 60.7% (High), indicating improvement in **Timeliness** (front-end).
- **ED LOS** decreased for *discharged* patients by 20.5% (High) and for *admitted* patients by 15.9% (Medium), indicating improvement in **Timeliness** (end-to-end).
- **LWBS** decreased by 54.8% (High), indicating improvement in **Safety**.
- **Admissions, Discharges:** Changes in these indicators reflect the altered case mix, not a performance consequence of the work adaptation.

RAZ implementation was associated with faster arrival-to-provider (Timeliness↑), shorter ED LOS for discharged and admitted patients (Timeliness↑), and lower LWBS (Safety↑). The evidence points to relief of a front-end intake bottleneck through vertical flow and early provider evaluation

5.3. Point-of-Care Testing (POCT)

Baron et al. (2022) introduced bedside molecular point-of-care testing (POCT) for COVID-19 in the Emergency Department of Saint-Louis Hospital, Paris, during France's second COVID-19 wave (18 October–30 November 2020). Before this change, nasopharyngeal samples were sent to the laboratory, which added hours of delay and aggravated bottlenecks. With the new ID NOW COVID-19 test, results became available in roughly 5–13 minutes and could be run directly by ED nurses and physicians in the department. To evaluate the effect, they compared two consecutive periods: a pre-implementation window (18 Oct–3 Nov) when testing remained lab-based, and an implementation window (4–30 Nov) when rapid POCT was used in the ED. The study's aim was to see whether faster confirmation could shorten ED length of stay, help prevent overcrowding, speed bed allocation, and ultimately improve patient management.

Similarly, Mortazavi et al. (2022) ran a retrospective cohort at Skåne University Hospital in Sweden during the same pandemic wave, bringing diagnostic testing into the ED in two steps: first adding rapid antigen tests, then adding point-of-care rapid RT-PCR (VitaPCR) so results were available at the bedside within minutes rather than after a 12–24 h core-lab delay. This makes the case directly comparable to Baron because it implements the same underlying work adaptation—moving diagnostic testing into the ED to shorten turnaround and support earlier disposition decisions. In Baron's paper, faster confirmation is primarily evaluated through its effect on crowding, but they also argue that it should enable safer, more appropriate placement, although that aspect is not directly measured. Mortazavi's study observes this "safer, more appropriate placement" by tracking the immediate placement decisions after ED—where admitted patients go on first placement—and whether an early intrahospital transfer (≤5 days) is needed to correct that decision. They also examine diagnostic coverage by tracking the proportion of patients discharged home who were classified as "Not tested", which decreases markedly across the study periods. Taken together, the two studies let us treat POCT as one coherent case: Baron supplies the front-end timeliness mechanism, while Mortazavi provides evidence on how this mechanism translates into placement, transfer patterns, and diagnostic coverage. The consolidated performance indicators from both studies are summarized in Table 5.3

Table 5.3: Performance Indicators for POCT Implementation

Author & Date	Performance Indicator	Reported Change (numerical / statistical)	Direction
Baron et al., 2022	ED length of stay (LOS)	Median LOS: 276 min $\rightarrow$ 208 min (-24.6%); $p < 0.0001$	Decrease (High)
Baron et al., 2022	Patients discharged <4h	38.3% $\rightarrow$ 61.3% (60.1%); $p < 0.0001$	Increase (High)
Baron et al., 2022	Time to result	Median: 261 min $\rightarrow$ 112 min (-57.1%); $p < 0.0001$	Decrease (High)
Baron et al., 2022	Patients per day	81 $\rightarrow$ 61 (-24.7%); p = 0.0001	Decrease (High)
Baron et al., 2022	Hourly occupancy rate	15.5 $\rightarrow$ 11.0 (-29%); p = 0.0004	Decrease (High)
Mortazavi et al., 2022	Diagnostic Coverage (Discharged Home classified as "Not tested")	26.2% $\rightarrow$ 8.1% (-69.1%)	Decrease (High)
Mortavazi et al., 2022	Targeted Admissions (ED-Neg $\rightarrow$ non-COVID ward)	32.9% $\rightarrow$ 59.8% (+81%)	Increase (Very High)
Mortazavi et al., 2022	Intrahospital Transfers	50% $\rightarrow$ 34% (-32%) $p < 0.0001$	Decrease (High)

Below is the classification of the performance indicators from Baron et al. (2022) and Mortazavi et al. (2022) that were not defined in the previous two cases:

- **Patients Discharged within 4 Hours (ED LOS < 4h) [Timeliness | Outcome]**: proportion of ED patients whose visit from registration to final disposition decision is completed within four hours; reflects success in meeting the NHS four-hour benchmark and reducing delays and crowding in the ED (Timeliness); records whether the ED visit is finished within the benchmark window as the end result of care (Outcome).
- **Time to Result [Timeliness | Process]**: interval between ED registration and release

of the SARS-CoV-2 test result; measures diagnostic turnaround time and delay before definitive information is available to guide disposition decisions (Timeliness); represents a diagnostic step within the actions during care delivery (Process).

- **Patients per Day [Efficiency | Process]**: number of patients seen in the ED per day; can give a basic view of throughput and workload in relation to available capacity (Efficiency); counts the volume of patients moving through the actions during care delivery (Process).
- **Hourly Occupancy Rate [Efficiency | Process]**: median number of patients present in the ED per hour; reflects how intensively ED space, staff, and other resources are used relative to demand (Efficiency); summarizes the concurrent patient load within the actions during care delivery (Process).
- **Diagnostic Coverage (Discharged home classified as "Not tested") [Effectiveness | Outcome]**: proportion of patients discharged home who leave the ED without a recorded SARS-CoV-2 test result; when this proportion decreases, more discharges are backed by a documented positive or negative result instead of symptoms alone, indicating more evidence-based and appropriate disposition and infection-control decisions (Effectiveness), with secondary benefits for Safety by reducing the likelihood of unrecognized infections leaving the ED or being sent to inappropriate wards; records whether discharge decisions are supported by confirmed diagnostic information as the end result of care (Outcome).
- **Targeted Admissions (ED-Neg → non-COVID ward) [Effectiveness | Outcome]**: proportion of ED-tested patients with a negative SARS-CoV-2 result who are admitted directly to an appropriate non-COVID ward on first placement; reflects how well rapid test results are used to match patients to the correct ward for their underlying condition and to avoid unnecessary placement in COVID isolation wards (Effectiveness), with secondary benefits for Safety by lowering avoidable exposure to COVID wards and reducing the need for later corrective transfers; records whether initial admission decisions place patients in an appropriate ward as the end result of care (Outcome).
- **Intrahospital Transfers [Safety | Outcome]**: movement of patients between hospital wards within the first five days after admission; from a hospital-wide perspective, high transfer rates mostly reflect inefficient use of beds and staff time, as resources are spent relocating patients instead of delivering care (Efficiency). However, in this ED-focused analysis, the indicator is used primarily as a safety-relevant outcome (Safety), because unnecessary transfers can disrupt continuity of care and expose patients to additional handovers and procedural risks (e.g. falls, medication errors, hospital-acquired infections), capturing downstream consequences of initial ED placement decisions as the end result of care (Outcome).

Faster front-end diagnosis drives end-to-end flow

Baron et al. (2022) report that introducing POCT in the ED sharply reduced time to result (from 261 to 112 minutes), so diagnostic confirmation happened much earlier in the ED stay instead of after a long central-lab delay. With results available sooner, teams could reach admission and discharge decisions earlier, rather than waiting on external laboratory processing. Consistent with this, median ED LOS fell from 276 to 208 minutes (−24.6%).

Meeting the 4-hour benchmark

Baron et al. (2022) frame this as a recognized quality and performance indicator (citing NHS benchmarks). The increase in patients leaving under 4 hours (ED LOS < 4h or patients discharged < 4h) is explicitly linked to reducing overcrowding and improving patient-management efficiency. Conceptually, it tracks the ED LOS result: if overall LOS shifts down, more patients naturally clear the 4-hour threshold. This improvement is also plausibly enabled by the shorter time-to-result with POCT, which accelerates disposition decisions.

Reduced crowding

Baron et al. (2022) also report that crowding eased during the POCT period. Alongside the shorter ED LOS and faster time to result, the median hourly occupancy rate (median number of patients present in the ED per hour) fell from 15.5 to 11.0. This pattern is consistent with the flow changes described above: when diagnostic confirmation happens earlier and patients spend less time in the department, each patient occupies ED space for a shorter period, which in turn lowers occupancy. At the same time, the study period with POCT also had fewer patients per day (81 to 61), so reduced inflow likely contributed as well.

Contextual enrichment of common indicators

In addition to the indicators reported by Baron et al. (2022), Mortazavi et al. (2022) use the disposition information to derive two further, performance-relevant measures: diagnostic coverage and targeted admissions linked to POCT. On their own, variables such as “discharged home” or “admitted to hospital ward” do not say much about ED performance. When they are viewed together with the ED test result (positive or negative) or which ward patients are admitted (COVID or non-COVID wards) and compared across periods, however, they become much more informative: we can see how many patients leave the ED without being tested at all, and how often admitted patients are sent directly to an appropriate ward. In this way, common disposition indicators are contextually enriched to reveal effects of the implemented work adaptation.

The following two sections use this contextual enrichment to describe higher diagnostic coverage and enhanced targeted admissions. A more detailed analysis of these patterns is provided in Appendix C.

Higher Diagnostic Coverage

Mortazavi et al. (2022) report that the proportion of patients discharged home who were classified as “Not tested” dropped sharply, from 26.2% to 8.1%, a total reduction of 69.1%. Over time, this means that many more discharged patients left the ED with a documented positive or negative result rather than no test at all. In other words, the introduction of rapid testing increased diagnostic coverage at discharge: fewer patients were discharged without having been tested, more decisions were backed by a confirmed result instead of symptoms alone, and overall confidence in disposition decisions are increased.

Mortazavi et al. (2022) note that overall ratios between discharge and admission remained stable across the study periods, indicating that these patterns are unlikely to be explained by changes in patient volume, case mix, or pandemic phase alone, but are closely linked to the implemented work adaptation. There are additional indicators in their study that point to the same mechanism of increased coverage, but for this POCT case we include only this most immediately observable effect, the decline in untested discharges as a signal of higher diagnostic coverage, for performance assessment.

Enhanced Diagnostic Coverage Enabled Targeted Admissions

Mortazavi et al. (2022) show that higher diagnostic coverage does not only change who gets tested; it also improves how admitted patients are placed in the hospital. As more patients received a definitive diagnosis in the ED, the hospital could make quicker and more targeted decisions about which ward they should go to, strengthening the effectiveness and efficiency of patient disposition.

Before rapid tests were introduced, patients with suspected COVID-19 were often admitted to dedicated COVID-19 wards while they waited for delayed core-laboratory RT-PCR results, regardless of their eventual diagnosis. This pre-POCT situation is similar to the condition described in Baron et al., where long laboratory turnaround contributed to longer ED stays and pressure on isolation capacity. After the new algorithms were implemented, participants who received a negative test at the ED were much less likely to be placed in a COVID-19 ward on first admission, dropping from 34.5% of admissions to 14.7%. Instead, these test-negative patients were far more likely to be admitted directly to an appropriate non-COVID hospital ward ("Other"), increasing from 32.9% to 59.8%. This pattern indicates that the improved diagnostic capacity allowed the ED to act immediately on a reliable negative result and initiate targeted admission based on the patient's underlying medical need, using isolation wards only when they were truly required. Although this shift also has implications for hospital-wide efficiency and bed utilisation, in this thesis targeted admissions are mainly interpreted as an ED performance signal of more effective and safer initial disposition decisions, rather than as a direct efficiency gain.

Reduced Intrahospital Transfers

Following the introduction of the POCT, intrahospital transfers within the first five days after admission declined by 32% (50% $\rightarrow$ 34%). More specifically, transfers among participants who had received a negative test result at the ED decreased steadily across the three periods, dropping from about one-third of cases pre-POCT to below one-sixth after POCT. This substantial reduction indicates that the faster and more reliable diagnostic workflow enabled more accurate initial ward placement, reducing unnecessary patient movements and thereby improving patient safety in hospital operations.

Collectively, these findings can be translated into the following IF–THEN statement that summarize how the adaptation influences ED performance:

IF the ED implemented point-of-care testing (POCT) for SARS-CoV-2, **THEN:**

- **ED LOS** decreased by 24.6% (High), indicating improvement in **Timeliness** (end-to-end).
- **Patients discharged within 4 hours (ED LOS < 4h)** increased by 60.1% (High), indicating improvement in **Timeliness** (meeting the 4-hour benchmark).
- **Time to result** decreased by 57.1% (High), indicating improvement in **Timeliness** (diagnostics process).
- **Patients per day:** Changes in this indicator reflect altered ED inflow and case mix, not a performance consequence of the work adaptation.
- **Hourly occupancy rate** decreased by 29.0% (High), indicating improvement in **Efficiency**.
- **Diagnostic coverage (Discharged home classified as "Not tested")** decreased by 69.1% (High), indicating improvement in **Effectiveness**.
- **Targeted admissions (ED-Neg $\rightarrow$ non-COVID ward)** increased by 81.0% (Very High), indicating improvement in **Effectiveness**.
- **Intrahospital transfers** decreased by 32.0% (High), indicating improvement in **Safety**.

In conclusion, POCT implementation was associated with shorter ED LOS and higher proportions of patients leaving within 4 hours (Timeliness), faster time to result in the diagnostic process (Timeliness), lower hourly occupancy and fewer intrahospital transfers (Efficiency, Safety), and higher diagnostic coverage together with more targeted admissions to appropriate non-COVID wards (Effectiveness, Safety). The evidence points to relief of a diagnostic bottleneck within the ED throughput, where rapid, bedside SARS-CoV-2 testing enables earlier, better-informed disposition decisions that both reduce crowding and improve the safety of initial ward placement.

With these structured descriptions of performance change now assembled, next chapter takes them into Phases 2 and 3 of the PI–RA framework, mapping each work adaptation onto the ED workflow and deriving the performance change scores (PC_M and PC_B) together with an IOM-based quality-of-care read-out.

6

Resilient Performance and Quality Trade-Offs

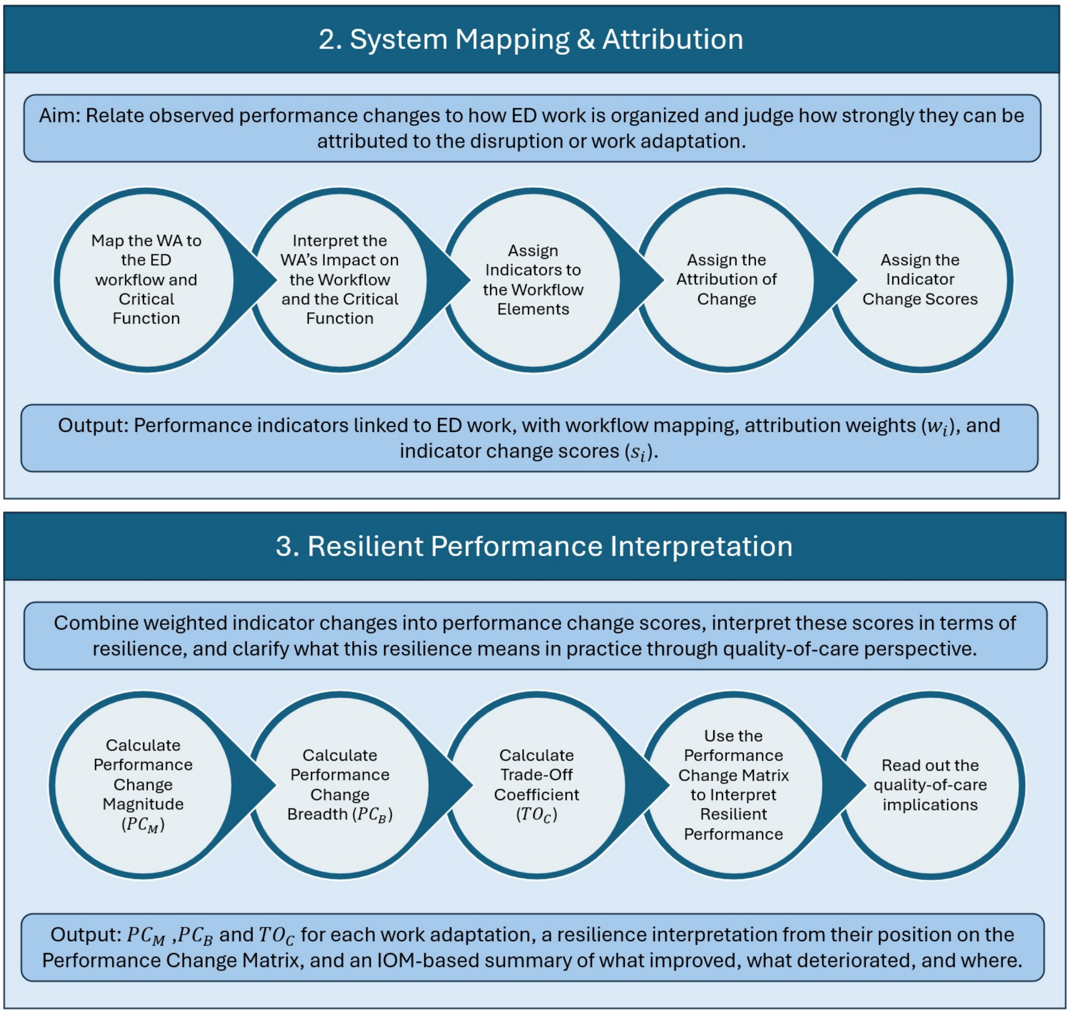

Figure 6.1: Phases 2 and 3 of PI-RA: System Mapping and Attribution & Resilient Performance Interpretation

This chapter operationalizes the second and third phases of the PI-RA framework for the example work adaptations introduced in Chapter 5 (see Figure 6.1). For each work adaptation, two structured outputs from Chapter 5 are taken forward. First, the IF–THEN statement summarizes the observable performance response: it lists which ED performance indicators (PIs) changed after implementation, the direction and magnitude category of each change, and the IOM quality dimensions these indicators speak to. Second, the accompanying case summary describes the mechanisms and performance implications of the adaptation in words: how the workflow was changed, where bottlenecks or relief points were created, and what this means for the ED's functioning. Together, these two elements provide both a numerical "what changed" and a narrative "how and why it changed".

Phase 2, *System mapping and attribution*, uses the ED workflow and critical-function model from Sub-question 1 to anchor these Chapter-5 outputs in the system. The IF–THEN statement is used to keep the analysis focused on the indicators that actually changed and on the parts of the workflow they measure (for example, changes in diagnostic waiting times direct attention to the diagnostics segment, not to initial assessment). The mechanism and implication narrative is then used to interpret how the adaptation is expected to alter flows and pressures across phases. On this basis, the analysis (i) maps the work adaptation to the relevant workflow phase(s) and critical-function element(s), (ii) assigns each indicator from the IF–THEN statement to the workflow location and critical-function component it reflects, (iii) judges how credibly the observed change in each indicator can be attributed to the adaptation (attribution weights w_i), and (iv) converts the direction and magnitude category from the IF–THEN statement into a signed indicator change score s_i. This conversion follows the generic magnitude-to-score mapping shown in Figure 6.2, so that a "High" or "Very high" change in the IF–THEN statement transparently translates into a larger absolute s_i.

Phase-3, *Resilient performance interpretation*, then aggregates these weighted indicator scores into performance-change measures: Performance Change Magnitude (PC_M), Performance Change Breadth (PC_B), and the trade-off coefficient (TO_C). PC_M summarizes the net direction and strength of performance change across all relevant indicators, PC_B expresses the net breadth of change in units of weighted indicators, and TO_C indicates how dominant improvements or deteriorations are across the weighted indicator set (i.e. how trade-off-heavy the overall pattern is). Together, PC_M and TO_C position each work adaptation on the Performance Change Matrix (with PC_M on the horizontal axis and TO_C on the vertical axis), while PC_B and the IOM-based quality-of-care narrative are used alongside the matrix to make explicit which aspects of care improved, which deteriorated, and where trade-offs occurred.

Applied to the three example cases from Chapter 5, these two phases show how the PI-RA framework turns structured performance-change descriptions into system-anchored measures of performance change and associated quality trade-offs. In doing so, this chapter demonstrates how resilient behaviour during disruption can be interpreted in a transparent and comparable way using ED performance indicators.

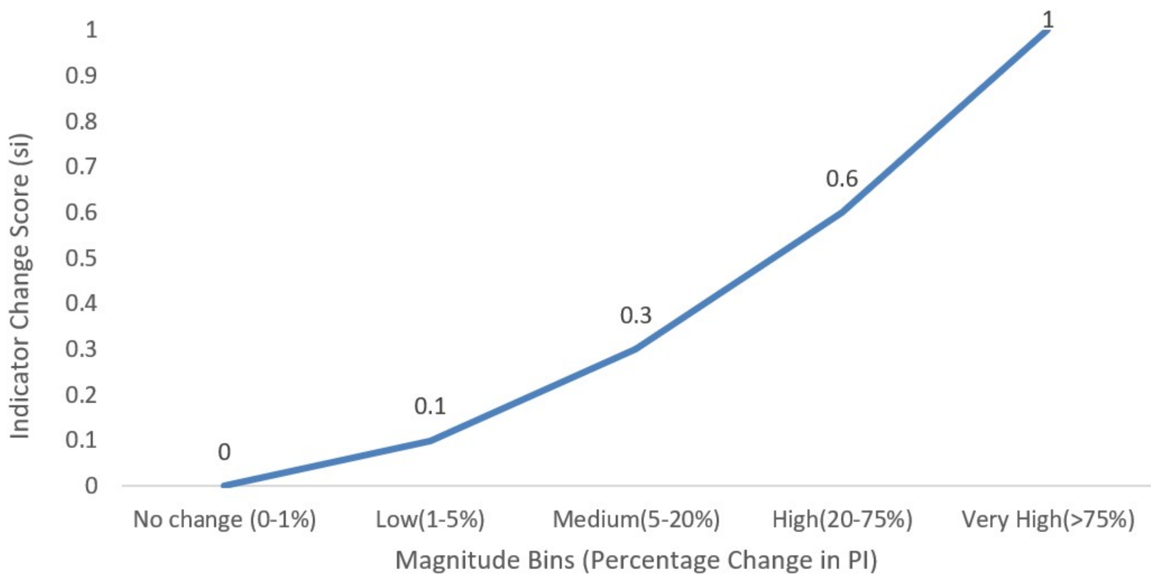

Figure 6.2: Mapping of PI Percentage-Change Magnitude Bins to Indicator Change Scores

6.1. Resilient Performance of the Enhanced Isolation Protocol for Fever or Respiratory Patients (FRPs)

Figure 6.3 highlights (red circled) the placement of this work adaptation on ED Workflow. The work adaptation reduces infection-contagion risk in Phase 1 by separating high-risk FRPs from NFRPs: FRPs are streamed into an isolation zone after triage, creating a capacity-limited queue before initial provider evaluation. This corresponds to the room-placement step, which is expected to be timely and efficient.

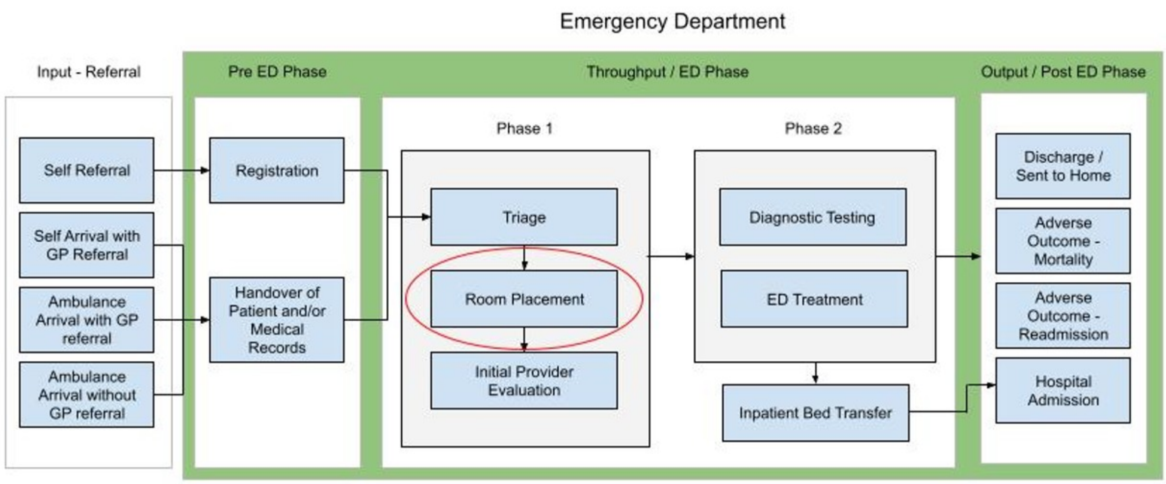

Figure 6.3: Placement of "Enhanced Isolation Protocol for Fever or Respiratory Patients (FRPs)" Work Adaptation on ED Workflow

To anchor the analysis in the following sections, we first restate the IF–THEN statement constructed in Chapter 5 for this work adaptation:

IF *patients with fever or respiratory symptoms were streamed to a capacity-limited isolation zone with restricted diagnostics,* ***THEN:***

- **ED Bed Occupancy Rate** decreased by 29.8% (High), indicating a improvement in **Efficiency**.
- **ED Length of Stay** decreased by 5.9% (Medium), indicating improvement in **Timeliness** (end-to-end).
- **Waiting Time** increased for FRPs by 87.5% (Very High) and for NFRPs by 13%, indicating deterioration in **Timeliness** (front-end).
- **LWBS** increased by 137% (Very High), indicating deterioration in **Safety**.
- **Disparity in WT and LWBS between FRPs and NFRPs** is 573.1% (Very High) for WT (FRPs +87.5% vs NFRPs +13.0%) and is 676.8% (Very High) for LWBS (FRPs +585.7% vs NFRPs +75.4%), indicating deterioration in **Equity**.
- **Infection contagion reduced**: It is the intended **safety** benefit however is not measured through any PI.
- **Admissions, Home Discharges, Transfers, Mortality:** Changes in these indicators reflect the altered case mix and routing, not a performance consequence of the work adaptation.

6.1.1. Impact on ED workflow and Critical Function

Mechanistically, the adaptation creates a local bottleneck in Phase 1. By streaming FRPs to isolation after triage, a capacity-limited queue forms before Initial Provider Evaluation, which reduces the timeliness of the rapid assessment function. This bottleneck increases LWBS among both FRPs (+587.7%) and NFRPs (+75.4%). The higher LWBS reduces the number of patients remaining in the ED, which in turn produces the observed decreases in BOR (−29.8%) and ED LOS (−5.9%), spanning the whole throughput phase (Phase 1 + Phase 2). However, the attribution of these reductions is *mixed*: the decreases in BOR and ED LOS reflect both the increase in LWBS and a decline in overall ED presentations in the post-COVID phase.

The large disparity in WT and LWBS between FRPs and NFRPs can be interpreted as an equity issue: different groups are affected differently. However, there is no systematic exclusion of any group—the ED is not preventing a group from getting a care. This disparity is not tied to the critical function of the ED and is therefore difficult to associate directly with system resilience.

Furthermore, to account for the intended safety benefit of this work adaptation, an explicit RIC (Reduced Infection Contagion) term is included in the scoring. This effect is not directly measured by a PI in the study, but it is central to the adaptation's purpose; omitting it would bias the assessment toward the observed increase in LWBS and ignore the safety impact.

