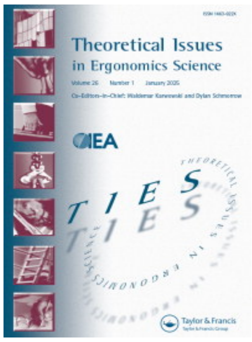

Taylor & Francis
Taylor & Francis Group

Theoretical Issues in Ergonomics Science

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/ttie20

The ironies of ‘human factors’

Erik Hollnagel & Sidney W. A. Dekker

To cite this article: Erik Hollnagel & Sidney W. A. Dekker (20 Dec 2024): The ironies of ‘human factors’, Theoretical Issues in Ergonomics Science, DOI: 10.1080/1463922X.2024.2443976

To link to this article: https://doi.org/10.1080/1463922X.2024.2443976

THEORETICAL ISSUES IN ERGONOMICS SCIENCE
https://doi.org/10.1080/1463922X.2024.2443976

RESEARCH ARTICLE

The ironies of 'human factors'

Erik Hollnagel[a] and Sidney W. A. Dekker[b]

[a]University of Southern Denmark, Odense, Denmark; [b]School of Humanities, Languages and Social Science, Griffith University, Brisbane, QLD, Australia

ABSTRACT
The term irony is here used in the sense pioneered in 1983 by Lisanne Bainbridge, to describe a solution which increases rather than reduces a problem. Bainbridge used the term in relation to automation, but it can be applied to other issues, particularly in how human factors engineering relies on training, procedures, design and automation as its main approaches to managing human variability. 'Human factors' tends to consider human agility or performance variability as a liability that should either be eliminated or brought under control. The paper encourages us to recognise that variability is an indispensable asset, without which few of the common human factors solutions would ever work.

ARTICLE HISTORY
Received 7 August 2024
Accepted 14 December 2024

KEYWORDS
Ironies; human factors; work-as-imagined; work-as-done; training; procedures; design; automation

Introduction

In1983 the British psychologist Lisanne Bainbridge published 'The ironies of automation', originally in the journal *Automatica* and soon after as a highly-cited book chapter (Bainbridge 1987). The first irony was that the more advanced a control system is, the more crucial the contribution of the human operator would be. Bainbridge noted that human supervisory controllers cannot really be unskilled in their monitoring of whether the automation is carrying out its work correctly, because they wouldn't know what they were looking at (or for), and would be useless in case they would need to take over (Billings 1997; Dekker and Woods 1999). The second irony was that the designer who tried to eliminate the operator still leaves the operator to do the tasks which the designer could not think how to automate—often without adequate support. The third irony was that the automatic control system was introduced because it could do the job more reliably, or cheaper, than the operator, yet the operator was asked to monitor that it worked effectively.

For the current paper we adopt Bainbridge's definition of an irony as a solution which expands rather than solves or eliminates a problem, thus making something worse rather than better. As Bainbridge did for automation, we do for human factors as a field in a broader sense—identifying and discussing its basic ironies. As we do so, we pull in more recent research on adaptive automation, cognitive systems, and resilience engineering, to explore how such ironies might shine through in the field's attempts to overcome its contemporary challenges.

CONTACT Sidney W. A. Dekker s.dekker@griffith.edu.au School of Humanities, Languages and Social Science, Griffith University, Macrossan Building N16, 170 Kessels Road, Nathan Campus, QLD 4111, Australia.

Fundamentally, the ironies of human factors coagulate around the widely-shared (if not always explicit) view of humans as parts or components of work systems. This reduction applies a physical mechanical analogy which goes back at least to Newton, and more specifically Julien Offray de la Mettrie in 1747 in his treatise 'Man a Machine', even though thinking of humans as intricate mechanisms or machines predates even de la Mettrie by millennia (cf. Leonardo da Vinci, or even earlier mentions of the Golem (Wiener 1964)). Humans, in this conception, are co-constitutive of the larger work system where they can be treated as if they were mechanical parts or sub-systems.

Humans are attributed a role both as a mechanical component themselves and as genuine parts of a larger socio-technical system or workplace. The temptation to think of humans as a machine only grew stronger and became nearly irresistible when digital computing machinery became a reality around the middle of the twentieth century.

