

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/153206>

How to cite:

Please refer to published version for the most recent bibliographic citation information.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Safety cases for digital health innovations: can they work?

Mark Sujan^{1,2} and Ibrahim Habli³

¹ Warwick Medical School, University of Warwick, Coventry, UK

² Human Factors Everywhere, Woking, UK

³ Department of Computer Science, University of York, York, UK

INTRODUCTION

Innovations in digital health, such as the introduction of smart infusion pumps, have the potential to improve patient safety, even though the evidence base remains weak.⁽¹⁾ Equally, however, new risks can be introduced, which might contribute to adverse events and patient harm.^(2, 3) The Healthcare Safety Investigation Branch (HSIB), which carries out independent investigations into patient safety concerns in the NHS in England, published in December 2020 its findings from a national investigation into the procurement, usability and adoption of smart infusion pumps.⁽⁴⁾ The report includes safety observations (suggested actions for wider learning and improvement) that suggest the use of (clinical) safety cases in order to demonstrate that patient safety risks have been addressed rigorously and proactively.

Safety cases are a common regulatory instrument used in UK safety-critical industries, as well as other countries such as Norway, Australia and New Zealand.⁽⁵⁾ Previous attempts at importing safety management practices from other industries to healthcare have not always delivered the anticipated benefits, e.g. the application of Failure Mode & Effects Analysis⁽⁶⁾ or the adoption of incident reporting systems.⁽⁷⁾ It is important to understand both the principles underlying such approaches and the context within which they were developed, as well as the unique cultural and institutional context that is specific to healthcare.^(8, 9)

In this paper, we review the thinking that underpins the use of safety cases across different safety-critical industries, and then reflect on their potential use for assuring the safety of digital health innovations. This builds in part on a previous review,⁽¹⁰⁾ but also considers the recent debate about the need for an evidence base for safety case adoption. We focus on digital health innovations, because safety cases are likely to be particularly relevant for software-based systems including, more recently, machine learning technologies, due to their increased complexity, fast pace of technological change and potential interactions with other systems. For example, the HSIB report looks specifically at the role of digital drug libraries used by smart infusion pumps.

Safety cases can work in healthcare, but they will require tailoring to account for the different regulatory landscape and the way patient safety is framed and evidenced.⁽¹⁰⁾ We suggest that safety cases might be put to best use (at least in the short term) as safety improvement tools, rather than as a regulatory (and mandatory) instrument.

WHAT ARE SAFETY CASES?

Safety cases form part of a proactive safety management approach. The purpose of a safety case is to *communicate* why a product, system or service is deemed acceptably safe for use in a particular environment. A safety case comprises two complementary components: (i) a structured and explicit *argument* that (ii) is supported by a body of *evidence*. The argument is usually risk-based, and is intended to demonstrate that all relevant risks have been understood and dealt with sufficiently. The evidence can come from diverse safety management activities, such as hazard and risk analyses, design specifications, testing and empirical evaluation. For

complex settings, such evidence is rarely self-evident, and hence the argument helps explain, appraise and challenge the extent to which the evidence is able to support the safety claims.

There are over one thousand medical device standards, several hundred of which are used for regulatory purposes. Many of the standards cover horizontal issues, e.g. electrical safety, i.e. they cover one specific hazard. These standards contain requirements, which – when followed – are intended to demonstrate compliance with regulations. While standards such as ISO 14971 (*Medical Devices – Application of Risk Management to Medical Devices*) take a risk-based approach, the lack of regulatory expectation for providing an explicit argument for how the body of evidence meets the regulatory requirements can reduce transparency and weaken confidence. As an analogy, the safety case can be thought of as the discussion in a research paper, as it explains and critically appraises the safety-related evidence and reflects on the limitations of the safety evidence and the safety activities that produced the evidence.

WHERE ARE SAFETY CASES BEING USED?

