Reliability Engineering and System Safety 140 (2015) 200–207

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

The development of safety cases for healthcare services: Practical experiences, opportunities and challenges

Mark Sujan [a,*], Peter Spurgeon [a], Matthew Cooke [a], Andy Weale [b], Philip Debenham [c], Steve Cross [d]

[a] Warwick Medical School, University of Warwick, Coventry CV4 7AL, UK
[b] North Bristol NHS Trust, Bristol BS10 5NB, UK
[c] Birmingham Children's Hospital NHS Foundation Trust, Birmingham B4 6NH, UK
[d] United Lincolnshire Hospitals NHS Trust, Lincoln LN2 5QY, UK

ARTICLE INFO

Article history:
Received 25 May 2014
Received in revised form
18 December 2014
Accepted 29 March 2015
Available online 8 April 2015

Keywords:
Patient safety
Healthcare
Safety case
Risk management
Safety assurance
Safety regulation

ABSTRACT

There has been growing interest in the concept of safety cases for medical devices and health information technology, but questions remain about how safety cases can be developed and used meaningfully in the safety management of healthcare services and processes. The paper presents two examples of the development and use of safety cases at a service level in healthcare. These first practical experiences at the service level suggest that safety cases might be a useful tool to support service improvement and communication of safety in healthcare. The paper argues that safety cases might be helpful in supporting healthcare organisations with the adoption of proactive and rigorous safety management practices. However, it is also important to consider the different level of maturity of safety management and regulatory oversight in healthcare. Adaptations to the purpose and use of safety cases might be required, complemented by the provision of education to both practitioners and regulators.

1. Introduction

Healthcare has made significant advances in the treatment of a wide range of diseases. At the same time, it is now widely acknowledged that modern healthcare systems may inflict preventable harm on patients [1,2]. A systematic review of the literature suggests that around one in ten patients admitted to hospitals around the world will suffer an adverse event, and that as many as half of these may be potentially preventable [3]. In addition to the impact on patients, this has significant financial implications for healthcare systems in terms of additional treatment costs and the cost of litigation [4,5].

As a result of this growing body of evidence of patient harm, healthcare organisations are making continuous progress in their attempts at understanding and managing risks to patient safety. For example, the use of Root Cause Analysis for the review of serious adverse events, and the use of incident reporting to learn from past experience are now mandatory requirements throughout the National Health Service (NHS) in the UK. Increasingly, organisations are also adopting proactive analytical methods such as Failure Mode, Effects and Criticality Analysis (FMECA) in order to anticipate and to assess risk. During the past few years FMECA or variations thereof have been used in healthcare to assess the risks associated with, for example, organ procurement and transplantation [6], intravenous drug infusions [7], and communication in the emergency care pathway [8]. However, it could be argued that in terms of maturity of safety management systems and practices, there may still be many valuable lessons to be learned from safety-critical industries in order to move healthcare organisations towards more proactive and more rigorous safety management practices.

In safety-critical industries, such as railways and nuclear energy, regulators require evidence from developers and operators that they have adopted a thorough and systematic process to understand proactively the risks associated with their systems and to control these risks appropriately. In the UK, these duties are often fulfilled through the use of safety cases [9]. The adoption of safety cases marked a shift from compliance-based to more goal-based regulatory approaches, where the regulator formulates goals, but the demonstration that these goals have been achieved, and how, is left to the manufacturer or operator of systems. Under a predominantly prescriptive regulatory regime, manufacturers and operators claim safety through the satisfaction of specific standards and technical requirements specified by the regulator. This approach has been criticised for prompting bureaucratic practices of safety management, where risks may not be properly understood, and for

* Corresponding author. Tel.: +44 24765 72941.
E-mail addresses: m-a.sujan@warwick.ac.uk (M. Sujan), p.c.spurgeon@warwick.ac.uk (P. Spurgeon), m.w.cooke@warwick.ac.uk (M. Cooke), andy.weale@nbt.nhs.uk (A. Weale), philip.debenham@bch.nhs.uk (P. Debenham), stevecross@googlemail.com (S. Cross).

http://dx.doi.org/10.1016/j.ress.2015.03.033

potentially hindering progress in industries that are driven by technological innovations [10–12]. Broadly speaking, the purpose of a safety case is to provide a structured argument, supported by a body of evidence that provides a compelling, comprehensible and valid case that a system is acceptably safe for a given application in a given context [13]. Safety cases can be structured in different ways, but key elements are a risk-based argument and the corresponding evidence. The risk-based argument is intended to demonstrate that all risks associated with a particular system have been identified, that appropriate risk controls have been put in place, and that there are appropriate processes in place to monitor the effectiveness of the risk controls and the safety performance of the system on an on-going basis. Safety cases are usually confidential, but there are publicly available safety cases (see for example the Safety Case Repository [14]). The history of safety cases and safety case legislation in high-risk industries in the UK has been closely linked to the occurrence of high profile accidents, such as the Windscale fire in 1957, which led to the Nuclear Installations Act 1959 (with revision in 1965) that introduced a licensing scheme along with the requirement for the production of a safety case. While there has been criticism of the safety case approach [15–17], the use of safety cases remains an accepted best practice in UK safety-critical industries as a means to provide rigour and structure to safety management activities.

