CLINICAL

Hot cheese: a processed Swiss cheese model

[1]Y Li, [2]H Thimbleby
[1]EPSRC Postdoctoral Researcher; [2]Professor of Computer Science, Computer Science, Swansea University, Swansea, Wales, UK

ABSTRACT James Reason's classic Swiss cheese model is a vivid and memorable way to visualise how patient harm happens only when all system defences fail. Although Reason's model has been criticised for its simplicity and static portrait of complex systems, its use has been growing, largely because of the direct clarity of its simple and memorable metaphor. A more general, more flexible and equally memorable model of accident causation in complex systems is needed. We present the hot cheese model, which is more realistic, particularly in portraying defence layers as dynamic and active – more defences may cause more hazards. The hot cheese model, being more flexible, encourages deeper discussion of incidents than the simpler Swiss cheese model permits.

KEYWORDS education, human error, incident causation model

DECLARATION OF INTERESTS No conflict of interest declared.

Correspondence to Y Li
Computer Science
Swansea University
Swansea SA2 8PP, UK

tel. **+44 (0)1792 602693 ext 4026**
e-mail **yunqiu.li@swansea.ac.uk**

INTRODUCTION

For many years, the graphic representation of James Reason's substantial insight into accident causation in complex systems, the Swiss cheese model (SCM),[1] has proven extremely effective and powerful. It visualises incidents as the result of the accumulation of multiple failures in defences (represented as the holes in slices of cheese) that unfortunately align, creating a 'hazard trajectory' that results in harm (Figure 1). The model is clear, insightful, memorable and justifiably a classic; it can be used not just in teaching, but also in helpfully focusing attention in accident investigations as well as in learning and planning defences against future harm. Despite some criticism,[2,3] the simple model has been widely taken up in risk analysis and risk management, especially in safety critical fields where human operators play an important role in incidents, for example, in aviation, nuclear, petrochemicals industries and, indeed, healthcare.

The recent authoritative *Berwick report*[4] has highlighted patient safety, but it is really a powerfully worded strategic report rather than a vehicle trying to provide mechanisms to think more clearly about safety. Although it has been instrumental from an organisational and systemic perspective on safety, the SCM is subject, as is any model, to some limitations. It is purely an abstract schematic. Even Reason himself has warned that 'the pendulum may have swung too far in our present attempts to track down possible errors and accident contributions that are widely separated in both time and place from the events themselves.'[1] Reason has also questioned the use of the model, quipping: 'Is Swiss cheese past its sell-by date?'[5] We recommend Carthey's excellent, brief, critical review of SCM and alternative healthcare safety models.[6]

In this paper, we argue that, as a simple visual model, the SCM has a crucial oversight: feature interaction. A familiar form of feature interaction is that Drug A is good and Drug B is good, but these two drugs interact causing some unwanted drug reaction. In the Detroit nuclear reactor meltdown (mentioned in more detail below) a defence introduced to make the reactor safer, namely a monitoring camera, fell off and contributed to the severity of the incident. Generally, new feature interactions are unforeseeable.

Feature interaction can be explicitly introduced in what we call the hot cheese model (HCM). The HCM is more active than the original SCM in terms of highlighting interaction between system defence layers and is just as memorable. Hot cheese allows the severity of the consequences of active failure and latent conditions to be clearly defined and explored. It also enables more discussion on how to best design, implement and execute some defence mechanisms that are more dynamic and active. A further advantage is that the HCM, because of its increased flexibility, is much better at stimulating and sustaining discussion, and hence encouraging deeper exploration of the relevant issues.

Interestingly, the original SCM does not discriminate between different sorts of defence failures: the holes in the cheese can equally represent errors (whether intentional or slips), deliberate violations or misconduct and omissions, such as neglect. In particular, the SCM does not address whether problems in patient safety are accidental, misguided, willful or reckless; it is concerned with how multiple defences – the slices of cheese – can be organised so that the system as a whole may be resilient against errors from whatever causes. Only by having a thorough discussion of an incident can the appropriate learning be achieved and, as we will argue, the SCM somewhat restricts the discussion compared with the HCM.

