

# **A Field Guide to** **Epistemic Failure**

---

**James Roy Dennis**

PortusSophia, LLC

*v0.8 — Sample Edition*

A Field Guide to Epistemic Failure

© 2026 James Roy Dennis. All rights reserved.

Published by PortusSophia, LLC — Maryland, USA

ORCID: 0009-0000-6689-8763

SSRN Author ID: 9648792

---

This is an early release draft. No part of this publication may be reproduced or transmitted in any form without permission from the author.

This book was developed through extensive AI-assisted dialogue, critique, and refinement. The framework, taxonomy, and conclusions presented herein remain the responsibility of the author.

# Contents

---

## **Reader's Guide**

Preface

Introduction

## **Part I — Foundations**

Chapter 1 — What Epistemic Failure Is

Chapter 2 — Knowledge as Constraint-Bearing Structure

Chapter 3 — Seeing Structure: Structural, Empirical, and Hybrid Knowledge

Chapter 4 — The Carrying Problem

Quick-Reference Taxonomy

## **Part II — Knowledge Survivability Failures**

Chapter 5 — Carrying Failure

Chapter 6 — Schema Degradation

Chapter 7 — Class Mismatch

Chapter 8 — Proxy Substitution

Chapter 9 — Context Collapse

Chapter 10 — Reconstruction Loops

## **Part III — Governance and Operational Failures**

Chapter 11 — Decision Recurrence

Chapter 12 — Legitimacy Inflation

Chapter 13 — Legitimacy Erosion

Chapter 14 — Corrective Exclusion

Chapter 15 — Signal Suppression

Chapter 16 — Cost Displacement

## **Part IV — Ethical and Existential Failures**

Chapter 17 — Responsibility Diffusion

Chapter 18 — Consequence Distance

Chapter 19 — Purpose Dissolution

Chapter 20 — What This Taxonomy Cannot Do

## **Part V — Application and Reflection**

Chapter 21 — Using the Taxonomy

Chapter 22 — The SEH Audit

Chapter 23 — Illustrative Applications

Chapter 24 — Future Work

Epilogue — Here and Now!

Acknowledgments

About the Author

Diagnostic Quick Reference

Appendix A — SEH Method SSOT

Appendix B — SEH Codebook v1.1

## **A Field Guide to Epistemic Failure — How Knowledge-Bearing Systems Break Down**

**James Roy Dennis**

---

---

### **Reader's Guide**

---

This book was written for people who have repeatedly encountered failures that existing explanations seem unable to fully describe.

You may have experienced projects that consumed extraordinary effort while producing surprisingly little progress. You may have watched organizations rediscover lessons they had already learned, repeat decisions that had already been made, or reconstruct knowledge that already existed somewhere within the system.

Most readers already recognize the symptoms.

The challenge is recognizing the underlying condition.

Different readers will approach this material from different directions.

AI practitioners may recognize patterns involving context loss, repeated reconstruction, governance friction, deployment failures, or the recurring cost of recovering knowledge that was previously available.

Governance and compliance professionals may recognize environments where documentation exists, controls are present, audits pass, and yet critical knowledge repeatedly fails to function at the point of use.

Organizational leaders and managers may recognize recurring initiatives, coordination failures, rediscovered lessons, and questions that seem permanently open despite years of discussion.

Researchers and systems thinkers may recognize deeper questions concerning how knowledge survives, transfers, degrades, accumulates, and fails within complex environments.

The chapters may be read sequentially, but they need not be.

Readers seeking practical recognition patterns may begin with the taxonomy chapters.

Readers interested in underlying theory may begin with Chapters 1 through 3.

Readers interested in methodology may consult the appendices.

The purpose of this book is not to convince you that the taxonomy is correct.

The purpose of this book is to provide distinctions that may help you recognize patterns that already exist within the systems you inhabit.

What you conclude after that recognition remains your responsibility.

### **Illustrative Use and Evidentiary Boundary**

The examples and applications in this book are illustrative rather than documentary.

They are used to clarify diagnostic patterns, not to make evidentiary claims about specific organizations, institutions, products, professions, or events.

Where a failure mode is discussed through an example, the example should be read as a structural illustration of the pattern under examination, not as a case study unless explicitly identified as such.

This distinction matters because the purpose of the taxonomy is diagnostic orientation, not accusation, adjudication, or institutional judgment.

---

### **Scope and Posture**

This book is diagnostic rather than prescriptive.

It attempts to identify and describe a class of failures associated with the functioning of knowledge-bearing systems. It does not attempt to provide universal remedies, optimization strategies, governance prescriptions, or organizational mandates.

The taxonomy presented here offers distinctions rather than verdicts.

It may help a reader recognize a condition. It cannot determine whether a condition is present in every circumstance. That judgment remains contextual and requires examination by those closest to the system being evaluated.

Readers will encounter observations, illustrations, and methodological artifacts throughout the text. These do not all carry identical evidentiary weight.

Some claims are structural observations derived from the properties of knowledge-bearing systems.

Some are illustrative descriptions intended to make patterns easier to recognize.

Some are supported by documented methodological exercises whose scope remains bounded to the materials examined.

Where uncertainty exists, uncertainty is intentionally preserved rather than concealed.

The goal is not certainty.

The goal is legibility.

---

## **A Note on Adjacent Traditions**

The framework presented in this book does not emerge in isolation.