There has been some recent interest in the potential application of safety cases in healthcare [18]. However, this is limited at present to the domain of medical devices and health informatics [19,20], such as the guidance for infusion pumps issued by the US Food & Drug Administration (FDA) [21] and guidance issued by the NHS Health and Social Care Information Centre (HSCIC) on Clinical Safety Cases for health informatics products [22,23]. A systematic review of the published literature suggests that there is a lack of published empirical evidence of safety case use in healthcare, in particular at the system or service level, where safety cases have not been used at all [9]. It is far from clear whether, and if so, how safety cases might be developed and put to reasonable use at the healthcare system or service level [18,24]. The organisational and institutional context in healthcare is very different from safety-critical industries: the maturity of safety management practices in healthcare is arguably at a lower level than in traditional safety-critical industries; there does not exist a notion of acceptable levels of risk within the healthcare community; and the regulatory regime is predominantly prescriptive, relying largely on inspection and routinely collected mortality data.

The aim of this paper is to describe and to reflect upon first practical experiences with the development and use of safety cases in healthcare at a service level. Two examples from a multi-site service improvement project in the NHS are presented. The paper argues that safety cases might be helpful in supporting healthcare organisations with the adoption of proactive and rigorous safety management practices. However, it is also important to consider the different professional, regulatory and cultural background in healthcare, which might require adaptations to the purpose and use of safety cases in this domain.

The next section provides background to the service improvement project within which the work was undertaken. Section 3 then describes two examples, where safety cases have been applied as a service improvement tool and as a safety communication tool. Section 4 provides a reflection on these practical experiences, and draws out lessons about opportunities and challenges for the use of safety cases in healthcare settings. Section 5 provides a conclusion with implications for research and for practice.

2. Safer Clinical Systems – a context for exploring the use of safety cases

Safer Clinical Systems (SCS) is a service improvement programme supported by the Health Foundation, and facilitated by a technical support team at Warwick University Medical School in conjunction with eight collaborating sites in the NHS. The programme has been running over five years since 2008, including an exploratory Phase I and an application Phase II (three years), which has been completed at the end of 2013. The specific aim of the work was to explore a new approach (in the context of healthcare) to improving patient safety. SCS combines methods from improvement science and safety engineering, and learning from high reliability organisations, to identify hazards and associated risks in a defined health delivery pathway through a structured process of analysis. Interventions aimed at eliminating, mitigating or controlling risks are then implemented and evaluated. The unique features of the SCS approach for the health sector could be described as:

- A systems perspective, incorporating the interdependencies, culture and context of influencing factors.
- A proactive and rigorous search for hazards and associated risks.
- An application of human factors thinking and specific tools as transferred from other industries.

The programme is delivered in four steps, which in this instance were managed by the Warwick technical support team to ensure that NHS sites progressed at a relatively similar rate. The four steps are shown in Table 1.

The SCS approach lends itself to providing some of the information typically used in safety cases to support claims about system safety. After a trial run with one study site during phase 1 of SCS, the subsequent phase 2 of SCS introduced study sites to the safety case concept putting emphasis on some key aspects from the literature [9]:

- Structured thinking and argument about hazards, risk and safety status.
- Integrating various sources of evidence as derived from the safety management activities.

Table 1
Overview of the SCS safety improvement approach.

Step	Description
I – Pathway definition	This step includes the identification and description of the care pathway, its boundaries and relationships with other departments and organisations, and an assessment of the organisational safety culture.
II – System diagnosis	This step includes mapping of the process, the analysis of activities using Hierarchical Task Analysis (HTA), and an assessment of vulnerabilities using FMECA. Specific reliability and performance measures are also established to serve as baseline for subsequent outcome assessments.
III – Option appraisal	This step includes the identification of a set of potential interventions linked to the hazard and risk information. Interventions are then subjected to an appraisal process to select candidate interventions for implementation.
IV – System improvement cycles	In this step the interventions that were selected as part of Step III are implemented with associated measures of reliability, risk reduction and outcomes.

- Making the implicit explicit about the safety of particular pathways.
- Offering a means of communication to various stakeholders.

The SCS programme represented a collective programme in which the eight NHS sites worked to a common time schedule and were introduced to the use of a number of analytic tools – from the human factors literature and practice. The tools were used in the analytic phase of SCS, and were aimed at enabling sites (many for the first time) to analyse their care delivery pathway, identify hazards within the pathway, assess the associated risks, and consider some form of intervention to eliminate, control or mitigate these risks. Within the collaborative framework of the SCS programme, it became a requirement on sites to present the information they had gathered to different stakeholders, such as other members of the clinical team, the chief executive and other senior managers of their organisation, other teams within the SCS programme, and more broadly to healthcare practitioners outside of their organisation.

Initially, collaborating sites were given some freedom, in terms of presentation, to link their SCS work to the safety case report. Over three iterations, a safety case report template was defined, with the following sections:

1. Title: the pathway or focus of the case.
2. Background: the information available to form our safety claim or evaluation of the current system.
3. Assurance statement: the pathway has been assessed and is described as safe to an agreed extent.
4. Risks and Control Mechanisms: a history of hazards, associated risks and control measures
5. The Residual Risks: description of significant uncontrolled risks.
6. Action: being taken, or needing to be taken.
7. Sources/confidence: an evaluation of the robustness of and confidence in the basis of the argument presented.