Finally, as stated in the case analysis and reflected in the IF–THEN rule, admissions, home discharges, transfers, and mortality do not capture the impact of this work adaptation. The adaptation acts in Phase 1, before initial provider evaluation. These indicators are (i) outcome measures more related to the care given; (ii) mortality could be affected by longer ED LOS and LWBS, but here LOS improved and no change in mortality is observed; and (iii) they require a clear contextual goal to meaningfully signal improvement or deterioration. For example, if the adaptation aimed to increase admission or disposition rates to reduce ED resource use, changes in these indicators would reflect Efficiency and resilience in action; without such a goal, there is no meaningful direction for s_i. Consistent with this, their w_i are set to 0, which avoids any mathematical issue in the scoring.

6.1.2. Performance Indicator Analysis

To make the connection between the IF–THEN statement and the indicator scores s_i explicit, Table 6.1 restates the mapping from the percentage-change magnitude categories (low, medium, high, very high) to the corresponding absolute score levels used for s_i. With the IF–THEN statement explicitly linked to the critical function and the ED workflow, the following Performance Indicator analysis is presented. A concise summary of the inputs for the PC_M, PC_B, and TO_C calculation are provided in Table 6.2.

Table 6.1: Magnitude bins and corresponding absolute indicator change scores.

Magnitude Bins (Percentage Change in PI)	Indicator Change Score (s_i Absolute Value)
Very High	1.0
High	0.6
Medium	0.3
Low	0.1
No change	0

ED Bed Occupancy Rate (BOR)

Classification: Process
IOM dimension: Efficiency
Link to ED workflow: Utilization of ED treatment/holding spaces across Phase 1 + Phase 2 (end-to-end throughput).
Link to critical function: Capacity to manage patient flow (whole throughput function).
Score: $s_i = +0.6$
Weight & rationale: $w_i = 0.5$. The observed ED BOR decrease has mixed attribution, as higher LWBS and a lower overall volume of ED presentations reduce throughput pressure, diluting direct attribution to the adaptation.

ED Length of Stay (LOS)

Classification: Process
IOM dimension: Timeliness
Link to ED workflow: Total time from arrival to disposition across Phase 1 + Phase 2 (end-to-end throughput).
Link to critical function: Capacity to manage patient flow (whole throughput function).
Score: $s_i = +0.3$
Weight & rationale: $w_i = 0.5$. The observed ED LOS decrease has mixed attribution, as higher LWBS and a lower overall volume of ED presentations reduce throughput pressure, diluting direct attribution to the adaptation.

Waiting Time (WT)

The study did not report an overall ED waiting time, so it is approximated by weighting the subgroup medians for FRPs and NFRPs by their respective shares in the ED population. Let p_{FRP} denote the proportion of FRPs; then

$$WT_{\text{overall}} \approx p_{\text{FRP}} \cdot WT_{\text{FRP}} + (1 - p_{\text{FRP}}) \cdot WT_{\text{NFRP}}.$$

Using the reported subgroup medians (FRP: $24.0 \to 45.0$ minutes; NFRP: $23.0 \to 26.0$ minutes) and FRP shares (pre: 24.2%, post: 22.3%), we obtain $WT_{\text{overall}} \approx 23.24$ minutes pre-COVID and 30.24 minutes post-COVID, yielding a change of $\approx +30\%$ (High)

Classification: Process
IOM dimension: Timeliness
Link to ED workflow: Queue time (front-end timeliness) in Phase 1 from Triage to Initial Provider Evaluation.
Link to critical function: Rapid assessment function.
Score: $s_i = -0.6$
Weight & rationale: $w_i = 1$. Direct node measure; the isolation stream creates a capacity-limited queue that directly lengthens WT.

Left Without Being Seen (LWBS)

Classification: Outcome
IOM dimension: Safety
Link to ED workflow: Exit from the system in Phase 1 before Initial Provider Evaluation
Link to critical function: Rapid assessment function.
Score: $s_i = -1.0$
Weight & rationale: $w_i = 1$. It is a direct consequence of the bottleneck created by the work adaptation at the node.

Discharged to Home

Classification: Outcome
IOM dimension: Efficiency (disposition mix; context-dependent)
Link to ED workflow: Output of the whole throughput, not a part of the throughput processes themselves.
Link to critical function: Capacity to manage patient flow.
Score: $s_i = 0.6$
Weight & rationale: $w_i = 0$. Not a performance consequence of this Phase-1 work adaptation.

Transfer to Another ED

Classification: Outcome
IOM dimension: Efficiency (disposition mix; context-dependent)
Link to ED workflow: Output of the whole throughput, not a part of the throughput processes themselves.
Link to critical function: Capacity to manage patient flow.
Score:$s_i = 0.6$
Weight & rationale: $w_i = 0$. Not a performance consequence of this Phase-1 work adaptation.

Admitted to Inpatient Care

Classification: Outcome
IOM dimension: Efficiency (disposition mix; context-dependent)
Link to ED workflow: Output of the whole throughput, not a part of the throughput processes themselves. It is connected to the inpatient transfer node.
Link to critical function: Capacity to manage patient flow.
Score: $s_i = 0.3$
Weight & rationale: $w_i = 0$. Not a performance consequence of this Phase-1 work adaptation.

Mortality

Classification: Outcome
IOM dimension: Safety
Link to ED workflow: Related to the effectiveness of Phase 2 (Diagnostic Testing + ED Treatment)
Link to critical function: Safety consequence of diagnosis and treatment functions
Score: $s_i = 0$.
Weight & rationale: $w_i = 0$. Not a performance consequence of this Phase-1 work adaptation.

Reduced Infection Contagion (RIC) — assumed safety benefit

IOM dimension: Safety
Link to ED workflow: Separation of patients in Phase 1 (Room Placement)
Link to critical function: Rapid assessment function — safety benefit during front-end flow by reducing infection contagion risk.
Score: Assumed Very High $\rightarrow s_i = +1.0$.
Weight & rationale: $w_i = 1$. Excluding this benefit would bias the resilience rating away from the adaptation's primary goal.

Table 6.2: Indicators, scores, and weights used to compute PC_M, PC_B, and TO_C for the FRP case.

Performance Indicator	% Change	Magnitude Bin	s_i	w_i
ED Bed Occupancy Rate (BOR)	-29.8%	High	$+0.6$	0.5
ED Length of Stay (LOS)	-5.9%	Medium	$+0.3$	0.5
Left Without Being Seen (LWBS)	$+137\%$	Very High	-1.0	1
Waiting Time (overall ED)	$+30\%$	High	-0.6	1
Discharged to Home	-11.6%	Medium	0.3	0
Transfer to Another ED	-28.9%	High	0.6	0
Admitted to Inpatient Care	$+10.5\%$	Medium	0.3	0
Mortality Rate	0%	No change	0	0
Reduced Infection Contagion (RIC)			$+1.0$	1

6.1.3. Performance Change Scores and Resilient Performance Interpretation

We now compute the work adaptation's performance-change measures using formulas (3.1), (3.2), and (3.3) defined in the Methodology section, using the indicator inputs for the FRP case given in Table 6.2.

Performance Change Magnitude (PC_M)

$$\begin{aligned}\sum_i w_i s_i &= (0.5)(+0.6) + (0.5)(+0.3) + (1)(-1.0) + (1)(-0.6) + (1)(+1.0)\\ &= 0.30 + 0.15 - 1.00 - 0.60 + 1.00 = -0.15,\\ \sum_i w_i &= 0.5 + 0.5 + 1 + 1 + 1 = 4.0,\\ PC_M &= \frac{-0.15}{4.0} = -\mathbf{0.0375}.\end{aligned}$$

Performance Change Breadth (PC_B)

$$\begin{aligned}\sum_i w_i \operatorname{sign}(s_i) &= (0.5)(+1) + (0.5)(+1) + (1)(-1) + (1)(-1) + (1)(+1)\\ &= 0.5 + 0.5 - 1 - 1 + 1 = 0,\\ PC_B &= \mathbf{0}.\end{aligned}$$

Trade-off Coefficient (TO_C)

$$TO_C = \frac{PC_B}{\sum_i w_i} = \frac{0}{4.0} = \mathbf{0}.$$

Resilient Performance Interpretation

When plotted on the Performance Change Matrix (Figure 6.4), the FRP adaptation lies almost at the origin, slightly to the left on the horizontal axis with a small negative performance change magnitude ($PC_M = -0.0375$) and on the vertical midline with a trade-off coefficient of zero ($TO_C = 0$). This places the FRP adaptation near the centre of the matrix but in the deteriorating half, on the horizontal midline between the two deterioration regions (Regions III and IV). The small negative value of PC_M indicates that, on average, ED performance deteriorates slightly compared to the pre-adaptation situation rather than improving. The unnormalised net breadth score is also zero ($PC_B = 0$), meaning that—after accounting for the relevance weights—the number of improving and deteriorating indicators balances out. In other words, gains and losses cancel out, and $TO_C = 0$ shows that neither improvements nor deteriorations dominate: the overall pattern is fully trade-off-driven. FRPs are therefore a trade-off-heavy work adaptation in which performance gains in some parts of the ED are offset by losses elsewhere, resulting in a small net deterioration rather than a clear net improvement.

On the case-specific resilience curve (Figure 6.5), this position on the matrix is shown as an illustrative trajectory in which performance drops away from P_0 after the disruption at t_1 and continues to move slightly downward once the FRP strategy is introduced around t_2. Given that $|PC_M|$ is very small, this interpretation should be treated with caution: such a slight net deterioration may partly reflect the chosen binning and scoring scheme rather than a strong substantive effect. The curve in Figure 6.5 is therefore not a reconstruction of the empirical time series, but a schematic representation of the net effect that can be inferred from the scores. Rather than implying that the curve will necessarily keep moving downwards, the safer conclusion is that the FRP adaptation does not succeed in bending the trajectory back towards the required operations level (or pre-disruption level) P_0. Taken together, the scores therefore suggest that FRPs, on their own, act as a maladaptive or failed adaptation: they change the configuration of ED work in response to COVID-19 but do not succeed in restoring performance onto an upward recovery trajectory.

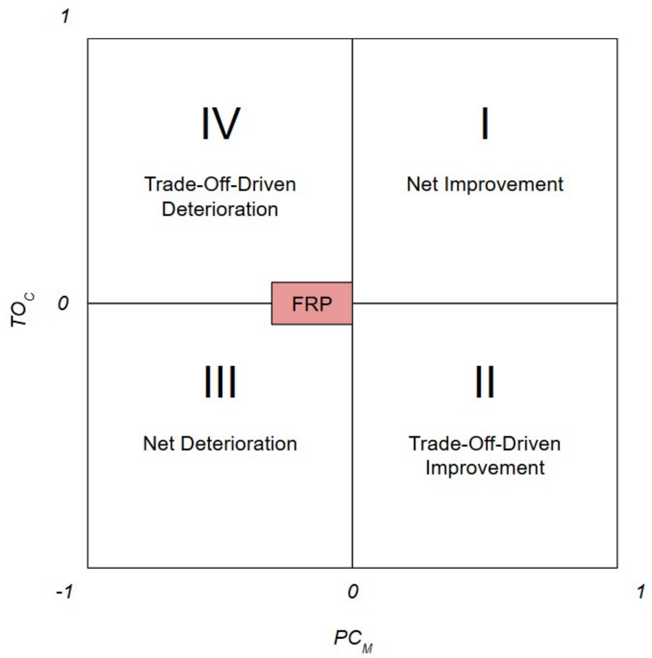

Figure 6.4: Placement of the FRP isolation strategy on the Performance Change Matrix ($PC_M = -0.0375$ & $TO_C = 0$)

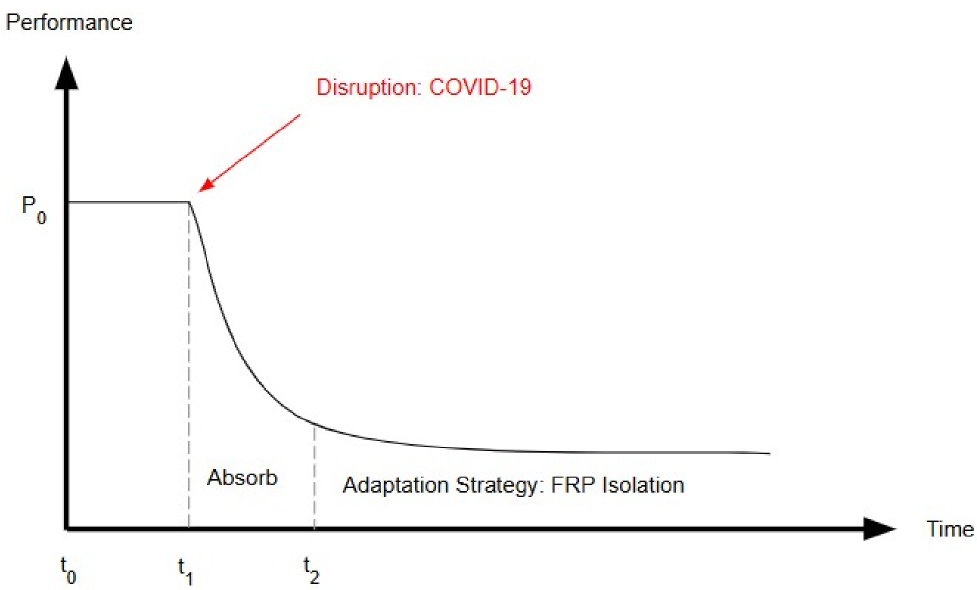

Figure 6.5: Illustrative resilience curve for the FRP isolation strategy.

6.1.4. IOM Dimension Trade-offs

Drawing on the analysis above, we observe a Timeliness deterioration at the front end and a Timeliness improvement end-to-end. We read the front-end deterioration—captured by Waiting Time—as more impactful than the end-to-end improvement. In parallel, there is a small Efficiency improvement due to reduced BOR.

For Safety, we assume that the benefits of reduced infection contagion are as impactful as the increased safety risks associated with patients who LWBS. This constitutes a safety trade-off for patients who remain in the ED and for those who LWBS, and it also relates to safety concerns linked to infection contagion outside the ED (in public and other institutions).

Building on this safety disparity, we conclude that Equity deteriorates while safety in the ED improves. This interpretation is consistent with the ED goal: there is no negative impact on the critical function; rather, there is an improved safety benefit within the ED. An alternative reading of this work adaptation would exclude the LWBS impact from the ED performance on the grounds that ED operations and relevant functions are not themselves deteriorated. Our interpretation, however, is that failing to provide the required care constitutes a deterioration of the ED's critical function. In this sense, the equity issue that arises from the disparity is already considered in the performance assessment; it simply lacks a direct tie to the critical function or a specific ED workflow node.

In conclusion, the FRP isolation strategy represents a maladaptive or failed adaptation in resilient-performance terms: despite reconfiguring the front-end workflow, the combined indicator pattern places the case near the origin on the negative side of the Performance Change Matrix, indicating that performance does not move into a clear recovery phase towards P_0. Within this region of the matrix, the quality-dimension trade-off reads as increased Safety in the ED versus deterioration of Timeliness at the front end, deterioration of Safety outside the ED, and deteriorated Equity. Consistent with this, the performance change breadth ($PC_B = 0$) indicates balanced gains and losses across the weighted indicators—no net breadth of improvement—in line with the observed trade-offs ($TO_C = 0$).

6.2. Resilient Performance of Implementing Rapid Assessment Zone (RAZ)

RAZ is modeled as a parallel front-end care path that begins with triage and is eligible for non-critical patients. In this path, patients receive initial provider evaluation and simple interventions in vertical bays and early orders are initiated. From RAZ, patients either (i) are discharged directly after reassessment when appropriate, or (ii) enter the main ED flow for Diagnostic Testing and ED Treatment when fuller care is required; critical patients bypass RAZ to a main-ED bed. This pathway is overlaid on the previously defined ED workflow and illustrated in Figure 6.6.

To anchor the analysis in the following sections, we first restate the IF–THEN statement constructed in Chapter 5 for this work adaptation:

- **Arrival-to-provider** decreased by 60.7% (High), indicating improvement in **Timeliness** (front-end).
- **ED LOS** decreased for *discharged* patients by 20.5% (High) and for *admitted* patients by 15.9% (Medium), indicating improvement in **Timeliness** (end-to-end).
- **LWBS** decreased by 54.8% (High), indicating improvement in **Safety**.

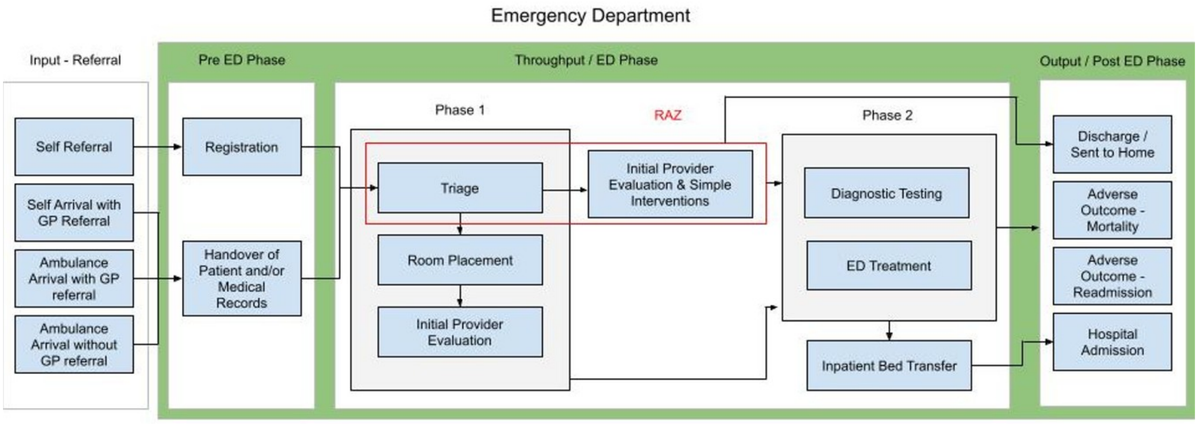

Figure 6.6: Implementation of Rapid Assessment Zone (RAZ) on ED Workflow

- **Admissions, Discharges:** Changes in these indicators reflect the altered case mix, not a performance consequence of the work adaptation.

6.2.1. Impact on ED Workflow and Critical Function

Mechanistically, the adaptation establishes a parallel front-end path in Phase 1. Eligible, non-critical patients receive initial provider evaluation in RAZ without waiting for Room Placement. This shortens the queue to first contact (Arrival-to-provider time 60.7%↓), strengthens rapid assessment function, and reduces dependence on bed availability, thereby supporting capacity to manage flow.

End-to-end Timeliness improves for both groups (ED LOS ↓). For discharged patients, earlier evaluation and disposition within the RAZ path shorten total time in the department. For admitted patients, earlier clinical decision-making at the front end moves them through the system more quickly despite back-end constraints due to patient boarding. These effects reflect bringing assessment forward in the visit and limiting unnecessary Room Placement for low-acuity patients, which preserves capacity for higher-acuity care.

Safety improves through LWBS ↓, as faster access to a clinician at arrival reduces the likelihood of patients leaving before evaluation or treatment.

Admissions, Discharges (case-mix). Post-implementation shifts in these proportions are interpreted as case-mix changes under the vertical-care model rather than performance consequences of the adaptation and are therefore noted but not scored.

6.2.2. Performance Indicator Analysis

To make the connection between the IF–THEN statement and the indicator scores s_i explicit, Table 6.3 restates the mapping from the percentage-change magnitude categories (low, medium, high, very high) to the corresponding absolute score levels used for s_i. With the IF–THEN statement explicitly linked to the critical function and the ED workflow, the following Performance Indicator analysis is presented. A concise summary of the inputs for the PC_M, PC_B, and TO_C calculation are provided in Table 6.4.

Table 6.3: Magnitude bins and corresponding absolute indicator change scores.

Magnitude Bins (Percentage Change in PI)	Indicator Change Score (s_i Absolute Value)
Very High	1.0
High	0.6
Medium	0.3
Low	0.1
No change	0

Arrival-to-Provider Time

Classification: Process
IOM dimension: Timeliness
Link to ED workflow: Time from arrival to initial provider evaluation in Phase 1 (front-end timeliness).
Link to critical function: Rapid assessment function.
Score: $s_i = +0.6$
Weight & rationale: $w_i = 1$. The adaptation operates directly at this node: eligible, non-critical patients receive the *initial provider evaluation in RAZ* without waiting for room placement, collapsing the queue to first contact and producing the observed high (60.7%) decrease.

ED Length of Stay (ED LOS)

The study did not report an overall ED length of stay, so it is approximated by weighting the subgroup medians for *discharged* and *admitted* patients by their respective shares in the ED population. Let p_{disch} denote the proportion of discharged; then

$$LOS_{\text{overall}} \approx p_{\text{disch}} \cdot LOS_{\text{disch}} + (1 - p_{\text{disch}}) \cdot LOS_{\text{adm}}.$$

Using the reported subgroup medians (**Discharged:** $205.0 \rightarrow 163.0$ minutes; **Admitted:** $395.0 \rightarrow 332.0$ minutes) and discharge shares (*pre:* 70.95%, *post:* 73.32%), we obtain $LOS_{\text{overall}} \approx \mathbf{260.2}$ minutes pre-COVID and $\mathbf{208.1}$ minutes post-COVID, yielding a change of $\approx \mathbf{-20.0}\%$ (High).

Classification: Process
IOM dimension: Timeliness
Link to ED workflow: Total time from arrival to disposition across Phase 1 + Phase 2 (end-to-end throughput).
Link to critical function: Capacity to manage patient flow (whole throughput function).
Score: $s_i = +0.6$
Weight & rationale: $w_i = 1$. The overall ED LOS decreased by 20% (High). Between the pre- and post-periods, total ED visits increased slightly while the case mix shifted toward more mid- to lower-acuity patients (more ESI 3–4, fewer ESI 2). Lower acuity would normally shorten LOS, whereas higher volume would be expected to lengthen LOS because of crowding. As these opposing background trends are modest relative to the observed LOS reduction, the net improvement is treated as primarily driven by the rapid-assessment-zone work adaptation, so no mixed-attribution penalty is applied.

Left Without Being Seen (LWBS)

Classification: Outcome
IOM dimension: Safety

Link to ED workflow: Exit from the system in Phase 1 before Initial Provider Evaluation
Link to critical function: Rapid assessment function.
Score: $s_i = +0.6$
Weight & rationale: $w_i = 1$. The reduction is a direct consequence of the RAZ path enabling faster first clinician contact at this node, lowering the likelihood that patients leave prior to evaluation.

Discharged to Home

Classification: Outcome
IOM dimension: Efficiency (disposition mix; context-dependent)
Link to ED workflow: Output of the whole throughput, not a part of the throughput processes themselves.
Link to critical function: Capacity to manage patient flow. **Score:** $s_i = 0.1$
Weight & rationale: $w_i = 0$. Not a performance consequence of this Phase-1 work adaptation.

Admitted to Inpatient Care

Classification: Outcome
IOM dimension: Efficiency (disposition mix; context-dependent)
Link to ED workflow: Output of the whole throughput, not a part of the throughput processes themselves. It is connected to the inpatient transfer node.
Link to critical function: Capacity to manage patient flow.
Score: $s_i = 0.1$
Weight & rationale: $w_i = 0$. Not a performance consequence of this Phase-1 work adaptation.

Table 6.4: Indicators, scores, and weights used to compute PC_M, PC_B, and TO_C for the RAZ case

Performance Indicator	% Change	Magnitude Bin	s_i	w_i
Arrival-to-Provider Time	-60.7%	High	$+0.6$	1
ED Length of Stay (overall ED)	-20.0%	High	$+0.6$	1
Left Without Being Seen (LWBS)	-54.8%	High	0.6	1
Discharged to Home	$+7.1\%$	Low	0.1	0
Admitted to Inpatient Care	-4.8%	Low	0.1	0

6.2.3. Performance Change Scores and Resilient Performance Interpretation

We now compute the work adaptation's performance-change measures using formulas (3.1), (3.2), and (3.3) defined in the Methodology section, with the inputs expressed in Table 6.4.

Performance Change Magnitude (PC_M)

$$\sum_i w_i s_i = (1)(+0.6) + (1)(+0.6) + (1)(+0.6) = 1.8,$$

$$\sum_i w_i = 1 + 1 + 1 = 3.0,$$

$$PC_M = \frac{1.8}{3.0} = +\mathbf{0.60}.$$

Performance Change Breadth (PC_B)

$$\sum_i w_i \operatorname{sign}(s_i) = (1)(+1) + (1)(+1) + (1)(+1) = 3,$$

$$PC_B = \mathbf{3}.$$

Trade-off Coefficient (TO_C)

$$TO_C = \frac{PC_B}{\sum_i w_i} = \frac{3}{3.0} = +\mathbf{1.0}.$$

Resilient Performance Interpretation

When plotted on the Performance Change Matrix (Figure 6.7), the RAZ adaptation lies in the upper-right quadrant (Region I: net improvement), with a performance change magnitude of $PC_M = +0.60$ and a trade-off coefficient of $TO_C = +1.0$. This placement indicates net improvement: performance improves overall, and all weighted indicators move in the desired direction. The relatively high positive value of PC_M reflects a substantial net gain in the performance of required operations, driven by shorter arrival-to-provider times, reduced ED LOS, and fewer LWBS events. The value $PC_B = 3$ indicates that, in weighted terms, the equivalent of three indicators improve on net. The accompanying $TO_C = +1.0$ confirms that these gains are not realized through trade-offs: all scored indicators improve and none deteriorate.

On the case-specific resilience curve (Figure 6.8), this position on the matrix is shown as an illustrative trajectory in which performance drops away from P_0 after the disruption at t_1, reaches a minimum during the absorption phase, and then turns upward once the RAZ strategy is introduced around t_2. The positive value of PC_M and the fact that all scored indicators improve (as reflected by $PC_B = 3$ and $TO_C = +1.0$) together indicate that, over the period in which RAZ is analyzed, the ED moves into a recovery segment of the resilience curve. Because pre-COVID performance levels and longer-term developments are unknown, the dashed curves in Figure 6.8 represent this uncertainty: the final performance level P_F may end up below or above P_0, and other concurrent work adaptations will also shape the eventual trajectory. The curve is therefore not a reconstruction of the empirical time series, but a schematic representation of the net effect captured by the scores: the RAZ adaptation helps to bend the performance trajectory upwards after the disruption, expressing resilient performance through an improvement in timeliness and safety without visible trade-offs in the measured indicators.