The comparison between humans and machines has always been made on technological premises, and was therefore rarely advantageous to humans. An early expression of that view was provided by Paul Fitts, undisputed pioneer of Human Factors Engineering, who argued that 'the final consideration which needs mention is the relative fallibility of a man to a machine' (Fitts 1951, p. 6). The notion that humans at best are fallible machines has been legitimate across the field ever since. Fitts also noted that 'We have been very much occupied in perfecting the machines and tools which the worker uses in the economic arts', and went on to say that 'we have hardly attempted to improve the worker himself' (Fitts 1951, p. iv). His interests led him to propose a rigorous method to compare humans and machines, which we now refer to as the Fitts' list:

> We begin with a brief analysis of the essential functions … We then consider the basic question: Which of these functions should be performed by human operators and which by machine elements? (Fitts 1951 p. x).

This approach later became known as the MABA-MABA list after the initials in (Men Are Better At)—(Machines Are Better At) (Dekker and Woods 2002; Sanders and McCormick 1992). The generic set of solutions and approaches to improving the worker, which Fitts was instrumental in developing, became known as Human Factors Engineering. Note that Fitts simply—and perhaps more honestly—called it Human Engineering. Human engineering from the very beginning needed ways to enhance and improve the human factor, not only to ensure a better fit between humans and the machines they had to work with, but to make humans perform more reliably like the machines or technological artefacts they were compared to. Fitts' three main remedies to engineer the human, from the beginning, were training, design, and automation.

Training was to be used to shape humans to meet the requirements of technology and was initially seen as an ideal and inexpensive solution, which unlike design and automation had a pedigree stretching back thousands of years. Training was also widely practiced by most industries and therefore familiar. Design was used to ensure that the requirements of machines corresponded to the 'natural' abilities of humans, as determined e.g. by the Fitts' list. And automation was the ultimate solution of replacing unreliable humans by reliable technology, thereby ensuring a 'perfect' match between system components. The three solutions were adopted from the start that have been used ever since, although in various proportions since they are not mutually independent.

The ironies of training

It did not take long before it became obvious that training was not the perfect way to engineer the human factor and overcome the problems Fitts wanted to solve. This was due in part to performance variability, and insufficient reliability although it was not described as an irony at the time. The Achilles heel of human factors, as a discipline, is that it has adopted the machine analogy as a universal explanatory and analytical principle for what a human is, and can and cannot do. Because of this technological bias, humans were from the very beginning seen as inefficient, variable, and unreliable. This view has led to the perceived need of human factors engineering as in the Fitts quote above. It is an analogy which still is widely, if quietly, accepted. Humans are seen as a liability in consequence of which the human factor *per se*, and specifically human performance variability, became problems which had to be addressed:

> We have been very much occupied in perfecting the machines and tools which the worker uses in the economic arts. We have hardly attempted to improve the worker himself (Fitts 1951, p. iv).

The premise, or foregone observation, was that the human operator was unreliable and inefficient in comparison to machines. These human imperfections made it impossible to fully exploit the potential offered by new technologies (Burnham 2009). Where this deficit could not be solved by putting in even more technology—particularly autonomous systems—training was the way to apply 'human factors.' Training can be portrayed as a 'Procrustean' approach (Hollnagel and Woods 1983), referring to the legendary Procrustes, a robber of Attica, who had an Inn on the sacred road between Athens and Eleusis (home of the Eleusian mysteries). The Inn had a single iron bed, where Procrustes invited everyone who passed by to spend the night. Procrustes became famous for making the visitors fit the bed, rather than the other way around. If visitors were shorter than the bed, Procrustes would stretch them until they were long enough, and if they were too long he would make them fit the bed by having parts of their legs lopped off.

In 1959, Taylor and Garvey (see Hollnagel and Woods 1983) used this analogy to criticise training when they wrote that two rather different human factor approaches could be distinguished in efforts to optimise the performance of man-machine systems. One sought to standardise performance by 'making humans shorter,' i.e. limit or constrain what they should do, i.e. using less than their full potential by limiting or constraining their performance. The other effort was aimed at 'making humans longer' i.e. extend or stretch human capabilities to meet task demands through additional specialised training. It has produced at least three ironies.

The idea that people should have just the competence needed to do the job, but no more because it would be a waste of effort, got a new lease on life with the introduction of *Scientific Management* widely known as *Taylorism* (Taylor 1911). Scientific management claimed that performance variability could—and certainly should—be reduced. Work reliability, quality and efficiency would be improved by ensuring that people did exactly what they were told to do; neither more nor less than requirements of the system demanded. The solution of 'making people shorter' resembles the third principle of Taylorism, listed below (Taylor 1911, p. 14):

1. Analyse tasks to determine most efficient performance;
2. Select people to achieve best match between task requirements and capabilities;
3. Train people to ensure the specified and required performance but nothing more and nothing less than that;
4. Insure compliance by economic incentives.

