

The Unintended Consequences of Health Information Technology Revisited

E. Coiera¹, J. Ash², M. Berg³

¹ Australian Institute of Health Innovation, Macquarie University, Australia

² Department of Medical Informatics and Clinical Epidemiology, School of Medicine, Oregon Health & Science University, USA

³ Principal, Advisory, KPMG LLP (US)

Summary

Introduction: The introduction of health information technology into clinical settings is associated with unintended negative consequences, some with the potential to lead to error and patient harm. As adoption rates soar, the impact of these hazards will increase.

Objective: Over the last decade, unintended consequences have received great attention in the medical informatics literature, and this paper seeks to identify the major themes that have emerged.

Results: Rich typologies of the causes of unintended consequences have been developed, along with a number of explanatory frameworks based on socio-technical systems theory. We however still have only limited data on the frequency and impact of these events, as most studies rely on data sets from incident reporting or patient chart reviews, rather than undertaking detailed observational studies. Such data are increasingly needed as more organizations implement health information technologies. When outcome studies have been done in different organizations, they reveal different outcomes for identical systems. From a theoretical perspective, recent advances in the emerging discipline of implementation science have much to offer in explaining the origin, and variability, of unintended consequences.

Conclusion: The dynamic nature of health care service organizations, and the rapid development and adoption of health information technologies means that unintended consequences are unlikely to disappear, and we therefore must commit to developing robust systems to detect and manage them.

Keywords

Safety, error, electronic health records, computer provider order entry, human-computer interaction, incident reporting.

Yearb Med Inform 2016:163-9

<http://dx.doi.org/10.15265/IY-2016-014>

Published online November 10, 2016

Introduction

The proposition that health information technology (HIT) is necessary but not sufficient to deliver modern, complex health care now seems undebatable. Today, we understand that HIT does not work in isolation, but is just one of many components that come together to create a genuinely complex and socio-technical health care system [1]. If we do not fit technology to workflow and to user, if we do not factor in the competing demands that clinicians must juggle as they use a technology, and if the technology is not fit for purpose, then no one should be surprised if it does little good, and indeed leads to patient harm or increased cost.

Such thinking is, however, relatively recent. It was only a little over a decade ago that the authors wrote a paper on the unintended consequences of HIT [2]. That paper challenged the existing dominant paradigm in informatics that emphasized technology architecture and design over human-computer interaction, and the collection and retrieval of data in patient records over the meaningful use of information by health professionals, patients, or consumers to improve decisions. The paper began by observing that while patient care information systems (PCISs) were being lauded as a core building block for a safer health care system, the reality was that they were creating unanticipated negative consequences:

“It is obvious that PCISs will ultimately be a necessary component of any high-quality health care delivery system. Yet, in our research in three different countries, we have each encountered many instances in which PCIS applications seemed to foster errors rather than reduce their likelihood.

In health care practices in the United States, Europe, and Australia alike, we have seen situations in which the system of people, technologies, organizational routines, and regulations that constitutes any health care practice seemed to be weakened rather than strengthened by the introduction of the PCIS application. In other words, we frequently observed instances in which the intended strengthening of one link in the chain of care actually leads unwittingly to a deletion or weakening of others.”

Such thinking did not emerge from nowhere. The authors drew together three streams of prior work to unpack the origins of the unintended negative consequences of technology. Ash drew on an important stream of research that explored the importance of people and their organizations to the successful implementation and use of technology [3]. Berg brought socio-technical theory to help understand the reasons why clinical decision-support systems were not always adopted or used successfully [4]. Coiera, prompted by observational studies of clinicians as they carried out their work, identified that it was human-to-human communication, rather than documentation, that was the primary information task in health care [5]. By ignoring everything in the ‘communication space’, informatics was building systems that did not actually fit the real needs of clinical practice.

For a discipline in its infancy still seeking legitimacy, the discussion of unintended consequences initially received a poor reaction in some quarters. The prevailing belief was that technology was clearly the key to building a safer, more efficient and evi-

dence-based health care system, and focusing on the 'rare' instances when something went wrong was 'talking down' the industry. Much has changed in the last decade. A steady stream of research from many independent groups around the world has described in growing detail the nature, extent, and consequences of technology-induced harm [6]. Adoption of HIT has skyrocketed to become the norm rather than the exception. At the clinical frontline, clinicians are only too aware of the difficult relationship they have with information systems. Clinical IT offers much, but extracts a price because of its complexity and difficulty of use, and looks nothing like the slick technologies available to consumers outside of health care. Those responsible for large health care IT projects are only too aware of the very real risks of project failure, cost overruns, patient harm, and clinical pushback. As adoption rates soar, the impact of these hazards is likely to increase as well [7].