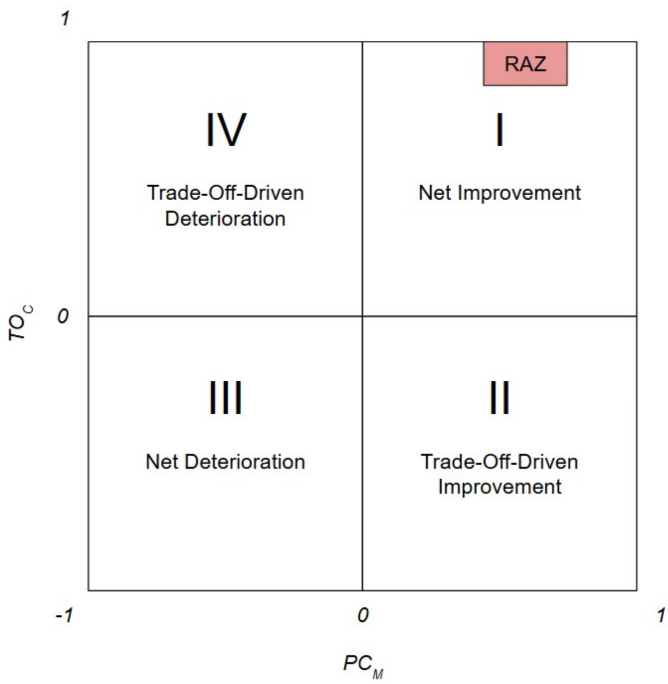

Figure 6.7: Placement of the RAZ on the Performance Change Matrix ($PC_M = +0.6$ & $TO_C = +1$)

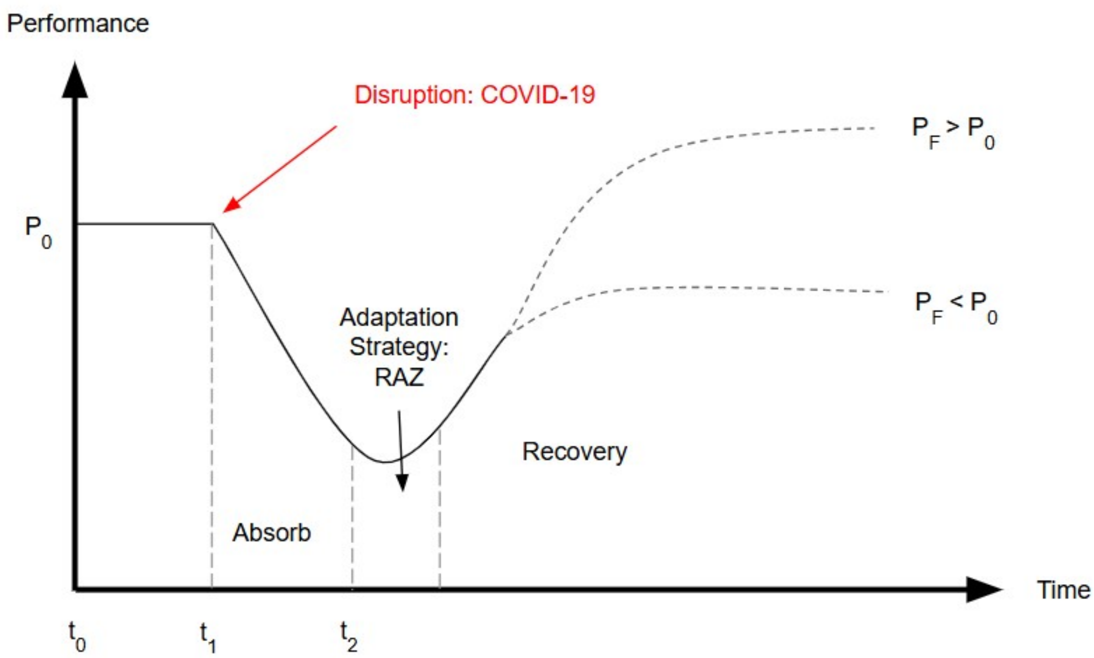

Figure 6.8: Illustrative resilience curve for the RAZ adaptation strategy.

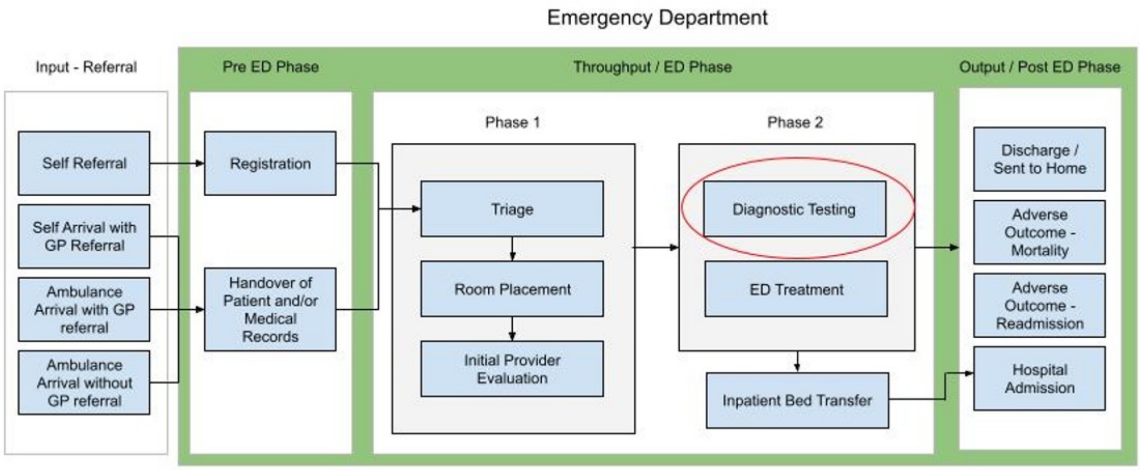

Figure 6.9: Implementation of Point of Care Testing (POCT) on ED Workflow

6.2.4. IOM Dimension Trade-offs

Drawing on the analysis above, we observe Timeliness improvement at the front end (Arrival-to-provider) and end-to-end Timeliness improvement (ED LOS for both discharged and admitted). These gains follow from performing the initial provider evaluation within RAZ as a parallel front-end path, shortening time to first clinician contact and accelerating progression to disposition for eligible, non-critical patients.

For Safety, performance improves through a substantial reduction in LWBS. Faster access to a clinician reduces the likelihood of patients leaving prior to evaluation, reinforcing the rapid-assessment element of the ED's critical function.

In conclusion, the Rapid Assessment Zone represents a successful adaptation in resilient-performance terms: by reconfiguring the front-end workflow, the combined indicator pattern places the case in the first region of the Performance Change Matrix, indicating that performance moves into a recovery phase towards P_0. Within this region of the matrix, the quality-dimension pattern reads as improved Timeliness at the front end and end-to-end, together with improved Safety through the reduction in LWBS, without deterioration in the other scored quality dimensions. Consistent with this, the positive performance change breadth ($PC_B = 3$) indicates a net breadth of improvement across the weighted indicators, in line with the absence of observable trade-offs ($TO_C = 1.0$).

6.3. Implementation of Point of Care Testing (POCT)

Figure 6.9 highlights (red circled) the placement of POCT on the ED workflow. In this model, the work adaptation modifies the diagnostic testing step in Phase 2 of the throughput / ED phase, where SARS-CoV-2 confirmation is performed within the ED. This step operationalizes the "correctly diagnosing" component of the ED's critical function and feeds into downstream admission and discharge decisions. With bedside molecular POCT replacing delayed core-laboratory RT-PCR, this diagnostic part of the workflow is expected to function effectively, by providing reliable results in time to support appropriate ward placement for each patient.

To anchor the analysis in the following sections, we first restate the IF–THEN statement constructed in Chapter 5 for this work adaptation:

IF the ED implemented point-of-care testing (POCT) for SARS-CoV-2, **THEN:**

- **ED LOS** decreased by 24.6% (High), indicating improvement in **Timeliness** (end-to-end).
- **Patients discharged within 4 hours (ED LOS < 4h)** increased by 60.1% (High), indicating improvement in **Timeliness** (meeting the 4-hour benchmark).
- **Time to result** decreased by 57.1% (High), indicating improvement in **Timeliness** (diagnostics process).
- **Patients per day:** Changes in this indicator reflect altered ED inflow and case mix, not a performance consequence of the work adaptation.
- **Hourly occupancy rate** decreased by 29.0% (High), indicating improvement in **Efficiency**.
- **Diagnostic coverage (Discharged home classified as "Not tested")** decreased by 69.1% (High), indicating improvement in **Effectiveness**.
- **Targeted admissions (ED-Neg → non-COVID ward)** increased by 81.0% (Very High), indicating improvement in **Effectiveness**.
- **Intrahospital transfers** decreased by 32.0% (High), indicating improvement in **Safety**.

6.3.1. Impact on ED Workflow and Critical Function

Mechanistically, the adaptation relieves a diagnostic bottleneck in Phase 2 of the throughput phase. By moving SARS-CoV-2 confirmation from delayed core-laboratory RT-PCR to rapid molecular POCT in the ED, test results are available much earlier in the visit. This shortens the time between assessment and a definitive diagnosis, allowing disposition decisions to be made sooner and contributing to the observed reductions in time to result, overall ED LOS, and the higher proportion of patients leaving within 4 hours. Together, these changes support the "correctly diagnosing" component of the ED's critical function and help maintain end-to-end flow by improving decision-making in the middle of the throughput pathway.

The decreases in hourly occupancy reflect that patients occupy ED space for a shorter period, which is consistent with faster diagnostic turnaround and earlier disposition. However, the attribution of this effect is mixed: the observed reduction in occupancy is driven both by these flow improvements (shorter LOS) and by a concurrent drop in patients per day during the POCT study period. Since part of the change is explained by lower inflow and altered case mix rather than by the work adaptation itself, only a portion of the occupancy reduction can be directly attributed to POCT.

Finally, the improvements in diagnostic coverage, targeted admissions, and intrahospital transfers show how the adapted diagnostic step influences the quality of downstream outcomes. Higher diagnostic coverage means that a much larger share of discharged patients leave the ED with a documented positive or negative test, rather than an unconfirmed suspicion of COVID-19, reducing the number of patients whose infection status remains unclear at the end of the visit. With POCT results available during the ED stay, clinicians can apply placement rules that direct test-negative patients straight to an appropriate non-COVID ward, which explains the increase in targeted admissions based on underlying medical need. Because more patients start in the correct ward from the outset, there is less need to correct earlier placement decisions once additional information becomes available, which is reflected in the substantial reduction in intrahospital transfers. In terms of the critical function, this pattern indicates more correct diagnoses being made within the ED and more appropriate initial dispositions, aligning with an effective ED (correctly diagnosing) and contributing to a safe ED by reducing avoidable adverse consequences of misplacement and unnecessary transfers.

6.3.2. Performance Indicator Analysis

To make the connection between the IF–THEN statement and the indicator scores s_i explicit, Table 6.5 restates the mapping from the percentage-change magnitude categories (low, medium, high, very high) to the corresponding absolute score levels used for s_i. With the IF–THEN statement explicitly linked to the critical function and the ED workflow, the following Performance Indicator analysis is presented. A concise summary of the inputs for the PC_M, PC_B, and TO_C calculation are provided in Table 6.6.

Table 6.5: Magnitude bins and corresponding absolute indicator change scores.

Magnitude Bins (Percentage Change in PI)	Indicator Change Score (s_i Absolute Value)
Very High	1.0
High	0.6
Medium	0.3
Low	0.1
No change	0

ED Length of Stay (ED LOS)

Classification: Process
IOM dimension: Timeliness
Link to ED workflow: Total time from arrival to disposition across Phase 1 + Phase 2 (end-to-end throughput).
Link to critical function: Capacity to manage patient flow (whole throughput function).
Score: $s_i = +0.6$.
Weight & rationale: $w_i = 1.0$. Even though overall ED presentations decreased during the study period, the LOS reduction remains after accounting for case-mix and volume effects and is therefore attributed primarily to POCT-driven faster diagnostic confirmation and earlier disposition.

Patients Discharged within 4 Hours (ED LOS < 4h)

Classification: Outcome
IOM dimension: Timeliness
Link to ED workflow: Completion of the full ED episode within the 4-hour benchmark from registration to exit.
Link to critical function: Capacity to keep whole throughput function within targeted time limits.
Score: $s_i = +0.6$.
Weight & rationale: $w_i = 0.0$. From an ED performance perspective, this indicator captures the same underlying effect as the reduction in overall ED LOS: improved end-to-end timeliness. Although the increase in patients discharged within 4 hours is a useful benchmark measure, it is essentially a thresholded transformation of LOS rather than an independent performance consequence. To avoid double-counting the same timeliness effect in the resilience scoring, this indicator is assigned zero weight.

Time to Result

Classification: Process
IOM dimension: Timeliness
Link to ED workflow: Diagnostic testing step in Phase 2; interval from registration to SARS-CoV-2 result availability.
Link to critical function: Correctly diagnosing, by providing timely diagnostic information.
Score: $s_i = +0.6$.
Weight & rationale: $w_i = 1.0$. Direct mechanistic effect of switching from delayed laboratory RT-PCR to rapid POCT.

Patients per Day

Classification: Process
IOM dimension: Efficiency (context-dependent)
Link to ED workflow: Overall daily ED volume across all phases.
Link to critical function: Reflects demand, not the functioning of the throughput processes.
Score: $s_i = 0.6$.
Weight & rationale: $w_i = 0.0$. Mainly driven by reduced inflow and case mix during the study period, not by the work adaptation.

Hourly Occupancy Rate

Classification: Process
IOM dimension: Efficiency
Link to ED workflow: Concurrent number of patients present per hour across Phase 1 + Phase 2.
Link to critical function: Capacity to manage patient flow and avoid crowding.
Score: $s_i = +0.6$.
Weight & rationale: $w_i = 0.5$. Decrease reflects both shorter LOS (adaptation effect) and fewer presentations (demand effect), so attribution is mixed.

Diagnostic Coverage (Discharged home classified as ``Not tested")

Classification: Outcome
IOM dimensions: Effectiveness; Safety (secondary)
Link to ED workflow: Disposition step for discharged patients; indicates whether they leave with a documented SARS-CoV-2 result.
Link to critical function: Correctly diagnosing and safely ending the ED episode.
Score: $s_i = +0.6$.
Weight & rationale: $w_i = 1.0$. Strong, directly attributable improvement in evidence-based discharge decisions under POCT.

Targeted Admissions (ED-Neg $\rightarrow$ non-COVID ward)

Classification: Outcome
IOM dimensions: Effectiveness; Safety
Link to ED workflow: Initial inpatient ward placement after ED; whether test-negative patients are sent directly to an appropriate non-COVID ward.
Link to critical function: Correctly diagnosing and safely initiating the subsequent care pathway.
Score: $s_i = +1.0$.
Weight & rationale: $w_i = 1.0$. Very large, directly attributable improvement in appropriateness of first-time ward placement.

Intrahospital Transfers

Classification: Outcome
IOM dimension: Safety (Efficiency secondary at hospital level)
Link to ED workflow: Ward movements within five days after admission; downstream consequence of initial placement decisions originating in the ED.
Link to critical function: Correctly diagnosing and safe ED care by avoiding harmful, unnecessary ward moves.
Score: $s_i = +0.6$.
Weight & rationale: $w_i = 1.0$. Substantial reduction closely linked to more accurate initial placement enabled by POCT.

Table 6.6: Indicators, scores, and weights used to compute PC_M, PC_B, and TO_C for the POCT case

Performance Indicator	% Change	Magnitude Bin	s_i	w_i
ED Length of Stay (ED LOS)	$\downarrow 24.6\%$	High	$+0.6$	1.0
Patients discharged within 4 hours (ED LOS < 4h)	$\uparrow 60.1\%$	High	0.6	0.0
Time to result	$\downarrow 57.1\%$	High	$+0.6$	1.0
Patients per day	$\downarrow 24.7\%$	High	0.6	0.0
Hourly occupancy rate	$\downarrow 29.0\%$	High	$+0.6$	0.5
Diagnostic coverage (Discharged home "Not tested")	$\downarrow 69.1\%$	High	$+0.6$	1.0
Targeted admissions (ED-Neg $\rightarrow$ non-COVID ward)	$\uparrow 81.0\%$	Very High	$+1.0$	1.0
Intrahospital transfers	$\downarrow 32.0\%$	High	$+0.6$	1.0

6.3.3. Performance Change Scores

We now compute the work adaptation's performance-change measures using formulas (3.1), (3.2), and (3.3) defined in the Methodology section, with the inputs expressed in Table 6.6.

Performance Change Magnitude (PC_M)

$$\begin{aligned}\sum_i w_i s_i &= (1)(+0.6) + (1)(+0.6) + (0.5)(+0.6) + (1)(+0.6) + (1)(+1.0) + (1)(+0.6)\\ &= 0.6 + 0.6 + 0.3 + 0.6 + 1.0 + 0.6 = 3.7,\end{aligned}$$

$$\sum_i w_i = 1 + 1 + 0.5 + 1 + 1 + 1 = 5.5,$$

$$PC_M = \frac{3.7}{5.5} \approx +\mathbf{0.67}.$$

Performance Change Breadth (PC_B)

$$\begin{aligned}\sum_i w_i\,\mathrm{sign}(s_i) &= (1)(+1)+(1)(+1)+(0.5)(+1)+(1)(+1)+(1)(+1)+(1)(+1)\\ &= 1+1+0.5+1+1+1 = 5.5,\\ PC_B &= \mathbf{5.5}.\end{aligned}$$

Trade-off Coefficient (TO_C)

$$TO_C = \frac{PC_B}{\sum_i w_i} = \frac{5.5}{5.5} = +\mathbf{1.0}.$$

Resilient Performance Interpretation

When plotted on the Performance Change Matrix, the POCT adaptation lies in the upper-right quadrant (Region I), with a performance change magnitude of $PC_M \approx +0.67$ and a trade-off coefficieny of $TO_C = +1.0$ (Figure 6.10). This placement indicates net improvement: performance improves overall, and all weighted indicators move in the desired direction. The relatively high positive value of PC_M reflects a strong net gain in the performance of required operations, driven by shorter overall ED LOS, faster time to result, lower hourly occupancy, better diagnostic coverage, more targeted admissions, and fewer intrahospital transfers. At the same time, $TO_C = +1.0$ indicates that this improvement does not rely on observable trade-offs in the scored indicators. In terms of the matrix regions, POCT therefore also sits in the desired part of the space where work adaptations strengthen ED performance broadly rather than narrowly or at the expense of other dimensions.

On the case-specific resilience curve (Figure 6.11), this position on the matrix is shown as an *illustrative* trajectory in which performance drops away from P_0 after the disruption at t_1, reaches a minimum during the absorption phase, and then turns upward once the POCT strategy is introduced around t_2. The positive value of PC_M and the broad improvement signalled by $PC_B = +5.5$ together indicate that, over the period in which POCT is active, the ED moves into a recovery segment of the resilience curve rather than remaining at a lowered plateau. Because pre-COVID performance levels and longer-term developments are unknown, the dashed curves in Figure 6.11 again represent this uncertainty: the final performance level P_F may end up below or above P_0, and other concurrent work adaptations will also shape the eventual trajectory. The curve is therefore not a reconstruction of the empirical time series, but a schematic representation of the net effect captured by the scores: the POCT adaptation helps to bend the performance trajectory upwards after the disruption, expressing resilient performance through broad improvements in timeliness, effectiveness, efficiency, and safety without visible trade-offs in the measured indicators.

Resilient Performance Interpretation

When plotted on the Performance Change Matrix (Figure 6.10), the POCT adaptation lies in the upper-right quadrant (Region I: net improvement), with a performance change magnitude of $PC_M \approx +0.67$ and a trade-off coefficient of $TO_C = +1.0$. This placement indicates net improvement: performance improves overall, and all weighted indicators move in the desired direction. The high positive value of PC_M reflects a strong net gain in the performance of required operations, driven by shorter overall ED LOS, faster time to result, lower hourly occupancy, better diagnostic coverage, more targeted admissions, and fewer intrahospital transfers. The value $PC_B = 5.5$ indicates that, in weighted terms, the equivalent of five and a half indicators improve on net, and $TO_C = +1.0$ confirms that these gains are not realized through trade-offs: all scored indicators improve and none deteriorate.

On the case-specific resilience curve (Figure 6.11), this position on the matrix is shown as an illustrative trajectory in which performance drops away from P_0 after the disruption at t_1, reaches a minimum during the absorption phase, and then turns upward once the POCT strategy is introduced around t_2. The positive value of PC_M and the fact that all scored indicators improve (as reflected by $PC_B = 5.5$ and $TO_C = +1.0$) together indicate that, over the period in which POCT is active, the ED moves into a recovery segment of the resilience curve. Because pre-COVID performance levels and longer-term developments are unknown, the dashed curves in Figure 6.11 again represent this uncertainty: the final performance level P_F may end up below or above P_0, and other concurrent work adaptations will also shape the eventual trajectory. The curve is therefore not a reconstruction of the empirical time series, but a schematic representation of the net effect captured by the scores: the POCT adaptation helps to bend the performance trajectory upwards after the disruption, expressing resilient performance through improvements in timeliness, effectiveness, efficiency, and safety without visible trade-offs in the measured indicators.

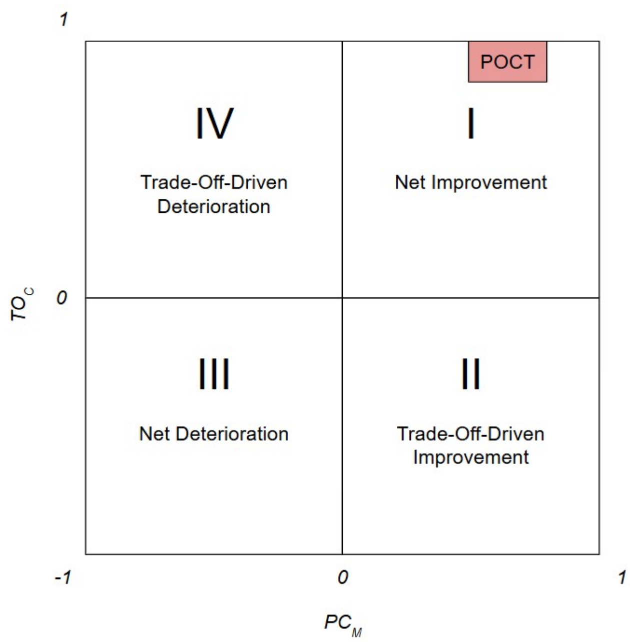

Figure 6.10: Placement of the POCT on the Performance Change Matrix ($PC_M = +0.67$ & $TO_C = +1$)

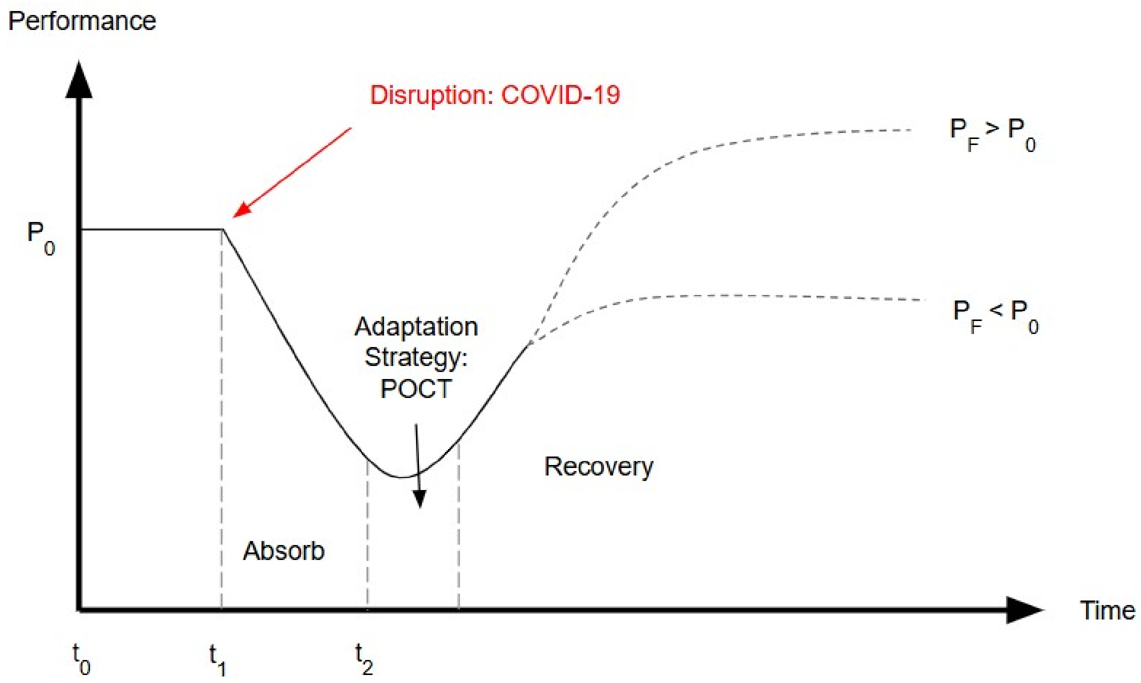

Figure 6.11: Illustrative resilience curve for the POCT implementation.

6.3.4. IOM Dimension Trade-offs

Drawing on the analysis above, we observe clear *Timeliness* gains under POCT. Time to result improves strongly at the diagnostic step, and this translates into shorter end-to-end ED LOS and a higher proportion of visits completed within 4 hours. Together, these indicators show that decisions about disposition are made earlier in the throughput phase, supporting timely flow through the ED. Even though "patients discharged within 4h" is not weighted in the PC_M calculation to avoid double-counting LOS effects, it still corroborates the same timeliness improvement from a benchmark perspective.

For *Effectiveness*, performance improves via higher diagnostic coverage and more targeted admissions. POCT ensures that many more discharges are backed by a documented test result, reducing uncertainty about COVID-19 status at ED exit. At the same time, negative results are used to place patients directly in appropriate non-COVID wards on first admission, indicating that the ED's diagnostic work more reliably supports correct downstream care decisions. These changes align with the "correctly diagnosing" component of the critical function.

Regarding *Safety*, the substantial reduction in intrahospital transfers suggests fewer harmful or unnecessary ward moves resulting from incorrect initial placement. Safety is also indirectly strengthened by the improved diagnostic coverage and targeted admissions, which reduce the chances that undetected infections or misplaced patients create additional risk. In contrast, *Efficiency* benefits are present but more modest and context-dependent: the lower hourly occupancy rate points to less crowding and better use of ED space, yet part of this effect is explained by reduced inflow (fewer patients per day), so only a fraction of the efficiency gain can be attributed directly to POCT.

In conclusion, the POCT adaptation represents a successful adaptation in resilient performance terms: by introducing a faster point-of-care testing step into the diagnostic process, the combined indicator pattern places the case in the first region of the Performance Change Matrix, indicating that performance moves into a recovery phase towards P_0. Within this region of the matrix, the quality-dimension pattern reads as improved *Timeliness*, *Effectiveness*, and *Safety* without deterioration in any other IOM dimension; the small, partly attributable improvement in *Efficiency* complements these gains. Consistent with this, the positive performance change breadth ($PC_B = 5.5$) indicates a net breadth of improvement across the weighted indicators, in line with the absence of observable trade-offs ($TO_C = 1.0$).

Chapter 6 synthesised the indicator analysis for each work adaptation into resilience readings: it calculated the performance-change measures (PC_M, PC_B, and TO_C), positioned each adaptation on the Performance Change Matrix, and unpacked the associated IOM quality-of-care trade-offs. Chapter 7 now turns to the discussion, bringing these case results together, reflecting on the methodological choices behind the framework, and drawing out implications for ED practice.

PART IV

Synthesis and Outlook

7

Discussion

Resilience was defined in this thesis as the intrinsic ability of a system to adjust its functioning before, during, and after disruptions in order to sustain required operations. Resilience engineering emphasis's the ability of systems to anticipate, absorb, adapt to, and recover from disruptions. In this thesis, the focus is on how the last three processes—absorption, adaptation, and recovery—are reflected in changes in performance over time. These processes are visualized in the resilience curve (see Figure 7.1), where system performance is plotted against time. The curve depicts a period of stable performance at a desired level P_0, a performance drop when a disruption occurs at t_1 (absorption), a subsequent phase in which the system reorganizes its functioning (adaptation), and finally a period of performance improvement back towards the desired level (recovery). Thus, sustaining required operations would ideally correspond to the resilience curve remaining close to the level P_0 despite the disruption introduced at time t_1. The more resilient the ED, the smaller and shorter the deviation from this P_0 would be, which would be seen as a shallower dip between t_1 and t_4.

Patriarca et al. (2017) describe resilience engineering as the study of what resilient performance looks like, how it can be measured or assessed, and how it can be improved. This thesis can be read as an attempt to operationalize that agenda for the emergency department: it specifies what resilient performance means in terms of sustaining required ED operations, proposes a way to measure and assess this using performance change scores derived from routine indicators, and offers initial reflections on how such assessments could inform improvement. The empirical material available for this work did not capture the entire resilience curve from pre-disruption baseline through the initial drop: the case studies report performance indicators only for periods before and after specific work adaptations, at a time when the COVID-19 disruption had already occurred and performance had already degraded. As a result, the thesis mainly observes the adaptation and recovery segments. Within this scope, the proposed framework quantifies how strongly, how widely, and with how much trade-offs ED performance changes after a work adaptation is introduced. In other words, it provides a structured way of reading how performance moves along the adaptation and recovery segments of the resilience curve while anticipation and the initial absorption phase remain unobserved in the available data. If time-resolved data from the pre-disruption and early disruption phases were available, the same logic could be used to characterize the initial drop and the full recovery trajectory. The performance change matrix introduced in the methodology already hints at this broader potential: some score combinations reflect clear deterioration in performance and can be interpreted as signatures of disruption impact and limited absorption capacity.