Training can also be used to 'make humans longer' i.e. to extend or stretch their capabilities and skills beyond what they would naturally or normally do through long and specialised training. This has become increasingly necessary as we build human-machine systems and work environments that no longer are intuitive despite attempts such as ecological interface design (Flach et al. 1996; Vicente 1999). Few people can naturally or intuitively fly a modern aircraft, or control a nuclear power plant, and perhaps not even drive a present-day EV without some additional 'stretch' of their knowledge, skills and capabilities. This realization has come to the fore recently in the two fatal accidents involving the latest version of the Boeing 737 (the MAX). The manufacturer's (and indeed, many operators' implicit) insistence that no training should be required for pilots transitioning from earlier 737 models became directly implicated in the two accidents. Pilots didn't stand much of a chance in the face of an automation take-over and nose-dive by hidden software (cf. Dekker and Woods 1999; Sarter, Woods, and Billings 1997). This software was, ironically, meant to smoothen out the difference between the MAX and previous generations for certification purposes, but pilots (and airlines) were never made aware of its existence. It denied them the chance to build out their variability so that they might have stretched their capabilities to meet its potentially fatal automation surprise (Defazio and Larsen 2020; Dekker, Layson, and Woods 2022; Herkert, Borenstein, and Miller 2020).

The irony of procedures

Procedures have, since the early years of HFE, been an essential solution (Brown, Moran, and Williams 1982; Degani, Heymann, and Shafto 1999; IFALPA 2005; Rasmussen and Jensen 1974). There is, however, always a difference what people actually do (Work-as-Done, or WAD) and what they are or were supposed to do (Work-as-Imagined or WAI), which is the opposite of the intended outcome (De Keyser, Decortis, and Van Daele 1988; Hollnagel 2012b):

- Work-as-imagined (WaI) represents the various assumptions, explicit or implicit, that people have about how their own work as well as the work of others' (e.g. co-workers or team mates) should be done and what others do, for instance how people at the sharp end think of people at the blunt end, and vice versa. WaI is related to the concept of requisite imagination. WaI represents what we at present are able to imagine about the future; about another time and place.
- Work-as-done (WaD) represents how something is actually done, either in a specific case or more routinely the typical or habitual way of carrying out a prescribed activity, and therefore represents what ought to have been imagined.
- There will for a number of reasons always be a difference between how work is 'imagined' or thought of and how work is actually done.

- The solution to this difference is to try to understand what determines how work is done and to find effective ways of managing that to keep the variability of WaD within acceptable limits. (But not by constraints and compliance, or, in other words, by making people 'shorter'.

Efforts to improve the uniformity of work in practice often rely on standard procedures, and it is generally taken for granted that the standards themselves (a form of WaI, after all) are complete and correct. Following the guidance provided by the standards to the letter is therefore assumed to compensate for human shortcomings and performance variability and result in work that is correct and flawless as long as work and workers comply with the standards. There is usually an insistence on compliance to standards, where non-compliance is a frequently used explanation when something has gone wrong. Of course, the irony is that there is ample evidence that non-compliance, particularly in non-standard situations, can be safer than sticking with the rules (Bieder and Bourrier 2013; Carim et al. 2016; Dekker 2001, 2003; McGinty 2008), if anything because there is no appropriate WaI script for the situation facing the humans: as in the MAX case cited above, nobody (at least nobody in the cockpit) could have 'imagined' the work that needed doing to save the day.

The ironies of design

To improve the performance characteristics of human-machine combinations, the choice is thus either trying to alter the human so that they fit the machine better, or modifying the machine to fit the human. This very dichotomy was of course the origin of 'human engineering' in the first place, departing from behaviourism in the 1940s: the world and its machines weren't fixed just for people to have to adapt to: they could be changed, modified and improved to fit human characteristics (Fitts 1951; Fitts and Jones 1947; Roscoe 1997). The logical alternative was to design the technology so that the requirements correspond to how people naturally perform, requiring neither too much, nor too little (Hollnagel 1993).