THE ORIGINAL SWISS CHEESE MODEL

The SCM has been an extraordinarily good education tool. In formal use, it has been considered useful because of its clarifying role in accident investigation and because it can help identify potential hazards (e.g. missing defences) before harm happens. Notably, instead of focusing blame or investigative efforts on the person who apparently committed the active failure that caused an incident – for example, a nurse who administrated an overdose to a patient – the SCM encourages investigators to view such actions just as one unsafe act in a chain of many, rather than the only cause.

The SCM helps us understand active failures and latent conditions (Figure 1). Active failures are generally committed by persons who perform their duty at the so-called sharp end[7] of the system (visualised to the left in Figure 1). Latent conditions are the potential contributing factors that lie dormant in the system and occur upstream at the more remote layers, called the blunt end.[7] These latent conditions can be organisational, contextual and diffuse in nature, or they may be design-related, to do with the system that people work in. Their consequences only become evident when they combine with other factors to breach system defences.

Here is an example. Administering the wrong drug is an active failure, but the latent condition might have been the confusingly similar names of two different drugs. In the SCM, both types of failure are represented the same way, as holes in defences. In particular, the SCM gains power by not differentiating between the 'sharp end' failures and the 'blunt end' failures. A hole is a hole, and it does not matter where it is; if an incident occurs, every defence failed, and seeking the root cause is misleading. Thus the SCM helps stop us focusing just on the sharp end active failures, such as the 'nurse pressing the wrong button'. This insight is nicely captured in the slogan 'system flaws, not character flaws'.[8] The SCM makes it clear that blame cannot be placed on problems involving just one slice of cheese; indeed, it would be a system design flaw to rely on only one defence for any critical process.

Anything from poorly designed policy, untimely training, under-staffing, the poor condition of the physical environment, communication deficiency, interruptions or bad medical device user interfaces can contribute to an accident – and sometimes magnify the consequences of an active failure.

A NEW HOT CHEESE MODEL

The SCM has successfully worked in the accident investigation culture for many years. Unfortunately, its graphical representation may be misleading: system defences are visualised somehow as a collection of passive, unchanging, independent features (namely, slices of cheese). The model
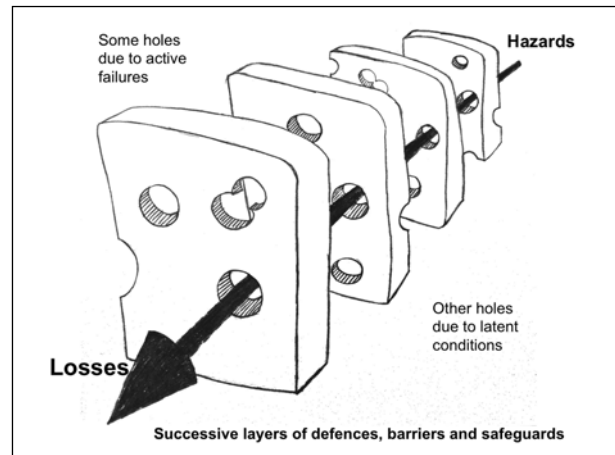

FIGURE 1 The original the Swiss cheese model. Note that our slices of cheese have been carefully drawn so they are not the same as each other – a common mistaken simplification in many representations of the model.

seems to imply that as long as enough layers of defences are in place, at least ones with non-aligning holes, the probability of an accident can be minimised.

The truth is, system defences are more active. There have been many cases when a defence layer was introduced in the hope of preventing errors and improving system performance, but only ended up provoking new errors and causing more harm. For example, in hospitals, the medication administration stage accounts for 26–32% of adult patient medication errors and 4–60% of paediatric patient medication errors.[9]

Evidence shows that Barcode Medication Administration (BCMA) systems can reduce pharmacy dispensing errors,[10] hence they have been strongly recommended for all hospitals in the USA. However, a study published in 2008 identified 15 types of BCMA-related workarounds and 31 separate probable causes of the identified workarounds, along with potential errors as a result of workaround occurrence.[11] In other words, the introduction of a defence – the barcoding – itself caused new problems.