Adjacent traditions including organizational learning, knowledge management, information science, governance studies, sociology of organizations, and related disciplines have long examined how knowledge is created, transferred, preserved, degraded, and applied within complex systems.

The contribution of this book is therefore not the claim that these fields are blind to knowledge as a structural phenomenon.

Rather, the contribution is the development of a diagnostic vocabulary intended for a specific class of failures that arise when knowledge-bearing structures cease to function effectively across time, transitions, interfaces, incentives, and conditions of use.

Readers familiar with these traditions will recognize areas of overlap.

That overlap is expected.

The taxonomy offered here is not intended to replace existing frameworks. It is intended to provide additional distinctions that may make certain patterns easier to recognize and discuss.

---

## **Relationship to Prior Work**

Several concepts discussed in this book emerged from earlier research concerning knowledge survivability, structural constraint, and the behavior of information under conditions of use.

Most directly, the work builds upon Schema Epistemics, a framework that treats knowledge not primarily as accumulated information but as constraint-bearing structure whose significance depends upon its continued capacity to function.

The present volume is not a restatement of that work.

It is an applied diagnostic extension.

Where Schema Epistemics focused primarily on the nature of knowledge-bearing structure, this book focuses on recognizable ways those structures break down.

The shift is from theory toward diagnosis.

From explanation toward recognition.

From architecture toward failure modes.

---

### **What This Book Cannot Do**

Before proceeding, one limitation should be stated clearly.

This taxonomy cannot determine whether a particular decision was correct.

It cannot determine whether a policy was wise.

It cannot determine whether a system is ethical.

It cannot determine whether a specific intervention will succeed.

It cannot resolve disagreement.

It cannot eliminate uncertainty.

What it can do is provide distinctions.

It can offer names for recurring patterns.

It can make certain forms of failure more visible.

It can make certain assumptions easier to examine.

And it can help readers ask a different question:

«Is what I am observing a failure of effort, a failure of process, a failure of technology, or a failure of knowledge to function?»

Everything that follows depends upon that question remaining open.

---

*Here and Now!*

### **How to Read This Book**

Not every reader arrives with the same question.

Some are attempting to understand why organizations repeatedly rediscover what they already knew.

Some are concerned with governance.

Some are concerned with technology.

Some are simply trying to understand why intelligent systems often produce avoidable failures.

This book was written for all of them.

The taxonomy does not require the reader to proceed in a strictly linear fashion. Different audiences may find different entry points more useful.

---

### *If You Work in Technology*

Begin with:

- Chapter 5 — Carrying Failure
- Chapter 6 — Schema Degradation
- Chapter 9 — Context Collapse
- Chapter 10 — Reconstruction Loops

These chapters address recurring challenges familiar to software teams, engineers, architects, and technical organizations.

Documentation survives.

Understanding does not.

Knowledge exists.

Availability fails.

Many technical readers will recognize these patterns immediately.

---

### *If You Work in Governance, Compliance, or Risk*

Begin with:

- Chapter 11 — Decision Recurrence
- Chapter 12 — Legitimacy Inflation
- Chapter 14 — Corrective Exclusion
- Chapter 15 — Signal Suppression
- Chapter 16 — Cost Displacement

These chapters focus on how organizations interact with knowledge once formal structures emerge.

The concern shifts from information itself to how institutions carry, evaluate, and respond to information.

---

### ***If You Lead Organizations***

Begin with:

- Chapter 16 — Cost Displacement
- Chapter 17 — Responsibility Diffusion
- Chapter 18 — Consequence Distance
- Chapter 19 — Purpose Dissolution

These chapters examine recurring tensions between authority, responsibility, consequence, and organizational purpose.

Many leadership challenges involve these relationships whether they are recognized explicitly or not.

---

### ***If You Are a Researcher***

Begin with:

- Chapter 3 — Structural, Empirical, and Hybrid Knowledge
- Chapter 20 — What This Taxonomy Cannot Do
- Chapter 22 — The SEH Audit
- Chapter 24 — Future Work

These chapters provide the methodological boundaries, limitations, and adjacent intellectual context underlying the framework.

The taxonomy should be understood as a diagnostic vocabulary rather than a completed theory.

---

### ***If You Are Simply Curious***

Read the book in order.

The sequence was intentionally constructed.

The progression moves from:

Knowledge survivability

to

Governance survivability

to

Responsibility survivability

Each stratum builds upon the previous one.

---

***What This Book Is***

This book is a field guide.

It provides distinctions intended to help observers recognize recurring patterns.

The framework is descriptive.

Not prescriptive.

The goal is not to determine what should happen.

The goal is to improve visibility regarding what appears to be happening.

---

***What This Book Is Not***

This book is not:

- a compliance framework
- a maturity model
- a certification program
- a management methodology
- a theory of consciousness
- a theory of truth
- a theory of governance

Nor does it attempt to prove that every pattern described here exists within every system.

The taxonomy offers vocabulary.

Observation remains the responsibility of the reader.

---

***A Suggested Posture***

Approach the taxonomy with curiosity rather than certainty.

The distinctions are most useful when treated as questions rather than conclusions.

Ask:

«Could this pattern be present?»

before asking:

«How do I prove it is present?»

The difference is small.

The consequences are not.

---

### *One Final Note*

Several chapters describe failures that may appear familiar.

That familiarity should not be mistaken for diagnosis.

Recognition is the beginning of inquiry.