The systematic study of Human-Machine Interaction (HMI) arose from the need to solve practical, technological problems, but the basis for the description was, as ever, the engineering view of humans as and machines. The technical and engineering fields had already developed a powerful vocabulary to describe how machines worked and it therefore seemed natural to apply the same vocabulary to how people worked—indeed to use this as a basis for modeling human performance. One consequence of this was the forced automaton analogy, for example to think of or describe human (cognitive) functioning by use of the information processing metaphor (Newell and Simon 1972; Wickens 1984), even though the automaton analogy can be found in practically every explanation of human performance including behaviourism and psychoanalytic theory (Watson 1978).

Some have helpfully conceived of design as telling stories about the future (Carroll and Campbell 1988; Roesler et al. 2001). Design is about shaping something, an artefact, a work process, a task or job, a machine, or a tool, that does not yet exist, in particular shaping how it is going to be used or how it should be used for the joint cognitive system to achieve the goals it intends to achieve. Designing an artefact cannot avoid also being Cognitive Task Design (Hollnagel 2003), but even if as we do so, telling stories about the future affects or changes the future. The specific problem here is that it proposes improvements to a situation and set of working conditions that are imprecisely known, and not precisely knowable (cf. WaI versus WaD).

An irony of design is that the inevitable differences WaI and WaI weaken the very basis from which the design is made. It is inevitable that a design cannot solve the problem it was supposed to address but possibly make it worse. By creating more uncertainty it can ironically increase rather than reduce the need for human performance variability to fill out the gaps between design (WaI) and reality (WaD). An extensive study of the necessary adaptations in the operating room by anaesthetic teams to accommodate new technologies for monitoring patients on cardiopulmonary bypass (Cook and Woods 1996) revealed how practitioners needed to engage in an extensive workaround before each procedure so that they were able at all to monitor (and respond to) rapid changes in blood pressure. The spontaneously emerging set-up procedure was elaborate:

> Identifying which transducer was connected to which physical channel was accomplished before the preanesthesia phase: Technicians plugged in the cables, assigned the labels to each channel, mechanically and electronically zeroed the transducers, assigned the channels to appear on the fixed-scale window, and finally called that window to the screen (Ibid, p. 600).

Practitioners took to the continuously reconfigured display so much that they started referring to it as the ‘normal’ screen, despite the fact that it was anything but, and that it needed significant set-up and configuration to produce—every time again.

The solution for human factors, at least in theory, is neither to force WaD to comply with WaI—as in the Zero Accident Vision (Zwetsloot et al. 2013) or, say, quality (Amalberti et al. 2005), nor is it to constrain WaI so that it corresponds to WaD as Scientific Management tried to do. A solution is rather to try to understand what determines how work is done and to find effective ways of managing that to keep the variability of WaD within acceptable limits (but not by using constraints and compliance to make humans ‘shorter’). This notwithstanding, there remains an irreducible difference between WaD and WaI, due to complexity, our limited imagination and the tradition to learn from what has gone wrong rather than what has gone well (Hollnagel 2008). The systems we design, of course, end up mostly functioning not because of their design or our ability to anticipate WAI, but because human performance variability compensates for design deficiencies and the limits of our imagination about the working conditions it will encounter (Cook and Woods 1996; Hollnagel 2012a).

The irony is that attempts at eliminating human performance variability—provided this was even possible—would not solve the problems of imprecision, of the nuances and ‘messy details’ of what actual practice requires, or of variability and lack of speed. Instead, it might only make them larger, and confront us with the realization that we cannot do without the ‘human factor’ which we have been trying to get rid of. Performance agility is necessary whatever the mode of operation, as the conceptual putty that fills out the cracks between WaI and WaD. Human factors as a field has been preoccupied with the legacy view of humans as a liability and source of error and variability. But again, without this variability, no system would be able to function. Performance agility is the putty for the inevitable discrepancy between WaI and WaD. Doubly ironic is that trying to constrain human performance by training or design, leads to an increase rather than a reduction of human performance variability.

The ironies of automation (redux)

An automaton can be described by a set of inputs, outputs, internal states and the corresponding state transitions; a classical example of this is the Turing machine. More formally, a finite automaton is a quintuple:

$$A = \left(We, O, S, \lambda, \delta\right)$$

where We is the set of inputs,

O is the set of outputs

S is the set of internal states

λ: S x We→ S is the set of rules for determining the next state, and

δ: S x We → O is the set of rules for determining the next output.