Perrow's classic book *Normal accidents* has many examples of multiple defence failures and interactions.[12] We mention just one. In the Fermi 1 nuclear reactor near Detroit in the US, a specially added defence, a safety device, fell off inside the reactor containment building and blocked a coolant drainage hole, resulting in temperatures that destroyed sensors. As a result, the reactor operators were working blind and guessing while the reactor suffered a meltdown.[13]

Similarly, multiple defence failures and interaction issues can be found in the use of medical devices. For instance, some manufacturers provide infusion pumps with a fast increasing/decreasing function, where the volume and the rate to be infused are entered via an arrow keypad. Such devices function similarly to an electric oven or
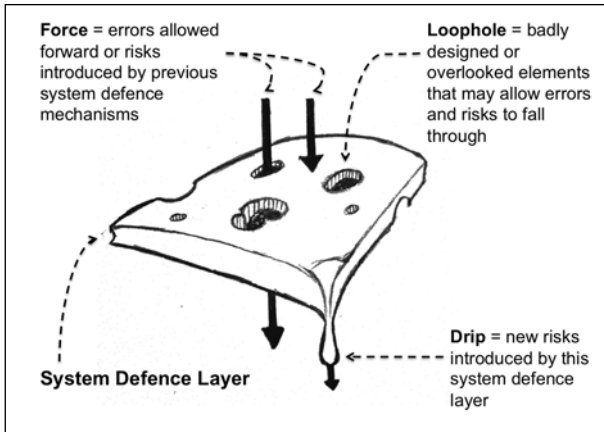
Force = errors allowed forward or risks introduced by previous system defence mechanisms

Loophole = badly designed or overlooked elements that may allow errors and risks to fall through

System Defence Layer

Drip = new risks introduced by this system defence layer

FIGURE 2 The hot cheese model.

central heating timer. The display changes faster, changing in larger steps, when the arrow button is pressed and held. Arguably, such a feature can help to improve the efficiency of programming infusion pumps. However, this is not a standard feature that has been implemented in every infusion pump model on the market. Nurses who are used to arrow keypads without fast increase/ decrease functions might have developed habits, such as relying on the tactile feedback from their fingers to count the number rather than looking at the display. This will cause potential overshoot or undershoot on a pump with the fast increase/decrease feature if they accidentally perform (what the pump considers to be) a press-and-hold rather than a press alone.

Three key contributors to incidents in HCM

In Figure 2, we present our new hot cheese model, the HCM. We argue that it is a more realistic representation of accident causation. Anything the original SCM can do, the HCM can do too, but with the added advantages of visualising some of the complexities we discussed above. Moreover, the HCM provokes 'heated' discussion, and thus a more useful and thorough exploration of incidents. We define some useful terms that are directly suggested by the HCM with simple examples:

Force: the combination of active errors allowed forward and risks introduced by system defence layers upstream. For example, a patient has forgotten to take his medication on Monday because he attended a wedding. Skipping medication may lead to health deterioration for the patient.

Loophole: a latent condition hidden in a layer that may allow force to fall through. For example, the pill case used by the patient is a simple compartmentalised storage box with a clear label of each type of medication and a reminder of daily medication. However, there isn't a history log built in. Although this pill case can help the patient to organise his medication, it cannot detect the above-mentioned error, nor can it assist the patient to recover from the error.

TABLE 1 Combinations of many possible consequences of active failure, latent conditions and defence layer design flaws

Severity level	Graphic representation	Explanation
0		Defence layer does not introduce new risk. Defence layer is designed with all possible errors and risks from upstream that it can encounter in mind. No error has slipped through previous defence layer
1		Defence layer does not introduce new risk. Defence layer is designed with all possible errors and risks from upstream that it can encounter in mind. Error has slipped through previous defence layer
2		Defence layer does not introduce new risk. Defence layer design has failed to cover some possible errors and risks from upstream. No error has slipped through previous defence layer
3		Defence layer does not introduce new risk. Defence layer design has failed to cover some possible errors and risks from upstream. Error has slipped through previous defence layer
4		Defence layer introduces new risk. Defence layer is designed with all possible errors and risks from upstream that it can encounter in mind. No error has slipped through previous defence layer
5		Defence layer introduces new risk. Defence layer is designed with all possible errors and risks from upstream that it can encounter in mind. Error has slipped through previous defence layer
6		Defence layer introduces new risk. Defence layer design has failed to cover some possible errors and risks from upstream. No error has slipped through previous defence layer
7		Defence layer introduces new risk. Defence layer design has failed to cover some possible errors and risks from upstream. Error has slipped through previous defence layer

