ASHGATE

Safety-I and Safety-II

The Past and Future of Safety Management

Erik Hollnagel

Safety–I and Safety–II

The Past and Future of Safety Management

ERIK HOLLNAGEL
University of Southern Denmark

ASHGATE

Published by
Ashgate Publishing Limited
Wey Court East
Union Road
Farnham
Surrey, GU9 7PT
England

Ashgate Publishing Company
110 Cherry Street
Suite 3-1
Burlington, VT 05401-3818
USA

www.ashgate.com

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

The Library of Congress has cataloged the printed edition as follows:
A catalogue record for this book is available from the Library of Congress.

ISBN 978 1 4724 2305 4 (hbk)
ISBN 978 1 4724 2308 5 (pbk)
ISBN 978 1 4724 2306 1 (ebk – PDF)
ISBN 978 1 4724 2307 8 (ebk – ePUB)

MIX
Paper from responsible sources
FSC® C013985
FSC
www.fsc.org

Printed in the United Kingdom by Henry Ling Limited,
at the Dorset Press, Dorchester, DT1 1HD

SAFETY–I AND SAFETY–II

Reviews for *Safety-I and Safety-II*

Much more than a technical book. Erik's work is a well documented journey into the multiple interactions between safety, work and human nature. A timely contribution to vindicate human beings and their variability from the one sided focus on the evils of human error. A groundbreaking look at 'the other story' that will certainly contribute to safer and more productive workplaces.

Dr Alejandro Morales, Mutual Seguridad, Chile

Safety needs a new maturity. We can no longer improve by simply doing what we have been doing, even by doing it better. Dr Hollnagel brings forth new distinctions, interpretations, and narratives that will allow safety to progress to new unforeseen levels. Safety–II is more than just incident and accident prevention. A must read for every safety professional.

Tom McDaniel, Global Manager Zero Harm and Human Performance, Siemens Energy, Inc., USA

Contents

List of Figures

List of Tables

Chapter 1
The Issues

The Need

'Safety' is a word that is used frequently and in many different contexts. Because it is used so often we all recognise it and we all believe that we know what it means – it is immediately meaningful. Because it is immediately meaningful to us, we take for granted that this is the case for others as well. Indeed, when we talk about safety we are rarely, if ever, met with the question 'What do you mean by that?' We therefore make the – unwarranted – inference that other people understand the word 'safety' in the same way that we do. The assumption that we all know and agree on what safety means is so widespread that many documents, standards, guidelines – and even doctoral theses (!) – do not even bother to provide a definition. A search for the etymology of safety, the origin of the word and the way in which its meanings has changed throughout history reveals that it seems to come from the Old French word *sauf*, which in turn comes from the Latin word *salvus*. The meaning of *sauf* is 'uninjured' or 'unharmed', while the meaning of *salvus* is 'uninjured', 'healthy', or 'safe'. (Going even farther back, the roots seem to be in the Latin word *solidus*, meaning 'solid', and the Greek word ηολοσ, meaning 'whole'.) The modern meaning of being safe, as in 'not being exposed to danger', dates from the late fourteenth century; while the use 'safe' as an adjective to characterise actions, as in 'free from risk', is first recorded in the 1580s.

A simple generic definition is that 'safety' means 'the absence of unwanted outcomes such as incidents or accidents', hence a reference to a condition of being safe. A more detailed generic definition could be that *safety is the system property or quality that is necessary and sufficient to ensure that the number of events*

that could be harmful to workers, the public, or the environment is acceptably low. Most people will probably agree to this definition without thinking twice about it. A second look, however, makes it clear that the definition is relatively vague because it depends on expressions such as 'harmful to workers' and 'acceptably low'. Yet because each of us finds these expressions meaningful, even though we do interpret them in our own way, we understand something when we encounter them – and naturally assume that others understand them in the same way. The vagueness of the definition therefore rarely leads to situations where the differences in interpretation are recognised.

Is it Safe?

Few who have seen John Schlesinger's 1976 film *Marathon Man* will ever forget the harrowing scene where the villain, Dr Szell, played by Sir Laurence Olivier, tortures the hero, Dustin Hoffman's character Babe, by probing his teeth. While testing existing cavities and even drilling a hole into a healthy tooth, Dr Szell keeps asking the question, 'Is it safe?' While the question is meaningless for Babe, it refers to whether it will be safe for Dr Szell to fetch a batch of stolen diamonds he has deposited in a bank in New York, or whether he risks being robbed by someone – specifically by Babe's brother, who works as an agent for a secret and mysterious US government agency. For Dr Szell, the meaning of 'safe' is whether something will turn out badly when he carries out his plan, specifically whether the diamonds are likely to be stolen. The meaning is therefore the conventional one, namely whether there is a risk that something will go wrong – whether the planned action will fail rather than succeed. But the question could also have been posed differently, namely whether the planned action, the complicated scheme to recover the diamonds, will succeed rather than fail. Here we find the basic juxtaposition between failure and success, where the presence of one precludes the other. But since the absence (negation) of failure is not the same as success, just as the absence (negation) of success is not the same as failure, it does make a difference whether the focus is on one or the other.