It was this framework that formed the basis of the production of the two examples presented below.

3. Examples

3.1. Example 1: safe shared care of patients undergoing surgery on a renal unit

The first example is drawn from work that was undertaken on the Renal Unit at a hospital within North Bristol NHS Trust. The Renal Unit serves a population of 1.57 million patients. It admits around 450 patients for surgery annually. The project aimed to improve the safety of shared care arrangements between the renal medicine team and the surgical team for patients with Established Renal Failure (ERF). The main drivers for the project were previous adverse events. Root Cause Analysis carried out on these adverse events prior to the project exposed inadequacies in the communication between the two clinical teams. The work was undertaken from February 2012 to December 2013.

3.1.1. Service description

Patients with ERF have diverse clinical needs, relating to dialysis, electrolyte and fluid management and appropriate prescribing. As such, these patients remain the primary responsibility of the renal team with pre- and post-operative surgical care delivered on the renal ward, and surgeons visiting rather than having primary responsibility for the patient. Multiple informal unstructured handovers occur at various stages of patient care. The corresponding high-level patient pathway is illustrated in Fig. 1.

3.1.2. Purpose of the safety case

In this organisation the predominant approach to safety management was reactive with a focus on harm rather than risk. Improvement activities were usually triggered by the review of adverse events leading to specific improvement interventions. These were not normally based on an in-depth risk analysis. There are no regulatory requirements mandating the use of safety cases, and the project was an improvement project. As such, the role of safety cases in the project was as an improvement tool. The safety case concept was adopted in order to encourage a proactive mindset, and to provide structure to proactive safety management activities. The safety case reports and the associated evidence were also intended to be a document repository of knowledge, which should become embedded in the department. As there are no criteria for determining acceptability of risk, the aim of the project, and the high-level safety claim, was articulated modestly as: "*G0: Shared care arrangements between renal medicine teams and surgical teams for patients with ERF are safer than currently,* i.e. *the risk associated with the pathway is reduced*". This high-level claim is broken down into two sub-claims. One consists of a risk argument intended to demonstrate that the main risks have been reduced. The other consists of a confidence argument intended to demonstrate that there is sufficient rigour in the safety management activities that were carried out to produce the evidence. Fig. 2 provides a graphical representation of this high-level argument in Goal Structuring Notation (GSN).

3.1.3. Risk argument

The sub-goal for the risk argument is: "*G1: The key risks associated with the pathway have been reduced*". The argument proceeds over a partial safety management life cycle by demonstrating that all risks have been understood (G1.1), the key risks have been controlled to

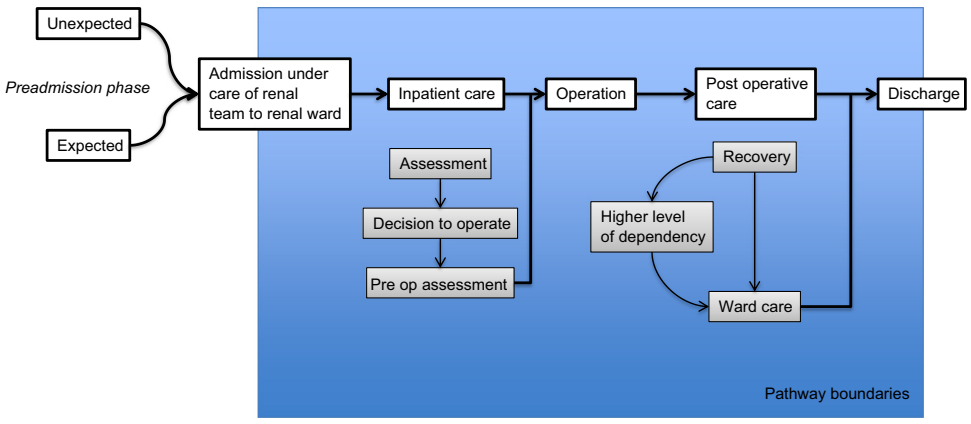

Fig. 1. High-level pathway for shared care arrangements.