Drip: new risks introduced by the current system defence layer due to design flaws. These risks may possibly turn into force, which layers downstream will have to deal with. For example, the manufacturer has introduced a sensor-enabled history log system to help the patient keep track of his medication. However, the list of skipped medication on Monday and the list of scheduled medication for Tuesday are displayed on the same page in similar font. There is a risk of the patient taking two days' medication at once 'as instructed'. This will cause an overdose, which in some case may be a much bigger hazard.

Crucially, all these definitions are readily visualised in the HCM. As a result, the HCM helps to illustrate eight different common cases of how an incident may be caused, as shown in Table 1.

When visualising the cause of an accident, force, loophole and drip are treated as three major contributors to accidents. It is possible to rank them based on the gravity of their roles once identified. For example, though force could be the most visible cause of an accident, if it can be previously identified and blocked by the current defence layer, the risk that it makes it to the sharp end will be minimal. Hence, we may grant it the lowest weight. Loophole does not count as an active error and does not introduce new force. However, just by overlooking possible active errors occurring in remote layers, it allows risks to fall through the defence with an unlikely chance to leave a trace. Therefore it can be assigned the medium weight. Drip may be considered with the greatest weight due to the fact that it is hidden within the system and provoking new forms of active failure and error that is previously unexpected. In addition, errors in this latter stage of system are far less likely to be intercepted and more likely to reach the sharp end than in any previous stage.

The original SCM can only represent four cases (which we cover as severity levels 0–3 in Table 1) out of the HCM's eight; it fails to portray a defence layer's active potential to provoke error. However, it is clear in the HCM that a drip generated by a defence layer could pose as a bigger threat to patient safety if used in a healthcare context, due to the fact that the error is unexpected and invisible. In addition, the combination of the existence of these three elements can cause severe harm to the patient, as shown in severity level seven.

Shape-shifting feature in HCM

The HCM also brings a dynamic aspect to the system defence layer, as in reality system loopholes can develop over time. For example, when performing complex tasks, people tend to find shortcuts and develop workarounds, and if a small piece of a system defence layer gradually melts and weakens as a result of these shortcuts and workarounds, it is just a matter of time for a loophole to form.

Furthermore, the HCM manifests the possibility and importance of enabling some 'shape-shifting' feature in downstream system defence layer so as to accommodate the needs of upstream ones. This may make it appealing to other stakeholders, such as procurers and evaluators.

A controversial study published in 2005 reveals the unexpected increased mortality rate (from 2.8% to 6.57%) in the Children's Hospital of Pittsburgh after it had purchased and implemented a commercially sold computerised physician order entry (CPOE) system.[14] The study points out that this new system was rapidly implemented across the hospital over only six days while the hospital was simultaneously instituting other system changes. Evidence also shows that insufficient training, untailored drug library design, lack of understanding of current hardware capacity and failure in end-user

involvement also contributed to the increased mortality rate. Clearly the hospital seems to have neglected the complexity of the system and how the mismatches between new procedures and legacy workflow might potentially cause medical errors.

Although there might have been the intention that the computer system would be a simple 'slice of cheese,' the slice failed, as suggested by the HCM. For example, the hospital placed the workstations yards from where the physicians worked, so whatever defences or other benefits it was intended to provide it had to at least offset the compromised patient attention the physician needed to use the computer in the first place.

The HCM allows a dynamic and active interpretation of complex system components and helps to see the system as a whole. It could be a powerful tool to provoke thinking and discussion, especially amongst medical device procurers, who may develop a more comprehensive framework to evaluate products even more thoroughly than they have already done, prior to making any purchase and implementation decisions.

