CRANFIELD
UNIVERSITY PRESS

Ashgate

The Field Guide to Human Error Investigations

SIDNEY DEKKER

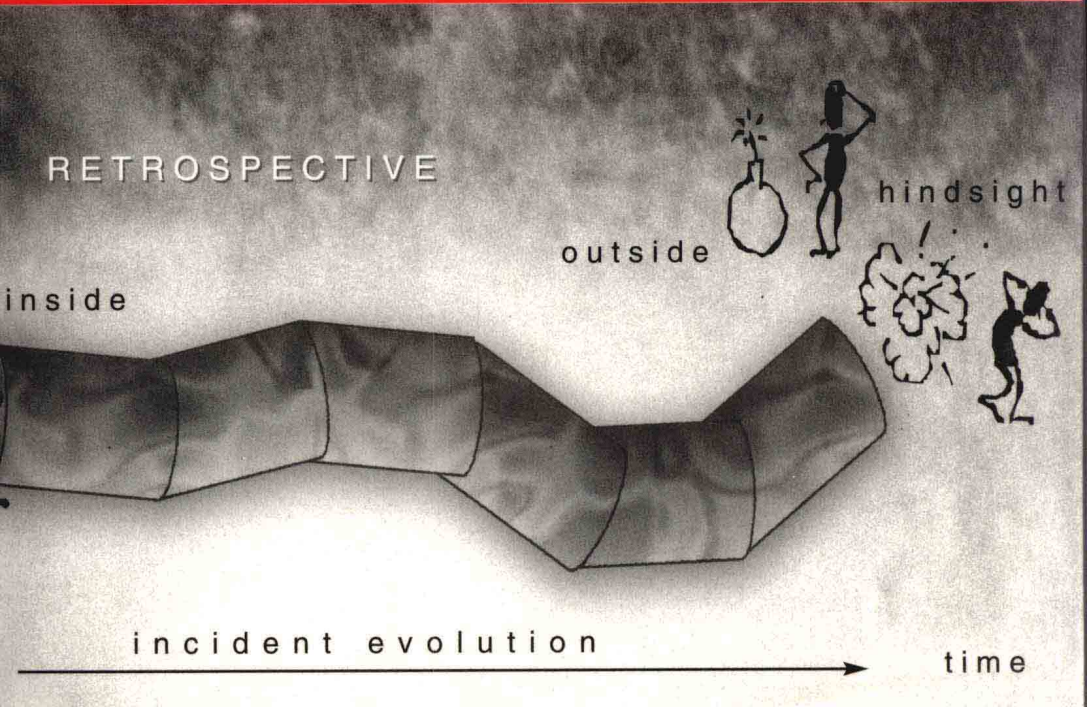

The Field Guide to Human Error Investigations

SIDNEY DEKKER
Linköping Institute of Technology, Sweden

Ashgate
Aldershot • Burlington USA • Singapore • Sydney

THE FIELD GUIDE TO HUMAN ERROR INVESTIGATIONS

Published by
Ashgate Publishing Limited
Gower House
Croft Road
Aldershot
Hants GU11 3HR
England

Ashgate Publishing Company
131 Main Street
Burlington, VT 05401-5600 USA

Ashgate website: http://www.ashgate.com

Published in association with Cranfield University Press

British Library Cataloguing in Publication Data
Dekker, Sidney
The field guide to human error investigations
1.System failures (Engineering) 2.Human engineering
3.Industrial accidents
I.Title
363.1'1

Library of Congress Control Number: 2001095458

ISBN 0 7546 1917 6 (HBK)
ISBN 0 7546 1924 9 (PBK)

Printed and bound in Great Britain by MPG Books Ltd, Bodmin, Cornwall

Preface

You are faced with an incident or accident that has a significant human contribution in it. What do you do? How do you make sense out of other people's controversial and puzzling assessments and actions? You basically have two options, and your choice determines the focus, questions, answers and ultimately the success of your probe, as well as the potential for progress on safety:

- You can see human error as the cause of a mishap. In this case "human error", under whatever label—loss of situation awareness, procedural violation, regulatory shortfalls, managerial deficiencies—is the conclusion to your investigation.
- You can see human error as the symptom of deeper trouble. In this case, human error is the starting point for your investigation. You will probe how human error is systematically connected to features of people's tools, tasks and operational/organizational environment.

The first is called the old view of human error, while the second—itself already 50 years in the making—is the new view of human error.

Table 0.1: Two views on human error

The old view of human error	The new view of human error
Human error is a cause of accidents	Human error is a symptom of trouble deeper inside a system
To explain failure, you must seek failure.	To explain failure, do not try to find where people went wrong.
You must find people's: inaccurate assessments, wrong decisions, bad judgments.	Instead, find how people's assessments and actions made sense at the time, given the circumstances that surrounded them.