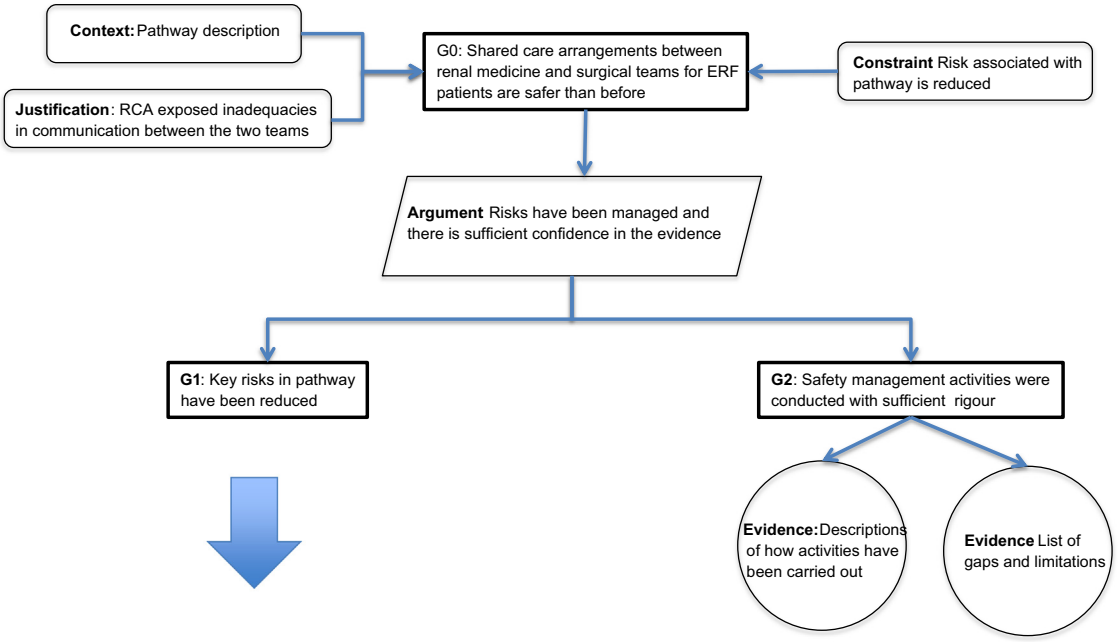

Fig. 2. High-level safety argument.

reasonable levels (G1.2), and that operational performance of the system has improved (G1.3). The argument remains partial, as there does not exist a mature safety management system [25] within the organisation that would ensure adequate safety policy, safety assurance and safety promotion. A simple graphical overview is given in Fig. 3

Sub-goal G1.1 is broken down into two claims to demonstrate that relevant hazards have been identified (G1.1.1) and that the risks associated with the hazards have been evaluated (G.1.1.2). The evidence to support both these sub-claims comes from previous experience (Root Cause Analysis of adverse events), Failure Mode, Effects and Criticality Analysis (FMECA) carried out on the process map, and Predictive Human Error Analysis (PHEA) carried out on a hierarchical task analysis (HTA) representation. The FMECA and PHEA provided risk scores based on expert consensus. In total, 99 hazards and associated risks were identified. The hazards were ordered according to the risk scores, and the six highest-ranking hazards were selected for further investigation. Table 2 provides a summary of these.

Sub-goal G1.2 (risk controls) has not been broken down further, but is supported directly by evidence from an options appraisal. The options appraisal consisted of a focus group session with a representative set of stakeholders in the pathway to identify and to select appropriate risk control interventions based on a number of dimensions, such as the amount of risk reduction likely to be achieved, cost of the intervention, time required to develop the intervention, ease with which measures for the intervention could be specified etc.

Finally, sub-goal G1.3 (operational performance) was broken down into two further sub-claims, namely that suitable measures have been identified for all key hazards (G.1.3.1), and that the measures over time demonstrate an improvement (G1.3.2). The evidence for G.1.3.1 comes from the options appraisal, where the measures were developed. The evidence for G1.3.2 comes from process audits carried out at regular intervals.

Fig. 4 presents an example of an audit carried out for Hazard 01. The hazard relates to situations where there is no documented review by a senior doctor pre-operatively. The absence of such a review can have many consequences, such as delays in treatment, unrecognised levels of risk, and incomplete knowledge of patient condition, which can result in unnecessary tests and interventions, patient harm and poor patient experience. Potential causes for the occurrence of the hazard include high levels of workload, low level of priority attached to this activity, and a lack of knowledge about the importance of the activity. It is also possible that the pre-operative assessment takes place, but remains undocumented. Interventions to control the risk associated with the hazard were: introduction of a formal policy and preoperative assessment clinic; and a change to the admissions forms to include a dedicated preoperative assessment box. The audit carried out over 46-week period shows that the reliability of this activity varies greatly from week to week, and that the interventions (introduced in week 31) have been unsuccessful thus far in in bringing about sustainable and high levels of reliability.

3.1.4. *Confidence argument*

The sub-goal for the confidence argument is: "*G2: The evidence produced by the safety management activities is sufficiently trustworthy*". The sufficiency criterion was determined qualitatively through consensus of the project team and through external expert review. This claim has not been broken down further. Instead it provides as evidence simple descriptions of how the different safety activities referred to in sub-goal G1 have been carried out. Limitations and gaps in the confidence of individual safety management activities are also noted.

As an implicit argument, there were two main strategies for ensuring trustworthiness of the evidence produced, which ran through most of the safety management activities. On the one hand, hazard identification, risk analysis and the design of risk controls all attempted to involve a broad and representative set of stakeholders in the pathway. On the other hand, all findings were continuously presented to, and reviewed and critiqued by external experts (external to the site) from the SCS Warwick technical support team over the duration of the project lifetime. Limitations included difficulties in engaging staff, and in bringing together representative sets of staff for the risk analysis and options appraisal focus groups.