A machine, or a program, can be described as a finite state automaton or a state machine in terms of the quintuple defined above (a set of inputs, outputs, internal states, and state transitions: all programming languages are actually based on that assumption). In order for the machine to work and to produce a predefined output, it must get a correct input. This means that the human user must respond in a way that corresponds to the predefined categories of inputs that the machine can recognise. If not, the machine, hence the joint system (Hollnagel and Woods 1983) cannot function appropriately; it will become ineffective or even unreliable (Sarter and Woods 1997).

Assume, for instance, that the automaton is in a state S_j. From that state it can, as defined by λ, progress to a predefined set of other states (S_k, …, S_m) or produce a predefined output (O_j) as defined by δ only if it gets the correct input (I_j). If the system in question is a joint human-machine system then the input to the automaton is provided by the human. The human must provide an input that the automaton can interpret. Any other response by the user will lead to one of two cases:

1) The input may not be understood by the automaton, i.e. transition functions λ and δ are not defined for the input. In this case the automaton may either do nothing or move to a default state.
2) The input may be misunderstood by the automaton, for instance if the semantics or the syntax of the response have not been rigorously defined. In this case the automaton may possibly malfunction, i.e. be forced to a state which has not been anticipated. This happens easily if the input was a physical control action or manipulation, e.g. like switching something on or off, opening a valve, etc.

If the Human-Machine System is to function properly, the human user must provide a response that falls within a limited set of pre-defined responses (the set of recognisable inputs). But the determination of the human user's response is a function of the information available. The output from the machine constitutes the input to the user, which itself is a function of the previous input i.e. what the user did. Current operations therefore depend on previous output from the machine, or on what has happened before. The content and structure of the machine's output must therefore be correctly understood by the user, i.e. correctly interpreted and mapped onto one of the predefined answer options (Sarter and Woods 1995). In order to do that, the designer has to consider the user as a finite state automaton, as a machine. Of course, people may interpret information in many different ways depending on the context. And there is no way in which we can possibly account for this infinity of interpretations (Work-as-Imagined or WAI will always be different from Work-as-Done or WAD) (Cilliers 2002). Work-as-Done remains a moving target because internal and external working conditions (demands, and resources) never are stable or fully predictable. The design requires that the user interprets the information in a limited number of (pre-) specified ways. Design, in other words, forces the human user to function as a finite state automaton (An everyday example of that is the menu-based user interfaces—visual or auditory—which helps restrict the user's degrees of freedom).

An issue of automation which Bainbridge did not address explicitly is the substitution principle (Hollnagel 1999). The substitution principle expresses the common assumption that artefacts are neutral in their effects and that their introduction into a system therefore only has intended and no unintended consequences. The basis for this principle is the concept of interchangeability (the same basis from which Fitts' or MABA-MABA lists operate). In general, however, substitutability only works when parts are not interacting and when there is no appreciable tear and wear. If parts are interacting, they constitute a system with (inter-)dependencies, which almost by definition invalidates the substitution assumption.

The irony is not necessarily that the automaton analogy is ineffectual as a basis for describing human performance. The irony, rather, is that the automaton analogy is also useless for describing machines in the context of human-machine systems where the functioning of the machine must be seen together with the functioning of a person (Hollnagel and Woods 2005; Roth, Bennett, and Woods 1987). It probably the case that whatever analogy we use for one (machine or human), we will have to use for the other (human or machine) as well. Because we want to retain some distinct human elements in the description of the user we are forced to apply the same elements to the description of the machine (as a kind of forced anthropomorphism). This was of course the reason behind the idea of 'joint cognitive systems:'

> A cognitive system produces intelligent action, that is, its behaviour is goal oriented, based on symbol manipulation and uses knowledge of the world (heuristic knowledge) for guidance Furthermore a cognitive system is adaptive and able to view a problem in more than one way. A cognitive system operates using knowledge about itself and the environment in the sense that it is able to plan and modify its actions on the basis of that knowledge (Hollnagel and Woods 1983, p. 589).

Conclusion: worker-as-imagined

Ironies remain in how human factors engineering relies on training, procedures, design and automation as its chief approaches to managing human variability. In addition to 'Work-as-Imagined', we have created a kind of 'Worker-as-Imagined,' who has to make up for the inevitable shortcomings of our own imagination as we design artifacts and attempt to automate more of the work that gets done with them. The Worker-as-Imagined is redolent with paradoxes: compliant for the most part but creative when necessary (though the second part barely has a formal role in design processes or decisions). 'Human factors' tends to consider human agility or performance variability as a liability that should either be eliminated or brought under control. The arguments in this paper might nudge us to once again recognise that variability is an indispensable asset, without which few of the common human factors solutions would ever work. To learn about its necessity, and appreciate the dynamic sacrifices and tradeoffs that get made in Work-as-Done all the time (Rasmussen 1997), the ability to somehow put oneself in the perspective of those who carry out the work is critical. Curiosity can replace judgment about what work is expected or 'should' be done (Havinga, Dekker, and Rae 2018). In recent research projects (Rae, Weber, and Dekker 2021), we have augmented the typical tools of cognitive work analysis and cognitive task design (Vicente 1999) with questions such as:

- I notice that… Help me understand why it makes (more) sense to work this way;
- What are the obstacles we're putting in your way to getting things done?
- What is the stupidest thing we're asking you to do to get this thing to work?

With answers to questions such as these, we've learned that a more compassionate human factors poise can emerge. It avoids the construction of a new project or design on some 'human-as-imagined:' for by the time it is done and implemented, all that is left for us is to hurl accusations at the user about why they couldn't be more perfect, as some indeed did in the wake of the Boeing 737 MAX crashes and in response to earlier sentinel accidents (NTSB 2009) which might have pointed to the eventual MAX disasters (Englehardt, Werhane, and Newton 2021; Tkakic 2019). This not only respects the humanity of everyone involved in the ergonomic enterprise; it also allows us early on to capture the systemic factors that contribute to human-machine breakdowns and failures (Dekker 2024).

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

Amalberti, R., Y. Auroy, D. Berwick, and P. Barach. 2005. "Five System Barriers to Achieving Ultrasafe Healthcare." *Annals of Internal Medicine* 142 (9): 756–764. https://doi.org/10.7326/0003-4819-142-9-200505030-00012.

Bainbridge, L. 1987. "Ironies of Automation." In *New Technology and Human Error*, edited by J. Rasmussen, K. Duncan, and J. Leplat, 271–283. Chichester, England: Wiley.

Bieder, C., and M. Bourrier, eds. 2013. *Trapping Safety into Rules: How Desirable or Avoidable is Proceduralization?* Farnham, UK: Ashgate Publishing Co.

Billings, C. E. 1997. *Aviation Automation: The Search for a Human-Centered Approach*. Mahwah, NJ: Lawrence Erlbaum Associates.

Brown, J. S., T. P. Moran, and M. D. Williams. 1982. *The Semantics of Procedures*. Palo Alto, CA: Xerox Palo Alto Research Center.

Burnham, J. C. 2009. *Accident Prone: A History of Technology, Psychology and Misfits of the Machine Age*. Chicago, IL: The University of Chicago Press.

Carim, G. C., T. A. Saurin, J. Havinga, A. Rae, S. W. A. Dekker, and É. Henriqson. 2016. "Using a Procedure Doesn't Mean following It: A Cognitive Systems Approach to How a Cockpit Manages Emergencies." *Safety Science* 89: 147–157. https://doi.org/10.1016/j.ssci.2016.06.008.

Carroll, J. M., and R. L. Campbell. 1988. "*Artifacts as Psychological Theories: The Case of Human-Computer Interaction*." *IBM Research Report RC 13454*

Cilliers, P. 2002. "Why we Cannot Know Complex Things Completely." *Emergence* 4 (1): 77–84. https://doi.org/10.1207/S15327000EM041&2-07.

Cook, R. I., and D. D. Woods. 1996. "Adapting to New Technology in the Operating Room." *Human Factors* 38 (4): 593–613. https://doi.org/10.1518/001872096778827224.

De Keyser, V., F. Decortis, and A. Van Daele. 1988. “The Approach of Francophone Ergonomy: Studying New Technologies.” In *The Meaning of Work and Technological Options*, edited by V. De Keyser, T. Qvale, B. Wilpert, and S. A. Ruiz-Quintallina, 148–163. New York, NY: Wiley.

Defazio, P. A., and R. Larsen. 2020. *The Design, Development and Certification of the Boeing 737 MAX (Final Committee Report)*. Washington, DC: The House Committee on Transportation and Infrastructure.

Degani, A., M. Heymann, and M. Shafto. 1999. “Formal Aspects of Procedures: The Problem of Sequential Correctness.” 43rd Annual Meeting of the Human Factors and Ergonomics Society, Houston, TX.

Dekker, S. W. A. 2001. “Follow the Procedure or Survive.” *Human Factors and Aerospace Safety* 1 (4): 381–385.