This Field Guide helps you reconstruct the human contribution to system failure according to the new view. In Part II, it presents a method for how to "reverse engineer" the evolving mindset of people who were caught up in a complex, unfolding situation. The Field Guide also wants to make you aware of the biases and difficulties in understanding past puzzling behavior—which is what Part I is about.

PART I OF THE FIELD GUIDE

The first six chapters of The Field Guide talk about the old view of human error—the problems it holds, the traps it represents, and the temptations that can make you fall into them. These chapters help you understand:

- The bad apple theory: why throwing out a few bad apples does not get rid of the underlying human error problem;
- Reactions to failure: why the surprising nature of failure makes you revert easily to the bad apple theory;
- That there is no such thing as a root or primary cause: accidents are the result of multiple factors—each necessary and only jointly sufficient;
- That large psychological labels may give you the illusion of understanding human error but that they hide more than they explain;
- Why human error cannot be explained by going into the mind alone. You have to understand the situation in which behavior took place;
- Why human factors data need to be left in the context from which they came: cherry picking and micro-matching robs data of its original meaning.

PART II OF THE FIELD GUIDE

The last seven chapters show you that human error is not necessarily something slippery or something hard to pin down. They show you how to concretely "reverse engineer" human error, like other components that need to be put back together in a mishap investigation. It shows how to rebuild systematic connections between human behavior and features of the tasks and tools that people worked with, and of the operational and organizational environment in which they carried out

their work. The Field Guide will encourage you to build a picture of:

- how a process and other circumstances unfolded around people;
- how people's assessments and actions evolved in parallel with their changing situation;
- how features of people's tools and tasks and organizational and operational environment influenced their assessments and actions.

The premise is that if you really understand the evolving situation in which people's behavior took place, you will understand the behavior that took place inside of it. Here is what the last seven chapters talk about:

- Human error as a symptom of deeper trouble. Connecting people's behavior with the circumstances that surrounded them points you to the sources of trouble and helps explain behavior;
- How and where to get human factors data: from historical sources, interviews and debriefings, and process recordings;
- A method for the reconstruction of people's unfolding mindset—this is the central part around which the rest of The Field Guide revolves;
- Patterns of failure: Directs you to various patterns of failure in complex dynamic worlds, including the contributions from new technology, the drift into failure through unremarkable repetition of seemingly innocuous acts, failures to adapt and adaptations that fail, and coordination breakdowns;
- Writing meaningful human factors recommendations;
- Learning from failure as ultimate goal of an investigation: failures represent opportunities for learning—opportunities that can fall by the wayside for a variety of reasons;
- Rules for in the rubble: these are steps for how to understand human error, wrapping up the most important lessons from the book.

Contents

PART I

The Old View of Human Error:

Human error is a cause of accidents

To explain failure,
investigations must seek failure

They must find people's inaccurate assessments, wrong decisions and bad judgments

1. The Bad Apple Theory

There are basically two ways of looking at human error. The first view could be called "the bad apple theory". It maintains that:

- Complex systems would be fine, were it not for the erratic behavior of some unreliable people (bad apples) in it;
- Human errors cause accidents: humans are the dominant contributor to more than two thirds of them;
- Failures come as unpleasant surprises. They are unexpected and do not belong in the system. Failures are introduced to the system only through the inherent unreliability of people.

This chapter is about the first view, and the following five are about the problems and confusion that lie at its root.

Every now and again, nation-wide debates about the death penalty rage in the United States. Studies find a system fraught with vulnerabilities and error. Some states halt proceedings altogether; others scramble to invest more in countermeasures against executions of the innocent.

The debate is a window on people's beliefs about the sources of error. Says one protagonist: "The system of protecting the rights of accused is good. It's the people who are administering it who need improvement: The judges that make mistakes and don't permit evidence to be introduced. We also need improvement of the defense attorneys."[1] The system is basically safe, but it contains bad apples. Countermeasures against miscarriages of justice begin with them. Get rid of them, retrain them, discipline them.

But what is the practice of employing the least experienced, least skilled, least paid public defenders in many death penalty cases other than systemic? What are the rules for judges' permission of evidence other than systemic? What is the ambiguous nature of evidence other than inherent to a system that often relies on eyewitness accounts to make or break a case?

Each debate about error reveals two possibilities. Error is either the result of a bad apple, where disastrous outcomes could have been

avoided if somebody had paid a bit more attention or made a little more effort. In this view, we wonder how we can cope with the unreliability of the human element in our systems.