Not the end of it.

The taxonomy was written to support observation.

What follows remains yours.

---

## **Preface**

---

This book began with a question.

Not a philosophical question.

Not a technical question.

A practical one.

Why do intelligent people, capable organizations, and sophisticated systems repeatedly lose things they already possess?

The question appeared in different forms.

A team rediscovered a solution it had already found.

An organization revisited a decision it had already made.

A governance process recreated work that had already been completed.

A system accumulated information while losing understanding.

The specifics changed.

The pattern remained.

At first these appeared to be separate problems.

Documentation problems.

Communication problems.

Governance problems.

Process problems.

Management problems.

Technology problems.

The more closely they were examined, the less separate they appeared.

Different environments.

Different participants.

Different consequences.

Similar recurrence.

The same questions continued appearing.

How does knowledge survive?

How does it disappear?

How does it become unavailable while remaining present?

How does a system lose context without losing information?

How does correction become increasingly difficult despite increasing awareness?

These questions eventually led to the distinctions presented throughout this book.

---

The framework that follows was not developed from a single discipline.

Nor was it developed from a single institution, organization, technology, or field of study.

The observations emerged through years of work across software systems, healthcare systems, governance environments, organizational structures, research efforts, and knowledge-intensive domains.

The result is neither a management methodology nor a compliance framework.

It is a field guide.

Its purpose is descriptive.

Not prescriptive.

The goal is not to determine what should happen.

The goal is to provide language for describing patterns that often appear difficult to discuss without shared vocabulary.

---

Several readers may recognize similarities to existing traditions.

That recognition is expected.

Questions surrounding knowledge, learning, organizational memory, governance, legitimacy, and institutional adaptation have been explored by many disciplines.

This book does not claim to replace those traditions.

Nor does it claim novelty for every distinction presented.

Its contribution, if one exists, lies in the particular arrangement of observations and the vocabulary used to describe them.

The framework should therefore be evaluated by its usefulness rather than by its originality alone.

---

A second clarification is equally important.

The taxonomy presented in these pages is not intended as a mechanism for assigning blame.

Epistemic failures frequently emerge without incompetence.

Without negligence.

Without malice.

Without bad intent.

Many arise naturally from scale, complexity, specialization, time, growth, success, and ordinary human limitation.

The framework describes patterns.

It does not assign moral character.

---

Throughout the manuscript, a recurring boundary is maintained.

*Recognition is not diagnosis.*

*Diagnosis is not judgment.*

*Judgment is not action.*

The distinctions provided here should be understood as tools for inquiry rather than instruments of certainty.

They are intended to support observation.

Not replace it.

Readers should feel free to reject any distinction that fails to illuminate reality as they encounter it.

No vocabulary deserves authority independent of usefulness.

---

The title of this book includes the phrase Field Guide intentionally.

A field guide does not tell observers what they must see.

It helps them describe what they notice.

The observations remain theirs.

The judgment remains theirs.

The responsibility remains theirs.

The same posture applies here.

---

If the distinctions prove useful, carry them.

If they do not, set them aside.

Either outcome is acceptable.

The framework seeks neither agreement nor adherence.

Only the possibility that certain recurring patterns, once named, may become easier to recognize.

For the purposes of this book, that possibility is enough.

---

## **Introduction**

---

### **The System Already Knew**

The meeting lasted three hours.

The participants were intelligent.

Prepared.

Experienced.

The issue under discussion was important.

Reports had been gathered.

Historical information had been collected.

Multiple stakeholders had been consulted.

Concerns were raised.

Alternatives were evaluated.

Recommendations were debated.

By the end of the meeting, the group reached a conclusion.

Several people left satisfied.

Several remained uncertain.

One participant sat quietly.

Not because they disagreed.

Because they remembered.

The organization had already reached the same conclusion years earlier.

---

The earlier decision had not been hidden.

The documentation existed.

The reports existed.

The recommendations existed.

The rationale existed.

At some point the knowledge stopped functioning.

The organization did not forget in the ordinary sense.

The information survived.

The understanding did not.

The system possessed the answer.

The system could no longer access the answer.

Weeks of effort were spent recovering what had already been known.

The organization experienced diligence.

The system experienced recurrence.

---

Many readers will recognize some version of this story.

The setting changes.

The pattern remains.

A software team recreates an architectural solution previously implemented.

A hospital revisits a workflow problem previously solved.

A governance body reopens a decision already examined.

A research group rediscovers a finding already established.

An institution commissions a study that reaches conclusions contained in an earlier study.

The details differ.

The experience feels familiar.

A recurring thought appears:

«Haven't we already done this?»

Often the answer is yes.

---

The observation is easy to dismiss.

Organizations are complex.

People leave.

Conditions change.

Information becomes difficult to locate.

All of these explanations are valid.

None fully explain the recurrence.

The deeper question remains.

Why do systems repeatedly lose access to knowledge they already possess?

Not information alone.

Understanding.

Context.

Rationale.

Correction.

Purpose.

The issue appears across domains.

Across industries.

Across technologies.

Across institutions.

Across scales.

The recurrence suggests something larger than isolated mistakes.

---

This book began as an attempt to describe that recurrence.

Not solve it.

Describe it.

The distinction matters.

Many frameworks begin with solutions.

Optimization strategies.

Best practices.

Maturity models.

Recommendations.

Prescriptions.

This project began elsewhere.