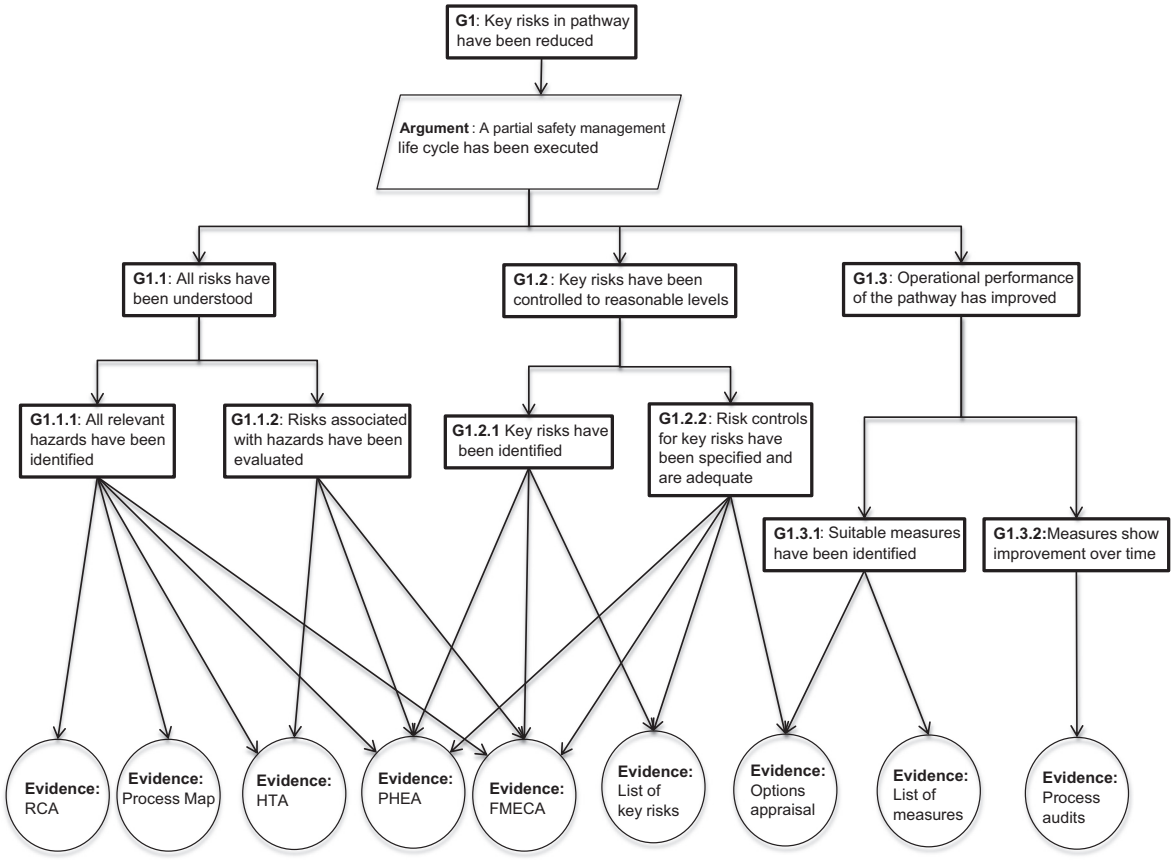

Fig. 3. Risk argument.

Table 2
Main hazards identified (example 1).

Hazard ID	Description
01	No documented medical review by senior doctor pre-operatively
02	No documented surgical plan pre-operatively
03	Appropriate haemodialysis not organised
04	Correct surgeon, operation and safety behaviours in theatre not documented
05	Requirements for safe discharge not met
06	Documented surgical review not provided post-operatively

3.1.5. *Lessons about risk*

The residual risk in the pathway following the improvement activities remains high. While certain performance measures have improved, others have not despite the safety management activities that have been carried out. Reasons for this may be the prototypical nature of many of the interventions that require further development iterations, as well as difficulties experienced in engaging all staff. This also implies that some of the qualitative evidence based on expert judgement and consensus is not as trustworthy as would be required to produce a strong and convincing argument. The safety case also acknowledges that there are many risks that have been identified, but have not been included in the risk control stage.

3.2. *Example 2: safe transfer of responsibility and accountability of medical services (daytime) to the Hospital at Night service*

The second example is drawn from work that was conducted at Birmingham Children's hospital, which serves both the local community, for standard care, as well as providing highly specialised tertiary level services. The focus was the clinical handover from all medical daytime services to a Hospital at Night (HaN) team. The main drivers for the project were the Trust Board focus on handover as a specific area for safety improvement, and the learning derived from previous incident reports that suggested that 33% of incidents reported occurred between 6 p.m. and 6 a.m., a period where the HaN pathway is active within the hospital. The work was undertaken between February 2012 and December 2013.

3.2.1. *Service description*

The handover from the day service covers all 13 medical specialties and, therefore, a large and wide spectrum of information, to a much smaller cross covering team of doctors, advanced nurse practitioners and senior nurse coordinators. The aim of the handover process is to facilitate prioritisation of workload and proactive monitoring of patients deemed to be at risk of deterioration. The size of the transfer content and inaccuracies results in a great deal of individual error recovery along the pathway of care. As a consequence, there can be many wasteful requests for clinical reviews of patients by the HaN team to re-establish the level of risk. The high-level handover pathway is described in Fig. 5.