A review by the Health Foundation describes safety case practices across six industries: automotive, civil aviation, defence, nuclear, petrochemical and railways.⁽⁵⁾ Safety cases are used widely across these safety-critical industries, particularly in the UK (see a review⁽¹⁰⁾ for further details). The UK nuclear industry adopted safety cases in 1965, following the Windscale fire accident in 1957. Accidents were major drivers for the adoption of safety cases also in other industries, such as offshore oil and gas production (Piper Alpha oil platform explosion 1988) and railways (e.g. King's Cross escalator fire 1987; Clapham main line derailment 1988). In the automotive domain, the increased complexity of interconnected electronics and software components was reflected in a requirement for an automotive safety case specified in the international standard on automotive functional safety (ISO 26262).

In healthcare, the application of safety cases has been limited. In 2010 (draft version, then finalised in 2014) the Food and Drug Administration (FDA) in the United States issued guidance to manufacturers of infusion pumps that recommends the use of an assurance (safety) case as part of the pre-market notification 510(k) submission route. This was triggered by high numbers of reported incidents involving such devices. However, the impact on adverse event rates has not been evaluated since.

In England, NHS Digital issued two risk management standards for health information technology, which specify safety assurance requirements and practices including the development of clinical safety cases for both manufactures and health organisations (referred to as *DCB 0129* and *DCB 0160*, respectively). Although compliance with these requirements is mandated by NHS England, the standards are only enforced for systems that directly connect to the national infrastructure.

WHY DO SAFETY-CRITICAL INDUSTRIES DO SAFETY CASES?

The use of safety cases is usually part of a regulatory approach that is known as “goal-based” as opposed to the more traditional prescriptive regulatory approach. Prescriptive regulation sets out in standards detailed requirements for which risks need to be controlled, and how. Such prescriptive standards are based on past experiences and work well for established and well-understood systems. However, in settings where there is a fast pace of technological innovation and change, prescriptive standards quickly become outdated and might even hinder innovation. Goal-based approaches are more flexible, because they only specify what needs to be achieved, but leave open how this is done. If there are applicable standards, which are deemed relevant, there is still the expectation that these are complied with. Otherwise, a good argument needs to be provided for why the standards are not followed.

Part of the regulatory requirements in a safety case regime is the duty to demonstrate that risks have been reduced as low as reasonably practicable (ALARP), or similar wording with the same intent. This means that operators of hazardous systems need to consider all reasonable ways of reducing risk, even if these are not prescribed in existing standards.

In the literature, a range of different reasons for why industries have adopted safety cases can be found.⁽¹¹⁾ Among these are expectations that safety cases:

- promote structured risk assessment and management;
- tell the story of a system's safety to a wider and diverse readership;
- show how high-level safety requirements are implemented in the detailed design;
- establish confidence in safety;
- stimulate critical thought around safety;
- explain safety evidence; and
- focus regulatory inspection.

DO SAFETY CASES IMPROVE SAFETY?

Even though safety cases have been used across diverse industries for many years, there is a lack of conclusive evidence that the use of safety cases improves outcomes.⁽¹²⁾ There are two reasons for this. First, safety cases are used traditionally for high-hazard settings, where the focus is on high-severity, low-frequency events, i.e. the rare, but catastrophic failure of a system, such as the loss of an aircraft. Given the low frequency of such events, it is difficult to provide meaningful statistical data about the impact of a regulatory instrument, such as safety cases.⁽¹³⁾ Second, the practice of safety cases is very varied, and it is frequently not explicitly articulated what kinds of benefits safety cases might have (see above) and how these are achieved.⁽¹²⁾ Consequently, the adoption of safety cases is usually based on a face-validity principle, i.e. regulators and industry act on the assumption that it is a good idea to use safety cases.

Critics point to this lack of evidence as well as to the fact that high-profile accidents continue to happen in countries that require safety cases.^(14, 15) A frequently used example is the catastrophic loss of a Royal Air Force Nimrod aircraft in Afghanistan in 2006. The independent Haddon-Cave review⁽¹⁶⁾ highlighted significant weaknesses in safety case practices as part of wider criticisms of poor risk management systems across the different organisations that were involved in the design, operation and assessment of the aircraft. Such a culture was found to undermine the intended value of a safety case leading to a “tick-box” and compliance-driven approach to safety.

MAKING SAFETY CASES WORK IN HEALTHCARE

In the NHS in England, the clinical safety case concept promoted by NHS Digital is suggested for wider use in the HSIB safety observation. However, bearing in mind the complexity and contested nature of safety case practices in safety-critical industries there is a danger that in healthcare the concept will be misunderstood, misused and ultimately fail to make care safer.