What Are Unintended Consequences?

In general, unintended consequences are outcomes that have not been anticipated. These consequences can be either desirable or undesirable, positive or negative. Although there can be positive unanticipated consequences of HIT, most interest is in unanticipated negative consequences. Negative consequences for patients can be minor, such as having a duplicate blood test, but extend through to injuries or even death, for example through incorrect or delayed treatment [8]. As with many things, what appears negative to one stakeholder group (the additional time taken by providers to enter structured data) may appear positive to another (administrators who see improvements in billing through more detailed records) [9].

In our 2004 paper, we described for the first time a broad categorization scheme for the negative unintended consequences of HIT:

- The first category, *errors in the process of entering and retrieving information*, included:

- 1) The existence of human-computer interfaces not suitable for health care's highly interruptive context, and;
 - 2) Increased cognitive load from overemphasizing the need for structured data.
- The second category, *errors in the communication and coordination process* included:
 - 1) The misrepresentation of health care work as a linear process, leading to inflexibility, workarounds, problems with transfers, and;
 - 2) The misrepresentation of communication as information transfer, rather than interactive sense-making weakened communication actions, resulting in loss of feedback, decision support overload, and the need for constant human diligence to catch errors.

Since 2004, numerous studies have described unintended consequences in different settings such as residential aged care homes [10], and with different HIT beyond EHRs such as bar code medication administration [11, 12], health information exchange [13], hands free communication devices [14], and speech recognition [15]. Not all negative consequences occur because of error, and we now know that delays in the process of diagnosis or starting treatment can be caused by HIT problems, and these are just as capable of leading to patient harm [16]. The literature on unintended consequences has thus continued to produce reports of new types of HIT problem, and understanding how each relates to the other has not always been straightforward. Complicating matters further, while some negative consequences cannot happen without the use of an electronic system (computer menu pick list errors), others existed in the pre-computer world (writing in the wrong patient's notes), but might become more likely with automation [17].

Perhaps the most comprehensive attempt at integrating our knowledge of the different types is the classification scheme developed by Magrabi, Coiera, and colleagues, which utilizes patient safety concepts, and has been driven by analysis of incident reports from Australia, the US, and the UK [16, 18, 19]. The scheme has a primary axis which separates unintended consequences into

those that have a primary genesis in machine errors (such as poorly designed user interfaces or computer network downtimes) or in human-initiated errors (such as a work-around). The classification now includes large-scale events where potentially many hundred or indeed thousands of patients might be affected because of an error (such as letters with patient laboratory results being sent to the wrong address) [19]. A more specialized taxonomy of medication error types associated with the use of Computerised Physician Order Entry (CPOE) has also been developed, based on medication errors reported via incident reporting systems [20].

Unintended Consequences of Clinical Decision-Support Systems

Clinical decision support systems (CDSSs) take many forms, and support a wide variety of tasks. Unintended consequences of CDSSs have mainly been reported in electronic prescribing [21] and medication administration [12]. CPOE systems have been documented in many studies to both reduce existing errors, as well as to introduce new errors [22]. These new problems include duplication of orders and selection errors where a user might pick the wrong drug, dose, dose frequency, or formulation from a drop-down menu.

Some CPOE problems arise because of inconsistency between different systems in the type and content of alerts generated for similar prescribing scenarios [20, 23]. Introducing even plausible design elements into CPOEs such as 'hard stops' to prevent clearly inappropriate drug combinations from being ordered can lead to new problems. A hard stop alert to prevent warfarin and trimethoprim-sulfamethoxazole being ordered seemed to effectively change prescribing behaviors but also led to clinically important treatment delays for patients, who needed immediate therapy [24].

Qualitative research by Ash et al. has identified nine types of CPOE unintended consequences which include: 1) workflow issues, 2) new kinds of errors, 3) changes

in communication patterns and practices, 4) more/new work for clinicians, 5) never ending system demands, 6) changes in the power structure, 7) overdependence on the technology, 8) emotions, and 9) paper persistence [9, 25]. These “new kinds of errors” are of particular interest because of their potential to do the greatest harm to patients, and have since been labeled by some as “e-iatrogenesis” [26], although patient harm can come from any of the classes identified.