The original SCM makes it seem that defences are either solid cheese or holes. The HCM makes it clear that the cheese slice may be hot and flexible and that new holes may form. Imagine a man needs to operate a chainsaw, which if it accidentally comes in contact with his legs will lead to trauma. The operator takes reasonable precautions and wears protective trousers, which are intended to help stop the chainsaw should the chain touch them. However, it is not guaranteed that no injury would occur. The trousers are designed to quickly tear up and clog the chainsaw; they do not actually provide physical protection like armour and the operator would have to be very lucky to escape without injury – ironically, the protective trousers end up with holes in them. (Perhaps you can't have your cheese and heat it?)

HOT CHEESE METAPHORS
Fondue pot

The HCM is also a platform and foundation for discussion of incident causation and relevant fields, such as incident reporting and learning. If we consider a system defence layer as a slice of hot and melting cheese, then the current incident reporting and learning system has got to be the fondue pot placed beneath those layers (see Figure 3) to gather everything that managed to fall through the system, so as to visualise the accumulation of active failures and system design flaws, as well as to generate evidence to feedback to the design of the system and to facilitate future learning that benefits everyone.

In healthcare, the existence of such a mechanism can play a crucial role in preventing an incident from recurring. Incident reporting and learning systems are established and
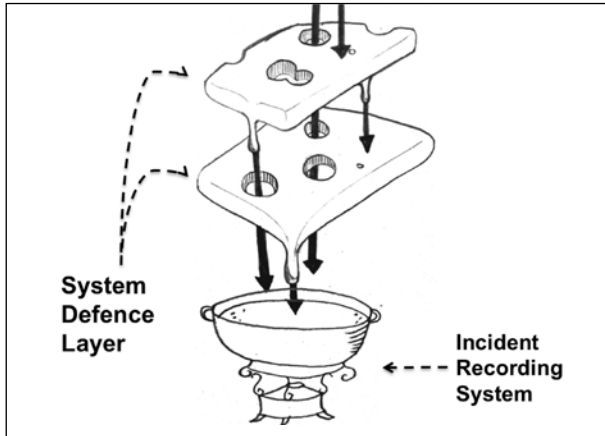
Y Li, H Thimbleby

FIGURE 3 Current incident reporting and recording system in HCM metaphors.

reinforced by regulatory agencies all around the globe, such as the US FDA's Medical Device Report (MDR) regulation and the formal UK National Reporting and Learning System (NRLS). The Medicines and Healthcare Products Regulatory Agency in the UK also host a reporting system that deals with medical device-related incidents.

Despite the well-known and well-advertised strengths and benefits of incident reporting and recording systems, under-reporting is known as a common problem. Possible barriers are suggested by several studies, such as inaccessibility and complexity. A collaborative hospital study states that a quarter of their participants did not know how to access an incident form, and more than 40% of consultants and registrars had never completed a report.[15] Other factors identified are cultural issues such as fear of punitive action, discrimination at the workplace and legal ramifications. From an individual practitioner's point of view, it is hard to report and record if an incident or a slip is simply unnoticed or unrecognisable. On top of that, lack of clarity regarding what should be reported and how the reports might lead to improvement in the existing systems are themselves responsible for poor reporting.

Rubbery cheese

In order to enable future learning, detailed and structured data should be provided via incident reporting, so to facilitate the reconstruction of the incident and to answer the questions of 'what, how and why'. Evidence shows that relying on the users to report an incident when it happens is, although theoretically promising, sometimes ineffective and inaccurate.

Under-reporting aside, incidents that do get reported often are poorly documented. A good-quality incident report should lend itself for detailed analysis of the chain of events that lead to the incident.[16] However, incidents that were immediate and often witnessed are better reported than those that had gradual development and multiple contributing factors.[15] In addition, current incident reporting and learning systems are not strong
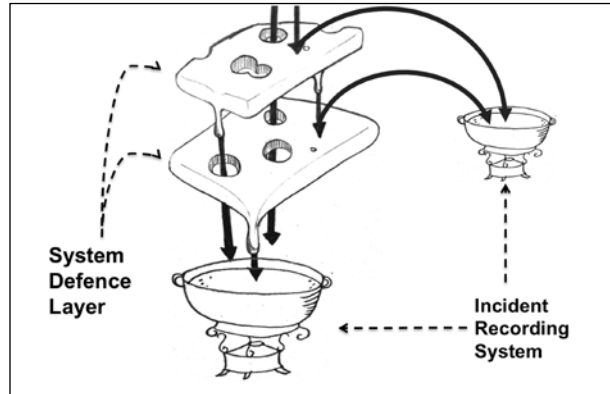