Or errors are the inevitable by-product of people doing the best they can in systems that themselves contain multiple subtle vulnerabilities; systems where risks and safety threats are not always the same; systems whose conditions shift and change over time. These systems themselves are inherent contradictions between operational efficiency on the one hand and safety (for example: protecting the rights of the accused) on the other. In this view, errors are symptoms of trouble deeper inside a system. Like debates about human error, investigations into human error mishaps face the choice. The choice between the bad apple theory in one of its many versions, or what has become known as the new view of human error.

A Boeing 747 Jumbo Jet crashed upon attempting to take-off from a runway that was under construction and being converted into a taxiway. The weather at the time was terrible—a typhoon was about to hit the particular island: winds were high and visibility low. The runway under construction was close and parallel to the intended runway, and bore all the markings, lights and indications of a real runway. This while it had been used as a taxiway for quite a while and was going to be officially converted at midnight the next day—ironically only hours after the accident. Pilots had complained about potential confusion for years, saying that by not indicating that the runway was not really a runway, the airport authorities were "setting a trap for a dark and stormy night". The chief of the country's aviation administration, however, claimed that "runways, signs and lights were up to international requirements" and that "it was clear that human error had led to the disaster." Human error, in other words, was simply the cause, and that was that. There was no deeper trouble of which the error was a symptom.

The ultimate goal of an investigation is to learn from failure. The road towards learning—the road taken by most investigations—is paved with intentions to follow the new view. Investigators intend to find the systemic vulnerabilities behind individual errors. They want to address the error-producing conditions that, if left in place, will repeat the same basic pattern of failure.

In practice, however, investigations often return disguised versions of the bad apple theory—in both findings and recommendations. They sort through the rubble of a mishap to:

- Find evidence for erratic, wrong or inappropriate behavior;
- Bring to light people's bad decisions; inaccurate assessments; deviations from written guidance;
- Single out particularly ill-performing practitioners.

Investigations often end up concluding how front-line operators failed to notice certain data, or did not adhere to procedures that appeared relevant after the fact. They recommend the demotion or retraining of particular individuals; the tightening of procedures or oversight. The reasons for regression into the bad apple theory are many. For example:

- Resource constraints on investigations. Findings may need to be produced in a few months time, and money is limited;
- Reactions to failure, which make it difficult not to be judgmental about seemingly bad performance;
- The hindsight bias, which confuses our reality with the one that surrounded the people we investigate;
- Political distaste of deeper probing into sources of failure, which may de facto limit access to certain data or discourage certain kinds of recommendations;
- Limited human factors knowledge on part of investigators. While wanting to probe the deeper sources behind human errors, investigators may not really know where or how to look.

In one way or another, The Field Guide will try to deal with these reasons. It will then present an approach for how to do a human error investigation—something for which there is no clear guidance today.

UNRELIABLE PEOPLE IN BASICALLY SAFE SYSTEMS

This chapter discusses the bad apple theory of human error. In this view on human error, progress on safety is driven by one unifying idea:

COMPLEX SYSTEMS ARE BASICALLY SAFE

THEY NEED TO BE PROTECTED FROM UNRELIABLE PEOPLE

Charges are brought against the pilots who flew a VIP jet with a malfunction in its pitch control system (which makes the plane go up or down). Severe oscillations during descent killed seven of their unstrapped passengers in the back. Significant in the sequence of events was that the pilots "ignored" the relevant alert light in the cockpit as a false alarm, and that they had not switched on the fasten seatbelt sign from the top of descent, as recommended by jet's procedures. The pilot oversights were captured on video, shot by one of the passengers who died not much later. The pilots, wearing seatbelts, survived the upset.[2]

To protect safe systems from the vagaries of human behavior, recommendations typically propose to:

- Tighten procedures and close regulatory gaps. This reduces the bandwidth in which people operate. It leaves less room for error.
- Introduce more technology to monitor or replace human work. If machines do the work, then humans can no longer make errors doing it. And if machines monitor human work, they can snuff out any erratic human behavior.
- Make sure that defective practitioners (the bad apples) do not contribute to system breakdown again. Put them on "administrative leave"; demote them to a lower status; educate or pressure them to behave better next time; instill some fear in them and their peers by taking them to court or reprimanding them.

In this view of human error, investigations can safely conclude with the label "human error"—by whatever name (for example: ignoring a warning light, violating a procedure). Such a conclusion and its implications supposedly get to the causes of system failure.