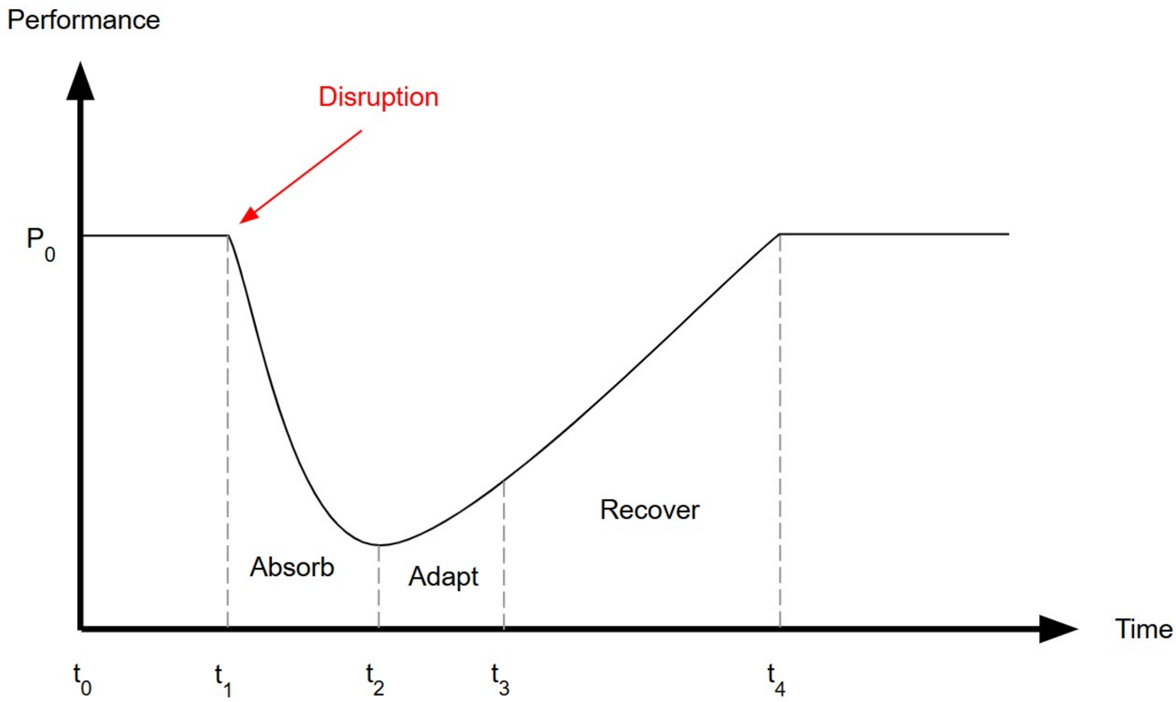

Figure 7.1: Resilience Curve: System performance trajectory across disruption, absorption, adaptation, and recovery

The first sub-question addressed what these "required operations" are in the context of an ED. By reconstructing a generic ED workflow and formulating an explicit critical function, the study described the system whose resilience is being assessed. The workflow breaks down the main phases and transitions of emergency care, while the critical function captures what the ED must achieve for patients to be considered adequately served. Together, these definitions clarify which parts of ED functioning matter most for performance and resilience assessment, and provide the reference against which any disruption or work adaptation is interpreted in the rest of the thesis.

The second sub-question then focused on how to describe changes in ED performance in a consistent way. In answering this sub-question, the first phase of the PI-RA framework organized performance indicators by IOM quality dimension and Donabedian category, and added a standardized treatment of direction and magnitude of change. International case studies of ED work adaptations during COVID-19 were processed by first summarizing the adaptation, then classifying the reported indicators, translating percentage changes into magnitude bins, and interpreting the mechanisms and context. This turns a heterogeneous set of indicators into a comparable description of how a work adaptation affects different aspects of ED functioning. It also makes explicit where trade-offs appear, for example when improvements in timeliness coincide with deterioration in equity or efficiency.

The third sub-question brought these elements together to assess resilient performance more directly. Building on the first phase, the second phase of PI-RA mapped indicator changes onto the ED workflow and critical function and assigned relevance weights based on how credibly each change could be attributed to the work adaptation. The third phase of PI-RA then aggregated these weighted indicator changes into the performance change magnitude, performance change breadth, and trade-off coefficient scores and positioned them on the

Performance Change Matrix and along the resilience curve. In this way, the scores summarize how strongly the required operations improved or deteriorated, how broadly the effects of a work adaptation spread across indicators, and how much trade-off is involved to achieve these changes. Applied to the three cases, they distinguish between adaptations that lead to net positive changes and one where improvements were accompanied by deterioration elsewhere, providing a compact view of how the ED performed under strain in relation to its defined critical function.

Taken together, these three sub-questions move from defining the system and its required operations, to organizing case evidence on how performance changes, to reading those changes as expressions of resilient performance in the adaptation and recovery parts of the resilience curve.

The remainder of this chapter is organised accordingly. Section 7.1 compares the three work adaptations in terms of their effects on ED performance and their position on the resilience curve. Section 7.2 reflects on the methodological choices and limitations of the PI-RA framework, including the treatment of performance change magnitude and breadth. Section 7.3 then discusses how the framework can be used in practice and further developed, **BLA BLA BLABLA BLA BLA BLA BLA - COME BACK HERE**

7.1. Cross-case comparison of the three work adaptations

This section compares the three work adaptations along three dimensions: their location in the ED workflow, their impact on quality-of-care patterns across the IOM dimensions, and their performance change profiles, summarized by the magnitude, net breadth, and trade-off scores PC_M, PC_B, and TO_C. Together, these dimensions show where in the patient journey each adaptation intervenes, which aspects of quality it affects, and how strongly and with what trade-off pattern it shifts the ED's position on the resilience curve as represented in the Performance Change Matrix. Table 7.1 summarizes this information. The following subsections then unpack each of these comparison dimensions in more detail to show how different ways of adjusting ED functioning during disruption translate into different forms of resilient performance.

7.1.1. Location in the ED workflow and type of intervention

Placing the three work adaptations on the ED workflow shows that they adjust different parts of the patient journey, and therefore influence different elements of the ED's critical function. The enhanced isolation protocol for FRPs acts in the front-end of the workflow, in Phase 1, immediately after triage, by inserting an additional placement step for patients with fever or respiratory symptoms. RAZ also acts in the front-end, but instead of adding a new queue it creates an alternative assessment path for eligible, non-critical patients. POCT, in contrast, is introduced in Phase 2 at the diagnostic step in the middle of the throughput, with effects that propagate into the back-end of the workflow when patients are admitted or discharged.

This positioning matters for how the adaptations affect the critical function. The FRP protocol modifies how quickly and where patients can be placed in a room at the front-end, which directly interacts with the rapid assessment function and with the way flow is managed when there is limited isolation capacity. Because the isolation area has fixed, restricted capacity, the extra step mainly acts as a filter and can easily become a bottleneck. RAZ modifies the same phase of the workflow but in a different way: it moves the first provider assessment earlier for a large subset of patients and partly decouples that rapid assessment from physical bed availability, which directly supports the rapid assessment function and front-end flow management. POCT does not change the front-end flow; instead, it strengthens the correct diagnosis function in the middle of the throughput by providing earlier and more complete test results, which then influence back-end disposition and ward placement decisions.

Viewed together, the three cases illustrate how two adaptations that are both located at the front end of the workflow (FRP protocol and RAZ) can have very different consequences because one inserts a capacity-limited safety buffer while the other creates additional assessment capacity. They also show that a diagnostic intervention further downstream (POCT) can improve not only diagnostic effectiveness but also confidence of disposition decisions later in the process.

Mapping the adaptations to phases and workflow elements helps to link them to specific functions (rapid assessment, correct diagnosis and flow management) and makes it easier to connect performance indicators to the parts of the workflow they represent. When changes in indicators are interpreted in that way, they say something about how much the adaptation helps the system to sustain or improve its required operations. The FRP case also shows the limits of this connection: the indicators clearly reveal a strong inequality between FRPs and non-FRPs, but this inequality is not directly anchored in a specific workflow element or critical function failure, so it remains difficult to translate it into a clear improvement or deterioration of resilient performance, even though it is an existing quality trade-off.

Table 7.1: Cross-case summary of the three work adaptations

Work adaptation	Workflow location	Critical functions targeted	IOM pattern (gains / losses)	PC_M	PC_B	TO_C	Resilient Performance Assessment
POCT	Phase 2 diagnostics (middle of throughput); affects back-end disposition and ward placement	Correct diagnosis; flow management	Timeliness ↑; Effectiveness ↑; Safety ↑; Efficiency ↑; no observed deterioration in other measured dimensions	+0.67	5.5	1	Moves the ED onto a recovery trajectory without any observed trade-offs in quality
RAZ	Phase 1 front-end assessment path after triage	Rapid assessment; front-end flow management	Timeliness ↑; Safety ↑; no detectable deterioration in other measured dimensions	+1.00	3	1	Moves the ED onto a recovery trajectory without any observed trade-offs in quality
FRP isolation protocol	Phase 1 front-end isolation / room-placement process after triage	Rapid assessment for FRPs; front-end flow management; infection-control emphasis	Intended infection safety ↑ inside ED (not directly measured); some end-to-end Timeliness and Efficiency ↑; front-end Timeliness ↓; Safety for LWBS patients ↓	-0.0375	0	0	Fails to moves the ED onto a recovery trajectory and introduces trade-offs in quality

7.1.2. Quality trade-offs across IOM dimensions

In the FRP isolation protocol, the intention is clearly to strengthen patient safety in the context of a highly infectious disease. The protocol separates suspected FRPs from other patients and keeps them in a controlled area, which is meant to reduce contagion risk inside the ED. However, by routing many more fever or respiratory-symptom patients through the same isolation process without a corresponding increase in isolation capacity, the protocol creates a bottleneck that brings this form of safety at the cost of other dimensions. Timeliness deteriorates for FRPs, who wait longer before being seen. Some end-to-end timelines and efficiency are gained due to decreased load due to increased LWBS. Equity is compromised, as FRPs and non-FRPs experience very different waiting times and LWBS probabilities. The trade-off pattern is therefore one where a presumed gain in infection safety is balanced by worse front end timeliness and equity. At the same time, equity was not explicitly included in the ED critical function and therefore did not enter the scoring directly, so this inequity appears mainly as an additional quality concern alongside the quantified assessment of resilient performance. This trade-off-heavy pattern is reflected in the trade-off coefficient ($TO_C = 0$), which shows that improvements and deteriorations balance each other across the scored indicators.

RAZ is designed explicitly around timeliness: bringing the first provider contact forward for eligible patients, regardless of bed availability. This primarily strengthens timeliness for non-critical patients at the front-end and reduces the risk that they leave without being seen, which also provides safety benefits. RAZ therefore functions as a straightforward example of a work adaptation that yields timeliness gains, with an accompanying safety benefit through fewer LWBS, without detectable losses in the other measured IOM dimensions. This aligns with the maximum trade-off coefficient ($TO_C = 1.0$), indicating that none of the scored indicators deteriorate.

By providing earlier and more reliable diagnostic information, POCT reduces diagnostic uncertainty at the middle of the throughput and supports more appropriate ward placement decisions at the back-end. In IOM terms, this strengthens effectiveness (more patients receive the test and are placed in the right ward on the first attempt) and safety (fewer intrahospital transfers). Timeliness improves because ED LOS shortens and results are available sooner. Efficiency gains appear in the form of reduced hourly occupancy; although part of this reduction is influenced by lower inflow, it still suggests a more efficient use of ED capacity under the new diagnostic process. POCT therefore functions as a straightforward example of a work adaptation that yields gains in timeliness, effectiveness, safety and efficiency, without detectable losses in the other IOM dimensions. The same pattern is visible in the trade-off coefficient ($TO_C = 1.0$), again indicating that all scored indicators improve without observable trade-offs.

7.1.3. Resilient performance profiles across the three work adaptations

This subsection compares the three work adaptations in terms of their resilient performance: how strongly they change the required operations, how many weighted indicators are involved in that net change, and how much trade-off is required to achieve it. To support this cross-case reading, Figures 7.3-7.2–7.4 compile, for each adaptation, its placement on the Performance Change Matrix together with the illustrative resilience curve already discussed in Chapter 6.

Across the three work adaptations, these placements reveal two distinct patterns of resilient performance. The FRP isolation strategy represents a response that does not bring the ED into recovery: its point on the matrix lies close to the origin but on the deteriorating side (Figure 7.2a), with a small negative performance change magnitude ($PC_M = -0.0375$), zero net breadth ($PC_B = 0$), and a trade-off coefficient of zero ($TO_C = 0$). Taken together, these scores indicate that the required operations remain in a degraded post-disruption state, with small improvements in some indicators offset by deteriorations elsewhere. Because $|PC_M|$ is very small, this pattern should be interpreted cautiously – the slight net deterioration may partly reflect the chosen binning and scoring scheme – but it is clear that the FRP strategy does not shift performance back towards the required operations level P_0 (Figure 7.2b). By contrast, the RAZ and POCT adaptations are associated with clear net improvements in the required operations. Their performance change magnitudes are clearly positive ($PC_M \approx +0.60$ for RAZ and $PC_M \approx +0.67$ for POCT), their net breadth scores are positive ($PC_B = 3$ for RAZ and $PC_B = 5.5$ for POCT), and both have a maximum trade-off coefficient ($TO_C = +1.0$), meaning that all scored indicators move in the desired direction. In the resilience curves, this corresponds to trajectories where performance turns upward after the adaptation is introduced, rather than remaining on a lowered plateau (panels (b) of Figures 7.3 and 7.4).

Because PC_M summaries the net change between two measurement windows, these magnitudes are implicitly time-bound. In the RAZ case, the before–after comparison spans two six-month periods, whereas in the POCT case it spans two much shorter windows of a few weeks during a single COVID-19 wave. If one were to translate these net changes into an explicit rate of change over time, a similar or slightly higher PC_M over a much shorter interval would correspond to a steeper recovery segment on the resilience curve. At the same time, the available data do not provide daily or monthly trajectories, so it is not possible to determine whether RAZ might have produced similarly strong improvements over a shorter interval after implementation. The point here is not to rank RAZ and POCT by recovery speed, but to highlight that the performance change measures are tied to the length of their observation windows, and that making this time dimension explicit would be a natural next step when using such scores to reconstruct resilience curves more directly. Section 7.3 sketches how this idea could be developed into an approach for constructing empirical resilience curves from performance-change scores and, ultimately, for quantifying resilience more directly.

Along the breadth dimension, the three cases show different profiles. For FRP, $PC_B = 0$ indicates that, in weighted terms, the net number of improving and deteriorating indicators is balanced, consistent with $TO_C = 0$ and the trade-off-heavy quality-of-care pattern described earlier. For RAZ and POCT, the positive breadth scores ($PC_B = 3$ and $PC_B = 5.5$) show that several indicators improve on net, while $TO_C = +1.0$ confirms that these gains are not realized through observable trade-offs: none of the scored indicators deteriorate. Comparing these two cases, the higher PC_B for POCT suggests that it improves a broader set of required operations than RAZ. However, this difference should be treated cautiously. First, PC_B is calculated over a relatively small, study-specific indicator set that does not capture all aspects of ED functioning. Second, the POCT profile combines results from two related case episodes, which increases the apparent breadth of impact. In this sense, the comparison between $PC_B = 3$ and $PC_B = 5.5$ is best seen as illustrative of how the score can be used, rather than as a definitive statement about which adaptation is broader in practice. More generally, because PC_B is expressed in weighted indicator units, it provides a sense of how extensive the net improvement is across the measured aspects of ED functioning, but cross-case differences in absolute breadth should be interpreted as indicative rather than precise here. In practice, the joint use of PC_M, PC_B, and TO_C offers a compact summary of these aspects of resilient performance—the direction and strength of change, the breadth of impact, and the reliance on trade-offs—and can in principle be applied to a wide range of changes in ED functioning, including both disruptive event itself and the consequent work adaptations, provided the indicator set and weights are chosen carefully.

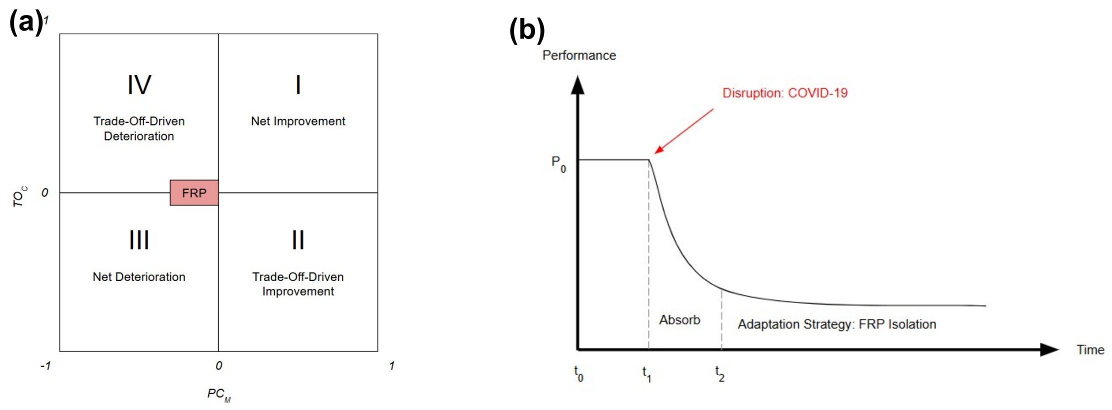

Figure 7.2: Placement of the FRP isolation strategy on the Performance Change Matrix (a) and illustrative resilience curve for the FRP isolation strategy (b).

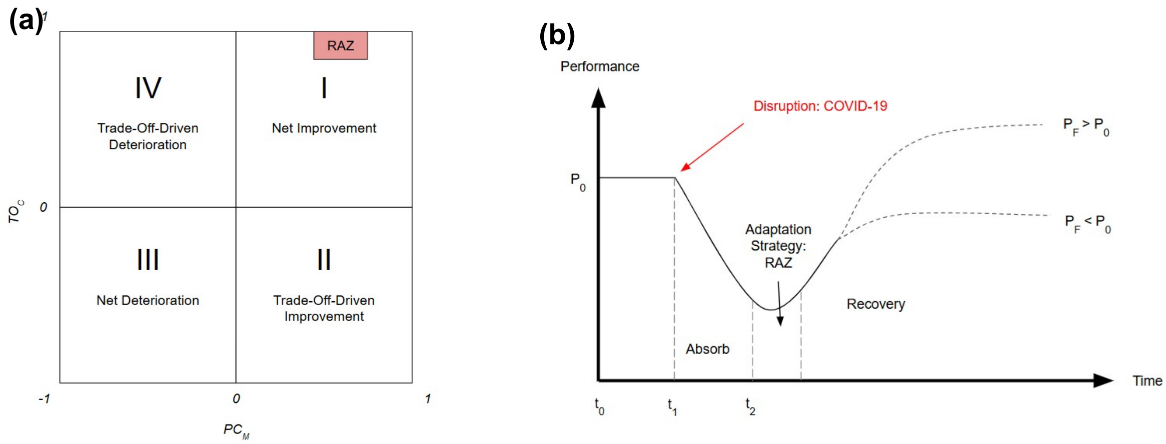

Figure 7.3: Placement of the RAZ on the Performance Change Matrix (a) and illustrative resilience curve for the RAZ (b).

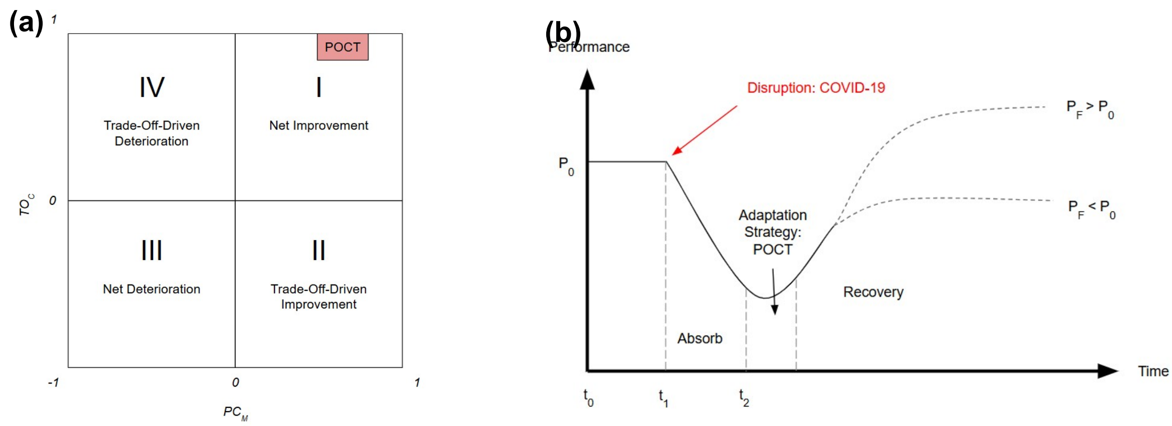

Figure 7.4: Placement of the POCT implementation on the Performance Change Matrix (a) and illustrative resilience curve for the POCT implementation (b).

7.2. Methodological Reflections: Strengths and Limitations

The cross-case comparison showed that the framework can distinguish between different patterns of resilient performance, but it also relies on several methodological choices that need refinement. In particular, the way indicators are weighted, how magnitudes of change are binned, and how classifications such as Donabedian category and IOM dimension are used all shape the performance-change measures PC_M, PC_B, and TO_C. This section reflects on these choices, what they enabled, and where they limit the interpretation of the results.

7.2.1. Indicator weighting and mixed attribution

The performance-change measures PC_M, PC_B, and TO_C all depend on how strongly each indicator is assumed to be influenced by the work adaptation, In the current application, this influence is expressed through a simple relevance weight w_i taking values 0, 0.5, or 1. Indicators with weight 1 are treated as fully attributable to the adaptation: changes in these indicators are interpreted as direct impact of the work adaptation. Indicators with weight 0.5 are judged to be partly driven by the adaptation but also by other factors. Indicators with weight 0 are treated as not being meaningfully affected by the work adaptation itself; their changes are interpreted as mainly driven by other external factors and therefore do not contribute to the performance-change measures.

This step is necessary because some indicators can be influenced by other factors or mechanisms not affected by the work adaptation. For example, in the POCT case, the reduction in hourly occupancy is affected by fewer patients presenting during the study period. In the FRP case, the change in occupancy reflects both the protocol and broader demand fluctuations. Without some adjustment for mixed attribution, the framework would either attribute all of these changes to the work adaptation or ignore them completely, both of which would be misleading.

At the same time, the $0/0.5/1$ scheme is based on informed judgement rather than data. It reflects a reasoned view about causal pathways in the ED workflow, but different analysts might make different calls for the same indicator. The approach is transparent and easy to apply, but it leaves open the question of how much of an observed change is really due to the adaptation and how much is due to other factors.

To address this mixed-attribution problem, the thesis proposes a multi variable regression-based refinement as explained in Section 8.3.2.

7.2.2. Magnitude bins and uncertainty about ``how big is big?"

A second methodological choice concerns how the framework treats the size of changes in indicators. There is no accepted way to say whether, for example, a 15% reduction in ED LOS corresponds to “a small” or “a large” amount of performance gain, or how that should be compared with a 40% reduction in waiting time or a 20% reduction in ED bed occupancy rate. The raw percentages are easy to compute, but they are not directly comparable across indicators that measure different constructs and have different practical implications.

In principle, it would be possible to calculate performance change scores directly from these observed percentage changes, treating a 25% reduction in ED LOS, a 40% reduction in waiting time and a 500% increase in LWBS rate as commensurable inputs. In practice, this would be misleading, because the same percentage change can correspond to very different performance effects across indicators. Some indicators can vary widely without altering the usefulness or safety of care, while smaller changes in others may already represent substantial gains or losses.

To avoid over-interpreting these raw percentages, the analysis uses five qualitative magnitude bins, ranging from “very small” to “very large”, defined by percentage thresholds that are chosen. Each observed percentage change is first mapped into one of these bins and then converted into a standardized magnitude score for use in PC_M. The bins therefore serve two functions. First, they spread uncertainty across labeled categories instead of pretending that there is a precise, linear mapping from percentage change to “amount of performance gain”. Second, they keep the performance change scores readable by preventing very large percentage changes in individual indicators from dominating the overall result.

At the same time, this binning approach is a clear limitation. The thresholds are judgement-based, and the same cut-offs are applied to all indicators, even though the practical meaning of a “medium” change in LOS is not the same as a “medium” change in LWBS or occupancy. Because the bins are not empirically derived, different analysts could reasonably choose different thresholds, and the resulting performance change scores might change. Around the bin cut-offs, even very small changes in the observed percentages can also push an indicator from one bin to the next, so minor fluctuations may flip the assigned magnitude category and alter PC_M. In principle, this could be mitigated by defining a continuous mapping from percentage change to magnitude instead of discrete bins, but such an approach is not further explored here. This makes it difficult to compare PC_M values in a precise way across work adaptations or across sites, and it is one of the reasons why we treated the low PC_M value for the FRP case with caution.

At present, this thesis does not offer a definitive solution to this problem. The magnitude bins are best seen as a pragmatic starting point rather than a finished method: they make the framework usable with the available data, but they do not resolve the underlying question of how big a given change in each indicator “ought” to be considered. Future work could test alternative threshold sets, calibrate indicator-specific bins using expert judgement, outcome-based models or simulation studies, or develop continuous mappings from percentage change to magnitude scores. More fundamentally, the current framework still lacks a strong quantitative link between the composite scores and the absolute scale of performance change: because raw percentage changes are first collapsed into qualitative bins, PC_M mainly acts as an ordinal index or coefficient rather than a direct measure of effect size. For applications that aim to reconstruct empirical resilience curves or compare effect sizes across settings, it would be necessary either to calibrate the bins against absolute outcome changes or to design an aggregation scheme that preserves more information about the size of the underlying changes. Until such a link is established, the magnitude dimension in the resilience scores should be interpreted as a relative indication of the magnitude of improvement or deterioration, rather than as a precise estimate of how much performance has changed.

7.2.3. Donabedian classification

In practice, most of the interpretive leverage in this thesis came from the IOM quality dimensions, from how indicator changes were positioned on the Performance Change Matrix and the resilience curves, and from the location of indicators in the ED workflow. Within this broader framing, the Donabedian classification (structure, process, outcome) played a supporting role.

First, it emphasized that structural indicators describe what the ED has rather than what it does. Structural features such as numbers of beds, physical layout or basic staffing models do not change from day to day and will typically receive an s_i score of 0. For that reason, they will usually not contribute directly to the performance change scores PC_M and PC_B. However, if a disruption or work adaptation explicitly alters structural features—for example by adding isolation rooms or reconfiguring triage space—structural indicators could in principle be included in the scoring as additional signals of how the ED's capacity to perform has changed. This logic creates a complication for any future attempt to express performance change per unit time: because structural features usually change once at implementation rather than gradually, they cannot be mapped as a daily trajectory in the same way as process indicators. In the three case studies analyzed here, some structural changes were implemented, but they were not reflected in the reported performance indicators, so this angle could not be explored further.