With observation.

Before deciding what should happen, it seemed useful to develop language for describing what appeared to be happening.

That objective shaped everything that followed.

---

Over time certain patterns became increasingly difficult to ignore.

Knowledge existed but remained unavailable.

Representations survived while context disappeared.

Metrics replaced objectives.

Decisions repeatedly returned.

Signals failed to travel.

Responsibility became difficult to locate.

Costs moved without disappearing.

Purpose weakened while activity continued.

The patterns appeared distinct.

They also appeared related.

Each represented a different way in which a system could lose the ability to effectively carry, interpret, apply, or respond to knowledge.

Eventually those patterns acquired names.

The chapters that follow describe them.

---

A caution is necessary.

The taxonomy presented in this book is not intended as proof.

It does not establish that any particular failure mode exists within any particular system.

It does not determine who is responsible.

It does not identify universal solutions.

The framework provides distinctions.

Observation remains necessary.

Judgment remains contextual.

Action remains separate.

This boundary will appear repeatedly throughout the book because it is easy to forget.

*Recognition is not diagnosis.*

*Diagnosis is not judgment.*

*Judgment is not action.*

The taxonomy should be understood within those limits.

---

Readers familiar with organizational learning, knowledge management, governance studies, systems thinking, information science, or related disciplines will encounter familiar themes.

That familiarity is expected.

This book does not claim that other traditions are blind to these concerns.

They are not.

The contribution, if one exists, lies in assembling a specific vocabulary around recurring epistemic failure patterns and exploring their relationships to one another.

Whether that vocabulary proves useful is a question best answered through use rather than assertion.

---

The title Field Guide was chosen intentionally.

A field guide does not tell observers what they must see.

It helps them describe what they notice.

The same posture applies here.

The pages that follow offer distinctions.

Not verdicts.

Questions.

Not conclusions.

Possibilities.

Not certainties.

---

The observation that initiated this work remains surprisingly simple.

A system may possess knowledge.

A system may preserve information.

A system may maintain documentation.

A system may retain records.

And still lose the ability to effectively use what it already knows.

The consequences can be expensive.

The pattern can be difficult to describe.

The recurrence can be difficult to ignore.

This book is an attempt to provide language for that condition.

Nothing more ambitious is required.

The rest belongs to observation.

---

# Part I

## *Foundations*

### Chapter 1

## **What Epistemic Failure Is — and Why Existing Frameworks Struggle to Recognize It**

---

You have already experienced epistemic failure.

You may not have called it that.

You called it a failed project.

A bad meeting.

A recurring problem.

A system that kept breaking in the same place.

You called it bureaucracy, miscommunication, organizational dysfunction, technical debt, poor leadership, inadequate process, or insufficient training.

Some of those diagnoses were correct.

Many were not.

What you were experiencing—beneath the visible event, beneath the personnel conflict, process gap, policy failure, or technology issue—may have been something more structural.

Something that existing frameworks often observe indirectly but do not always isolate directly.

A failure in the functioning of knowledge itself.

Not the absence of knowledge.

Not the absence of intelligence.

Not the absence of effort.

A failure in the structures responsible for carrying knowledge through the conditions of its own use.

This book gives that condition a vocabulary.

Not because the phenomenon is entirely new.

But because the distinctions required to recognize it consistently are often missing.

Before the failure modes can be examined, however, a more fundamental question must be addressed:

Why are these failures so difficult to recognize in the first place?

---

### *The Frameworks We Already Have*

Modern organizations are not without diagnostic tools.

Engineering provides Failure Mode and Effects Analysis (FMEA), root cause analysis, fault trees, reliability engineering, and a wide range of approaches for understanding why systems fail.

Governance disciplines provide audit frameworks, compliance structures, accountability mechanisms, and control architectures.

Organizational theory provides models of communication breakdown, coordination failure, institutional drift, and learning dysfunction.

AI governance and safety research increasingly examine alignment failures, deployment risks, adversarial behaviors, oversight limitations, and unintended consequences.

Each of these traditions contributes something valuable.

Each sees part of the landscape.

None should be dismissed.

At the same time, each tradition tends to focus on a particular class of failure.

Physical systems.

Technical systems.

Organizational systems.

Behavioral systems.

Governance systems.

What often remains difficult to isolate is a specific category of breakdown occurring within the knowledge-bearing structures upon which those systems depend.

The distinction is subtle.

A system may possess knowledge and still fail to function.

A system may generate knowledge and still fail to retain it.

A system may document knowledge and still fail to make it available where it is needed.

When these conditions occur, the visible symptoms frequently appear elsewhere.

The failure becomes visible as coordination overhead, repeated work, recurring mistakes, compliance friction, rediscovery, context reconstruction, or decision instability.

The underlying epistemic condition remains difficult to name.

This book focuses on that condition.

---

### *A Note on Adjacent Traditions*

The framework presented here does not claim that other disciplines are blind to knowledge as a structural phenomenon.

Organizational learning, knowledge management, information science, institutional theory, and related traditions have spent decades studying how knowledge is created, transferred, preserved, degraded, and applied.

Readers familiar with the work of Argyris and Schön, Nonaka, Fricker, and others will recognize areas of overlap.

That overlap is expected.

The contribution of this book is narrower.

Its purpose is not to replace those traditions.