On a less dramatic level there are commonly used expressions such as 'have a safe flight' or 'drive safely back' and 'you will be safe here'. The meaning of the first expression is a hope that a journey by plane will take place without any unwanted or unexpected events, that it will be successful and that you will land in Frankfurt – or wherever – as expected. That flying is safe is demonstrated by the fact that on 12 February 2013 it was four years since the last fatal crash in the US, a record unmatched since propeller planes gave way to the jet age more than half a century ago. (It did not last. On 6 July Asiana flight 214 landed short of the runway in San Francisco, killing three passengers and injuring dozens.) The meaning of the second expression, 'drive safely back', is, again, a hope that you will be able to drive home and arrive without any incidents or problems (but not necessarily without having been exposed to any harm). And the meaning of the third expression, 'you will be safe here', is that if you stay here, in my house or home, then nothing bad will happen to you.

What we mean in general by 'being safe' is that the outcome of whatever is being done will be as expected. In other words, that things will go right, that the actions or activities we undertake will meet with success. But strangely enough, that is not how we assess or measure safety. We do not count the tasks where people succeed and the instances when things work. In many cases we have no idea at all about how often something goes right, how often we have completed a flight without incidents, or driven from one place to another without any problems. But we do know, or at least have a good idea of, how many times something has gone wrong, whether it was a delay, the luggage lost, a near miss with another car, a small collision, or other problems. In other words, we know how often we have had an accident (or incident, etc.), but we do not know how often we have not!

The same discrepancy of focus can be found in the management of safety, whether in the deliberate sense of a safety management system or just in the way that we go about what we do in our daily lives. The focus is usually on preventing or avoiding that something goes wrong, rather than on ensuring that something goes right. While it would seem reasonable, if not outright logical, to focus on a positive outcome qua a positive outcome rather than on the absence of a negative outcome, the professional and

everyday practice of safety seems to think otherwise. Why this is so is explained in Chapter 3.

The Need for Certainty

One of the main reasons for the dominant interpretation of safety as the absence of harm is that humans, individually and collectively, have a practical need to *be* free from harm as well as a psychological need to *feel* free from harm. We need to *be* free from harm because unexpected adverse outcomes can prevent us from carrying out work as planned and from achieving the intended objectives. Hazards and risks are a hindrance for everyday life and for the stability of society and enterprises, as well as individual undertakings. We need to *be* free from harm in order to survive. But we also need to *feel* free from harm because a constant preoccupation or concern with what might go wrong is psychologically harmful – in addition to the fact that it prevents us from focusing on the activities at hand, whether they be work or leisure. There are many kinds of doubt, uncertainty and worries and, while some of them cannot easily be relieved the doubt about why something has gone wrong can – or at least we presume that is the case. Whenever something happens that we cannot explain, in particular if it was accompanied by unwanted outcomes, we try willingly or unwillingly to find some kind of explanation, preferably 'rational' but if need be 'irrational'. The philosopher Friedrich Wilhelm Nietzsche (1844–1900) described it thus:

> To trace something unfamiliar back to something familiar is at once a relief, a comfort and a satisfaction, while it also produces a feeling of power. The unfamiliar involves danger, anxiety and care – the fundamental instinct is to get rid of these painful circumstances. First principle – any explanation is better than none at all.

There are thus both practical and psychological reasons for focusing on things that have gone wrong or may go wrong. There is a practical need to make sure that our plans and activities are free from failure and breakdowns and to develop the practical means to ensure that. But we also have a psychological need for

certainty, to feel that we know what has happened – and also what may happen – and to believe that we can do something about it, that we can master or manage it. Indeed, long before Nietzsche, Ibn Hazm (944–1064), who is considered one of the leading thinkers of the Muslim world, noted that the chief motive of all human actions is the desire to avoid anxiety. This semi-pathological need for certainty creates a preference for clear and simple explanations, expressed in terms that are easy to understand and that we feel comfortable with – which in turn means equally simple methods. It is a natural consequence of these needs that the focus traditionally has been on what I will call the 'negative' side of safety, i.e., on things that go wrong.