Second, the distinction between process and outcome indicators offers a useful angle on how indicators relate to the ED's critical function. Process indicators measure what happens inside the ED: they describe the sequence and timing of activities, the way resources are used, and how patients move through the system. Because they describe the functioning of the ED itself, they can be read more or less directly as performance signals and, when interpreted against the critical function, as signals of resilient performance in action: they show how the system is coping with demand and disruption in real time. Outcome indicators, by contrast, reflect the end results of the care process. These end states can be adverse or not (for example, complications, returns, or intrahospital transfers), which naturally connects some outcome indicators to the safety dimension of care. In addition, by contextualizing outcome indicators, they can also be used to say something about the effectiveness of the care that was provided, whether the relevant intervention is at the front end or the back end of the ED.

The POCT case illustrates this contextualization. There, outcome indicators such as admissions and discharges home were not just reported as crude totals, but were broken down by ward destination and testing status. Admissions were distinguished by which ward patients were admitted to, which turned the admission figures into a signal about correct placement and thus the effectiveness of diagnostic decisions. Discharges home were separated into patients who were tested and those who were not, providing a proxy for how thorough the diagnostic work-up had been before sending patients home. This example shows how outcome indicators, once enriched with clinical and pathway context, can be read as evidence about both the safety and the effectiveness of the care delivered, rather than as stand-alone counts. Similar contextualization could, in principle, be applied to other outcome indicators and to some process indicators as well.

Taken together, this suggests a more structured way of thinking about Donabedian categories in the assessment of resilient performance. Process indicators primarily reflect how the ED carries out its work in practice and therefore provide direct information about performance and resilience in action. Outcome indicators primarily reflect the consequences of that work and, when interpreted with appropriate context, provide information about the safety and effectiveness of the care. Both types contribute to understanding resilient performance, but they do so in different ways and at different points along the ED workflow.

7.2.4. IOM quality dimensions and the ED critical function

The IOM quality dimensions line up closely with how the ED critical function was defined in this thesis. The critical function – rapid assessment, correct diagnosis, initiating treatment and managing flow – already embeds three IOM dimensions as core aims. Rapid assessment and managing flow correspond directly to timeliness, while correct diagnosis and initiating treatment correspond to effectiveness. Safety appears as a non-negotiable boundary condition on all of these functions: the ED must avoid adverse events while it works quickly and effectively. In the empirical material, efficiency tends to show up more as a by-product than an explicit goal via bed occupancy, which reflect how tightly capacity is used when the critical function is under strain.

The relatively low representation of efficiency in the analysis is partly a consequence of how it was defined. Efficiency was defined as avoiding the waste of resources, where "resources" were mainly understood as beds, staff availability and physical capacity, not time. Conceptually, time can also be seen as a resource, and under that view any timeliness gain would also count as an efficiency gain: if patients spend less time in the ED and staff spend less time per patient, the same capacity is used more efficiently. However, if timeliness improvements had been systematically double-counted as efficiency improvements, almost every positive change in timeliness would have produced a parallel efficiency signal. This would have diluted the analysis and risked undermining the distinct role of the timeliness dimension. For this reason, the thesis treats timeliness as the primary dimension for time-related effects and uses efficiency more narrowly for capacity-related indicators.

Equity and patient-centredness do not line up with the critical function, but the FRP case shows how equity issues can still surface when indicator movements are examined by subgroup. Even though the observed inequity between FRPs and non-FRPs does not map onto the critical function, it remains a quality concern that sits alongside, rather than within, the assessment of resilient performance.. Patient-centredness is not represented in the analyzed cases, but given the way the critical function is framed, it is also not expected to impact assessment of resilient performance.

7.2.5. Locating work adaptations on the ED workflow and critical function

Locating the work adaptations on the ED workflow and linking them to the critical function turned out to be a key methodological step. It provided a concrete way to describe how each adaptation interacted with the system (which phase and workflow elements were affected); and to anchor performance indicators in those same parts of the system. This made it possible to interpret PI changes as more than abstract numbers: improvements or deteriorations could be read as gains or losses in rapid assessment, correct diagnosis, initiating treatment or managing flow, and then aggregated into performance change scores that summarize the magnitude and breadth of these changes for the required operations. In principle, an even more detailed, site-specific workflow map for each ED could deepen these links further by capturing configuration-specific pathways and local variations in how activities are organized, thereby allowing an even tighter connection between particular indicators and the processes around them. Developing such ED-specific workflow models was outside the scope of this thesis. Instead, the workflow map was kept at a generic level so that it would capture the ED critical function in a way that is recognizable and applicable across different departments, providing a common backbone for assessment of resilient performance while leaving room for future studies to add local detail where needed.

7.2.6. Reflection on PC_M, PC_B, and TO_C

The Performance Indicator Resilience Assessment (PI-RA) framework ultimately rests on three numerical measures: the performance change magnitude PC_M, the net breadth of change PC_B, and the trade-off coefficient TO_C. Taken together, these scores summarize three complementary aspects of how ED functioning changes between two measurement windows: how strongly the required operations improve or deteriorate on average (PC_M), how much improvements outnumber deteriorations (or vice versa) across the indicator set reflecting the net breadth of the impact (PC_B), and how far the net breadth of change is driven by improvements versus deteriorations (TO_C).

For PC_M, the main strength is that it compresses heterogeneous percentage changes into a single, readable index of net performance gain or loss. By mapping raw percentage changes into judgement-based magnitude bins and then computing a relevance-weighted mean, PC_M avoids having individual indicators with very large percentage swings dominate the overall result. Because the ED's required operations span several quality dimensions, this construction also makes it possible to combine timeliness-, safety-, effectiveness- and efficiency-related indicators into one index of "required operations" performance, which is essential for assessing the ED's critical function as a whole rather than dimension by dimension. At the same time, this construction means that PC_M should be interpreted as an ordinal indicator rather than a precise effect size. Because the bins are defined by thresholds and are applied uniformly across indicators, the score reflects a relative direction and strength of change rather than the exact magnitude of improvement in minutes of LOS, percentage points of LWBS, or other concrete units. Moreover, dividing by the sum of non-zero relevance weights normalizes PC_M to a common range across cases, which is helpful for comparison and for plotting the scores on the Performance Change Matrix, but also further distances the score from the absolute scale of change: for example, a single indicator improving by 70% and ten indicators each improving by 70% will yield the same PC_M, because the score averages the weighted magnitudes rather than accumulating their total size. A related limitation is that PC_M treats all included IOM quality dimensions as equally important as long as their indicators can be traced back to the critical function. Equity and patient-centredness were excluded on that basis, but within the remaining dimensions safety would in principle deserve more weight than, for example, efficiency. In practice, however, reductions in timeliness, effectiveness or efficiency can all undermine safety, so differentiating these dimensions in terms of importance is not only difficult but may be infeasible; at the level of PC_M, they are therefore handled symmetrically. The same symmetry also applies to improvements and deteriorations: a deterioration of a given magnitude moves the score down as much as an equal-sized improvement moves it up, even though real-world value judgements may be more loss-averse. In this sense, PC_M behaves more like a direction-and-intensity coefficient than a direct measure of how much performance has changed in concrete units. As with PC_B and TO_C, one way to retain more information about absolute magnitude would be to complement the normalized PC_M with an unnormalized companion score, so that both the direction-and-intensity coefficient and the underlying total size of change remain visible.

Unlike for the magnitude dimension, where normalization currently obscures the absolute size of change, the breadth of performance change is already captured through a complementary pair of scores, PC_B and TO_C. By construction, PC_B is an absolute breadth measure: it sums the relevance weights of improving indicators and subtracts the relevance weights of deteriorating indicators, so that the score (in weighted indicator units) captures the net breadth of improvement across the measured aspects of ED performance. This makes PC_B useful for seeing how extensive the net impact is in terms of affected indicators. The trade-off coefficient TO_C then adds a second lens by normalizing this net breadth by the total relevance weight, providing a ratio-type measure that indicates whether improvements dominate, deteriorations dominate, or gains and losses largely balance each other. Scores of TO_C close to $+1$ indicate improvement without observable trade-offs, values near zero indicate trade-off-heavy patterns, and negative values indicate that deteriorations dominate. Because PC_B and TO_C are built from the signs of the indicator change scores rather than their magnitudes, they deliberately separate breadth and trade-off information from the intensity information that is already summarized in PC_M. In that sense, PC_M and the pair (PC_B, T_{OC}) capture two complementary dimensions of the same underlying data: PC_M answers how strong the net change is, while PC_B and TO_C show how that change is distributed over improvements and deteriorations. This separation is also what makes it meaningful to display PC_M and TO_C on the orthogonal axes of the Performance Change Matrix in the next subsection, where their combined position visualizes both the direction-and-intensity of performance change and the extent to which it is driven by trade-offs.

Because all three scores are derived from the same indicator set, any change in which indicators are included or how they are weighted will propagate directly into PC_M, PC_B, and TO_C. The interpretation of these scores is therefore only as robust as the underlying indicator selection and weighting decisions. Taken together, the reflections above suggest that PC_M, PC_B, and TO_C are most useful as compact summary tools to describe how ED performance changes around disruptions and accompanying work adaptations, and *to distinguish different qualitative profiles of resilient performance.* Their main value lies in organizing and visualizing the underlying indicator changes in a consistent numerical form, so that these numerical profiles can be interpreted alongside the IOM dimension patterns, the workflow and critical-function mapping, and the case narratives developed in this study to obtain a more complete picture of resilient performance.

7.2.7. Reflection on Performance Change Matrix

The Performance Change Matrix combines the performance change magnitude PC_M and the trade-off coefficient TO_C into a two-dimensional representation of how an ED's required operations change between two measurement windows. The horizontal axis captures the direction-and-intensity of net change through PC_M, while the vertical axis captures how dominant improvements or deteriorations are across the weighted indicator set through TO_C. This way, the matrix separates the question of *how strong* the net change is from the question of *how trade-off-heavy* that change is. In this sense, the Performance Change Matrix does not simply re-plot two scaled versions of the same information, but combines structurally different aspects of performance change—direction-and-intensity along one axis and reliance on trade-offs along the other—while at the same time providing a compact visualization that makes the numerical scores more readable and supports cross-case comparison and classification of performance-change profiles beyond the three work adaptations analyzed here.

At the same time, the matrix should be interpreted as a single before–after summary of resilient performance. Each point is derived from one pre-adaptation and one post-adaptation measurement window, so the matrix only reflects the net change between these two periods and does not show how quickly the system moved towards that point, whether recovery occurred in a short burst or gradually, or whether performance overshot and then leveled off. In that sense, the matrix provides a structured way to interpret how the resilience curve is likely to shift, but it does not reveal the detailed trajectory of that curve. Moreover, because both PC_M and TO_C are normalized, ordinal indices (as discussed above), they do not convey absolute performance levels or absolute breadth of impact. The Performance Change Matrix is therefore most reliable for coarse, categorical comparisons: it can distinguish adaptations that end in net deterioration from those that end in net improvement, and those that rely on heavy trade-offs from those that do not. By contrast, small horizontal differences between points within the same region should not be interpreted as evidence that one adaptation improves performance “more” than another; for example, a slightly higher PC_M for POCT than for RAZ does not justify ranking POCT as the better intervention in absolute terms or in terms of recovery speed. However, the clear difference in sign and magnitude between the FRP profile and the RAZ/POCT profiles does support a more robust conclusion that RAZ and POCT can be regarded as performance-improving adaptations relative to the FRP strategy, even if the exact size of that advantage cannot be quantified yet.

One could in principle extend the representation into three dimensions by adding PC_B as a third axis to distinguish narrow from broad impacts within each region, but developing and testing such a visualization falls outside the scope of this thesis.

A further implication of the way PC_M and TO_C are constructed is that the quadrant labeling of the matrix is most robust for points that lie clearly away from the axes. When PC_M or TO_C are close to zero, small changes in the indicator set, relevance weights or bin assignments may shift a work adaptation from a “net improvement” to a “net deterioration” quadrant, or from a “trade-off-free” to a “trade-off-driven” quadrant. Near these boundaries, the matrix should therefore be interpreted together with the underlying indicator table and the IOM dimension patterns, rather than as a hard categorical verdict.

Overall, the Performance Change Matrix can therefore be seen as a structured way of translating the coefficient scores PC_M and TO_C into qualitative profiles of resilient performance. The four regions summarize whether the ED’s required operations stabilize in a degraded state, continue to drift downwards despite an adaptation, or move back into a recovery trajectory in which net improvements outweigh losses and trade-offs are limited. Given the present data and the construction of PC_M, this link to the resilience curve remains conceptual, but it still provides a useful language for comparing how different disruptions and work adaptations reshape ED performance.

7.3. Future use of the PI–RA framework

So far, the PI–RA framework has been applied to case studies where specific work adaptations were introduced in response to a disruption (in this case, COVID-19). This use already shows what the framework can do with relatively limited data: it allows us to assess whether a given adaptation manages to put performance “back on track” or not, whether this improvement involves trade-offs, and, if so, what kind of trade-offs these are. Because each performance indicator is mapped to IOM quality dimensions, we can see whether gains come mainly in safety, timeliness, effectiveness, or efficiency. Because the indicators are also mapped to the ED workflow and to the critical function model, we can locate where in the process the adaptation acts and what knock-on effects it has on other parts of the system. Taken together, this already gives a structured way of profiling resilient performance: we can read how an adaptation changes the required operations, what this implies for healthcare quality, and which parts of the ED are most affected.

The same logic can be used to analyse the *impact of the disruption itself*, before any deliberate work adaptation is in place. In such pre-adaptation windows, the scores mainly describe how the disruption pushes the ED away from its baseline functioning. Negative values of PC_M indicate net deterioration of the required operations relative to the pre-disruption period, while PC_B shows how broadly this deterioration is spread across the measured aspects of ED performance. The trade-off coefficient TO_C would, in principle, show whether deterioration in some indicators coincides with improvement in others. For a disruption phase, however, the trade-off interpretation is different from the work-adaptation case: we do not expect a disruption to create "good" trade-offs in the sense of deliberate improvements elsewhere. At the same time, we also do not always expect $T_{OC} = -1$, because a disruption can relieve pressure on some parts of the system even while it harms others. For instance, if a failure earlier in the care chain prevents certain patient groups from reaching the ED at all, crowding pressures may fall and ED LOS or bed occupancy may temporarily improve, even though the overall situation is clearly undesirable. TO_C is useful here because it can make such patterns visible: it shows whether deterioration is truly system-wide or whether some parts of the ED experience reduced pressure while others are overstretched. This perspective also connects to the notion of robustness in resilience engineering, where robustness is understood as the system's ability to absorb or withstand disruptions and thus as the absorption capability component of resilience (Yang et al., 2023).

The Performance Change Matrix was constructed with this dual perspective in mind. Regions III and IV do not only represent failed or harmful work adaptations, but also cover the situation in which no adaptation is present and we are simply observing the disruption's impact on the required operations. In that sense, the matrix can place both the disruption phase and any subsequent adaptations in the same space, making it easier to compare how far the ED is pushed away from its baseline and how far it is later pulled back.

These two applications—profiling the disruption and profiling the adaptation—are most informative when they are analyzed together across time. If performance indicators are measured and recorded at regular intervals (for example, weekly rather than daily, to smooth out case-mix and arrival-rate fluctuations), the PI–RA framework can be used to build a richer picture of resilient performance over the whole episode. With repeated measurements, we do not only know that performance changed between a single "before" and "after" point; we can also approximate how quickly the disruption impact unfolds, how fast recovery sets in once an adaptation is introduced, how deep performance drops at its lowest point, and at what level it stabilises afterwards. Because PC_M is a normalised, ordinal index, it will still not give the exact numerical size of the performance drop by itself, but knowing the baseline and post-episode PI levels provides more information about the depth of the valley and the height of the recovery.

A key strength of this way of using the framework is that scoring happens at indicator level, with each indicator mapped both to specific segments of the ED workflow and to the IOM quality dimensions. By examining which indicators contribute most to negative or positive values of PC_M and PC_B, deterioration or improvement can be traced back to concrete parts of the process and to particular dimensions such as safety, timeliness, effectiveness, or efficiency. This makes it possible to see not only *how much* performance is lost or regained during a disruption episode, but also *where* in the system those changes occur and which aspects of quality are most affected. Such profiles can inform which kinds of interventions are most promising (for example, whether an ED would benefit more from redesigning intake, reallocating diagnostic capacity, or addressing bed availability) and then provide a way to compare their impact using the same scoring scheme.

Finally, the PI–RA framework does not have to be limited to COVID-19 or even to disruption-and-adaptation pairs. Any change that affects ED performance and can be analyzed: policy reforms, operational changes, the introduction of new technology or new clinical pathways, and so on. For each such change, the framework can show how the required operations shift, which parts of the workflow are most affected, which quality dimensions move in the desired or undesired direction, whether improvements come with trade-offs, and whether the change appears worth sustaining from a resilience and quality-of-care perspective.

Data Implications

To apply the framework, sites should begin with the initial set of performance indicators provided in Appendix A. The set was selected for its relevance to the ED critical function, strength of evidence in the literature, and frequency of use in practice. It is not exhaustive but offers a coherent baseline that is readily implementable and comparable across settings. If regression-based weighting is pursued, accompanying covariates should be recorded to adjust for non-adaptation effects.

7.4. Extending performance indicators through contextual enrichment

The framework in this thesis works with standard indicators such as ED LOS, waiting time, LWBS and admissions. On their own, these indicators already support resilience assessment when they are mapped to the ED workflow and critical function. However, the POCT example and the additional material about it in Appendix C show that these indicators can be made much more informative by enriching them with contextual information about patient subgroups and care pathways. In practice, this means splitting standard indicators by clinically or organizationally relevant categories: for example, tested versus non-tested patients, ward A versus ward B, admitted versus discharged, or other groups defined by risk status, triage category or diagnostic profile.

This kind of enrichment serves several purposes. First, it makes visible which subgroups benefit from, or are disadvantaged by, a work adaptation or a disruption. The inequality pattern in the FRP case, and the difference between admitted and discharged patients in the RAZ LOS results, already hint at this: disaggregating indicators shows that not all patients experience the same effects. Second, when outcome indicators such as admissions and discharges are enriched by destination, testing status or diagnosis, they begin to reveal the effectiveness of the diagnostic and placement processes rather than just their volume. In the POCT case, distinguishing discharges by whether patients were tested, and admissions by ward destination, turned routine disposition figures into signals about correct diagnosis and the safety associated with the end results. Finally, similar contextual enrichments could be used to explore other questions, such as how disruptions or new policies affect particular vulnerability groups or how changes propagate through specific pathways in the ED.

Appendix C walks through the POCT enrichment in detail as an example of how such enriched indicators can be analyzed with the framework developed in this thesis. The main point for the methodological reflection is that contextual enrichment does not replace the existing indicators or scoring; it extends them. By adding subgroup and pathway information to routine measures, the same framework can say more about who gains or loses from a given change, how effective key processes are, and how safety and equity are distributed across the ED during both work adaptations and disruptions.

8

Conclusion

8.1. Answers to the research questions

This section answers the three sub-research questions (SQ1–SQ3) in order to address the main research question of this thesis: "How can observed changes in ED performance indicators be translated into a transparent assessment of resilient performance, and what do these assessments reveal about quality of ED care?" For each sub-question, it briefly summarizes the key findings and how they were obtained. Together, these answers provide the basis for the final, integrated answer to the main research question.

8.1.1. SQ1 -- ED workflow & critical function

In order to answer SQ1, "What are the ED's critical function and workflow?" this thesis first positioned Dutch EDs within a strongly gatekept acute care chain. Non-urgent complaints are filtered by general practitioners and out-of-hours cooperatives, while Emergency Medical Services perform protocol-guided assessment, treatment and transport decisions in the field. As a result, most clinical triage, urgency assignment, and much of the registration and documentation work are handled before ED entry, supported by standardized digital data exchange between GPs, EMS and hospitals. Because these input and Pre-ED activities are largely governed by primary care, EMS and national IT infrastructure rather than by the ED itself, they were excluded from performance assessment. Consequently, the ED's direct responsibility in this thesis was located in the throughput phases (Phase 1 and Phase 2) and in the post-ED/output phase, because these reflect the actions taken during ED care and the immediate outcomes of that care. The workflow steps and phase boundaries used in the analysis are shown in Figure 8.1. Within this workflow, the ED's critical function was defined as four tightly connected elements: rapid assessment, correct diagnosis, initiating treatment, and managing patient flow so that these activities can be sustained under variable demand and disruption. The analysis showed that this critical function is primarily carried by the throughput phases: rapid assessment is concentrated in Phase 1, diagnostic accuracy and treatment in Phase 2, while flow management links both phases to the output step where patients leave the ED. This clarified where "required operations" actually take place in the ED workflow and provided the basis for later performance and resilience assessment.

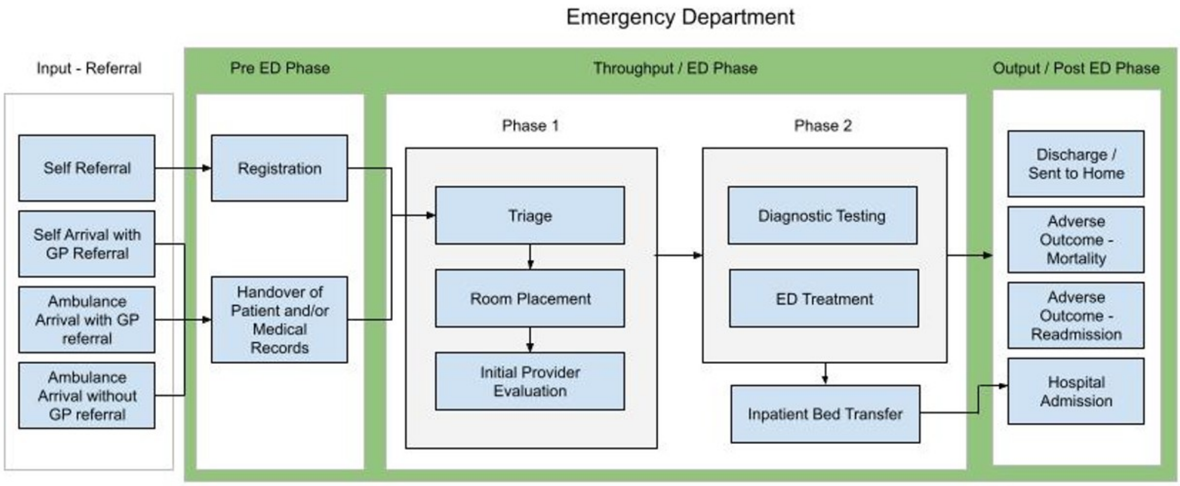

Figure 8.1: ED Workflow

8.1.2. SQ2 -- Impact of work adaptations on ED performance

In order to answer SQ2, “How can ED work adaptations in response to disruptions (e.g. COVID-19) be analyzed systematically to extract structured evidence and insights on performance?”, this thesis operationalized Phase 1 of the PI–RA framework (see Figure 8.2). Three international empirical cases were selected that described concrete organisational or process-level work adaptations in the ED and reported quantitative performance indicators before and after implementation: an enhanced isolation protocol for fever or respiratory patients (FRPs), the introduction of a Rapid Assessment Zone (RAZ), and the implementation of point-of-care testing (POCT) for SARS-CoV-2. For each case, the reported indicators and their pre/post values were extracted and the adaptation was summarized with a clear performance focus.

The Phase-1 steps were then applied in sequence. Each performance indicator was classified using the Donabedian structure–process–outcome scheme and mapped to one or more IOM quality dimensions, so that heterogeneous measures such as waiting times, length of stay, diagnostic coverage and left-without-being-seen rates could be interpreted on a common footing. The direction and magnitude of change for every indicator were coded into discrete bins (No change, Low, Medium, High, Very High) with an explicit sign for improvement or deterioration. Finally, these coded changes were condensed into a single IF–THEN statement per work adaptation that lists which indicators changed, by how much, and what those changes imply for different quality dimensions.

This procedure showed that ED work adaptations can be translated from diverse case-study reports into a standardized, comparable description of performance change. The resulting tables and IF–THEN statements capture both improvements and deteriorations across multiple dimensions and form the structured performance-change input that is carried forward into the system-mapping and resilient performance assessment in SQ3.

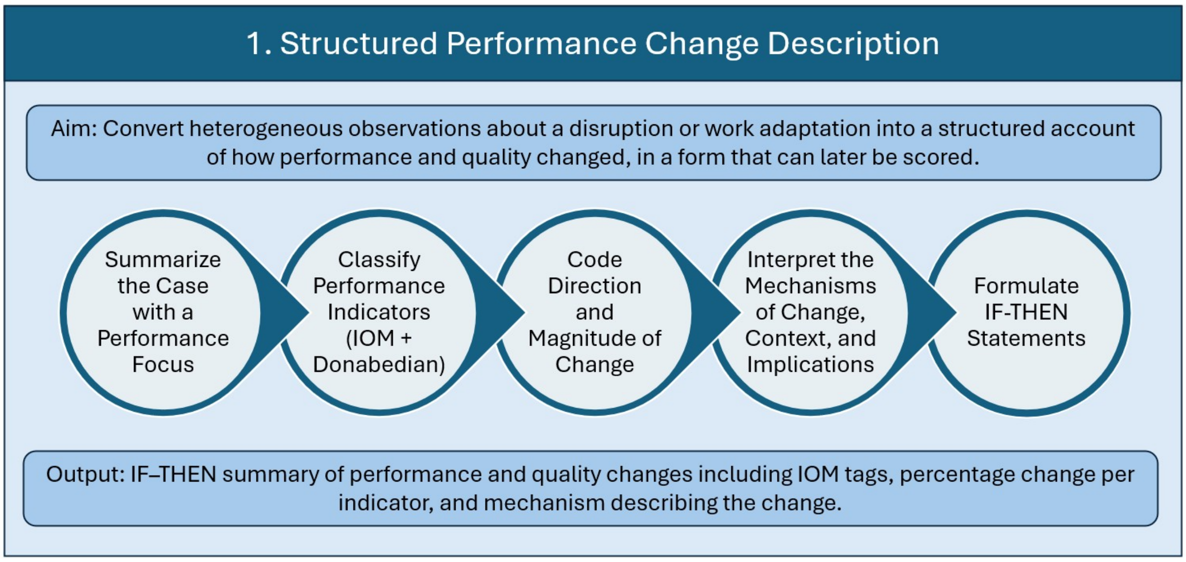

Figure 8.2: Phase 1 of PI-RA: Structured Performance Change Description

8.1.3. SQ3 -- Resilient performance and trade-offs

In order to answer SQ3, “How can the structured evidence and insights on ED performance be used to assess the resilient performance of different work adaptations?”, this thesis applied the second and third phases of the PI–RA framework. Building on the workflow and critical function from SQ1 and the structured performance-change descriptions from SQ2, Phase 2 first mapped each work adaptation onto the ED workflow and critical function. For each case, the analysis identified where in the workflow the adaptation acted, how it was expected to alter flows and bottlenecks, and which elements of the critical function it primarily affected. The performance indicators from the IF–THEN statements were then assigned to those workflow locations and critical-function components, while retaining their IOM and Donabedian classifications. On this basis, each indicator received an attribution weight reflecting how directly its observed change could be linked to the adaptation, and its coded direction and magnitude were converted into a signed indicator change score. Together, these steps anchored the indicator changes in the ED system and expressed them on a common numeric scale.

Phase 3 then translated these weighted indicator scores into explicit measures of performance change and trade-offs. Using the attribution weights and indicator scores, the analysis calculated the performance change magnitude PC_M as a relevance-weighted mean, summarizing the net direction and strength of change in the required operations. The performance change breadth PC_B captured how many weighted indicators improved or deteriorated on net, while the trade-off coefficient TO_C normalised this breadth to show whether improvements dominated, deteriorations dominated, or gains and losses largely balanced each other. Each work adaptation was positioned on the Performance Change Matrix using PC_M (horizontal axis) and TO_C (vertical axis), and this placement was interpreted alongside PC_B, the IOM quality tags, and an illustrative resilience curve. In this way, the framework provided both a compact numerical profile and a narrative account of how performance changed and which quality dimensions were traded off.