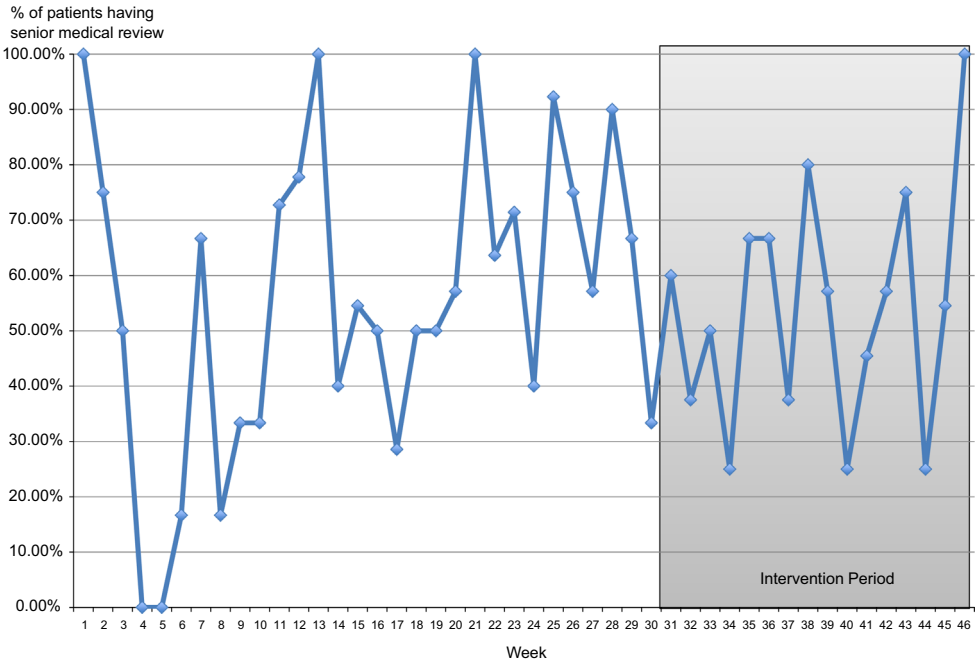

Fig. 4. Audit data of senior medical review.

Fig. 5. High-level pathway for handover to Hospital at Night team.

3.2.2. Purpose of the safety case

Each day service provides a written handover sheet, but in varying formats, multiple locations and no audit trail of amendments. No formal controls existed to check the material, and the HaN team added amendments by hand. The organisation was seeking an approach to enhance the safety of the handover process (hence participation in the SCS programme) and was seeking to establish with the Trust Board and Executive Management Team the challenges faced within the handover process. Although there was no prior experience of the safety case methodology, it was used here as a communication tool about safety to the Board – a statement of the current risk position. The safety case report does not attempt to demonstrate acceptable safety of the handover pathway; rather it claims that key vulnerabilities have been understood, and supports this through a comprehensive exposition of risks. The safety case report follows, however, a similar structure as the safety case report developed in example 1 (see Fig. 2). This entails an argument about the level of risk, and an argument about the level of confidence in the safety management activities that provide the evidence for the risk argument.

3.2.3. Risk argument

The risk argument in the current version of the safety case report provides a modest claim and corresponding evidence that the risks associated with the HaN handover pathway have been understood. This is comparable to sub-goal G1.1 in example 1. As in that case, the evidence to support this claim comes from previous experience as well as from the techniques used within the SCS programme (Process Mapping, HTA, FMECA, PHEA). The FMECA and PHEA provided risk scores based on expert consensus. In total, 123 hazards and associated risks were identified. The hazards were ordered according to the risk scores. Review of the hazards and associated risk scores led to the identification of three high-risk areas within the pathway, see Table 3. Existing risk mitigation measures were also identified.

3.2.4. Confidence argument

The confidence argument followed the same strategy as that described in example 1. Safety management activities were undertaken collaboratively with staff working in the handover pathway. Focus groups with staff were held, and input from individuals with knowledge of the handover system at the organisation was sought. Identical patterns of high-risk areas were identified independently across the different specialties.

The project sought the engagement of executive, senior management, and both senior and junior clinical staff. A stakeholder analysis informed the development of the Communication Strategy, which has guided the project in engaging key groups of staff. However, involvement of junior doctors was low.

3.2.5. Lessons about risk

There is a high-level assumption that the current handover pathway is safe; in that harm is not routinely caused to a large number of patients as a direct result of handover. However, there remains a high level of residual risk, and the system may not be as safe as outcome data implies. The risk analysis has demonstrated that within the whole HaN system there is a high rate of error production that is mitigated by high error recovery rate at multiple steps along the pathway. This overreliance on the individual leaves the handover process fallible to performance influencing factors.