Its purpose is to introduce a diagnostic vocabulary for a specific family of failures that become visible when knowledge-bearing structures stop functioning effectively across time, transitions, interfaces, incentives, and conditions of use.

The distinction is one of emphasis rather than exclusivity.

The question asked here is:

«What happens when knowledge itself ceases to function as a carrying structure?»

---

### ***What Epistemic Failure Is***

Epistemic failure is not the same thing as being wrong.

A team may make an incorrect prediction, test it, revise its understanding, and improve its model of reality.

Nothing in that sequence necessarily constitutes epistemic failure.

The knowledge-bearing structure functioned.

The system learned.

Epistemic failure is not ignorance.

A team may recognize that it lacks information, investigate, acquire new understanding, and proceed accordingly.

Again, the knowledge-bearing structure functioned.

The absence of knowledge was recognized and addressed.

Epistemic failure refers to something more specific.

It occurs when the structures responsible for carrying meaning, constraint, orientation, and prior understanding cease to function effectively.

Knowledge may still exist.

Documentation may still exist.

Records may still exist.

Experts may still exist.

Data may still exist.

Yet the knowledge no longer functions where it is needed.

It is unavailable.

Unrecoverable.

Unusable.

Or repeatedly reconstructed at significant cost.

The failure lies not in whether knowledge exists.

The failure lies in whether knowledge functions.

---

### ***The Difference Between Possession and Function***

Many systems possess more knowledge than they can effectively use.

Organizations often maintain archives, repositories, policies, procedures, governance records, postmortems, lessons learned databases, audit histories, technical documentation, and institutional memory.

Possession is not the same thing as function.

A lesson learned that cannot be located at the point of need is functionally absent.

A governance determination that must be rediscovered every audit cycle is functionally absent.

A technical rationale that survives only in the memory of one departing expert is functionally absent.

The information exists.

The knowledge does not function.

This distinction appears repeatedly throughout the failure modes that follow.

The taxonomy is not primarily concerned with whether knowledge exists.

It is concerned with whether knowledge continues to function under the conditions imposed upon it.

---

### *A Different Kind of Cost*

When knowledge ceases to function, systems often compensate without realizing they are compensating.

Teams reconstruct context.

Organizations repeat analysis.

Governance processes revisit settled questions.

Projects repeat prior mistakes.

Technical systems repeatedly rebuild information that was already available.

The resulting costs are frequently attributed elsewhere.

Management.

Process.

Technology.

Training.

Communication.

Resources.

Sometimes those explanations are correct.

Sometimes they describe symptoms rather than causes.

The possibility explored in this book is that a meaningful portion of those costs originate from failures in the carrying, preservation, transfer, availability, and survivability of knowledge itself.

The failure mode is epistemic.

The cost merely appears elsewhere.

---

### *What This Book Does*

The chapters that follow describe fifteen recurring patterns through which knowledge-bearing structures break down.

Some occur at the level of information mechanics.

Some emerge within governance and operational systems.

Some appear only when questions of legitimacy, correction, authority, and interpretation become involved.

The taxonomy is not offered as a complete theory of knowledge.

It is not offered as a replacement for existing frameworks.

It is not offered as a prescription.

It is offered as a vocabulary.

A set of distinctions.

A way of seeing.

Whether those distinctions prove useful remains for the reader to determine.

The intended outcome is not agreement.

The intended outcome is recognition.

Because once a pattern can be recognized, it can at least be examined.

And examination is where every meaningful diagnosis begins.

---

Chapter 2

## Knowledge as Constraint-Bearing Structure

---

You already know that knowledge is not merely information.

You have encountered information that could not guide action.

Reports that were accurate but useless.

Documentation that existed but was never consulted.

Lessons learned that were carefully recorded and repeatedly rediscovered.

Systems that contained everything required to make a sound decision and still produced poor outcomes.

The distinction matters.

Information is not knowledge.

And knowledge is not simply information organized more effectively.

This chapter introduces the central framework underlying the taxonomy presented in this book:

Knowledge functions as a constraint-bearing structure.[^schema-epistemics]

Its significance lies not merely in what it contains but in what it enables.

---

### *What a Structure Does*

Structures are not defined solely by their contents.

They are defined by their function.

A bridge is not defined by steel, concrete, or cables.

A bridge is defined by its capacity to carry load across a span.

Remove that capacity and the structure ceases to function as a bridge, regardless of whether the materials remain.

Knowledge behaves similarly.

Knowledge is not defined by the quantity of information it contains.

Knowledge is defined by its capacity to carry meaning, constraint, orientation, and prior understanding through the conditions imposed upon it.

When that carrying capacity degrades, the information may remain while the knowledge ceases to function.

This distinction is operational rather than philosophical.

A functioning knowledge structure reduces uncertainty, narrows decision space, preserves prior learning, and enables coherent action.

A non-functioning knowledge structure may contain identical information while accomplishing none of those things.

---

### ***Constraints as Knowledge***

The word "constraint" requires clarification.

In everyday language, constraints are often understood as limitations.

Here the term is used more precisely.

A constraint is a condition that narrows the range of possible interpretations, actions, configurations, or outcomes available to a system.

Knowledge functions by carrying constraints.

It tells a system what has already been learned.

It identifies questions that no longer need to be asked.

It preserves distinctions that prevent repeated mistakes.

It reduces the number of possibilities requiring active evaluation.

A team that knows a requirement is architecturally necessary does not need to rediscover that necessity each time a project begins.