What Have We Learned?

The first stage in the evolution of our understanding of unintended consequences necessarily focused on empirical studies which provided examples of the harms that might arise from HIT and the different circumstances and errors that led to those harms. Since our 2004 paper, over 50 new papers discussing the unintended consequences of HIT have been added to Medline. Only about half focus on outcomes. In 2012 when the Institute of Medicine study about HIT safety was published, only a handful of studies had analyzed HIT safety events [27]. While important, few of these studies tell us much about *likelihood* and *impact*. In other words, just because harms can arise in principle does not mean that they are common events, nor that they cause significant problems. The next stage in the research endeavor has thus been to identify both the frequency of HIT-related harm events, and their cost to the health system and to patients.

The Epidemiology of Unintended Consequences

The great strength of voluntary incident report studies is the richness of events that they contain, and the often-invaluable descriptions of the circumstances that lead to an event [28]. While it is tempting, especially with large incident data sets, to draw conclusions about the probability of events, any such analyses are statistically flawed because the frequency of incident reports do not reflect the true incidence of events in the world [29]. Incident reports typically contain

events that are somehow notable to those writing the reports, and so they will under-report minor events that seem unimportant or that have been previously reported. To estimate true frequencies of events, we need to conduct prospective studies, which count events in an unbiased and reproducible way.

One source of such data comes from software companies, which can provide estimates of the likely number of defects in a working clinical system. Unfortunately such data are rarely available to the public. One now old dataset was generated from three releases of a major commercial USA medical record system and is illustrative [30, 31]. That system contained 188 separate software components across 6500 files. Release 1 had defects in 58 of 180 components, seven of which were discovered post release. Release 2 had 64 defects in 185 components, with five discovered post release. Release 3 showed a numerical improvement in quality, with only 40 of 188 components being defective, but still six were only discovered post release.

True harm rates can also be estimated by chart reviews, which randomly select patient records for assessment. A CPOE pediatric medication error study analyzed 104 errors detected during 352 randomly selected pediatric admissions to a teaching hospital [32]. Of these, 7 serious errors were computer-related. The authors concluded that “serious pediatric computer-related errors are uncommon (3.6 errors per 1000 patient-days)”. The challenge with chart reviews however is that they can only count what has been recorded. Chart reviews share several of the limitations of incident reports: clinicians cannot record events that they did not see, may incorrectly fail to ascribe causation of an event to HIT, or may not record events at all.

Prospective studies that directly observe clinical work provide us probably the richest and most accurate estimates of the true frequency of HIT-related events and harms. Using a standardized workflow measurement tool and error categorization scheme, Westbrook and colleagues undertook detailed analyses of medication error rates before and after the introduction of a CPOE system in two hospitals [33–34]. Use of a CPOE resulted in a decline in medication errors from 6.25 errors per admission to

2.12 at one hospital, and from 3.62 to 1.46 errors at the second hospital. This decrease was driven by a large reduction in unclear, illegal, and incomplete orders. Serious errors decreased by 44% (0.25 per admission to 0.14) compared to the control wards (17% reduction; 0.30–0.25). The study is interesting from the perspective of unintended consequences because post-implementation the CPOE system introduced new classes of error. It was estimated that about 40% of the residual error rates of electronic prescribing were system- rather than prescriber-related. The authors noted that system-related errors are thus very frequent, yet few are routinely detected, reinforcing the inadequacy of incident reporting or chart review as a way of measuring true incident frequency.

Most clinical organizations are unlikely to have the resources to undertake such detailed observations, which means that routine estimates of IT-related events and harms must depend on alternate approaches. A promising approach is to use IT to monitor such events. Real-time monitoring of clinical information systems has the potential to identify system downtimes and periods of network congestion that may risk patient care [35]. Statistical process control models can identify errors in patient records, such as high rates of duplicate or missing orders, unexpected delays in orders, and special trigger rules can identify events that might be risky, such as the ordering and deletion of a medication order in rapid succession [36].

Several studies have attempted to assess the impact of HIT-related harms on patient outcomes, including death rates. Three US hospitals in Pittsburgh, PA, Seattle, WA, and Palo Alto, CA implemented the same EHR and CPOE [37–39]. Five months after implementation the mortality rate in the Pittsburgh paediatric intensive care unit (ICU) increased from 2.8% to 6.6%, while there was a non-significant change in Seattle 13 months after implementation from 4.2% to 3.5%, a finding subsequently replicated in other institutions [40]. Hospital-wide, there was a significant decrease in mortality in Palo Alto. The disparity in patient outcomes probably reflects the socio-technical nature of computer systems and was most likely due to differences in implementation processes, including differences in clinical workflow,

speed of implementation and staff training [41]. The latter papers offer insight into successful implementation processes. These studies as a group offer examples of how outcomes vary depending on the journey taken between conception and implementation of HIT.