An Australian report identifies five key criteria for successful safety case regimes:⁽¹³⁾ (i) an established risk or hazard management framework; (ii) a legal requirement to make the case to the regulator; (iii) a competent and independent regulator; (iv) workforce involvement; and (v) a general duty of care imposed on the operator. It is clear that most health systems do not currently meet these success criteria, not least because much of the patient safety improvement work is driven by outcomes (reactive) rather than by consideration of risk in processes and systems (proactive), while regulators also do not provide incentives for reducing risk as such.⁽¹⁷⁾

Bearing in mind the differences between safety-critical industries and healthcare, the Health Foundation convened in 2013 a multi-professional working group to investigate the potential use

of safety cases in healthcare,(18) and the findings remain highly relevant. The report suggested that the health sector might benefit from the use of safety cases because they provide a structure for proactively assessing risk, they can have a positive impact on safety culture, and because they bring together and synthesise a range of information and evidence relating to a particular service. These benefits might best be realised when safety cases are used as part of service improvement or as part of an assurance process. For example, a review of clinical safety cases submitted to NHS Digital found that many organisations were struggling to define the functionality of health information technology and how it integrates into their local clinical context.(19) This is reflected by the HSIB report, which suggests that the investigated organisations lacked an understanding of how smart pump functionality might differ from current practice, who the users of the smart pumps were, how smart pumps would interface and interact with other IT systems, and what risks might need to be addressed. Irrespective of regulatory requirements, the use of clinical safety cases could support organisations in considering more adequately the scope of change that comes with the adoption of digital health technologies, and making explicit their risk position so that risks do not go undetected or undocumented. This is illustrated in Table 1 based on the reference investigations described in the HSIB report (4).

Table 1: A clinical safety case could help make an organisation's risk position explicit: smart infusion pump example

Stage of safety assurance process	Smart infusion pump example
Scope definition	Consideration of smart infusion pumps as a system that includes drug libraries developed by the hospital team, and that interfaces with other IT systems. Understanding of who the users are (e.g. doctors as well as nurses). Appreciation of local context and working practices.
Hazard identification and risk analysis	Consideration of wider system hazards rather than exclusive focus on technical smart pump failures, e.g. doctors getting confused about drug concentrations due to use of non-standardised drug libraries.
Risk control	Traceability between identified hazards, associated risks and corresponding risk controls put in place, or justification for why risk controls not put in place, e.g. reasons for why non-standardised drug libraries are used and how risks arising from this are going to be controlled.
Post-deployment monitoring	Routine audit can be used to check and challenge assumptions, e.g. routine audit of smart pump event log can provide insights into whether there are more frequent overrides than anticipated.
Modification	Consideration of risks associated with modification or lack of modification of the system, e.g. plans for regular and timely update of smart pump drug libraries over the existing IT network.

However, in order to facilitate and achieve successful adoption of safety case practices in healthcare, suggestions for the use of clinical safety cases (as in the HSIB safety observations) need to be underpinned by additional work and changes to the (patient) safety management infrastructure:

- *Evidence base:* Researchers need to articulate the mechanisms by which safety cases can improve outcomes and build a persuasive evidence base about benefits and the conditions

that create the most fertile ground for using safety cases. In this respect, healthcare might be better placed than other industries, because the rigorous evaluation of complex interventions has gained a lot of traction in recent years, and because (sadly) adverse events happen at a rate that is more amenable to statistical analysis.

- *Capability*: Safety experts and patient safety specialists need to identify the level of training and support that healthcare staff and regulators require in order to support and to implement a safety case approach. In England, bodies such as NHS Digital and Health Education England should consider how capability can be built at scale. NHS Digital offer courses, but these might not scale up, and there are few publicly available examples of clinical safety cases. Health Education England have developed the national patient safety syllabus, which includes consideration of proactive safety management and safety cases, but questions remain about how the syllabus can be implemented and delivered across the NHS. Internationally, the patient safety curriculum developed by the World Health Organisation could potentially be a vehicle, but does not currently include safety cases.
- *Criteria for risk reduction*: Health systems should develop and adopt a healthcare-specific notion of acceptable levels of risk, and a framework that can be used in the decision-making process about the management of risk. In safety-critical industries decisions about risk reduction are based on the ALARP principle, but the health systems face different challenges, such as the duty to provide care to an ageing population with complex health needs while at the same time being bound by a budget set by the government. There is a need for a broader dialogue around the criteria based on which healthcare organisations should manage the trade-off between risk reduction and cost, and to inform their evaluation of whether services are acceptably safe.