A governance process that carries prior determinations forward does not need to re-litigate settled questions every review cycle.

A system that successfully carries constraints preserves effort.

A system that fails to carry constraints repeatedly spends effort recovering what it already possessed.

This distinction becomes central throughout the failure modes described later in the book.

---

### ***Storage Is Not Carrying***

Many systems confuse storage with preservation.

The confusion is understandable.

Stored information appears preserved.

But storage and carrying are not the same function.

A document stored in an archive may remain inaccessible at the moment it is needed.

A database may contain critical information that no process can effectively retrieve.

An organization may possess extensive documentation that no longer influences decisions.

The information exists.

The carrying function has failed.

Carrying refers to the active preservation of knowledge's functional capacity across time, transitions, interfaces, and conditions of use.

For knowledge to be carried successfully, meaning must survive.

Constraints must survive.

Orientation must survive.

Availability must survive.

When carrying succeeds, prior learning remains available without requiring reconstruction.

When carrying fails, systems repeatedly recover knowledge they already possessed.

The resulting cost often appears as inefficiency, coordination overhead, rediscovery, or repeated analysis.

Its origin is epistemic.

---

### ***Survivability as a Test***

Knowledge structures are not evaluated solely at the moment they are created.

They are evaluated across the conditions they must survive.

A governance decision that functions correctly for one week but disappears six months later has survived poorly.

A technical rationale that remains available only while its original author remains employed has survived poorly.

A compliance determination that must be reconstructed every audit cycle has survived poorly.

The question is not merely whether knowledge existed.

The question is whether it remained functional under stress.

Several recurring stresses appear throughout complex systems:

- Time
- Organizational transition
- Scale
- Interface change
- Incentive pressure
- Context loss

A knowledge structure that cannot survive these pressures has not necessarily failed to exist.

It has failed to function.

That distinction forms the foundation of the taxonomy that follows.

---

### *Knowledge Beyond Human Systems*

Although this book focuses primarily on human and organizational knowledge-bearing systems, similar patterns appear in other domains.

Biological systems carry information through inheritance, adaptation, and selection.

Chemical systems preserve constraints through stable molecular relationships.

Astronomical systems carry information through structure, composition, and observable history.

The present work does not claim that these domains are identical.

Nor does it attempt to reduce them to a single explanatory framework.

The comparison is narrower.

In each case, information becomes meaningful only insofar as it survives the conditions acting upon it.

The concern of this book is the human expression of that problem:

How knowledge survives.

How it transfers.

How it degrades.

How it accumulates.

And how it fails.

---

***Relationship to Schema Epistemics***

The framework presented here emerged in part from earlier work in Schema Epistemics.

Schema Epistemics treats knowledge not primarily as accumulated information but as constraint-bearing structure whose significance depends upon its continued capacity to function.

The present volume extends that perspective into diagnosis.

Rather than asking what knowledge is, the taxonomy asks what happens when knowledge-bearing structures stop functioning effectively.

The shift is practical.

The objective is not explanation for its own sake.

The objective is recognition.

Because once a failure becomes visible, it can at least be examined.

And examination remains the beginning of meaningful diagnosis.

[^schema-epistemics]: Schema Epistemics is a related line of inquiry that treats knowledge not merely as information, but as constraint-bearing structure; its relationship to this taxonomy remains exploratory, and neither framework depends upon acceptance of the other.

---

Chapter 3

## Seeing Structure: Structural, Empirical, and Hybrid Knowledge

---

Before a failure can be recognized, something else must occur first.

The observer must distinguish between different kinds of knowledge.

Many disagreements within governance, engineering, compliance, research, and organizational decision-making do not arise because people possess different information.

They arise because people are treating different kinds of knowledge as though they were the same thing.

One participant is asking a structural question.

Another is asking an empirical question.

A third is asking both simultaneously.

The resulting disagreement often appears substantive.

In reality, the participants may be operating within different epistemic categories.

This chapter introduces a distinction that will be used throughout the remainder of the book.

Not because it resolves disagreement.

Because it helps reveal where disagreement originates.

---

### *A Recognition Example*

A security team debates whether a cloud deployment should require multi-factor authentication.

One participant argues that the requirement is obvious.

Another argues that effectiveness must be demonstrated.

A third argues that neither discussion matters until implementation details are specified.

The conversation continues.

Documentation expands.

Meetings accumulate.

Positions harden.

Yet the disagreement persists.

Why?

Because the participants are answering different questions.

One is discussing architecture.

One is discussing outcomes.

One is discussing both.

The disagreement appears technical.

Its origin is epistemic.

---

### *Structural Knowledge*

Structural knowledge concerns what can be determined through architecture, organization, dependency, sequence, relationship, or design.

Its validity does not depend upon observing an outcome.

It can often be evaluated before deployment, execution, or operation.

Examples include:

- Segregation of duties
- Governance structures
- Approval workflows
- Prerequisite conditions
- Dependency relationships
- Architectural constraints

A requirement such as:

«"Risk assessment must occur before deployment."»

is structural.

The requirement may or may not be followed.

Its classification does not depend upon whether compliance is achieved.

The structure exists independently of the outcome.

Structural knowledge narrows possibility space before observation occurs.

---

### *Empirical Knowledge*

Empirical knowledge concerns observation, measurement, verification, and attained outcomes.

Its validity depends upon conditions that must be observed.