FIGURE 4 Rubbery hot cheese and incident recording.

on facilitating the documentation of 'unremarkable errors',[17] that is, errors and near misses that have not led to patient harm and are not reported or remarked upon; much the same point as Perrow made in 1999.[12]

The quantity of incident reporting is lower than it should be in an ideal world, then. Meanwhile, the quality of actual reports has been hampering incident analyses for years.[18] To improve things, we need mechanisms that enable data capturing and monitoring within every system defence layer, so 'unremarkable errors' and near misses, along with the gradual development of an incident and its contributing factors, can be made tangible. In other words, we need our hot cheese to be rubbery, so forces and drips that didn't made it through a defence can be bounced into the recording system as well, instead of disappearing without a trace (Figure 4).

Interaction logging and monitoring in high-risk interactive medical devices is a good example of how the defence layers become rubbery. An analysis carried out on logs from 58 infusion pumps revealed that, contrary to what was being emphasised in training, infusions were being stopped by simply opening the pump door, instead of pressing the stop button before opening the door.[19] When the door is open, a pump with infusion bag can only rely on a plastic safety clip to block the fluid in the line. If the clip breaks, the bag can empty itself into the patient in seconds. This is a dangerous violation and unsafe practice that puts the patient at a risk of over-infusion. However, since it has never really made it to the sharp end of the system, nobody may have noticed it.

Interaction monitoring, logging and reconstruction are well received and implemented in other high-risk industries, such as aviation. However, it is yet to be seen in healthcare systems and medical device design.

Unfortunately, the HCM does not solve the reporting problem. Human error only causes losses because nothing notices and blocks or intercepts it; in particular, no human notices the error before it turns into harm. And an error (or the harm it causes) has to be noticed before

it can even be reported; Koppel provides further discussion or missed errors,[20] and particularly the impact this has on evaluating the effectiveness on computer healthcare systems (see also our discussion of Han's study above[14]).

EDUCATIONAL POTENTIAL OF HCM

No qualitative model of error can be perfect. No matter what model it is, the ultimate goal is to encourage insightful discussion about issues and limitations of the design of safety systems, as well as to encourage broader thinking in incident analysis after the fact. What both cheese models do is provide a shared language, as well as a worked-out way of conceptualising what would have been blind spots in thinking. In particular, it is arguable that the HCM encourages deeper thinking because of the far wider range of interactions it can portray; literally because we can easily see (as in Table 1) what would have been invisible 'holes' in our thinking, we can consciously think through what to do.[21]

The original SCM has shortcomings if taken too literally. Who, for instance, has seen thin slices of Emmental (or is it Gruyère?) stacked on edge successfully, as in all the Swiss cheese diagrams? Arguably, the model needs turning on its side for a more realistic effect. We can heat up the cheese to make it boil and have the holes bubbling through in a dynamic and unpredictable way. And when Swiss cheese is hot, everyone can enjoy a good cheese fondue… This is the sort of motivation that makes teaching enjoyable! Indeed, the possibilities for cheesy humour make the idea of hot cheese and the fondue pot helpful and very memorable, particularly in educational settings.

CONCLUSIONS

We have presented hot cheese as a new recipe for understanding and talking about error reduction, creatively covering more issues than the SCM. However, like the SCM, the HCM is not intended to be a rigorous scientific model as such, but a flexible model that promotes the view of accidents as a combination of active errors and a failure of barriers, along with the view of barriers themselves as an active feature that changes over time and may also provoke new forms of errors. The HCM has the flexibility to highlight and discuss quite deep issues, and educators will find its flexibility (and its scope for humour and creative engagement with students) easy to mould to their specific teaching needs. We discussed the need of making the hot cheese rubbery, to capture upstream errors, near misses and the gradual development of an incident. We believe the greater flexibility of hot cheese lends itself to even more stimulating – and thoughtful – discussions of human error, without losing any of the advantages of Reason's original and powerful model.