Safety as a Dynamic Non-event

One alternative to focusing on unwanted outcomes, which in a very real sense is what safety management does, is, curiously, to focus on what does *not* happen – or rather to focus on what we normally pay no attention to. In an article in *California Management Review* in 1987, professor Karl Weick famously introduced the idea of reliability as a dynamic non-event:

> Reliability is dynamic in the sense that it is an ongoing condition in which problems are momentarily under control due to compensating changes in components. Reliability is invisible in at least two ways. First, people often don't know how many mistakes they could have made but didn't, which means they have at best only a crude idea of what produces reliability and how reliable they are. [...] Reliability is also invisible in the sense that reliable outcomes are constant, which means there is nothing to pay attention to.

This has often been paraphrased to define safety as 'a dynamic non-event', and this paraphrase will be used throughout this book – even though it may be a slight misinterpretation. This is consistent with the understanding of safety as 'the freedom from unacceptable risk' in the sense that a system is safe when nothing untoward happens, when there is nothing that goes wrong. The 'freedom' from the unacceptable risk is precisely the non-event, although it is a little paradoxical to talk about something that is not there. The meaning of 'dynamic' is that the outcome – the

non-event – cannot be guaranteed. In other words, we cannot be sure that nothing will happen. It is not a condition of the system that can be established and then left alone without requiring any further attention. Quite the contrary, it is a condition that must constantly be monitored and managed.

Although the definition of safety as a dynamic non-event is very clever, it introduces the small problem of how to count or even notice or detect a non-event. A non-event is by definition something that does not happen or has not happened. Every evening I could, for instance, rightly ask myself how many non-events I have had during the day? How many times was I not injured at work or did not cause harm at work? How many times did I not say or do something wrong or make a mistake? How many cyclists or pedestrians – or cats or dogs – did I not hit when I drove home from work? But I never do, and I guess that no one ever does.

This problem is not just frivolous but actually real and serious. Consider, for instance, an issue such as traffic safety. Every year the traffic safety numbers are provided in terms of how many people were killed in traffic accidents, either directly or as the result of their injuries. And for a number of years, the trend has been that each year the number of dead has been smaller than the year before (see Figure 1.1). Since traffic safety has adopted the goal of zero traffic deaths, that is a development in the right direction. When I read my daily newspaper, I can see how many people were killed in the Danish traffic in the preceding 24-hour period and how many have been killed in the year so far. (Today, 3 August 2013, the total for the preceding 24-hour period is zero, and the total for the year to date is only 94.) But I cannot find out how many people were *not* killed in traffic. Nobody makes a count of that or has the statistics, perhaps because we take for granted that this is the normal outcome and we therefore concentrate on the opposite cases. But knowing how many were not killed is important because it is necessary to know how serious the problem is. We want to ensure that people can drive safely from A to B. We do want to ensure the non-event, but the question is whether it is best done by preventing the 'bad' events or the traffic deaths (which is what we do and which is how we count) or whether it is best done by furthering the 'good' events – which

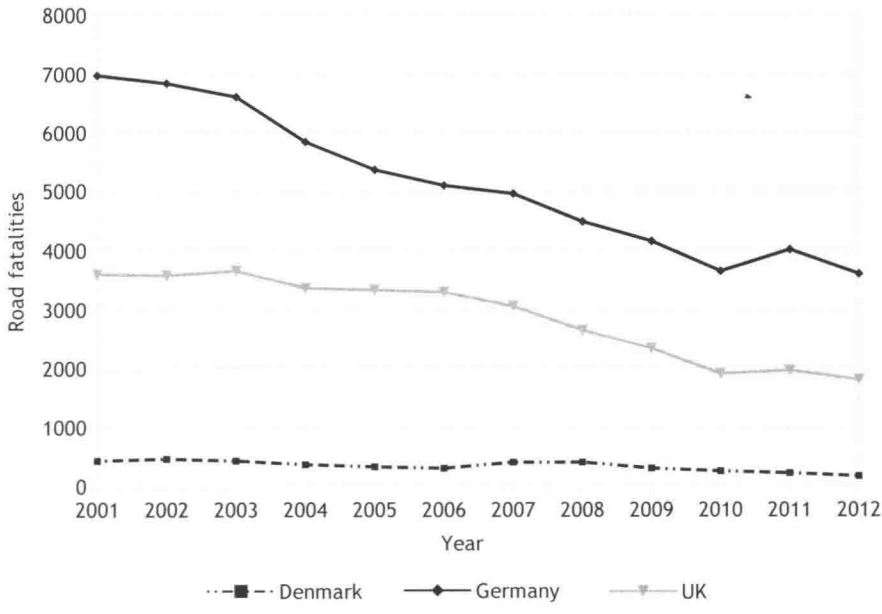

Figure 1.1 Traffic accident statistics for Germany, the UK, and Denmark

is also what we do, but which is not something that we count. Two examples can be used to illustrate these problems.