CONCLUSION

The national investigation into smart infusion pumps suggests the use of safety cases, which is an accepted practice in UK safety-critical industries. Safety cases can support the safe adoption of digital health innovations, but any such suggestion needs to be underpinned by far-reaching structural changes. These include the rigorous evaluation of safety case practices and their impact on outcomes, the scaling up of education and capability around proactive patient safety management practices, and the establishment of an agreed framework for how to make and justify decisions about patient safety risks.

REFERENCES

1. Agboola SO, Bates DW, Kvedar JC. Digital health and patient safety. JAMA : the journal of the American Medical Association. 2016;315:1697-8.
2. Furniss D, Dean Franklin B, Blandford A. The devil is in the detail: How a closed-loop documentation system for IV infusion administration contributes to and compromises patient safety. Health Informatics Journal. 2019;1460458219839574.
3. Sujan M, Scott P, Cresswell K. Digital health and patient safety: Technology is not a magic wand. Health Informatics Journal. 2019.
4. Healthcare Safety Investigation Branch. Procurement, usability and adoption of 'smart infusion' pumps. Farnham: Healthcare Safety Investigation Branch; 2020.
5. Bloomfield R, Chozos N, Embrey D, Henderson J, Kelly T, Koornneef F, et al. Using safety cases in industry and healthcare. London: Health Foundation; 2012.
6. Dean Franklin B, Shebl NA, Barber N. Failure mode and effects analysis: too little for too much? BMJ quality & safety. 2012;21(7):607-11.
7. Macrae C. The problem with incident reporting. BMJ quality & safety. 2015;25:71-5.
8. Sutcliffe KM, Paine L, Pronovost PJ. Re-examining high reliability: actively organising for safety. BMJ quality & safety. 2017;26:248-51.
9. Dixon-Woods M, Martin G, Tarrant C, Bion J, Goeschel C, Pronovost P, et al. Safer Clinical Systems: Evaluation Findings. London: Health Foundation; 2014.

10. Sujan MA, Habli I, Kelly TP, Pozzi S, Johnson CW. Should healthcare providers do safety cases? Lessons from a cross-industry review of safety case practices. *Safety Science*. 2016;84:181-9.
11. Graydon PJ. The Many Conflicting Visions of 'Safety Case'. 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W); Denver, USA: IEEE; 2017. p. 103-4.
12. Graydon MS. Towards efficacy hypotheses for safety cases. 16th European Dependable Computing Conference (EDCC); Munich, Germany: IEEE; 2020. p. 51-8.
13. Hopkins A. WP 87 - Explaining "Safety Case". Canberra: National Research Centre for OHS Regulation; 2012.
14. Leveson N. The use of safety cases in certification and regulation. *Journal of System Safety*. 2011;47(6).
15. Steinzor R. Lessons from the North Sea: Should "Safety Cases" come to America? . *Boston College Environmental Affairs Law Review*. 2011;38(2):417-44.
16. Haddon-Cave C. The Nimrod Review: An independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 aircraft XV230 in Afghanistan in 2006. London: The Stationary Office; 2009.
17. Sujan MA, Habli I, Kelly TP, Gühnemann A, Pozzi S, Johnson CW. How can health care organisations make and justify decisions about risk reduction? Lessons from a cross-industry review and a health care stakeholder consensus development process. *Reliability Engineering & System Safety*. 2017;161:1-11.
18. Health Foundation. Exploring the potential use of safety cases in health care. London: Health Foundation; 2014.
19. Habli I, White S, Sujan M, Harrison S, Ugarte M. What is the safety case for health IT? A study of assurance practices in England. *Safety Science*. 2018;110:324-35.