AN ILLUSION OF PROGRESS ON SAFETY

The shortcomings of the bad apple theory are severe and deep. Progress on safety based on this view is often a short-lived illusion. For example, focusing on individual failures does not take away the underlying problem. Removing "defective" practitioners (throwing out the bad apples) fails to remove the potential for the errors they made.

As it turns out, the VIP jet aircraft had been flying for a long time with a malfunctioning pitch feel system ('Oh that light? Yeah, that's been on for four months now'). These pilots inherited a systemic problem from the airline that operated the VIP jet, and from the organization charged with its maintenance.

In other words, trying to change your people by setting examples, or changing the make-up of your operational workforce by removing bad apples, has litte long-term effect if the basic conditions that people work under are left unamended.

Adding more procedures

Adding or enforcing existing procedures does not guarantee compliance. A typical reaction to failure is procedural overspecification—patching observed holes in an operation with increasingly detailed or tightly targeted rules, that respond specifically to just the latest incident. Is this a good investment in safety? It may seem like it, but by inserting more, more detailed, or more conditioned rules, procedural overspecification is likely to widen the gap between procedures and practice, rather than narrow it. Rules will increasingly grow at odds with the context-dependent and changing nature of practice.

The reality is that mismatches between written guidance and operational practice always exist. Think about the work-to-rule strike, a form of industrial action historically employed by air traffic controllers, or customs officials, or other professions deeply embedded in rules and regulations. What does it mean? It means that if people don't want to or cannot go on strike, they say to one another: "Let's follow all the rules for a change!" Systems come to a grinding halt. Gridlock is the result. Follow the letter of the law, and the work will not get done. It is as good as, or better than, going on strike.

Seatbelt sign on from top of descent in a VIP jet? The layout of furniture in these machines and the way in which their passengers are pressured to make good use of their time by meeting, planning, working, discussing, does everything to discourage people from strapping in any earlier than strictly necessary. Pilots can blink the light all they want, you could understand that over time it may become pointless to switch it on from 41,000 feet on down.

And who typically employs the pilot of a VIP jet? The person in the back. So guess who can tell whom what to do. And why have the light on only from the top of descent? This is hypocritical—only in the VIP jet upset discussed here was that relevant because loss of control occurred during descent. But other incidents with in-flight deaths have occurred during cruise. Procedures are insensitive to this kind of natural variability.

New procedures can also get buried in masses of regulatory paperwork. Mismatches between procedures and practice grow not necessarily because of people's conscious non-adherence but because of the amount and increasingly tight constraints of procedures.

The vice president of a large airline commented recently how he had seen various of his senior colleagues retire over the past few years. Almost all had told him how they had gotten tired of updating their aircraft operating manuals with new procedures that came out—one after the other—often for no other reason than to close just the next gap that had been revealed in the latest little incident. Faced with a growing pile of paper in their mailboxes, they had just not bothered. Yet these captains all retired alive and probably flew very safely during their last few years.

Adding a bit more technology

More technology does not remove the potential for human error, but relocates or changes it.

A warning light does not solve a human error problem, it creates new ones. What is this light for? How do we respond to it? What do we do to make it go away? It lit up yesterday and meant nothing. Why listen to it today?

What is a warning light, really? It is a threshold crossing device: it starts blinking when some electronic or electromechanical threshold is exceeded. If particular values stay below the threshold, the light is out.

If they go above, the light comes on. But what is its significance? After all, the aircraft has been flying well and behaving normally, even with the light on.

WHY IS THE BAD APPLE THEORY POPULAR?

Cheap and easy

So why would anyone adhere to the bad apple theory of human error? There are many reasons. One is that it is a relatively straightforward approach to dealing with safety. It is simple to understand and simple, and relatively cheap, to implement. The bad apple theory suggests that failure is an aberration, a temporary hiccup in an otherwise smoothly performing, safe operation. Nothing more fundamental, or more expensive, needs to be changed.

A patient died in an Argentine hospital because of an experimental US drug, administered to him and many fellow patients. The event was part of a clinical trial of a yet unapproved medicine eventually destined for the North American market. To many, the case was only the latest emblem of a disparity where Western nations use poorer, less scrupulous, relatively underinformed and healthcare-deprived medical testing grounds in the Second and Third World. But the drug manufacturer was quick to point out that "the case was an aberration" and emphasized how the "supervisory and quality assurance systems all worked effectively". The system, in other words, was safe—it simply needed to be cleansed of its bad apples. The hospital fired the doctors involved and prosecuters were sent after them with murder charges.[3]

Saving face

In the aftermath of failure, pressure can exist to save public image. Taking out defective practitioners is always a good start to saving face. It tells people that the mishap is not a systemic problem, but just a local glitch in an otherwise smooth operation.