Dekker, S. W. A. 2003. “Failure to Adapt or Adaptations That Fail: Contrasting Models on Procedures and Safety.” *Applied Ergonomics* 34 (3): 233–238. <Go to ISI>://WOS:000183149300004 https://doi.org/10.1016/S0003-6870(03)00031-0.

Dekker, S. W. A. 2024. “Forgiveness as Morally Serious Response to Errors in Healthcare: A Narrative Review.” *Journal of Hospital Administration* 13 (2): 52–58. https://doi.org/10.5430/jha.v13n2p52.

Dekker, S. W. A., M. D. Layson, and D. D. Woods. 2022. “Repentance as Rebuke: Betrayal and Moral Injury in Safety Engineering.” *Science and Engineering Ethics* 28 (6): 56–56. https://doi.org/10.1007/s11948-022-00412-2.

Dekker, S. W. A., and D. D. Woods. 1999. “To Intervene or Not to Intervene: The Dilemma of Management by Exception.” *Cognition, Technology & Work*, 1 (2): 86–96. https://doi.org/10.1007/s101110050035.

Dekker, S. W. A., and D. D. Woods. 2002. “MABA-MABA or Abracadabra? Progress on Human-Automation Co-Ordination.” *Cognition, Technology & Work* 4 (4): 240–244. https://doi.org/10.1007/s101110200022.

Englehardt, E., P. H. Werhane, and L. H. Newton. 2021. “Leadership, Engineering and Ethical Clashes at Boeing.” *Science and Engineering Ethics* 27 (1): 12–12. https://doi.org/10.1007/s11948-021-00285-x.

Fitts, P. M. 1951. *Human Engineering for an Effective Air Navigation and Traffic Control System*. Washington, DC: National Research Council.

Fitts, P. M., and R. E. Jones. 1947. “*Analysis of Factors Contributing to 460 “Pilot Error” Experiences in Operating Aircraft Controls (TSEAA94-12)*.” Aero Medical Laboratory, Air Material Command, Wright-Patterson Air Force Base.

Flach, J. M., P. A. Hancock, J. Caird, and K. J. Vicente. 1996. *An Ecological Approach to Human-Machine Systems I: A Global Perspective*. Mahwah, NJ: Lawrence Erlbaum Associates.

Havinga, J., S. W. A. Dekker, and A. J. Rae. 2018. “Everyday Work Investigations for Safety.” *Theoretical Issues in Ergonomics Science* 19 (2): 213–228. https://doi.org/10.1080/1463922X.2017.1356394.

Herkert, J., J. Borenstein, and K. Miller. 2020. “The Boeing 737 MAX: Lessons for Engineering Ethics.” *Science and Engineering Ethics* 26 (6): 2957–2974. https://doi.org/10.1007/s11948-020-00252-y.

Hollnagel, E. 1993. *Human Reliability Analysis: Context and Control*. San Diego, CA: Academic Press.

Hollnagel, E. 1999. “From Function Allocation to Function Congruence.” In *Coping with Computers in the Cockpit*, edited by S. W. A. Dekker and E. Hollnagel, 29–53. Aldershot, UK: Ashgate.

Hollnagel, E. 2003. *Handbook of cognitive task design*. Lawrence Erlbaum Publishers. http://www.loc.gov/catdir/enhancements/fy0709/2003042401-d.html

Hollnagel, E. 2008. “Investigation as an Impediment to Learning.” In *Resilience Engineering Perspectives: Remaining Sensitive to the Possibility of Failure*, edited by E. Hollnagel, C. P. Nemeth, and S. W. A. Dekker, 259–268. Farnham, UK: Ashgate Publishing Co.

Hollnagel, E. 2012a. *FRAM the Functional Resonance Analysis Method: Modelling Complex Socio-Technical Systems*. Boca Raton, FL: CRC Press.

Hollnagel, E. 2012b. “Resilience Engineering and the Systemic View of Safety Work: Why Work-as-Done is Not the Same as Work-as-Imagined.” In *Gestaltung Nachhaltiger Arbeitssysteme: Wege Zur Gesunden, Effizienten Und Sicheren Arbeit*, edited by G. f. A. e. V. (GfA), 58. Arbeitswissenschaftlichen Kongress GfA (Hrsg.), Kassel, Germany.

Hollnagel, E., and D. D. Woods. 1983. "Cognitive Systems Engineering: New Wine in New Bottles." *International Journal of Man-Machine Studies* 18 (6): 583–600. <Go to ISI>://WOS:A1983QV54000005 https://doi.org/10.1016/S0020-7373(83)80034-0.