Signals Passed at Danger

The first example is a train accident that took place in Buizingen, Belgium, on 15 February 2010. Two trains, carrying 250–300 people, collided in snowy conditions during the morning rush hour. The trains apparently collided 'laterally' (i.e., sideways rather than head-on) at a set of points at the exit of Halle station. Eighteen people were killed and 162 injured, with major damage to the tracks as well. The investigation found that one of the trains had passed a red signal without stopping (a situation that happens so often that it has acquired its own name: SPAD or Signal Passed At Danger), and that this could be a contributing cause to the collision, although not the only one.

Further investigation revealed that there were 130 SPAD events in Belgium in 2012, of which one third were serious. (In 2005 there

were 68 SPADs, but the number had been increasing each year.) But it was also estimated that there were about 13,000,000 cases per year of trains stopping at a red signal, which means that the probability of a SPAD is 10^{-5}. A value of 10^{-5} means that the system is safe, although not ultrasafe. However, for activities where humans are involved, 10^{-5} is not unacceptable. (The probability of an accident calculated from the same numbers is 7.7×10^{-8}, which is probably as good as it gets.) In this case it was possible to find, or rather to estimate, how many times the activity succeeded, and thereby get an idea about how serious the event was – not in terms of its outcomes, which were serious and even tragic, but in terms of its occurrence.

In relation to the accident at Buizingen, the fact that 13,000,000 trains stopped at a red signal does not mean that they all stopped in the same way. An elevator is a purely mechanical system that will stop in the same way whenever it reaches the floor it is going to. (Although even here there may be variability due to load, wear and tear, adjustments, maintenance, or other conditions.) But a train is a human–machine system, which means that it is stopped by the train engineer rather than by a mechanism. The way in which a train stops is therefore variable, and it is important to know the variability in order to understand how it is done – and how it can fail. By analogy, try to look at how a driver brakes when the car comes to a red light (but do it from the pavement, not while driving yourself). The way in which a car brakes depends on the load, the driver, the weather, the traffic conditions, etc. It can be, and is, done in many different ways, and all of them usually achieve their goal.

Switching from Left to Right

An interesting, although rather unique, case of a situation where the number of non-events was known with certainty was 'Dagen H' in Sweden. On this day, Sunday, 3 September 1967, Sweden changed from driving on the left-hand side to driving on the right-hand side. In order to accomplish that, all non-essential traffic was banned from the roads from 01:00 to 06:00. Any vehicle on the road during that time, for instance the fire brigade, ambulances, the police and other official vehicles, had to follow special rules.

All vehicles had to stop completely at 04:50, then carefully change to the other side of the road and remain there until 05:00, when they were allowed to proceed. (In the major cities, the driving ban was considerably longer in order for working crews to have sufficient time to reconfigure intersections.)

Because there was no traffic, or only so little traffic that it in principle could be monitored and accounted for, it can be said with certainty that there were no non-events during the change. Or at least the number of non-events was countable, had anyone bothered to count them. And since there were no non-events, since cars were not allowed to drive, there could not be any events either, i.e., no collisions between cars. (In absolute terms this lasted only for the ten minutes from 04:50 to 05:00, but in practice it lasted for the five hours between 01:00 and 06:00.)

Even when people agree that safety is a dynamic non-event, the practice of safety management is to count the events, i.e., the number of accidents, incidents, and so forth. By doing that we know how many events there have been, but not how many non-events. We may, however, easily turn the tables, by defining safety as a dynamic *event*. The event is now that an activity succeeds or goes well (that we come home safely, that the plane lands on time, etc.), and we are obviously safe when that happens. The non-event consequently becomes the situation when this does *not* happen, i.e., when things go wrong. We can count the non-events, i.e., the non-successes or failures, just as we have usually done. But we can now also count the events, the number of things that go right, at least if we make the effort.

The Measurement Problem

In order to know that we are safe – not just subjectively or psychologically, but also objectively or practically – industry and society need some way of demonstrating the presence of safety. In practice this means that there must be some way of quantifying safety.

Strictly speaking, it must be possible to confirm the presence of safety by means of *intersubjective verification*. To the extent that safety is an external, public phenomenon, the way in which it is experienced and described by one individual must correspond