Applied to the three COVID-19 work adaptations, this procedure yielded distinct resilient-performance profiles. For the enhanced isolation protocol for fever or respiratory patients (FRPs), the scores clustered near the origin with a slightly negative $PC_M = 0.0375$, $PC_B = 0$

and $TO_C = 0$, indicating that small efficiency and timeliness gains were offset by deteriorations in front-end timeliness and safety, leaving the ED in a degraded, trade-off-heavy post-disruption state rather than on a clear recovery path. In contrast, the Rapid Assessment Zone (RAZ) and point-of-care testing (POCT) adaptations produced clearly positive PC_M values ($PC_M = +0.60$ for RAZ and $PC_M \approx +0.67$ for POCT), positive breadth scores ($PC_B = 3$ and $PC_B = 5.5$), and a maximum trade-off coefficient $TO_C = +1$, showing that all weighted indicators moved in the desired direction. These adaptations correspond to upward movements on the resilience curve in which required operations recover and improve—RAZ mainly through better timeliness and safety in early throughput, and POCT through broader gains in timeliness, effectiveness, safety and efficiency. Taken together, the second and third phases of PI–RA thus show how structured indicator evidence can be converted into a transparent assessment of resilient performance, distinguishing maladaptive, trade-off-heavy responses from work adaptations that genuinely support recovery without observable losses in other quality dimensions.

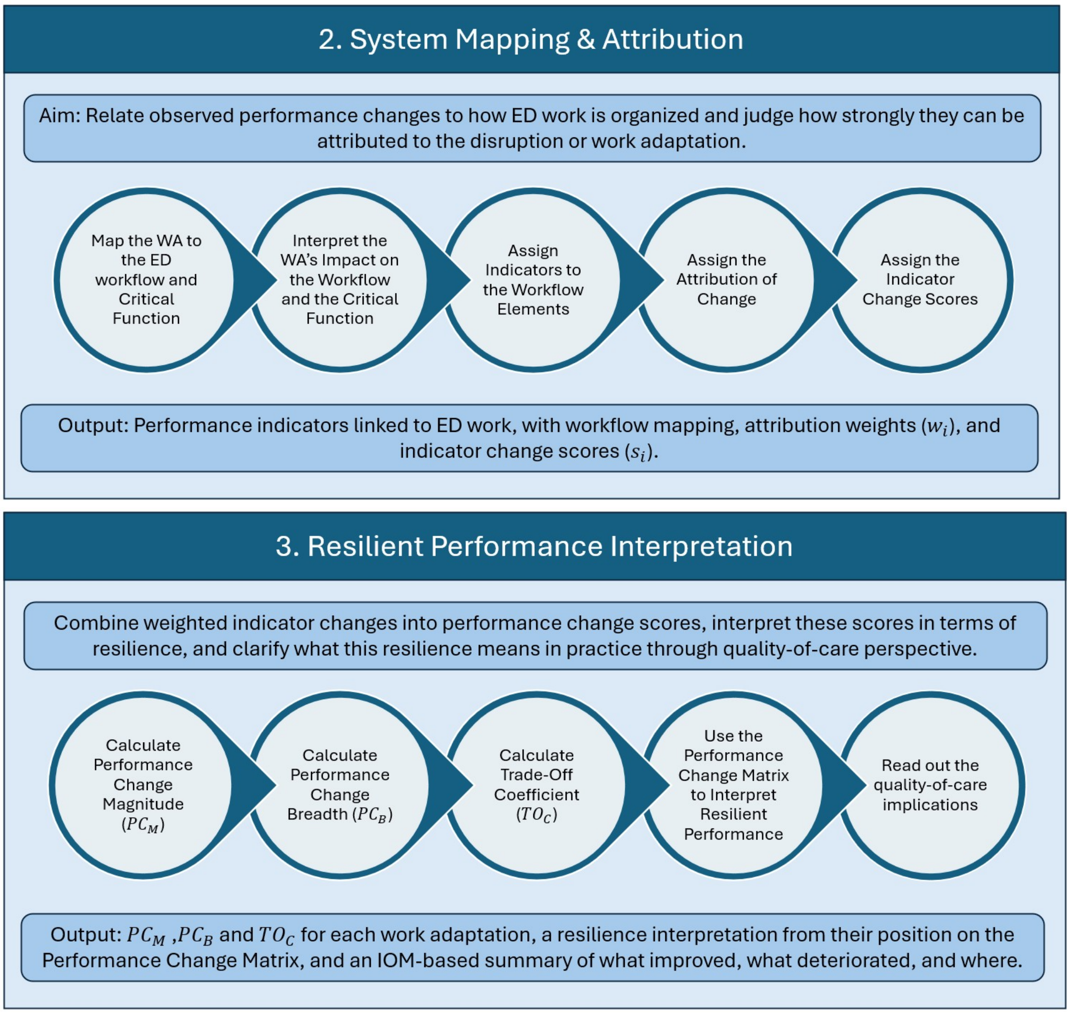

Figure 8.3: Phases 2 and 3 of PI-RA: System Mapping and Attribution & Resilient Performance Interpretation

8.1.4. Answer to the Main Research Question

The main research question asked how observed changes in ED performance indicators can be translated into a transparent assessment of resilient performance, and what such assessments reveal about the quality of ED care. This thesis shows that the PI–RA framework provides this translation. By (i) framing the ED workflow and critical function as the reference for “required operations”, (ii) structuring observed indicator changes and linking them to Donabedian and IOM quality dimensions, and (iii) aggregating them into the performance change scores PC_M and PC_B and the trade-off coefficient TO_C on the Performance Change Matrix, PI–RA turns heterogeneous indicator data into explicit resilient-performance profiles. These profiles make visible whether a work adaptation supports recovery or leads to degradation, and which dimensions of quality are gained or sacrificed, thereby answering the main research question.

8.2. Contributions

Conceptually, this thesis contributes by making the ED’s “required operations” explicit in a Dutch, strongly gatekept acute-care chain. It defines the ED’s critical function as rapidly assessing, correctly diagnosing, initiating treatment, and managing flow, and anchors these elements in a concrete ED workflow that focuses on the throughput and immediate output phases. By linking this critical function to the Institute of Medicine (IOM) quality dimensions, the thesis shows how routine performance indicators can be read as statements about how well the ED continues to perform its core task under disruption, rather than as isolated metrics.

Methodologically, the thesis develops the PI–RA framework as a coherent way of turning heterogeneous performance indicators into a structured assessment of resilient performance. PI–RA combines (i) a common coding grammar for indicator changes (Donabedian type, IOM dimension, direction and magnitude of change, and IF–THEN summaries), (ii) a system-mapping and attribution step that locates indicators on the ED workflow and judges how strongly they are affected by a given disruption or work adaptation, and (iii) composite performance-change scores that summarize the net direction, breadth, and trade-off balance of change (PC_M, PC_B, and TO_C) together with the Performance Change Matrix. In this way, the framework addresses the gap between qualitative resilience concepts and quantitative ED data by offering a transparent, empirically grounded way to profile how required operations change during disruption and adaptation, while keeping the underlying indicators and assumptions visible. The notion of contextual enrichment, introduced in the POCT case and Appendix A, further illustrates how standard indicators can be extended with clinically meaningful subgroups to make their contribution to resilience assessment more informative.

Practically, the thesis applies this framework to three COVID-19 work adaptations and demonstrates how EDs can use their own indicator data to evaluate where specific adaptations succeed, where they rely on trade-offs, and where they leave gaps. The FRP isolation protocol is shown to produce a trade-off-heavy, near-stagnant performance profile, whereas the Rapid Assessment Zone and point-of-care testing correspond to clearer improvement profiles with limited observable trade-offs. The initial indicator set in Appendix A and the worked examples in Chapters 5–7 provide a concrete template for practitioners who wish to map their own adaptations onto the ED workflow, score performance changes, and interpret the results in terms of resilience and quality of care.

8.3. Future research

This section outlines two complementary directions for future work. First, Section 8.3.1 sketches how PI–RA could be extended towards an empirical resilience curve and quantitative resilience measurement; second, Section 8.3.2 proposes more data-driven approaches to relevance weight estimation within PI–RA.

8.3.1. Towards Empirical Resilience Curve and Resilience Measurement

A first avenue for further work is to extend PI–RA from static performance profiles to a time-varying resilience curve and, in principle, a quantitative resilience index. Up to this point, the thesis has used the PI–RA framework mainly for resilient performance profiling: for each COVID-19 work adaptation, the indicators were scored once in a pre-adaptation and once in a post-adaptation window, and the resulting PC_M, PC_B and TO_C scores were interpreted as a structured summary of how the ED's performance changed and what kinds of quality-of-care trade-offs were involved. This already addresses a gap identified in the resilience engineering literature, namely that many studies describe resilience qualitatively while giving little guidance on how to express performance change and recovery in quantitative terms. Yang et al. (2023) explicitly point out that quantitative resilience assessment is still underdeveloped compared with the large body of qualitative work, and that resilience in engineering should be seen as a way of characterizing how system performance changes and recovers under disruptive conditions. In their notation, resilience can be represented as a triplet $R = (D, F, P)$, where D denotes the disruption, F the system functionality, and P the performance metrics used to measure that functionality over time.

This view aligns closely with the choices made in this thesis. In the empirical chapters, COVID-19 is treated as the main disruption; the functionality corresponds to the ED workflow and critical function modelled in Chapter 4.1; and performance is captured through PIs that measure timeliness, safety, efficiency and effectiveness that can be mapped both to this critical function and to specific segments of the ED workflow. Moreover, Yang et al. (2023) emphasize that any quantitative resilience model ultimately depends on having a performance profile over time: their framework "can only quantify the system resilience to the disruptions whose impacts on system performance can be measurable". The PI–RA framework can be seen as a starting point for a way of making such impacts measurable in a complex socio-technical setting like an ED; by mapping heterogeneous indicators to IOM quality dimensions and critical ED processes and combining them into interpretable composite scores.

Against this background, it is natural to ask whether PI–RA can be extended from static profiling to something closer to what Yang et al. (2023) term a quantitative resilience assessment. They summarize existing metrics into four broad types: metrics based on performance over a time period, metrics based on performance at a time instant, probabilistic metrics, and multiple-indicator metrics. In this thesis, resilience is understood as the behaviour of ED performance over the full disruption and recovery period: how far performance is pushed away from its baseline, how long this loss persists, and to what level it is restored.

Within this first type, Yang et al. (2023) describe two closely related formulations. One is "the ratio between actual and desired performance during the period"; the other is "the ratio of performance loss and desired performance during the period". Both refer to an underlying performance curve over time. Let the desired (baseline) performance during a disruption window $[T_0, T_f]$ be denoted $P_{\text{des}}(t)$ and the actual performance be $P(t)$. The first formulation can then be written as a ratio of areas under these curves:

$$R_{\text{attained}} = \frac{\displaystyle\int_{T_0}^{T_f} P(t)\,\mathrm{d}t}{\displaystyle\int_{T_0}^{T_f} P_{\text{des}}(t)\,\mathrm{d}t}. \tag{8.1}$$

This quantity, here denoted R_{attained}, expresses how much of the desired level of operations was actually delivered over the whole episode. If the system maintained baseline performance throughout, the numerator and denominator are equal and $R_{\text{attained}} = 1$. If performance drops and only partially recovers, the area under the actual curve is smaller and the ratio falls below one, indicating the fraction of required operations that were effectively sustained.

The second formulation can be interpreted as the fraction of desired operations that were *not* delivered during the episode, that is, the normalized area of performance loss. Algebraically, the two metrics are complements: expanding the numerator of the loss expression shows that

$$R_{\text{lost}} = 1 - \frac{\displaystyle\int_{T_0}^{T_f} P(t)\,\mathrm{d}t}{\displaystyle\int_{T_0}^{T_f} P_{\text{des}}(t)\,\mathrm{d}t} = 1 - R_{\text{attained}}. \tag{8.2}$$

Figure 8.4 schematically illustrates this relationship. The upper horizontal line represents the desired baseline $P_{\text{des}}(t)$ over the disruption window, while the broken curve depicts the actual performance $P(t)$, with a drop after the disruptive event and a subsequent recovery. The area under $P(t)$ (green area) corresponds to the numerator of R_{attained}, whereas the area between $P_{\text{des}}(t)$ and $P(t)$ (red area) represents the integrated loss in the numerator of R_{lost}. Together these two shaded regions fill the total area under the baseline curve, so that R_{attained} and R_{lost} indeed sum to one.

For the purposes of this thesis, these formulations are important not because they introduce new mathematics, but because they clarify what is missing from the current PI–RA analysis and what a next step could look like. PI–RA already provides a structured, multidimensional description of how ED performance changes between two periods and how those changes involve trade-offs across IOM quality dimensions, summarized in the scores PC_M, PC_B and TO_C. Yang et al. (2023) show that, if one can extend this description into a performance trajectory $P(t)$ over a disruption episode and specify an appropriate desired curve $P_{\text{des}}(t)$, then resilience can, in principle, be quantified as an area-based ratio such as R_{attained} or $1 - R_{\text{lost}}$.

Appendix D provides an exploratory illustration of what such an extension might look like. For the RAZ case, it constructs a hypothetical time series of PI–RA snapshots and aggregates the resulting PC_M scores over multiple periods to see whether performance would, in principle, return to its original level. This example also highlights why the current PI–RA scores cannot yet be used directly as $P(t)$: the binning of magnitudes (the "how big is big" problem) and the ordinal nature of P_{CM}, P_{CB} and TO_C mean that they do not represent a cardinal net performance level. If future work can develop a more quantitative mapping from indicator changes to such a level, the same logic illustrated in Appendix D could be used to construct empirical resilience curves and area-based resilience measures.

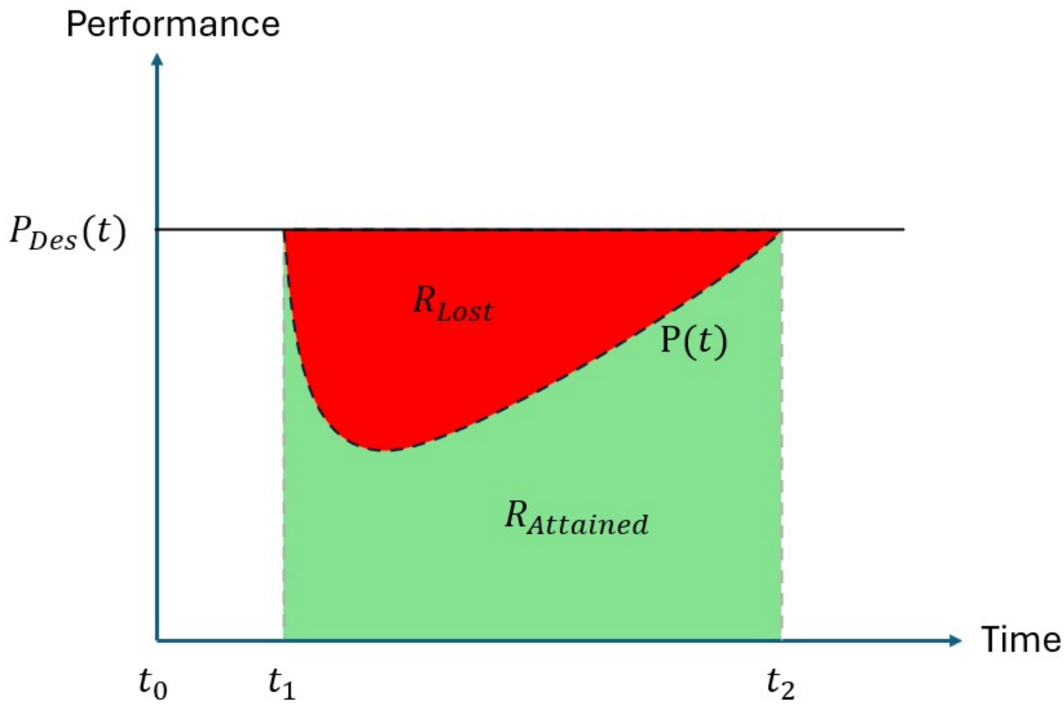

Figure 8.4: Schematic illustration of R_{attained} and R_{lost} as complementary area-based resilience measures. The desired performance $P_{\text{des}}(t)$ defines the total reference area; the area under the actual performance curve $P(t)$ corresponds to R_{attained} (green area), while the area between $P_{\text{des}}(t)$ and $P(t)$ represents R_{lost} (red area).

8.3.2. Multi Variable Regression Analysis for Weight Determination

To address the weighting problem discussed in Section 7.2.1, a key direction for future research is to explore how multivariable regression could be used to derive more data-driven relevance weights. Such work would need to identify which covariates (e.g. crowding measures, case-mix variables, seasonal patterns) are most relevant for different types of work adaptations, disruptions, or policy changes, and how these should be incorporated into a consistent modeling strategy. The overall aim would be to move from judgement-based 0/0.5/1 weights to empirically grounded estimates of how much of the observed change in a given indicator can plausibly be attributed to the intervention being studied.

In the POCT case, for example, the published evaluation reports that mean ED LOS dropped from 276 to 208 minutes after POCT was introduced, a crude reduction of about 68 minutes. In a regression model that adjusts for crowding and case mix, the independent effect of POCT on LOS is estimated at about 56.478 minutes. Taking the ratio of the model-estimated effect to the observed change gives a weight of approximately $56.478/68 \approx 0.955$. Interpreted as a relevance weight, this suggests that roughly 95% of the observed LOS reduction can be attributed to POCT, with the remaining 5% reflecting other influences.

Generalizing this idea, the relevance weight for an indicator could be defined as the proportion of the observed change that is explained by the work adaptation in a covariate-adjusted regression model. This would turn the current coarse weights (0, 0.5, 1) into continuous weights between 0 and 1. Because the weights would be derived from a specified regression model, the assumptions behind them would be visible; other analysts could try alternative models or sets of covariates and see how much the resulting weights change. If the weight stays roughly the same across reasonable model choices, this would increase confidence that the indicator is closely linked to the adaptation. If it moves around a lot, it would signal that the attribution is sensitive to modelling choices and should be treated with more caution.

References

Aaronson, E. L. et al. (2015). "Emergency department quality and safety indicators in resource-limited settings: an environmental survey". In: *International Journal of Emergency Medicine* 8.39, p. 7.

Agency for Healthcare Research and Quality (2015). *Six Domains of Health Care Quality*. https://www.ahrq.gov/talkingquality/measures/six-domains.html. Accessed: 2025-05-07.

Ambulancezorg Nederland (2016). *Landelijk Protocol Ambulancezorg (LPA) 8.1*. URL: https://www.ambulancezorg.nl/themas/kwaliteit-van-zorg/protocollen-en-richtlijnen/landelijk-protocol-ambulancezorg.

— (2023). *Onderwijsboek protocollen LPA9*. Zwolle: Ambulancezorg Nederland.

Anderson, Janet E. et al. (2016). "Implementing resilience engineering for healthcare quality improvement using the CARE model: a feasibility study protocol". In: *Pilot and Feasibility Studies* 2.61. DOI: 10.1186/s40814-016-0103-x. URL: https://doi.org/10.1186/s40814-016-0103-x.

AORTA-LSP (VZVZ) (n.d.). *Spoedverwijzing van HAP naar SEH*. URL: https://www.aorta-lsp.nl/zorgtoepassingen/spoedverwijzing-van-hap-naar-seh.

Arcuri, Rodrigo et al. (2022). "On the brink of disruption: Applying Resilience Engineering to anticipate system performance under crisis". In: *Applied Ergonomics* 99, p. 103632. DOI: 10.1016/j.apergo.2021.103632. URL: https://doi.org/10.1016/j.apergo.2021.103632.

Asplin, B. R. et al. (2003). "A Conceptual Model of Emergency Department Crowding". In: *Annals of Emergency Medicine* 42.2, pp. 173–180.

Backus, B. E. et al. (2020). "Organization of prehospital care in the Netherlands: a perspective article". In: *European Journal of Emergency Medicine* 27.6, pp. 398–399.

Baron, Audrey et al. (2022). "Impact of Fast SARS-CoV-2 Molecular Point-Of-Care Testing on Patients' Length of Stay in an Emergency Department". In: *Microbiology Spectrum* 10.4, e00636–22. DOI: 10.1128/spectrum.00636-22.

Berg, Siv Hilde et al. (2018). "Methodological strategies in resilient health care studies: An integrative review". In: *Safety Science* 110, pp. 300–312. DOI: 10.1016/j.ssci.2018.08.025. URL: https://doi.org/10.1016/j.ssci.2018.08.025.

BIG-register (2020a). *Maatregelen minister: voormalig verpleegkundigen en artsen mogen aan de slag en uitstel herregistratie*. Gepubliceerd 27 oktober 2020. URL: https://www.bigregister.nl/actueel/nieuws/2020/03/18/deur-open-voor-voormalig-verpleegkundigen-en-artsen-en-uitstel-herregistratie.

— (2020b). *Uitstel herregistratie: wat betekent dit*. Gepubliceerd 18 maart 2020. URL: https://www.bigregister.nl/actueel/nieuws/2020/03/18/uitstel-herregistratie.

Bos, N. et al. (2015). *SPOED MOET GOED: Indicatoren en normen voor zes spoedzorgindicaties*. URL: https://www.nivel.nl/sites/default/files/bestanden/Ketenbrede-Kwaliteitsindicatorensets-Acute-Zorg.pdf.

Brand, Crispijn L. van den et al. (2014). "Fracture prevalence during an unusual period of snow and ice in the Netherlands". In: *International Journal of Emergency Medicine* 7, p. 17. DOI: 10.1186/1865-1380-7-17.

Cho, Minsu et al. (2020). “Process Mining-Supported Emergency Room Process Performance Indicators”. In: *International Journal of Environmental Research and Public Health* 17.17, p. 6290. DOI: 10.3390/ijerph17176290. URL: https://doi.org/10.3390/ijerph17176290.

Chuang, Sheuwen et al. (2020). “Measurement of resilience potential: Development of a resilience assessment grid for emergency departments”. In: *PLOS ONE* 15.9, e0239472. DOI: 10.1371/journal.pone.0239472. URL: https://doi.org/10.1371/journal.pone.0239472.

Claassen, L. et al. (2025). “Characteristics of Dutch ED patients and their journey through the acute care chain: A province-wide flash-mob study”. In: *PLOS ONE* 20.4, e0318510.

Donabedian, A. (2005). “Evaluating the Quality of Medical Care”. In: *The Milbank Quarterly* 83.4, pp. 691–729.

Drabecki, Mariusz et al. (2023). “Multi-criteria assignment problems for optimising the emergency medical services (EMS), considering non-homogeneous speciality of the emergency departments and EMS crews”. In: *Scientific Reports* 13, p. 7496. DOI: 10.1038/s41598-023-33831-7. URL: https://doi.org/10.1038/s41598-023-33831-7.

Faber, Jayne et al. (2023). “Creating a Rapid Assessment Zone with Limited Emergency Department Capacity Decreases Patients Leaving Without Being Seen: A Quality Improvement Initiative”. In: *Journal of Emergency Nursing* 49.1, pp. 86–98. DOI: 10.1016/j.jen.2022.10.002.

Fenton, Joshua J. et al. (2012). “The Cost of Satisfaction: A National Study of Patient Satisfaction, Health Care Utilization, Expenditures, and Mortality”. In: *Archives of Internal Medicine* 172.5, pp. 405–411. DOI: 10.1001/archinternmed.2011.1662.

Feral-Pierssens, Anne-Laure et al. (2024). “Redirection of low-acuity emergency department patients to nearby medical clinics using an electronic medical support system: effects on emergency department performance indicators”. In: *BMC Emergency Medicine* 24.166. DOI: 10.1186/s12873-024-01080-0. URL: https://doi.org/10.1186/s12873-024-01080-0.

Garrone, M. (2011). “Prehospital ultrasound as the evolution of the Franco-German model of prehospital EMS”. In: *Critical Ultrasound Journal* 3, pp. 141–147.

Gezondheidsraad (2012). *De basis moet goed! Kwaliteit bij een Basis Spoedeisende Hulp binnen een regionaal netwerk*. Den Haag: Gezondheidsraad.

Greaney, J. (2009). *Developing Key Performance Indicators to Monitor Healthcare Quality*. URL: https://publications.scss.tcd.ie/theses/diss/2009/TCD-SCSS-DISSERTATION-2009-006.pdf.

Guasconi, M. et al. (2022). “Handover methods between local emergency medical services and Accident and Emergency: is there a gold standard? A scoping review”. In: *Acta Biomedica* 93.4, e2022288.

Hollnagel, Erik (2010). “How Resilient Is Your Organisation? An Introduction to the Resilience Analysis Grid (RAG)”. In: pp. 1–6. URL: https://www.researchgate.net/publication/281803346_How_Resilient_Is_Your_Organisation_An_Introduction_to_the_Resilience_Analysis_Grid_RAG.

Jerant, Anthony et al. (2019). “Association of Clinician Denial of Patient Requests With Patient Satisfaction”. In: *Journal of General Internal Medicine* 34.7, pp. 1373–1379. DOI: 10.1007/s11606-019-04985-9.

Jones, Peter et al. (2014). “What makes a good healthcare quality indicator? A systematic review and validation study”. In: *Emergency Medicine Australasia* 26.2, pp. 113–124. DOI: 10.1111/1742-6723.12195. URL: https://doi.org/10.1111/1742-6723.12195.

Källberg, Ann-Sofie et al. (2015). “Contributing factors to errors in Swedish emergency departments”. In: *International Emergency Nursing* 23.2, pp. 156–161. DOI: 10.1016/j.ienj.2014.10.002. URL: https://doi.org/10.1016/j.ienj.2014.10.002.

Kim, Dohyung et al. (2022). “Effect of fever or respiratory symptoms on leaving without being seen during the COVID-19 pandemic in South Korea”. In: *Clinical and Experimental Emergency Medicine* 9.1, pp. 1–9. DOI: 10.15441/ceem.21.105.

Kraaijvanger, N. et al. (2016). “Self-referrals in a Dutch Emergency Department: how appropriate are they?” In: *European Journal of Emergency Medicine* 23.3, pp. 194–202.

Lieshout, Arno P. W. van et al. (2010). “Peak incidence of distal radius fractures due to ice skating on natural ice in the Netherlands”. In: *Strategies in Trauma and Limb Reconstruction* 5, pp. 65–69. DOI: 10.1007/s11751-010-0087-7.

Linden, M. Christien van der et al. (2023). “Association between COVID-19 surge and emergency department patient flow and experience”. In: *International Emergency Nursing* 66, p. 101241. DOI: 10.1016/j.ienj.2022.101241.

Madsen, M. et al. (2015). “The level of evidence for emergency department performance indicators: systematic review”. In: *European Journal of Emergency Medicine* 22.4, pp. 298–305.

Makrides, T. et al. (2022). “From stretcher bearer to practitioner: A brief narrative review of the history of the Anglo-American paramedic system”. In: *Australasian Emergency Care* 25, pp. 347–353.

Mehrolhassani, M. H., A. Behzadi, and E. Asadipour (2025). “Key performance indicators in emergency department simulation: a scoping review”. In: *Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine* 33.15.

Minister van Volksgezondheid, Welzijn en Sport (VWS) (2020a). *Infectieziektenbestrijding*. Kamerstuk 25 295, nr. 249; brief aan de Voorzitter van de Tweede Kamer der Staten-Generaal, 15 april 2020. URL: https://zoek.officielebekendmakingen.nl/kst-25295-249.html.

— (2020b). *Regeling 2019-nCoV (Staatscourant nr. 6800)*. Regeling van 28 januari 2020, gepubliceerd 31 januari 2020. URL: https://zoek.officielebekendmakingen.nl/stcrt-2020-6800.pdf.