Examples include:

- Accuracy rates
- Performance metrics
- Response times
- Error rates
- Threshold attainment
- Measured effectiveness

A requirement such as:

«"System accuracy must exceed 95%."»

cannot be evaluated through architecture alone.

Measurement is required.

The requirement becomes meaningful only through observation.

Empirical knowledge narrows possibility space after observation occurs.

---

### *Hybrid Knowledge*

Many real-world requirements contain both structural and empirical components.

These requirements are neither purely architectural nor purely observational.

They combine both.

Examples include:

- Monitoring programs that require demonstrated effectiveness
- Security controls with specified validation standards
- Governance processes tied to measurable outcomes
- Operational controls requiring ongoing verification

A requirement such as:

«"Implement monitoring processes and validate their effectiveness."»

contains both structural and empirical elements.

The process must exist.

The process must also be shown to function.

Neither component alone is sufficient.

Hybrid knowledge emerges at the intersection of structure and observation.

---

### ***Why This Distinction Matters***

The distinction is useful because failures often occur when these categories become confused.

Structural questions become treated as empirical questions.

Empirical questions become treated as structural questions.

Hybrid questions are forced into one category or the other.

The resulting confusion produces recurring patterns:

- Endless debates
- Repeated analysis
- Misaligned governance discussions
- Conflicting audit interpretations
- Decision paralysis
- Classification instability

The problem is not necessarily disagreement.

The problem is category confusion.

A question cannot be answered coherently if its type has not first been recognized.

---

### ***The Structural–Empirical–Hybrid Audit***

During development of this framework, a structured classification audit was conducted using a predefined coding sheet, operational codebook, explicit decision rules, boundary clarifications, confidence ratings, dispute flags, reliability procedures, sensitivity analyses, and governance controls over method evolution.

The purpose of the audit was not to validate this book.

Nor was it intended to establish the correctness of the taxonomy presented here.

Its purpose was narrower.

The audit examined how selected governance requirements could be classified under explicit Structural–Empirical–Hybrid criteria.

The resulting classifications provide bounded methodological evidence regarding the analyzed corpus and support transparent review of classification decisions.

The audit should not be interpreted as proof of external effectiveness, validation of the Field Guide as a whole, or evidence that the framework necessarily generalizes beyond the materials examined.

Its contribution is methodological rather than confirmatory.

It makes classification decisions inspectable.

Nothing more.

Nothing less.

---

### *Why Classification Matters*

The distinction introduced here is not itself one of the fifteen failure modes.

It is an observational tool.

A lens.

A way of examining what kind of knowledge a system is attempting to carry.

Some failures occur because knowledge disappears.

Others occur because knowledge survives but becomes inaccessible.

Others emerge because different categories of knowledge are treated as interchangeable.

The taxonomy that follows builds upon these distinctions.

Not because every reader must agree with them.

But because recognizing the type of knowledge involved often makes subsequent failures easier to see.

And recognition remains the first step toward diagnosis.

---

## Chapter 4

# The Carrying Problem

---

Knowledge is often discussed as though it were an asset.

Organizations speak of acquiring it.

Protecting it.

Managing it.

Leveraging it.

Accumulating it.

The language is understandable.

Assets are valuable.

Knowledge is valuable.

The comparison appears natural.

Yet the comparison conceals something important.

Assets can retain value while remaining unused.

Knowledge cannot.

Knowledge derives its significance from function.

When it ceases to function, much of its value disappears with it.

This chapter examines the carrying problem:

How knowledge remains available across time, transitions, interfaces, and conditions of use.

And what happens when it does not.

---

### *A Recognition Example*

A team spends six months solving a difficult governance problem.

Stakeholders meet.

Requirements are analyzed.

Tradeoffs are documented.

Alternatives are evaluated.

A decision is reached.

The decision works.

Eighteen months later, most of the original participants have moved on.

A new team encounters the same question.

The documentation exists somewhere.

The rationale exists somewhere.

The lessons learned exist somewhere.

No one can find them.

No one trusts them.

No one understands them.

The analysis begins again.

The organization experiences the event as new.

The system experiences it as repetition.

The knowledge existed.

The carrying failed.

---

***Knowledge Must Travel***

Knowledge rarely remains where it is created.

Organizations change.

People leave.

Systems evolve.

Responsibilities shift.

Interfaces change.

Technologies are replaced.

Governance structures reorganize.

The knowledge that guides action must therefore move.

It must cross boundaries.

It must survive transitions.

It must remain intelligible to individuals who were not present when it originated.

Knowledge that cannot survive these conditions becomes increasingly dependent upon reconstruction.

The cost of reconstruction is often hidden.

Yet it accumulates.

Repeated meetings.

Repeated analysis.

Repeated audits.

Repeated mistakes.

Repeated explanations.

Repeated decisions.

The visible cost appears operational.

The underlying condition is epistemic.

---

### ***Reconstruction Is Not Free***

Organizations frequently treat reconstruction as normal.

The team gathers again.

The analysis is repeated.

The context is rebuilt.

The rationale is rediscovered.

The decision is revisited.

Sometimes this is necessary.

Conditions change.

New information appears.

Reevaluation may be appropriate.

But many reconstruction cycles occur for a different reason.

Not because the decision became invalid.

Not because the environment changed.

Not because the original reasoning failed.

The reconstruction occurs because the knowledge failed to remain available.

The organization pays again for work it has already performed.