Hollnagel, E., and D. D. Woods. 2005. *Joint cognitive systems: Foundations of cognitive systems engineering*. Taylor & Francis. http://www.loc.gov/catdir/enhancements/fy0647/2004064949-d.html

IFALPA. 2005. "*Design Deficiency Makes Primus II Transponder Mitigation Procedures a "Vital Action*." Safety Bulletin, Issue. IFALPA.

McGinty, S. 2008. *Fire in the Night: The Piper Alpha Disaster*. London, UK: Pan Books.

Newell, A., and H. A. Simon. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

NTSB. 2009. "*U.S. Summary Comments on Draft Final Report of Aircraft Accident Turkish Airlines Flight 1951*." *Boeing 737-800, TC-JGE February 25, 2009, Amsterdam Schiphol Airport (EHAM).*

Rae, A., D. Weber, and S. Dekker. 2021. "Work as Planned, as Done and as Desired." In *Inside Hazardous Technological Systems: Methodological Foundations, Challenges and Future Directions* edited by K. P. Gould and C. Macrae, 115–132. Routledge. https://doi.org/10.1201/9780429281587-8.

Rasmussen, J. 1997. "Risk Management in a Dynamic Society: A Modelling Problem." *Safety Science* 27 (2-3): 183–213. https://doi.org/10.1016/S0925-7535(97)00052-0.

Rasmussen, J., and A. Jensen. 1974. "Mental Procedures in Real-Life Tasks: A Case Study of Electronic Trouble Shooting." *Ergonomics* 17 (3): 293–307. https://doi.org/10.1080/00140137408931355.

Roesler, A., M. Feil, A. Puskeiller, D. D. Woods, and D. Tinapple. 2001. *Design is Telling Stories about the Future*. Columbus, OH: Cognitive Systems Engineering Laboratory, The Ohio State University.

Roscoe, S. N. 1997. "The Adolescence of Engineering Psychology." In *Human Factors History Monograph Series*, edited by S. M. Casey, Vol. 1, 1–9. Santa Monica, CA: Human Factors and Ergonomics Society.

Roth, E. M., K. B. Bennett, and D. D. Woods. 1987. "Human Interaction with an "intelligent" Machine." *International Journal of Man-Machine Studies* 27 (5-6): 479–525. https://doi.org/10.1016/S0020-7373(87)80012-3.

Sanders, M. S., and E. J. McCormick. 1992. *Ergonomics in Engineering and Design*. 7th ed. New York, NY: McGraw-Hill.

Sarter, N. B., and D. D. Woods. 1995. "How in the World Did we Get into That Mode?" Mode Error and Awareness in Supervisory Control." *Human Factors: The Journal of the Human Factors and Ergonomics Society* 37 (1): 5–19. <Go to ISI>://WOS:A1995RL73500002 https://doi.org/10.1518/001872095779049516.

Sarter, N. B., and D. D. Woods. 1997. "Teamplay with a Powerful and Independent Agent: A Corpus of Operational Experiences and Automation Surprises on the Airbus A320." *Human Factors* 39: 553–569.

Sarter, N. B., D. D. Woods, and C. E. Billings. 1997. "Automation Surprises." In *Handbook of Human Factors and Ergonomics*, edited by G. Salvendy, 2nd ed., 1926–1943. New York, NY: Wiley.

Taylor, F. W. 1911. *The Principles of Scientific Management*. New York, NY: Harper & Brothers.

Tkakic, M. 2019. "Crash Course: How Boeing's Managerial Revolution Created the 737MAX Disaster." *New Republic* 105 (10): 5–31.

Vicente, K. J. 1999. *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Lawrence Erlbaum Associates. http://www.loc.gov/catdir/enhancements/fy0740/98031670-d.html

Watson, R. I. 1978. *The Great Psychologists*. 4th ed. Philadelphia, PA: Lippincott.

Wickens, C. D. 1984. *Engineering Psychology and Human Performance*. Columbus, OH: Merrill.

Wiener, N. 1964. *God and Golem*. Cambridge, MA: M.I.T. Press.

Zwetsloot, Gerard I. J. M., Markku Aaltonen, Jean-Luc Wybo, Jorma Saari, Pete Kines, and Rik Op De Beeck. 2013. "The Case for Research into the Zero Accident Vision." *Safety Science* 58: 41–48. https://doi.org/10.1016/j.ssci.2013.01.026.