— (2021). *Aanvulling op “Tijdelijk beleidskader voor het waarborgen acute zorg in de COVID-19 pandemie” voor de IC-afhankelijke kritiek planbare zorg*. Richtlijn; Rijksoverheid. 19 november 2021, Nederland.

Minister voor Medische Zorg (2020). *Acute zorg*. Kamerstuk 29 247, nr. 317; brief aan de Voorzitter van de Tweede Kamer, 23 oktober 2020. URL: https://zoek.officielebekendmakingen.nl/kst-29247-317.pdf.

— (2021). *Informatie- en Communicatietechnologie (ICT) in de Zorg*. Kamerstuk 27529, nr. 262; brief aan de Voorzitter van de Tweede Kamer der Staten-Generaal, 30 april 2021. URL: https://zoek.officielebekendmakingen.nl/kst-27529-262.html.

Moes, F. B. et al. (2019). ““Strangers in the ER”: Quality Indicators and Third Party Interference in Dutch Emergency Care”. In: *Journal of Evaluation in Clinical Practice* 25, pp. 390–397.

Mortazavi, Susanne E. et al. (2022). “A retrospective cohort study of the effect of SARS-CoV-2 point of care rapid RT-PCR at the Emergency Department on targeted admission”. In: *BMC Infectious Diseases* 22, p. 536. DOI: 10.1186/s12879-022-07497-x.

Nederlands Huisartsen Genootschap (NHG) (2022). *Directive on Data Exchange in Acute Care*. URL: https://www.nhg.org.

Nederlandse Zorgautoriteit (2023). *Advies bekostiging acute zorg*. URL: https://puc.overheid.nl/nza/doc/PUC_706556_22/.

Nemeth, Christopher (2019). “Resilience Engineering, Safety, and Implications for Pediatric Care”. In: *Current Treatment Options in Pediatrics* 5.2, pp. 102–110. DOI: `10.1007/s40746-019-00154-7`. URL: `https://doi.org/10.1007/s40746-019-00154-7`.

Noel, G. et al. (2018). “Which indicators to include in a crowding scale in an emergency department? A national French Delphi study”. In: *European Journal of Emergency Medicine* 25.4, pp. 257–263.

O’Connor, Rory D., Dennis G. Barten, and Gideon H. P. Latten (2021). “Preparations of Dutch emergency departments for the COVID-19 pandemic: A questionnaire-based study”. In: *PLOS ONE* 16.9, e0256982. DOI: `10.1371/journal.pone.0256982`.

OECD/European Observatory on Health Systems and Policies (2023). *Netherlands: Country Health Profile 2023*. State of Health in the EU, OECD Publishing, Paris. DOI: `10.1787/19f79f7d-en`.

Patriarca, Riccardo et al. (2017). “A paradigm shift to enhance patient safety in healthcare, a resilience engineering approach: scoping review of available evidence”. In: *International Journal of Healthcare Technology and Management* 16.3/4, pp. 319–343.

Rief, M. et al. (2023). “Physician utilization in prehospital emergency medical services in Europe: an overview and comparison”. In: *Emergencias* 35, pp. 125–135.

Rooijen, M. Reitsma-van, A. Brabers, and J. de Jong (2013). “Selectie aan de poort: onterechte zelfverwijzers op de SEH terugdringen”. In: *TSG: Tijdschrift voor Gezondheidswetenschappen* 91.1, pp. 41–43.

Rutten, Martijn et al. (2017). “Patient and care characteristics of self-referrals treated by the general practitioner cooperative at emergency-care-access-points in the Netherlands”. In: *BMC Family Practice* 18, p. 62. DOI: `10.1186/s12875-017-0633-1`.

Safi, Mariam et al. (2022). “The application of resilience assessment grid in healthcare: A scoping review”. In: *PLOS ONE* 17.11, e0277289. DOI: `10.1371/journal.pone.0277289`. URL: `https://doi.org/10.1371/journal.pone.0277289`.

Schmutz, T. et al. (2023). “No waiting lying in a corridor: a quality improvement initiative in an emergency department”. In: *BMJ Open Quality* 12, e002431.

Al-Shaqsi, S. (2010). “Models of International Emergency Medical Service (EMS) Systems”. In: *Oman Medical Journal* 25.4, pp. 320–323.

Soldatenkova, A. et al. (2023). “Emergency department performance assessment using administrative data: A managerial framework”. In: *PLOS ONE* 18.11, p. 22.

Sørup, C. M., P. Jacobsen, and J. L. Forberg (2013). “Evaluation of emergency department performance: a systematic review on recommended performance and quality-in-care measures”. In: *Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine* 21.62, p. 14.

Stang, A. S. et al. (2015). “Crowding Measures Associated With the Quality of Emergency Department Care: A Systematic Review”. In: *Academic Emergency Medicine* 22.6, pp. 643–656.

Toh, L. et al. (2024). “Evaluation of unplanned reattendances to the pediatric emergency department – a five-year study”. In: *BMC Pediatrics* 24, p. 662.

Valk, J. de et al. (2014). “Self-referred patients at the Emergency Department: patient characteristics, motivations, and willingness to make a copayment”. In: *International Journal of Emergency Medicine* 7.30, pp. 1–6.

VZinfo.nl – Volksgezondheid en Zorg (2025). *Acute zorg—Gebruik: SEH-bezoeken*. URL: `https://www.vzinfo.nl/acute-zorg/gebruik/seh`.

Watt, Alison, Gyuchan Thomas Jun, and Patrick Waterson (2019). “Resilience in the blood transfusion process: Everyday and long-term adaptations to ‘normal’ work”. In: *Safety Sci-*

ence 120, pp. 498–506. DOI: 10.1016/j.ssci.2019.07.028. URL: https://doi.org/10.1016/j.ssci.2019.07.028.

Werkgroep Kwaliteitsindeling SEH (2009). *Spoedeisende hulp: Vanuit een stevige basis*. URL: https://www.rijksoverheid.nl/documenten/rapporten/2009/06/16/spoedeisende-hulp-kwaliteitsindeling.

Winslow, Rosalie (2020). "Failing the metric but saving lives: The protocolization of sepsis treatment through quality measurement". In: *Social Science & Medicine* 253, p. 112982. DOI: 10.1016/j.socscimed.2020.112982. URL: https://doi.org/10.1016/j.socscimed.2020.112982.

World Health Organization, Organisation for Economic Co-operation and Development, and World Bank Group (2018). *Delivering Quality Health Services: A Global Imperative for Universal Health Coverage*. Geneva: World Health Organization. ISBN: 9789241513906. URL: https://documents.worldbank.org/en/publication/documents-reports/documentdetail/482771530290792652/delivering-quality-health-services-a-global-imperative-for-universal-health-coverage.

Yang, Ming, Hao Sun, and Sunyue Geng (2023). "On the quantitative resilience assessment of complex engineered systems". In: *Process Safety and Environmental Protection* 174, pp. 941–950. DOI: 10.1016/j.psep.2023.05.019. URL: https://doi.org/10.1016/j.psep.2023.05.019.

A

Perforamnce Indicators

The paper by Madsen et al. (2015) presents a systematic review evaluating emergency department (ED) performance indicators using the Oxford Centre for Evidence-Based Medicine (OCEBM) Levels of Evidence. To prioritize candidate indicators, we refer to these OCEBM evidence levels to characterize the likely strength and reliability of the underlying evidence. The levels form a hierarchy of study designs: Level 2b corresponds, for example, to cohort studies—either retrospective (cohorts defined from existing data to examine past outcomes) or prospective (cohorts defined now and followed into the future). Levels placed higher in the hierarchy (numerically closer to 1) generally reflect evidence that is more systematic and reproducible than lower levels (e.g., case series or expert opinion). The levels used in this study are shown in Table A.1.

Table A.1: Oxford Evidence Levels as used in Madsen et al. (2015)

Two-rater article Oxford Evidence Level consensus	Type of study
1b. Prospective cohort study with good follow-up	Primary research
2a. Systematic review of 2b and better studies	Systematic review
2b. Retrospective cohort study or prospective cohort with poor follow-up	Primary research
3a. Systematic review of 3b and better studies	Systematic review
3b. Nonconsecutive cohort study, or very limited population	Primary research
4. Case series	Primary research
5. Expert opinion	Primary research

Importantly, we do not equate a higher evidence level with an automatically "best" performance indicator for ED evaluation. Instead, the evidence level increases our confidence that an indicator has been studied with more robust methods. We therefore mainly weigh indicators by the strength of their evidence base, with evidence levels of 2b and above given more weight in the selection.

Using these evidence levels, Madsen et al. (2015) reviewed 127 articles and extracted 202 individual performance indicators, which they grouped into five broad categories: process, outcome, satisfaction, structural/organizational, and equity. Each indicator was assessed based on the number of studies supporting it, as well as the strength of the supporting evidence. The result was a comprehensive ranking of indicators according to both their frequency of use and their empirical support, as shown in Table A.2.

		Level of evidence (*n*)							
Indicator group	Category	1b	2a	2b	3a	3b	4	5	Total
Patient satisfaction	Satisfaction	2		5	1	11	3		22
Standard of care treatment	Process	1	1	14	1	1	3		21
Correct diagnosis	Process			14		3	1		18
ED occupancy/crowding	Structural	2	1	12		1	1		17
Time to treatment	Process		2	12		3		1	18
ED LOS/wait	Process	1	1	12	1	2			17
ED returns	Outcome			12		1			13
LWBS	Outcome	1		4	2	4			11
Time to diagnosis	Process			7		1			8
Effectiveness	Process			4			2	1	7
Mortality	Outcome			4				1	5
Time to pain management	Process			4		1			5
Peer-assessed physician performance	Satisfaction						4		4
Satisfaction: provider communication	Satisfaction					4			4
Admission rate	Structural			3					3
Rate of complaints	Satisfaction			1		1			2
Resources	Structural					2			2
Provider satisfaction	Satisfaction					2			2
Satisfaction: waiting time	Satisfaction			1		1			2
Auxiliary services	Satisfaction			1			1		2
Satisfaction: attitude and respect	Satisfaction					2			2
Physician workload	Process	1							1
Time diagnostics to treatment	Process			1					1
Quality of care summary measure	Process						1		1
Severity of illness	Structural			1					1
Time to diagnosis	Process			1					1
Patient participation in own care	Satisfaction						1		1
Satisfaction: pain control	Satisfaction					1			1
Satisfaction: hygiene	Satisfaction			1					1
Satisfaction: A&E environment	Satisfaction					1			1
Appropriate treatment	Process			1					1
Triage vs. time to see provider	Process						1		1
Staff safety	Satisfaction						1		1
Consultant sign-off	Process			1					1
Satisfaction: accompanying persons	Satisfaction						1		1
Ambulance diversion	Outcome			1					1
Satisfaction: discharge info	Satisfaction			1					1
Opportunity cost of crowding	Structural					1			1
Grand total		8	5	118	5	43	20	3	202

Table A.2: Two hundred and two performance indicators, grouped and categorized according to ‘a guide to medical care administration’ by the American public health association, and level of evidence (Madsen et al., 2015)

This review provides valuable insight into which indicator categories have the strongest empirical support. Among the indicators assessed with evidence levels of at least 2b, the following types were found to have the most consistent backing:

- **Patient Satisfaction**
- **Standard of care treatment**
- **Correct diagnosis**
- **ED occupancy/crowding**
- **Time to treatment**
- **ED length of atay (LOS) and wait times**
- **ED returns**

While patient satisfaction received the highest number of supporting studies, they are mainly measured through questionnaires or patient feedback. Hence this category will not be focused further for indicator selection.

Regarding standard of care and correct diagnosis indicators indicator categories, Madsen et al. categorize these as process indicators and provides the explanation as ” Process indicators assess what the provider did for the patient, including treatment adherence to clinical guidelines and quality measures that use physician peer review to determine the degree of adherence to a standard of care” Our study could not identify any papers that provide a defined set of indicators for standard of care, and Madsen et al. (2015) do not offer further clarification on what kind of indicators fall under this category.

For correct diagnosis indicators, Madsen et al. (2015) again do not provide further elaboration, and we could not find a set of such indicators in other studies. However, as will be discussed in Section A.1.2 under “Other Indicators”, some indicators reflect the correct diagnosis measure directly or indirectly, and we aim to address diagnostic accuracy using those. In addition, this thesis uses contextual enrichment of routine data (see Appendix C), where diagnostic outcomes are examined more closely, to provide a complementary way of identifying potential diagnostic errors beyond the indicators listed in this appendix.

Crowding-related indicators also emerged from Madsen et al.’s review with strong empirical backing. Some studies categorize crowding as a structural indicator, as it often reflects broader systemic constraints, such as limited inpatient capacity or ED space. This thesis builds on those findings and uses crowding measures, due to their well-established link to quality of care and patient outcomes (Stang et al., 2015; Noel et al., 2018; Soldatenkova et al., 2023).

Finally, the remaining indicators with the strongest empirical support (e.g., time to treatment, LOS and wait times, ED returns, LWBS) are included in the indicator sets described in the Crowding Indicators and Other Indicators sections.

A.1. Selection of Performance Indicators

This Appendix begins by presenting crowding-related indicators, which multiple studies have shown to be the most empirically grounded measures of ED performance (Madsen et al., 2015; Stang et al., 2015; Noel et al., 2018; Soldatenkova et al., 2023). It then introduces additional indicators drawn from studies that offer the strongest evidence and are most aligned with the selection rationale of this thesis (Sørup et al., 2013; Soldatenkova et al., 2023). And finally maps these indicators to IOM dimensions and categorizes them under the Donabedian structure of structure-process-output.

A.1.1. Crowding Indicators

Among the studies that propose a well-supported set of ED crowding indicators, two were selected for in-depth analysis: Noel et al. (2018) and Stang et al. (2015).

Noel et al. (2018) conducted a Delphi study with 23 French ED clinicians, most of whom had over 10 years of emergency care experience and nearly half of whom worked in academic hospitals. These experts were selected based on their prior involvement in ED crowding–related workgroups or publications, ensuring both practical and conceptual familiarity with the issue. The goal of the study was to identify a core set of performance indicators relevant to daily ED operations. The result was a focused list of 15 indicators, organized into input, throughput, and output categories. These indicators are generic, directly tied to capacity to manage flow function, and reflect clinically plausible, operationally realistic measures.

Stang et al. (2015) conducted a systematic review of 32 peer-reviewed studies to identify crowding measures. From this review, they extracted 15 crowding measures and grouped them within a multi-layered structure. Broad indicators—such as ED occupancy or number of patients in the waiting room—were linked to diverse downstream outcomes (e.g., time to antibiotics, mortality, adverse events), depending on study context. This layered and indirect structure, while useful for synthesizing empirical evidence, is less suited for direct operational application. Many of the observed relationships are study-specific and require contextual interpretation, making the indicators harder to standardize for real-time or unit-level use.

However, Stang et al. (2015) highlight three indicators that were frequently rated as high-quality crowding measures suitable for use across different emergency care settings: ED occupancy rate, number of patients in the waiting room, and number of admitted patients awaiting inpatient beds. The first, ED occupancy rate, reflects the percentage of ED beds occupied at a given time. It offers more granular insight than a simple patient count and is categorized as a throughput measure. The second, number of patients in the waiting room, captures the volume of patients awaiting entry into the ED workflow which is categorized as an input measure. Although this is conceptually captured by the indicators “not seen by the triage nurse” and “not seen by a doctor,” these more detailed measures may not be available in the selected ED datasets. Therefore, this indicator is retained to avoid the risk of missing this aspect of crowding. The third indicator, number of admitted patients awaiting inpatient beds, is already represented in the existing set as ”Number of patients awaiting boarding”. However, this phrasing is clearer, so it will replace “number of patients awaiting boarding” in the final proposed indicator set.

One indicator from Noel et al. (2018) named “number of patients present in the UHCD over 24 hours” was excluded from this study due to contextual inapplicability. The UHCD (Unité d’Hospitalisation de Courte Durée) is a short-stay observation unit specific to French EDs and is not a standard component of Dutch emergency departments. The complete list of 16 crowding indicators selected for this study, organized by ED workflow phase, is presented in Table A.3 below. The study classified “Number of admissions over the past 24h,” “Number of patients not seen by the triage nurse,” “Number of patients not seen by a doctor,” “Time to be seen by a doctor,” and “Number of patients in the waiting room” as input indicators—understandably, since they reflect patients who have not yet progressed through ED care. However, given the workflow and critical-function definitions used in this thesis, these indicators actually measure throughput performance. Only number of admissions over the past 24h is retained as an input indicator, as it is an operational volume metric that signals inflow to the ED throughput phase. On its own—without relating it to capacity or contextual shifts that change the number of ED patients—it does not provide performance insight.

A.1.2. Other Indicators

Sørup et al. (2013) conducted a systematic review to identify the most relevant performance measures for evaluating emergency department (ED) performance at a macro level, reflecting an overall departmental performance level. The study was motivated by the lack of consensus on which indicators best reflect quality, efficiency, and sustainability in ED settings. Using the PRISMA guidelines, the authors reviewed 14 eligible peer-reviewed review articles from a pool of 1,314 database hits. The selected literature focused on performance indicators applicable across various ED contexts, rather than condition-specific or micro-level measures.

The study extracted a total of 55 performance indicators, which they organized into three categories: patient-related, employee-related, and operational, according to Traberg’s classification. Notably, very few employee-related indicators were reported in the reviewed literature. Instead, the analysis revealed that most emphasis was placed on operational indicators, partic-

Table A.3: Crowding Related Performance Indicators and Corresponding ED Workflow Phase

Performance Indicator	ED Workflow Phase
Number of admissions over the past 24h	Input
Number of patients not seen by the triage nurse	Throughput
Number of patients not seen by a doctor	Throughput
Time to be seen by a doctor	Throughput
Number of patients in the waiting room	Throughput
Patients' average length of stay	Throughput
Number of patients older than 75 years old	Throughput
Number of patients present	Throughput
Number of patients per doctor	Throughput
Number of patients per nurse	Throughput
Number of patients on a gurney or in the corridors	Throughput
ED occupancy rate	Throughput
Number of admitted patients awaiting inpatient beds	Output
Number of transfers for lack of bed over the last 3 days	Output
Average boarding time	Output
Number of patients hospitalized over the last 3 days	Output

ularly time-based metrics, and on patient-related indicators. This emphasis was evident both in the frequency with which these measures appeared in the literature and in the discussion, where the authors highlight their central role in evaluating ED performance.

Among the most frequently cited indicators across the reviewed literature were time-related measures. These indicators are already represented in the studies by Noel et al. (2018) and Stang et al. (2015), and have been included in the crowding indicators selected earlier. In addition, Sørup et al. (2013) emphasized indicators such as patients leaving without being seen (LWBS), unplanned re-attendance within 72 hours, mortality/morbidity, and unintended incidents, which they described as patient related measures. This interpretation should not be confused with the IOM definition of patient-centered care as "aligning care with patient preferences and values." Instead, these indicators can be considered patient related in the sense that they affect patients in the most consequential ways: while an extended ED stay may be frustrating but tolerable, having to return with unresolved issues, experiencing adverse outcomes, or facing mortality represents a far more critical impact on patients.

Sørup et al. (2013) also note that unintended incidents is an indicator that does not provide value if not accompanied by a qualitative description; hence, it will be excluded from this study. LWBS, unplanned re-attendance within 72 hours, and mortality/morbidity are treated as output indicators due to their clear relevance for evaluating patient outcomes. Unplanned re-attendance within 72 hours is considered an indirect measure of missed diagnosis or inadequate treatment (Sørup et al., 2013).

These are framed as essential for assessing patient safety and experience in the ED and are therefore added to the indicator set. Furthermore, these indicators were also supported by Madsen et al. (2015) and Soldatenkova et al. (2023), though they were not on the list of the previously selected crowding indicators.

In addition to the papers used for performance indicator selection discussed above, the indicator selection and mapping work by Soldatenkova et al. (2023), which aimed to develop a quality dashboard for emergency care, provides supporting confirmation for the indicators chosen in this study. Their study systematically compiled and analyzed 224 indicators from 30 existing indicator sets (total of 977 raw performance measures), focusing on simple, generalizable metrics that can be derived from routinely collected data. The authors explicitly excluded condition-specific or highly specialized indicators, which aligns with the direction of this thesis. All of the most frequently cited indicators in their review are already covered in the current selection, further confirming the validity of the chosen indicator set. Furthermore, their strong emphasis on crowding and prolonged lead times as key indicators of poor ED performance further reinforces the validity of the performance metrics selected in this study, and directly supports the second part of the critical function defined in Section 4.1, namely the capacity to manage flow through the department.

However, the indicator “mortality/morbidity” included from Sørup et al. (2013) is interpreted more precisely as “Death in ED” in Soldatenkova et al. (2023), and this is the version that will be used in this study. The three indicators included from this section, along with their respective ED workflow phase categorization, are presented in Table A.4.

Table A.4: Additional Performance Indicators and Corresponding ED Workflow Phase

Performance Indicator	**ED Workflow Phase**
Patients Left Without Being Seen (LWBS)	Output
Death in ED	Output
Unplanned Re-attendance Within 72 Hours	Output

A.2. Mapping to IOM Dimensions

The set of indicators compiled in this appendix does not capture the full breadth of ED performance. Instead, it brings together those measures that are already widely reported in the literature, are conceptually aligned with the critical function defined in this thesis, and are likely to be available in routine data. As such, the mapping presented here should be regarded as a starting point rather than a complete catalogue. Future research is needed to identify, refine, and validate additional indicators—ideally in collaboration with ED clinicians, hospital management, and other healthcare experts—to cover aspects of quality and resilience that are not yet well represented.

The selected performance indicators were mapped directly to the Institute of Medicine’s (IOM) quality dimensions, using the IOM’s own definitions as the primary guide. To complement this, Donabedian’s structure–process–outcome model was applied to clarify how each indicator operates in practice. Where needed, insights from the literature were used to support classification choices, and explanatory notes were added in the last column of the table to document the rationale behind specific classifications.

The full classification of indicators according to these dimensions is presented in Table A.5. This table includes both the literature-based indicators discussed earlier in this appendix and indicators that were analyzed empirically in this thesis that were not included in this selection.

Table A.5: Classification of performance indicators according to the IOM quality dimensions and Donabedian model

Performance Indicator	IOM Dimension	Donabedian	Notes
Number of Admissions over the past 24h	Efficiency	Structure, Process	Operational Volume Metric – Can be associated with efficiency when analyzed in relation to capacity.
Number of patients not seen by the triage nurse	Timely	Process	Even though timeliness is defined as "Reducing waits and sometimes harmful delays for both those who receive and those who give care", this indicator represents those patients who are waiting and can undergo these harmful delays.
Number of patients not seen by a doctor	Timely	Process	Same point as above.
Number of patients in the waiting room	Timely	Process	Same point as above.
Time to seen by a doctor	Timely	Process	Exact match with IOM Timely Dimension.
Patients' average length of stay (ED LOS)	Timely	Process	Exact match with IOM Timely Dimension.
Number of patients older than 75 years old	Equitable	Structure	According to Noel et al. (2018), patients over 75 contribute significantly to ED workload and typically have longer stays. This reflects a tradeoff between equity and efficiency: reductions in older patient admissions may improve efficiency metrics, but not necessarily due to improved health — rather due to systemic adaptations (e.g., post-COVID changes). From an IOM perspective, such efficiency gains may come at the cost of reduced equity.

Continued on next page

Performance Indicator	IOM Dimension	Donabedian	Notes
Number of patients present	Efficiency	Structure, Process	Operational Volume Metric – Can be associated with efficiency when analyzed in relation to capacity.
Number of patients per doctor	Efficiency	Structure	Efficiency is defined as “Avoiding waste of resources”, and staffing is directly related to resources.
Number of patients per nurse	Efficiency	Structure	Same point as above.
Number of patients on a gurney or in the corridors	Patient-Centered, Safety, Efficiency	Structure	Schmutz et al. (2023) highlight that patients lying in corridors face several risks, including lack of verbal and physical privacy, incomplete physical examinations, limited surveillance, increased patient aggressiveness, risk of psychiatric patients absconding, fragmented nursing care, inadequate staffing, and reduced patient satisfaction. These concerns closely align with the Patient-Centeredness, Safety, and Efficiency dimensions of care, as they reflect compromised patient experience, increased potential for harm, and inadequade resources. *(Or it is just efficiency and the other affected dimensions need to be measured with other indicators - its hard to fully seperate direct causes and what the PI directly meaasures)*
ED occupancy rate	Efficiency	Process	Efficiency is defined as “Avoiding waste of resources”, and this indicator is related to the resource of “beds” or “treatment spaces.”

Continued on next page

Performance Indicator	IOM Dimension	Donabedian	Notes
Unplanned re-attendance within 72 hours	Safety, Effective	Outcome	Toh et al. (2024) suggest that high unplanned re-attendance (UR) rates may reflect medical errors such as misdiagnosis or inappropriate management. This links the indicator to both the Safety dimension—"avoiding harm to patients from the care that is intended to help them"—and the Effectiveness dimension, which relates to providing services based on scientific knowledge. However, UR can also result from factors unrelated to ED performance, such as illness progression or patient behavior. This limitation should be considered when interpreting the indicator.
Number of patients awaiting boarding	Timely	Structure	Even though timeliness is defined as "Reducing waits and sometimes harmful delays for both those who receive and those who give care", this indicator represents those patients who are waiting and can undergo these harmful delays.
Number of transfers for lack of bed over the last 3 days	Efficiency	Outcome	The PI name directly resonates with the IOM Efficiency Definition.
Average boarding time	Timely	Process	Exact match with IOM Timely Dimension.
Number of patients hospitalized over the last 3 days	Efficiency	Outcome	Operational Volume Metric – Where changes in admission rates more clearly capture changes in ED resource use, thus it matches with IOM efficiency dimension.
Patients Left Without Being Seen (LWBS)	Safety	Outcome	Patients leaving the ED before treatment face increased risks of adverse events, particularly in high-risk populations requiring care.(Faber et al., 2023).

Continued on next page

Performance Indicator	IOM Dimension	Donabedian	Notes
Death in ED	Safety	Outcome	The IOM Safety dimension is defined as "Avoiding harm to patients from the care that is intended to help them" and this directly captures failure to avoid harm/help the patients.
Patients discharged < 4h	Timely	Outcome	This indicator was empirically analyzed in this thesis.
Time to result	Timely	Process	This indicator was empirically analyzed in this thesis. In the case study, "time to result" referred specifically to COVID test turnaround time, which is not relevant in current practice. However, it is retained here to highlight that time to result for other diagnostic tests performed in the ED should also be captured as a key timeliness measure.
Discharged out of ED	Efficiency	Outcome	Final disposition at ED exit; indicates throughput and bed-allocation consequences reflecting use of resources (Efficiency). When analyzed together with contextual information (see Appendix C for Contextual Enrichment), this indicator can provide more detailed insight to performance.
Transfer to another ED	Efficiency	Outcome	Final disposition at ED exit; indicates throughput and bed-allocation consequences reflecting use of resources (Efficiency). When analyzed together with contextual information (see Appendix C for Contextual Enrichment), this indicator can provide more detailed insight to performance.
Admitted to Inpatient Care	Efficiency	Outcome	Final disposition at ED exit; indicates throughput and bed-allocation consequences reflecting use of resources (Efficiency). When analyzed together with contextual information (see Appendix C for Contextual Enrichment), this indicator can provide more detailed insight to performance.

Continued on next page

Performance Indicator	IOM Dimension	Donabedian	Notes
Intrahospital Transfers	Safety, Efficiency	Outcome	Indicates movement of patients between hospital wards after admission; higher transfer rates can reflect inefficient use of beds and staff (Efficiency) and expose patients to additional handovers and procedural risks such as falls, medication errors, or hospital-acquired infections (Safety).