The Genesis of Unintended Consequences

We noted in our 2004 paper when HIT systems were implemented in a clinical setting that “we are confronted with a large socio-technical system in which many behaviors emerge out of the sociotechnical coupling” of humans and systems. By sociotechnical, we meant that technology is never independent of the context in which it is embedded, and that this context is a larger construct involving people and processes. Today implementation science researchers understand this perspective very well. We now see the implementation of HIT not as the injection of technology into a location, but as a process in which we mold together a unique bundle that includes technology, work processes, people, training, resources, culture, and more [42].

The socio-technical perspective brought with it a view that unintended consequences emerged out of this complexity, and that causation was rarely going to be simple and linear. From a patient safety perspective, much of what we call unintended consequences fits well within existing frameworks that apply to other domains in health care. Most patient safety problems do not arise because of single points of failure by a human, process, or technology. Rather, they emerge out of the interaction of many events, and many potential problems never cause actual harm because they are caught by system defenses early enough [43]. Patient safety studies and implementation science are both rich fields of research that can thus tell us much about why unintended consequences occur, and what we can do to minimize their likelihood and severity [44].

Recent work has tried to bring together what we know about HIT related harm and patient safety research. Three complementary sociotechnical frameworks have been proposed, each based on extensive fieldwork.

The first general model describes two frames of reference which, when they collide, can cause unsafe HIT situations. The first frame, the Practice Frame, is the user’s orientation and the second System Frame represents the technology implementation’s orientation. Each frame needs to adapt over time [45]. The second eight-dimensional sociotechnical model for HIT safety is more detailed and is designed to help us understand the risks related to the interplay of the different dimensions including workflow, people, and technology [46]. This model was used during development of the U.S. Office of the National Coordinator for HIT’s SAFER Guides addressing HIT safety [47]. A third sociotechnical framework takes an interactional perspective, looking at the different ways system components like workflows, culture, social interactions, and technologies can change each other [48].

The development of these more general frameworks is not without controversy. While the classification scheme of HIT harms developed by Magrabi is pragmatically designed as a tool to classify incident reports and not to explain their causes [16], the sociotechnical frameworks are theoretically oriented, and exist to explain. Some scholars suggest that because these frameworks are being developed at a time when our understanding of causation is still immature, using them as explicit, exclusive recipes for dealing with HIT harm is risky [49].

Cognitive Biases

Negative consequences arising from the use of technology have been long recognized in the literature on cognitive biases. The nature of human cognition, its limitations, and reliance on heuristics can lead humans to make erroneous decisions because of, or in spite of, the use of technology. Biases alter the way humans weight the importance of data when making a decision, for example leading to emphasis on data that was seen first or last, or that is most similar to past experience [42]. For example, clinicians and consumers can misinterpret data from information retrieval systems because they interpret any new information through the lens of their prior beliefs [50].

Recent progress in psychology suggests that many of the different biases associated with decision-making all arise from the same underlying mechanism. The decision by sampling model suggests that individuals preferentially work from samples of events that are most memorable, and that typically come from personal experience [51]. As any personal sample is typically small and unrepresentative of the true distribution, our decisions are often equally biased. Such distortions are common in human assessments of health risks, where individuals play down risks associated with behaviors such as smoking or exposure to HIV.

Of specific interest to decision support system use, automation biases or automation-induced complacency are increasingly recognized as an important source of decision error [52, 53]. For example, when using a decision support system, a user can make either errors of omission (they fail to perform actions because the system did not prompt them to take notice) or errors of commission (they did what the decision system told them to do, even when it contradicts their training and available data) [54].

It has been suggested that when humans delegate tasks to a computer system, they also shed task responsibility. Computer users may then take themselves out of the decision loop and develop an “out of loop unfamiliarity” with the system they are meant to be watching [55]. If an urgent event occurs, recovering from loop unfamiliarity requires additional time and cognitive resources to obtain the necessary understanding of all the variables required to make a decision, or situational awareness.