Table 3
High-risk areas and existing risk mitigation identified in the HaN handover pathway

Risk area	Description	Existing risk mitigation
Synthesis of information pre-handover	Failure to synthesise information into packaged clinical intelligence during handover process, e.g. identification of "at risk" patients.	Preparatory meetings before handover, but practices vary greatly.
Communication of information	Failure of information communication during handover preparation and handover process, e.g. not delivering up to date information.	Structured communication protocols, but this is not widely adopted and of variable quality when used.
Fidelity of written information	Failure of written information entry onto the handover documents, e.g. information inaccurate, incomplete or outdated.	No formal risk controls; HaN team members adopt workarounds using hand written annotations.

4. Discussion

The two examples described how the use of safety cases might support healthcare organisations in adopting a proactive, risk-driven approach to safety management at the service level. In both examples, the pathway had previously been tacitly assumed to be safe due to the low number of patients that were actually harmed. However, a large number of incident reports suggested that both pathways might not be as safe as the outcome data alone would lead to believe. The safety case approach supported the organisations in defining the pathways, understanding and ranking risks in a proactive way, whilst engaging the workforce in safety management activities. However, with such complex "legacy" systems that had evolved over many years, the number of risks uncovered was overwhelming, and the organisations found it challenging to make a safety claim about their pathways, because the work focused so much on the now explicitly uncovered and largely accepted risks. In particular, the organisations found it difficult to claim that the system had reached an acceptable level of safety because of the new understanding of the risk within the pathway. As a result, the safety claims and the safety cases were used as improvement tools and as communication tools providing an exposition of risk, rather than as approval or certification instruments.

Commonly, the high-level safety claim made in safety case reports of industrial systems states that the system under consideration is acceptably safe to operate for a given purpose and in a given environment [26]. In the UK the acceptability criterion is often determined through the application of the *As Low as Reasonably Practicable* (ALARP) principle [27]. However, in healthcare no comparable approach exists at the service level (there are standards for medical devices [28] and for health informatics products [22,23]). In fact, regulatory guidance often does not operate with the notion of risk as such, but rather specifies events that should not happen (so called "never events"), such as wrong site surgery or wrong route administration of chemotherapy [29]. This regulatory environment potentially presents a challenge to the successful and widespread use of safety cases. When healthcare organisations are "exposing" their levels of risk in a safety case, they need reassurance that they would not be regarded as less safe by the regulator, or even be penalised for this [18]. Ideally, the systematic exposition of the presence of risk should be regarded as an indication of the growing maturity and safety awareness of healthcare organisations.

Safety Management Systems (SMS) [25] are not well established in healthcare organisations, and the level of maturity of their safety management activities is arguably significantly lower than, for example, in aviation. Safety management in healthcare is still predominantly reactive. On the one hand, this provides an opportunity for the use of safety cases to contribute a framework and a structure to the safety management activities. On the other hand, the need for the provision of rigorous and comprehensive evidence may appear daunting and potentially overly resource intensive in a healthcare environment that is facing serious challenges in terms of staff shortages, high levels of workload, and often long waiting times for patients. This is exacerbated by the fact that there is little published evidence of the direct contribution of safety cases to improved safety performance [16,17]. There is a danger that safety case development might be regarded as a costly intervention without established evidence base—and hence not be promoted at senior management level.

Recent criticism of safety case use in industry following the Nimrod Review [15] into the loss of a Royal Air Force aircraft in Afghanistan in 2006 highlighted that safety cases should not be regarded as a panacea for successful safety management. The review criticised that the development of safety cases had degenerated into a "tick-box exercise" and that the culture had been one of paper safety at the expense of real safety. In a healthcare environment, it is easy to see similar challenges arising. NHS organisations are collecting a wealth of quality and safety data, such as data about serious untoward incidents ("never events") or data about common forms of harm (through the NHS Safety Thermometer). However, there is often a disconnect between safety managers and this type of data on the one hand, and front line staff working in clinical reality on the other hand. The Nimrod Review further criticised the lack of involvement of frontline operators. In healthcare, similar concerns have been raised in the literature, suggesting that there is a continuing lack of clinical engagement in safety improvement activities [30]. The engagement of junior doctors, who deliver a lot of the actual frontline medical care, might be particularly problematic, as they usually spend only a short amount of time in a particular environment before they rotate again. However, this group potentially has particular incentives to participate in safety improvement activities that could be exploited by contributing to their education and professional portfolios.

A review of safety case practices in industry found that the lack of experience and expertise with safety cases and the underlying safety management activities formed a significant threat to the successful adoption of the approach [9]. This will be particularly relevant in a healthcare context where many of the safety management activities that are well established in industry are largely unknown to practitioners. Trust Boards, patient safety and quality managers, and frontline staff all require some degree of education in order to engage with, contribute to and understand proactive approaches to safety management. This applies also to regulators, who are responsible for providing stimulus and guidance to healthcare organisations. Technical knowledge of the corresponding methods and tools is one aspect, but arguably a proactive orientation in itself is the key requirement. In the examples, the use of safety cases was intended to contribute towards achieving this mindset.