REFERENCES

1 Reason JT. *Managing the risks of organizational accidents*. Aldershot: Ashgate Publishing; 1997.
2 Shappell SA, Wiegmann DA. T*he human factors analysis and classification system – HFACS*. Washington DC: US Department of Transportation; 2000. p.2
3 Dekker S. *The field guide to human error investigations*. Aldershot: Ashgate Publishing; 2002.
4 Berwick D. *A promise to learn – a commitment to act: improving the safety of patients in England*. National Advisory Group on the Safety of Patients in England; 2013. Available from: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/226703/Berwick_Report.pdf
5 Reason J, Hollnagel E, Paries J. *Revisiting the «Swiss cheese» model of accidents*. Brussels: Eurocontrol; 2006. Available from: http://publish.eurocontrol.int/eec/gallery/content/public/document/eec/report/2006/017_Swiss_Cheese_Model.pdf
6 Carthey J. Understanding safety in healthcare: the system evolution, erosion and enhancement model. *J Public Health Res* 2013; 2:e25:144–9.
7 Cook RI, Woods DD. 'Operating at the sharp end: the complexity of human error' in Bogner MS, editor. *Human error in medicine*. Hillsdale, NJ: Lawrence Erlbaum Associates; 1994. p.255–310
8 Hinton-Walker P, Carlton G, Holden L et al. 'The intersection of patient safety and nursing research,' in Fitzpatrick JJ, Hinton-Walker P, editors. *Annual review of nursing research*. Vol. 24. New York: Springer Publishing; 2006. p.8–9
9 Walsh KE, Kaushal R, Chessare JB. How to avoid paediatric medication errors: a user's guide to the literature. *Arch Dis Child* 2005; 90:698–702. http://dx.doi.org/10.1136/adc.2003.048827

10 Poon EG, Cina JL, Churchill W et al. Medication dispensing errors and potential adverse drug events before and after implementing bar code technology in the pharmacy. *Ann Intern Med* 2006; 145:426–34.
11 Koppel R, Wetterneck T, Telles JL et al. Workarounds to barcode medication administration systems: their occurrences, causes, and threats to patient safety. *J Am Med Inform Assoc* 2008; 15:408–23. http://dx.doi.org/10.1197/jamia.M2616
12 Perrow C. *Normal accidents: living with high risk technologies*. Princeton, NJ: Princeton University Press; 1999.
13 Fuller JG. *We almost lost Detroit: this is not a novel*. New York: Readers Digest; 1975.
14 Han YY, Carcillo JA, Venkataraman ST et al. Unexpected increased mortality after implementation of a commercially sold computerized physician order entry system. *Pediatrics* 2005; 116:1506–12.
15 Evans SM, Berry JG, Smith BJ et al. Attitudes and barriers to incident reporting: a collaborative hospital study. *Qual Saf Health Care* 2006; 15:39–43.
16 Mahajan RP. Critical incident reporting and learning. *Br J Anaesth* 2010; 105:69–75. http://dx.doi.org/10.1093/bja/aeq133
17 Furniss D, Blandford A, Meyer A. Unremarkable errors: low-level disturbances in infusion pump use. *Proceedings British Computer Society HCI Conference*; 2010. Available from: http://discovery.ucl.ac.uk/1347246/2/2011_DFABAM_UnremarkErr_BHCI.pdf
18 Merry AF. Safety in anaesthesia: reporting incidents and learning from them. *Anaesthesia* 2008; 63:337–98. http://dx.doi.org/10.1111/j.1365-2044.2008.05517.x
19 Oladimeji P, Li Y, Cauchi A et al. Visualising medical device logs. Paper presented at 1st BCS Health Wales Workshop, 2011.
20 Koppel R. Is healthcare information technology based on evidence? *Yearb Med Inform* 2013; 8:7–12.
21 Kahnemann D. *Thinking, fast and slow*. London: Allen Lane; 2011.