B

Data Request Mail and Letter

Contact: Alp Engin Kuvrag | a.e.kuvrag@student.tudelft.nl

Requested Performance Indicators for Emergency Departments

The list below contains the performance indicators I aim to analyze as part of my master's thesis in Complex Systems Engineering and Management (CoSEM) at TU Delft. My thesis focuses on resilience engineering in healthcare, specifically examining how Emergency Departments **(Spoedeisende Hulp, SEH)** in the Netherlands managed trade-offs between healthcare quality dimensions during disruptions such as the COVID-19 pandemic. The goal is to map operational adaptations and policy changes to shifts in these performance indicators, in order to better understand and measure resilient performance.

Where possible, I am seeking data for these indicators from *2000 to the most recent available period*. Fine-grained data (e.g., **daily or weekly**) would be especially useful for the period **2019 onwards** to analyze the effects of *COVID-19-related policies and operational changes.* National-level data, regional data (e.g., from specific provinces), hospital-level data, or data from research universities would all be extremely helpful.

- Number of admissions of over the past 24h
- Number of patients not seen by the triage nurse
- Number of patients not seen by a doctor
- Time to seen by a doctor
- Number of patients in the waiting room
- Patients' average length of stay (ED LOS)
- Number of patients older than 75 years old
- Number of patients present
- Number of patients per doctor
- Number of patients per nurse
- Number of patients on a gurney or in the corridors
- Ed occupancy rate
- Unplanned re-attendence within 72 hours
- Number of patients awaiting boarding
- Average boarding time
- Number of patients hospitalized over the last 3 days
- Patients Left Without Being Seen (LWBS)
- Death in ED
- Number of transfers for lack of bed over the last 3 days

I understand that it may not be possible to provide data for all of these indicators. If only a **subset** of these indicators is available, that would still be extremely helpful. Similarly, if there are **similar or related indicators** that could serve as proxies for the measures listed above, I would be very grateful if those could be shared instead.

More information about the research is provided on the next page.

Contact: Alp Engin Kuvrag | a.e.kuvrag@student.tudelft.nl

Research Gap

Resilience engineering offers a promising framework for understanding how healthcare systems adapt under pressure but existing tools (e.g. Resilience Assessment Grid (RAG)) remain largely conceptual and lack empirical grounding. There is limited knowledge on *which performance indicators truly reflect resilience and how these indicators can be used to evaluate how well healthcare systems adjust and continue functioning* during disruptive events such as the COVID-19 pandemic.

Research Objective

The objective of this thesis is to develop a framework/model to assess the resilience of Emergency Departments by linking operational performance indicators to resilience capacities *(i.e., the ability to* ***respond*** *to disruptions effectively,* ***monitor*** *the current state of the system,* ***anticipate*** *future demands or challenges, and* ***learn*** *from past experiences to improve future strategies)*. By analyzing how EDs managed trade-offs among healthcare quality dimensions—as defined by the *Institute of Medicine's (IOM) widely recognized quality framework*—the study aims to identify which indicators can serve as meaningful proxies for resilience monitoring and provide insights to strengthen healthcare system resilience.

Main Research Question

How do emergency departments manage trade-offs among healthcare quality dimensions (i.e., safety, timeliness, effectiveness, efficiency, equity, and patient-centeredness) during service disruptions, and how does this enhance resilient performance?

Sub-Questions

1. **What is the critical function of Emergency Departments (EDs), and how can it be evaluated through performance indicators?**
   Defines the ED's *core role* and identifies *relevant indicators* for measuring performance.

2. **How does the performance of EDs change during disruptions (e.g., COVID-19)?**
   Examines how *operational adaptations* and *policy decisions* affected performance indicators and generated *trade-offs between healthcare quality dimensions.*

3. **How was resilient performance expressed in EDs during COVID-19, as reflected in adaptations and changes to performance indicators?**
   Interprets observed performance changes using the *resilience capacities* to understand how EDs maintained their *critical functioning*.

C

Contextual Enrichment

The studies used in this thesis use what is here termed contextual enrichment—for example, by separating FRPs from non-FRPs, or by distinguishing patients based on their COVID-19 testing status in the POCT case. In each of these examples, the choice of subgroups follows the logic of the intervention under study: when the intervention targets FRPs, it is meaningful to compare FRPs and non-FRPs; when the intervention concerns diagnostic testing, it is meaningful to compare groups defined by test status. Because these subgroup choices depend on the clinical and organizational context, this appendix does not propose a universal rule such as “for intervention X, always use subgroups Y and Z.” Instead, it makes the procedure explicit and introduces contextual enrichment as a term for this type of analysis. In future work, more systematic guidance on subgroup selection could be developed together with ED clinicians, hospital managers, and other experts.

As discussed in Chapter 7, the purpose of contextual enrichment is to make standard performance indicators more informative by combining them with clinically meaningful context. Contextual enrichment builds on the existing indicators and PI–RA scoring by providing more detailed views for specific subgroups. In this thesis, such disaggregation showed, for example, how work adaptations in the FRP (fever and respiratory patients) case affected different patient groups in different ways, raising questions about equality and safety, and how, in the POCT case, combining indicators with test status and ward destination made it possible to assess diagnostic coverage and the effectiveness and safety of placement decisions.

This appendix does not introduce additional study findings beyond those already presented in this thesis. Instead, it shows how the contextual enrichment used in the POCT case, based on Mortazavi et al. (2022), was carried out in practice and how the enriched indicators support the interpretations about diagnostic coverage, targeted admissions, and intrahospital transfers that are summarized in the main text.

Indicators and Subgroups in the Mortazavi Case

In the Mortazavi case, several routine indicators are re-analyzed through contextual enrichment. The base indicators are as follows:

- Discharge to home from the ED
- Hospital admission from the ED
- ED length of stay (ED LOS)

- Intrahospital transfers within the first five days after admission
- Destination ward at first admission

These indicators are then split into subgroups defined by the patients' COVID-19 testing status:

1. Positive test at the ED
2. Negative test at the ED
3. Positive test before admission
4. Not tested

The enriched discharge indicators are presented in Table C.1 (Discharge to home by test status and period). Enriched admission indicators are shown in Table C.2 (Hospital admissions by test status and period). The corresponding ED LOS and intrahospital transfer indicators, broken down by the same four test groups, are shown in Table C.3 (ED LOS by test status and period) and Table C.4 (Intrahospital transfers within five days by test status and period). Finally, Table C.5 (Targeted admissions by test status and ward type) displays the composition of first admissions to COVID wards, mixed wards, ICU, and other wards for each test-status group and study period.

Table C.1: Discharge to home

Sub-Groups	**Period 1 → Period 2**	**Period 2 → Period 3**	**Period 1 → Period 3**
Discharge to Home, Total	50.6% → 48.4% (–4.3%)	48.4% → 49.5% (+2.3%)	50.6% → 49.5% (–2.2%)
Positive test at the ED	2.6% → 5.1% (+96.2%)	5.1% → 6.9% (+35.3%)	2.6% → 6.9% (+165.4%)
Negative test at the ED	15.2% → 21.0% (+38.2%)	21.0% → 22.9% (+9.0%)	15.2% → 22.9% (+50.7%)
Positive test before admission	6.5% → 10.5% (+61.5%)	10.5% → 11.6% (+10.5%)	6.5% → 11.6% (+78.5%)
Not tested	26.2% → 11.8% (–55.0%)	11.8% → 8.1% (–31.4%)	26.2% → 8.1% (–69.1%)

Table C.2: Hospital admissions

Sub-Groups	Period 1 → Period 2	Period 2 → Period 3	Period 1 → Period 3
Hospital admissions, total	49.4% → 51.6% (+4.5%)	51.6% → 50.5% (–2.1%)	49.4% → 50.5% (+2.2%)
Positive test at the ED	6.4% → 8.6% (+34.4%)	8.6% → 10.4% (+20.9%)	6.4% → 10.4% (+62.5%)
Negative test at the ED	32.7% → 31.2% (–4.6%)	31.2% → 25.7% (–17.6%)3	32.7% → 25.7% (–21.4%)
Positive test before admission	5.2% → 9.0% (+73.1%)	9.0% → 12.6% (+40.0%)	5.2% → 12.6% (+142.3%)
Not tested	5.1% → 3.6% (–29.4%)	3.6% → 1.8% (–50.0%)	5.1% → 1.8% (–64.7%)

Table C.3: ED LOS

Sub-Groups	Period 1 → Period 2	Period 2 → Period 3	Period 1 → Period 3
ED LOS, total	383 min → 377 min (–1.6%)	377 min → 363 min (–3.7%)	383 min → 363 min (–5.2%)
Positive test at the ED	393 min → 365 min (–7.1%)	365 min → 350 min (–4.1%)	393 min → 350 min (–10.9%)
Negative test at the ED	430 min → 442 min (+2.8%)	442 min → 423 min (–4.3%)	430 min → 423 min (–1.6%)
Positive test before admission	296 min → 313 min (+5.7%)	313 min → 297 min (–5.1%)	296 min → 297 min (+0.3%)
Not tested	345 min → 243 min (–29.6%)	243 min → 253 min (+4.1%)	345 min → 253 min (–26.7%)

Table C.4: Intrahospital transfers in the first 5 days

Sub-Groups	Period 1 → Period 2	Period 2 → Period 3	Period 1 → Period 3
Intrahospital Transfers, Total	50.0% → 40.2% (–19.6%)	40.2% → 34.0% (–15.4%)	50.0% → 34.0% (–32.0%)
Positive test at the ED	9.1% → 8.8% (–3.3%)	8.8% → 9.3% (+5.7%)	9.1% → 9.3% (+2.2%)
Negative test at the ED	33.2% → 23.9% (–28.0%)	23.9% → 15.9% (–33.5%)	33.2% → 15.9% (–52.1%)
Positive test before admission	6.5% → 5.3% (–18.5%)	5.3% → 8.3% (+56.6%)	6.5% → 8.3% (+27.7%)
Not tested	0.8% → 2.2% (+175.0%)	2.2% → 0.5% (–77.3%)	0.8% → 0.5% (–37.5%)

	Period 1		Period 2		Period 3		Total	
	n=	% (95% CI)	n=	% (95% CI)	n=	% (95% CI)	n=	% (95% CI)
Positive test at the ED (N=257)								
All admissions	50	100 (92.9–100.0)	85	100 (95.7–100.0)	122	100 (96.9–100.0)	257	100 (98.5–100.0)
Covid-19 ward	41	82.0 (69.2–90.2)	62	72.9 (62.7–81.2)	101	82.8 (75.1–88.5)	204	79.4 (74.0–83.9)
Mixed Covid-19/internal medicine ward	6	12.0 (5.6–23.8)	17	20.0 (12.9–29.7)	17	13.9 (8.9–21.2)	40	15.6 (11.6–20.5)
ICU	1	2.0 (0.1–10.5)	2	2.4 (0.4–8.2)	1	0.8 (0.04–4.5)	4	1.6 (0.6–3.9)
Other	2	4.0 (0.7–13.5)	2	2.4 (0.4–8.2)	1	0.8 (0.04–4.5)	5	1.9 (0.8–4.5)
Missing data	0	0.0 (0.0–7.7)	2	2.4 (0.4–8.2)	2	1.6 (0.29–5.8)	4	1.6 (0.6–3.9)
Negative test at the ED (N=864)								
All admissions	255	100 (98.5–100.0)	308	100 (98.8–100.0)	301	100 (98.7–100.0)	864	100 (99.6–100.0)
Covid-19 ward	88	34.5 (28.9–40.5)	91	29.5 (24.7–34.9)	44	14.7 (11.1–19.1)	223	25.8 (23–28.9)
Mixed Covid-19/Internal Medicine ward	68	26.7 (21.6–32.4)	57	18.5 (14.6–23.2)	63	21.0 (16.8–26.0)	188	21.8 (19.2–24.7)
ICU	9	3.5 (1.9–6.6)	4	1.3 (0.5–3.3)	5	1.7 (0.7–3.8)	18	2.1 (1.3–3.3)
Other	84	32.9 (27.5–38.9)	144	46.8 (41.3–52.3)	180	59.8 (54.2–65.2)	408	47.2 (43.9–50.6)
Missing data	6	2.4 (1.1–5)	12	3.9 (2.2–6.7)	9	3.0 (1.6–5.6)	27	3.1 (2.2–4.5)
Positive test before admission (N=277)								
All admissions	41	100.0 (91.4–100.0)	89	100.0 (95.9–100.0)	147	100.0 (97.5–100.0)	277	100.0 (98.6–100.0)
Covid-19 ward	35	85.4 (71.6–93.1)	73	82.0 (72.8–88.6)	128	87.1 (80.7–91.6)	236	85.2 (80.5–88.9)
Mixed Covid-19/internal medicine ward	4	9.8 (3.9–22.5)	12	13.5 (7.9–22.1)	13	8.8 (5.2–14.5)	29	10.5 (7.4–14.6)
ICU	0	0.0 (0.0–8.6)	0	0.0 (0.0–4.1)	0	0.0 (0.0–2.5)	0	0.0 (0.0–1.4)
Other	1	2.4 (0.1–12.6)	2	2.2 (0.4–7.8)	2	1.4 (0.2–4.8)	5	1.8 (0.8–4.2)
Missing data	1	2.4 (0.1–12.6)	2	2.2 (0.4–7.8)	4	2.7 (1.1–6.8)	7	2.5 (1.2–5.1)
Not tested (N=89)								
All admissions	40	100 (91.2–100.0)	28	100 (87.9–100.0)	21	100 (84.5–100.0)	89	100 (95.9–100.0)
Covid-19 ward	2	5.0 (0.9–16.5)	3	10.7 (3.7–27.2)	1	4.8 (0.2–22.7)	6	6.7 (3.1–13.9)
Mixed Covid-19/internal medicine ward	8	20.0 (10.5–34.8)	5	17.9 (35.6–7.9)	5	23.8 (10.6–45.1)	18	20.2 (13.2–29.7)
ICU	0	0.0 (0.0–8.8)	1	3.6 (0.2–17.7)	0	0.0 (0.0–15.5)	1	1.1 (0.1–6.1)
Other	28	70.0 (54.6–81.9)	18	64.3 (45.8–79.3)	14	66.7 (45.4–82.8)	60	67.4 (57.1–76.3)
Missing data	2	5.0 (0.9–16.5)	1	3.6 (0.2–17.7)	1	4.8 (0.2–22.7)	4	4.5 (1.8–11.0)

Table C.5: Targeted admissions by test status and ward type (taken from (Mortazavi et al., 2022))

Diagnostic Coverage and Discharge Decisions

The enriched discharge indicators in Table C.1 make visible how diagnostic coverage at discharge changed over time. In Period 1, a large share of patients discharged home from the ED belonged to the "Not tested" subgroup (26.2%). Across the subsequent periods—during which RAD tests and then point-of-care RT-PCR were introduced—this proportion dropped to 11.8% in Period 2 and further to 8.1% in Period 3 (overall 26.2% → 8.1%). Over the same periods, the proportions of patients discharged with a documented test result increased: discharges with a positive test at the ED rose from 2.6% in Period 1 to 6.9% in Period 3, negative test at the ED from 15.2% to 22.9%, and positive test before admission from 6.5% to 11.6%. Taken together, these shifts indicate that a much larger share of patients leaves the ED with a confirmed diagnostic status rather than with unresolved uncertainty. This is what the main thesis refers to as higher diagnostic coverage at discharge: the same basic indicator ("discharged home") becomes an indicator of diagnostic thoroughness once it is enriched with test-status information.

Targeted Admissions, Transfers, and Placement Quality

A similar logic applies to hospital admissions and subsequent intrahospital transfers. Table C.2 shows that the overall proportion of patients admitted from the ED remains fairly stable from Period 1 to Period 3 (49.4% → 50.5%). The more relevant change for contextual enrichment appears in Table C.5, which links admissions to both test status and ward destination. Among patients with a negative test at the ED, the share admitted to COVID-19 wards decreases from 34.5% in Period 1 to 14.7% in Period 3, while the share admitted directly to non-COVID wards increases from 32.9% to 59.8%. At the same time, test-positive patients continue to be admitted predominantly to COVID or mixed wards. These patterns indicate that diagnostic information from RAD and VitaPCR is increasingly used to guide first admission decisions so that patients are placed in wards that better match their actual infection status. In other words, the enriched indicators in Table C.5 provide evidence of more targeted admissions.

The enriched intrahospital transfer indicators in Table C.4 reinforce this interpretation. Overall, the share of admitted patients who experienced at least one transfer within the first five days falls from 50.0% in Period 1 to 34.0% in Period 3 (a reduction of about one-third). This overall decline is driven in particular by patients with a negative ED test, whose transfer rate decreases from 33.2% to 15.9% over the same periods. In Mortazavi et al. (2022), intrahospital transfers are interpreted as a marker of non-targeted initial placement: transfers typically occur when patients are first admitted to a ward that does not match their final diagnosis or care needs. Thus, the reduction in total transfers after the introduction of RAD and VitaPCR is taken as evidence that more patients are placed correctly from the outset. From a resilience and safety perspective, fewer transfers mean fewer handovers and fewer opportunities for transfer-related errors or hospital-acquired complications. In this way, contextual enrichment—linking transfer rates to test status—allows the indicator “intrahospital transfers” to be interpreted as evidence about diagnostic and placement accuracy.

D

Example Case

This appendix provides an exploratory example of how the PI–RA scores might, in principle, be extended towards an empirical resilience curve, and why the current implementation of the framework is not yet suitable for that purpose. The aim is illustrative rather than conclusive: the example shows how one could construct a time series of performance-change snapshots for a single work adaptation, and what goes wrong if the resulting performance change magnitudes PC_M are treated as if they formed a cardinal performance trajectory $P(t)$ as required by the area-based resilience measures $R_{attained}$ and R_{lost} discussed in Section 7.3.

Synthetic monthly trajectories

The example is based on the Rapid Assessment Zone (RAZ) case. In the published evaluation, and in the main analysis in this thesis, the RAZ adaptation is assessed by comparing two six-month windows before and after implementation using five performance indicators. For the purposes of this illustration, the focus is restricted to the three indicators that were assigned a relevance weight $w_i = 1$: arrival-to-provider time (ATP), ED length of stay (ED LOS), and the percentage of patients who left without being seen (LWBS).

Because the original study reports only aggregate before/after values, the monthly trajectories in Table D.1 are synthetic. They were constructed as follows. First, the post-RAZ means for these three indicators were taken as approximating the desired baseline level of performance after the adaptation. These values were assigned to Months 1 and 12, so that the episode starts and ends at the same desired level. Second, the pre-RAZ means were placed at Month 7 to represent the worst point of the disruption immediately before the work adaptation. Third, the intermediate months were filled in by hand with a smooth deterioration–recovery pattern: ATP, ED LOS and LWBS are assumed to worsen gradually between Months 1 and 7 (reflecting the impact of the disruption before RAZ is fully in place) and then to improve over the remaining months as the adaptation takes effect. The resulting trajectories in Table D.1 form a plausible resilience-shaped pattern in which each indicator returns to its initial level by Month 12, but they are not meant to reproduce the empirical time series.

From monthly changes to PI-RA scores

Given the synthetic indicator levels in Table D.1, Table D.2 reports the month-to-month percentage changes for each indicator and the corresponding indicator change scores s_i obtained

Table D.1: Synthetic monthly trajectories for RAZ indicators used in the resilience-curve example

Performance Indicator	m1	m2	m3	m4	m5	m6	m7	m8	m9	m10	m11	m12
ATP (min)	11	15	18	21	24	26	28	25	22	18	14	11
ED LOS (min)	208	215	225	235	245	253	260	248	238	226	216	208
LWBS (%)	2.6	3.0	3.6	4.2	4.9	5.3	5.6	5.0	4.3	3.7	3.1	2.6

by applying the same magnitude bins and direction conventions as in the main PI–RA analysis. Improvements in ATP, ED LOS and LWBS are scored positively, deteriorations negatively, and changes are binned into low, medium and high magnitude categories as before. For this example, all three indicators are treated as equally relevant, with relevance weights $w_i = 1$.

Using these s_i values, Table D.3 then computes performance change magnitudes PC_M for each month-to-month period using formula (3.1) from Chapter 3, together with the associated breadth scores PC_B and trade-off coefficients TO_C from formulas (3.2) and (3.3). Because the synthetic trajectories were chosen such that all three indicators move in the same direction in each month, the breadth and trade-off scores are constant within the deterioration segment (Months 1–7) and within the recovery segment (Months 7–12): during deterioration, $PC_B = -3$ and $TO_C = -1.0$; during recovery, $PC_B = 3$ and $TO_C = +1.0$. In this appendix, these latter scores are reported for completeness but not used directly in the attempted construction of a resilience curve.

Table D.2: Monthly percentage changes and indicator change scores s_i for the synthetic RAZ example

Period	ATP (min) % change	ATP s_i	ED LOS (min) % change	ED LOS s_i	LWBS (%) % change	LWBS s_i
m1–m2	$+36.4$	-0.6	$+3.4$	-0.1	$+15.4$	-0.3
m2–m3	$+20.0$	-0.6	$+4.7$	-0.1	$+20.0$	-0.6
m3–m4	$+16.7$	-0.3	$+4.4$	-0.1	$+16.7$	-0.3
m4–m5	$+14.3$	-0.3	$+4.3$	-0.1	$+16.7$	-0.3
m5–m6	$+8.3$	-0.3	$+3.3$	-0.1	$+8.2$	-0.3
m6–m7	$+7.7$	-0.3	$+2.8$	-0.1	$+5.7$	-0.3
m7–m8	-10.7	$+0.3$	-4.6	$+0.1$	-10.7	$+0.3$
m8–m9	-12.0	$+0.3$	-4.0	$+0.1$	-14.0	$+0.3$
m9–m10	-18.2	$+0.3$	-5.0	$+0.3$	-14.0	$+0.3$
m10–m11	-22.2	$+0.6$	-4.4	$+0.1$	-16.2	$+0.3$
m11–m12	-21.4	$+0.6$	-3.7	$+0.1$	-16.1	$+0.3$

Why Accumulated PC_M is not a performance trajectory

To explore whether the monthly PC_M scores could be used to reconstruct a performance trajectory $P(t)$, we considers a deliberately simple mapping. The idea is to treat PC_M as an approximate discrete rate of change: if PC_M is negative over a given month, overall ED performance is assumed to move downward; if PC_M is positive, performance moves upward. Formally, we define a hypothetical discrete performance level P_m at the beginning of each

Table D.3: PC_M, PC_B, and TO_C for each monthly period in the synthetic RAZ example

	m1–m2	**m2–m3**	**m3–m4**	**m4–m5**	**m5–m6**	**m6–m7**	**m7–m8**	**m8–m9**	**m9–m10**	**m10–m11**	**m11–m12**
PC_M	-0.33	-0.43	-0.23	-0.23	-0.23	-0.23	$+0.23$	$+0.23$	$+0.30$	$+0.33$	$+0.33$
PC_B	-3	-3	-3	-3	-3	-3	$+3$	$+3$	$+3$	$+3$	$+3$
TO_C	-1.0	-1.0	-1.0	-1.0	-1.0	-1.0	$+1.0$	$+1.0$	$+1.0$	$+1.0$	$+1.0$

month m and set

$$P_{m+1} = P_m + PC_M^{(m,m+1)}, \tag{D.1}$$

Using this update rule, we construct a simple cumulative performance level. We set an arbitrary reference value $P_1 = 0$ at Month 1, and then add the monthly performance change magnitudes successively, where $PC_M^{(m,m+1)}$ denotes the performance change magnitude for the period from Month m to Month $m+1$ as reported in Table D.3. Starting from $P_1 = 0$ and iteratively applying (D.1) yields a sequence of levels $P_1, \ldots, P_{12}$. When these levels are plotted over time (Figure D.1), the resulting curve has the expected qualitative shape: it slopes downward during the first six to seven months, reaches a minimum when deterioration is greatest, and then slopes upward as the PC_M values turn positive. However, a striking feature of this construction is that the final point P_{12} does not return to the initial value $P_1 = 0$, even though by construction all three indicators in Table D.1 have returned exactly to their baseline levels in Month 12. In other words, this simple accumulated-PC_M trajectory shows a residual deficit long after the indicators themselves have recovered.

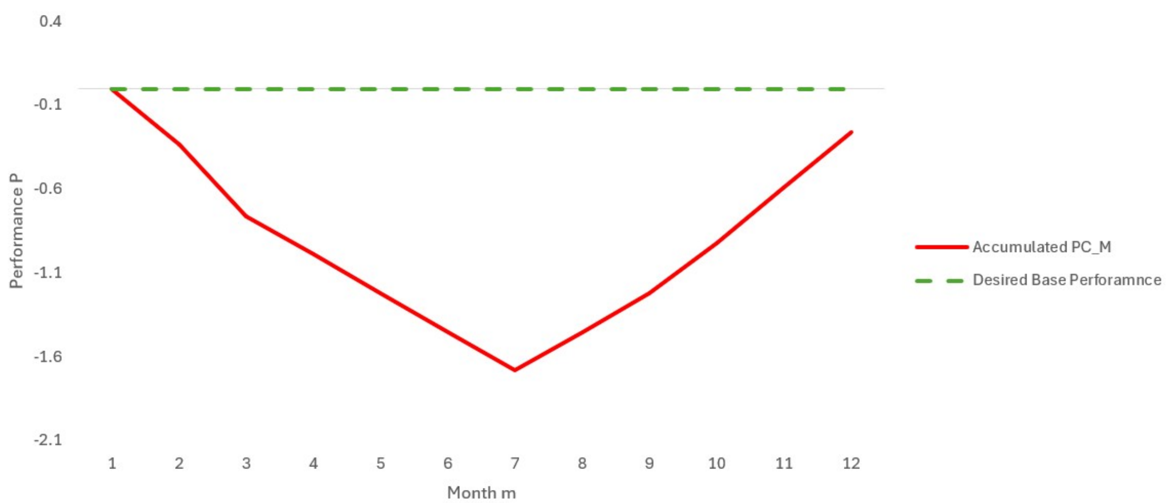

Figure D.1: Illustrative performance trajectory obtained by accumulating monthly PC_M scores. The curve shows a deterioration phase followed by recovery, but the final level P_{12} does not return to the initial level P_1, even though the underlying indicators in Table D.1 were constructed to return exactly to their baseline values.

The example in Figure D.1 therefore highlights an important limitation of the current PI–RA scores for empirical resilience measurement. While PC_M, PC_B and TO_C are well suited to profiling the direction, breadth and trade-off structure of performance change between two periods, they do not form a cardinal performance level $P(t)$ that can be integrated over time. They are designed as ordinal summaries of “how much better or worse things became on net”, not as exact increments that add up to zero whenever the underlying indicators return to their original values.

D.4 Implications for future work

The lesson from this appendix is not that the PI–RA framework cannot support empirical resilience measurement, but that an additional quantitative layer is required. To use area-based resilience metrics such as $R_{attained}$ and R_{lost}, one needs a performance function $P(t)$ on a cardinal scale, such that the area under $P(t)$ genuinely reflects how much of the desired ED function was maintained during the disruption episode. Simply accumulating the existing PC_M scores does not achieve this, because the binning of magnitudes and the ordinal nature of the scores mean that equal and opposite changes do not necessarily cancel out over time.

At the same time, the example also shows how the present thesis can serve as a starting point for developing such a quantitative mapping. Chapters 4–6 identify which functions are most relevant for the ED performance to aid relevant PI selection, how they relate to IOM quality dimensions, and how their changes should be interpreted as expressions of resilient performance. Future work could build on this foundation by replacing the current ordinal magnitude bins with more finely grained, empirically calibrated mappings from indicator changes to a cardinal performance level $P(t)$. Once such a mapping is available, the same logic as in this appendix—constructing a performance trajectory over time and comparing it to a desired baseline $P_{des}(t)$—can be used to derive empirical resilience curves and to compute area-based resilience indices in the sense of $R_{attained}$ or $1 - R_{lost}$.