The cost is real.

The category used to recognize the cost is often absent.

---

### ***The Difference Between Retrieval and Reconstruction***

The distinction is subtle but important.

Retrieval preserves prior effort.

Reconstruction repeats prior effort.

A system that successfully carries knowledge retrieves.

A system that fails to carry knowledge reconstructs.

The difference can be measured in:

- time
- attention
- labor
- coordination overhead
- governance friction
- opportunity cost

Most organizations track these costs indirectly.

Few track them as consequences of epistemic function.

As a result, recurring reconstruction is frequently accepted as unavoidable when it may instead indicate a carrying failure.

---

### ***Why Carrying Matters***

Carrying is not primarily about efficiency.

Efficiency is a consequence.

The deeper issue concerns continuity.

Knowledge that survives enables systems to build.

Knowledge that does not survive forces systems to begin again.

A team capable of retrieving prior understanding can move forward.

A team forced into reconstruction repeatedly returns to earlier stages of reasoning.

The distinction compounds over time.

Small failures become recurring costs.

Recurring costs become structural burdens.

Structural burdens eventually shape the behavior of entire systems.

What begins as an epistemic problem often becomes an organizational one.

---

### ***The Hidden Dependency***

Most systems assume carrying is occurring.

Governance assumes prior decisions remain available.

Compliance assumes controls remain understandable.

Engineering assumes rationale survives implementation.

Leadership assumes lessons learned remain learned.

The assumption is rarely examined.

Yet every one of these activities depends upon successful carrying.

When carrying fails, the dependent systems continue operating.

But they operate with increasing friction.

Increasing uncertainty.

Increasing reconstruction.

The symptoms become visible long before the dependency is recognized.

---

### ***The Central Question***

Throughout the remainder of this book, a recurring question will appear beneath many of the failure modes.

The question is simple.

When knowledge ceases to function, what happened?

Was it forgotten?

Was it degraded?

Was it displaced?

Was it substituted?

Was it misclassified?

Was it carried imperfectly?

Was it carried at all?

The failure modes that follow may be understood as different answers to that question.

Each describes a distinct way in which the carrying function can break down.

And once the carrying problem becomes visible, those failures become much easier to recognize.

---

*Here and Now!*

---

## Quick-Reference Taxonomy

### Morphology of Breakdown

What Happened to Representation?	Failure Mode
Not carried forward	Carrying Failure
Degraded over transfer	Schema Degradation
Applied to wrong class	Class Mismatch
Replaced by indicator	Proxy Substitution
Lost contextual frame	Context Collapse

What Happened to Representation?	Failure Mode
Repeated reconstruction	Reconstruction Loops
Repeated decision cycles	Decision Recurrence
Authority exceeds evidence	Legitimacy Inflation
Correction excluded	Corrective Exclusion
Signal blocked or muted	Signal Suppression
Cost moved elsewhere	Cost Displacement
Responsibility dispersed	Responsibility Diffusion
Consequence hidden by distance	Consequence Distance
Objective detached from purpose	Purpose Dissolution

### Failure Interaction Matrix

*The following interactions are illustrative, not exhaustive. Real systems typically exhibit multiple overlapping failure modes. The pairs below represent commonly observed relationships — not a complete specification.*

If you observe...	Also inspect...
Carrying Failure	Reconstruction Loops, Context Collapse
Proxy Substitution	Purpose Dissolution, Legitimacy Inflation
Signal Suppression	Corrective Exclusion
Cost Displacement	Responsibility Diffusion
Legitimacy Inflation	Signal Suppression
Decision Recurrence	Carrying Failure

The taxonomy chapters may be read as a sequence, but they also function as a practical recognition map. The table below offers a compact entry point before the failure-mode chapters begin.

Chapter	Failure Mode	Primary Diagnostic Question
5	<b>Carrying Failure</b>	Can the system access what it already knows?

Chapter	Failure Mode	Primary Diagnostic Question
6	<b>Schema Degradation</b>	Has the representation lost fidelity while remaining present?
7	<b>Class Mismatch</b>	Is valid knowledge being applied to the wrong class of problem?
8	<b>Proxy Substitution</b>	Has the indicator replaced the objective?
9	<b>Context Collapse</b>	Has information survived while interpretive context disappeared?
10	<b>Reconstruction Loops</b>	Is the system generating new understanding or recovering old understanding?
11	<b>Decision Recurrence</b>	Can the system preserve resolution once a decision has been made?
12	<b>Legitimacy Inflation</b>	Has trust grown beyond the warrant that supports it?
13	<b>Legitimacy Erosion</b>	Can evidence still generate appropriate trust?
14	<b>Corrective Exclusion</b>	Can valid correction still alter the system?
15	<b>Signal Suppression</b>	Can relevant signals reach system awareness?
16	<b>Cost Displacement</b>	Has the cost been reduced, or merely moved?
17	<b>Responsibility Diffusion</b>	Can the system identify who must act?
18	<b>Consequence Distance</b>	Can decision-makers perceive consequences with enough fidelity to learn?
19	<b>Purpose Dissolution</b>	Can the system still explain why it exists in a way that organizes action?

*The catalog continues with 19 additional failure modes across Parts II–V, including the Diagnostic Quick Reference at the back.*

*The full edition is available at [leanpub.com/fieldguideepistemic](http://leanpub.com/fieldguideepistemic)*

*Here and Now!*