5. Conclusions

In this paper we reflected on practical experiences with the development and use of safety cases at the service level in healthcare settings. Healthcare organisations require support with the adoption of proactive and systematic safety management practices that are a key feature of safety-critical industries. Safety cases might have the potential to provide structure and a proactive

mindset to healthcare organisations' safety management activities. However, it is unlikely that safety cases can be adopted in the same way as they are used in safety-critical industries due to the different professional, regulatory and cultural background in healthcare. Systems in healthcare have evolved over many years and many risks are accepted. There is no formal notion of what constitutes an acceptable level of risk. The maturity of safety management activities is comparably low, and the safety management approach is predominantly reactive. In the two examples described in the paper, we have outlined possible ways in which safety cases might be used in a healthcare context: as a service improvement tool and as a safety communication tool, rather than as an instrument to demonstrate that services are "acceptably safe".

There is, in addition, a need for empirical evidence of the impact of the adoption of safety cases on performance and outcomes, as well as on culture. This should be combined with education about proactive safety management to Trust Boards and regulators in particular, as these groups are best placed to initiate change.

Acknowledgements

We are grateful to the participants of the Safer Clinical Systems programme and their organisations. We acknowledge the input of the experts who were part of the Warwick based Technical Support Team. Prof Tim Kelly provided helpful suggestions for the development of the safety cases.

We are also grateful for the support of the Health Foundation (Registered Charity number: 286967), who commissioned the project.

References

[1] Department of Health. An organisation with a memory. London: The Stationery Office; 2000.
[2] Kohn LT, Corrigan JM, Donaldson MS. To err is human: building a safer health system. Washington: The National Academies Press; 2000.
[3] de Vries EN, Ramrattan MA, Smorenburg SM, Gouma DJ, Boermeester MA. The incidence and nature of in-hospital adverse events: a systematic review. Qual Saf Health Care 2008;17:216–23.
[4] Ovretveit J. Does improving quality save money?. London: Health Foundation; 2009.
[5] Vincent C, Neale G, Woloshynowych M. Adverse events in British hospitals: preliminary retrospective record review. BMJ 2001;322:517–9.
[6] Steinberger DM, Douglas SV, Kirschbaum MS. Use of failure mode and effects analysis for proactive identification of communication and handoff failures from organ procurement to transplantation. Prog Transpl 2009;19:208–14.
[7] Apkon M, Leonard J, Probst L, DeLizio L, Vitale R. Design of a safer approach to intravenous drug infusions: failure mode effects analysis. Qual Saf Health Care 2004;13:265–71.
[8] Sujan M, Spurgeon P, Inada-kim M, Rudd M, Fitton L, Horniblow S, et al. Clinical handover within the emergency care pathway and the potential risks of clinical handover failure (ECHO): primary research. Health Serv Deliv Res 2014:2.
[9] 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.
[10] Bishop P, Bloomfield R, Guerra S. The future of goal-based assurance cases. Proc Workshop Assur Cases 2004:390–5.
[11] Habli I, Kelly T. Process and product certification arguments: getting the balance right. ACM SIGBED Rev 2006;3:1–8.
[12] Hawkins R, Habli I, Kelly T, McDermid J. Assurance cases and prescriptive software safety certification: a comparative study. Saf Sci 2013;59:55–71.
[13] Maguire R. Safety cases and safety reports. Aldershot: Ashgate; 2006.
[14] Dependability Research Group University of Virginia. Safety Cases: Repository.
[15] 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 Stationery Office; 2009.
[16] Leveson N. The use of safety cases in certification and regulation. J Syst Saf 2011:47.
[17] Steinzor R. Lessons from the North Sea: should "Safety Cases" come to America? Boston Coll Environ Aff Law Rev 2011;38:417–44.
[18] Health Foundation. Exploring the potential use of safety cases in health care. London: Health Foundation; 2014.
[19] Sujan MA, Koornneef F, Chozos N, Pozzi S, Kelly T. Safety cases for medical devices and health IT: involving healthcare organisations in the assurance of safety. Health Inform J 2013;19:165–82.
[20] Sujan MA, Koornneef F, Voges U. Goal-based safety cases for medical devices: opportunities and challenges. Comput Saf Reliab Secur 2007:14–27.
[21] FDA. Total Product Life Cycle: Infusion Pump Premarket Notification 510 (k) Submissions. 2010.
[22] Health & Social Care Information Centre. Clinical Risk Management: Its Application in the Manufacture of Health IT Systems – Implementation Guidance. 2013.
[23] Health & Social Care Information Centre. Clinical Risk Management: Its Application in the Deployment and Use of Health IT Systems – Implementation Guidance. 2013.
[24] Sujan M, Harrison M, Steven A, Pearson P, Vernon S. Demonstration of safety in healthcare organisations. Comput Saf Reliab Secur 2006:219–32.
[25] International Civil Aviation Authority. Safety management manual. Montreal, Canada: ICAO; 2013.
[26] UK Ministry of Defence. Defence standard 00-56: safety management requirements for defence systems. London: The Stationery Office; 2007.
[27] Health and Safety Executive. Reducing Risk: Protecting People. 2001.
[28] International Organization for Standardization. ISO 14971:2012 medical devices: application of risk management to medical devices. Geneva: ISO; 2012.
[29] Department of Health. The never events policy framework. London: The Stationery Office; 2012.
[30] Health Foundation. Are clinicians engaged in quality improvement? A review of the literature on health professionals' views on quality improvement. London: Health Foundation; 2010.