Some evidence suggests that explicit training in automation bias has a short-term benefit only. Making individuals personally accountable for the consequences of their decisions however does seem to reduce automation bias. For example, if individuals are told that their actions are socially accountable, because the data of their performance are being recorded and will be shared with others, then more time is spent verifying the correctness of a decision support system’s suggestions leading to fewer errors [56].

Cognitive biases are thus predictable consequences of interaction with technology, and can be to some extent minimized by debiasing

designs, which recast decisions in a way that makes the appearance of biases less likely [57]. Although such biases are strictly not unanticipated technology effects, their very real impact on decision outcomes has meant that they are often considered alongside other causes of unintended consequences.

The Science of Unintended Consequences

We find ourselves at a challenging point in the development of HIT. The argument that information technology is essential for operating a safe, efficient and sustainable health system has been successfully made, and most nations are seeing large-scale implementation of clinical information systems across all sectors of health care. At the same time, we have yet to develop a deep safety culture to match [7]. HIT safety standards are rarely mandated and are still immature. There is weak reporting of HIT-related events so no nation is able at this point to accurately report on the true incidence or consequences of HIT related-harms.

The shift in recent years from home-grown and maintained clinical systems to commercially built, locally customized ones, also brings risks and benefits [58]. Mass production should see an increase in the quality of system construction and functionality. It could also however result in a poorer fit of technology to work, compared to earlier locally implemented and purpose built systems. Much thus remains to be done, both scientifically, to understand why and how often harms occur, how they can be avoided, and pragmatically, to implement this knowledge into technology, policy, and practice.

From a scientific point of view, much can still be borrowed from disciplines such as psychology, cognitive science, human factors, and safety science. In so far as HIT risks are associated with the ways humans think and react to circumstances, it would be foolish to ignore many decades of robust science. Health care however does bring its own special circumstances, and issues that loom large in our domain get less attention in other disciplines.

There has been nearly a decade and a half of research on the pervasive nature of interruption and multitasking in health care work and their patient safety implications [59]. While we recognize that humans must seek to minimize interruptions and actively manage their work when they are interrupted, such considerations do not seem to have made their way into HIT design. Clinicians are routinely interrupted by HIT-generated alerts, often unnecessarily. Rather than recognizing that clinicians might have to suspend their tasks because of interruption, HIT is designed on the unreasonable assumption that it should have the full and undivided attention of their users. We clearly need to better understand how interruption and multitasking lead to unintended consequences and we need to mitigate those risks through smarter HIT design [54].

Workflows in health care are also very different from those in more traditional safety critical industries like aerospace or power plants. Work patterns seem less linear, less designed, and patterns are more adaptive and emergent as clinicians juggle multiple competing demands under resource constraints. This no doubt leads to greater complexity in task structures, and we know from psychology that complexity leads to cognitive load and error [60]. HIT, as a well behaved actor in the socio-technical space that is health care, needs to contribute to complexity reduction and not make it worse [61]. We also need to understand more about the reasons why humans continue to generate adaptations in these complex work settings and what that means for technology design. The fact that humans create workarounds to subvert HIT, with unintended negative consequences, is now well documented [12], but that technology needs to be designed in a way that accepts the need for workarounds and supports that need is much less well understood [62].

The emergence of implementation science as a coherent discipline has been amongst the most significant ideas to emerge since we wrote our first paper on unintended consequences [63]. Implementation science takes what was once thought of as ‘mere’ pragmatics – the installation of a new process or technology – and reconceived it as a complex adaptive process, governed by

its own physics. Different outcomes were always anticipated when the same HIT system was installed in two different places, but implementation science helps us see that this is actually a crucial observation.

Once we understand that local context – the sociotechnical system – is almost unique for every implementation – then we understand that there is no silver bullet to ‘solving’ unintended consequences. It is also a dynamic problem given that context constantly changes. The relentless evolution in the nature of technology, work practices, and the behavior and skills of patients and their caregivers brings with it a similar evolution in unintended consequences. There are processes and designs that can minimize risk and catch unwanted events before they lead to harm, but we should not ever believe that we can predict well enough where, when, and how any specific event might occur.

A major perspective from implementation science is thus that we need to see the study of unintended consequences as local – what happened here might not happen there – which limits generalizability of specific events to other settings. However implementation science is developing general models that should guide how we adapt the process of HIT implementation by fitting technology to context, which should reduce unexpected surprises later on.

The final lesson from implementation science is that implementation never ends. It is not only the moment when a new technology is introduced. As an organization changes over time in role, the work that it does, and in the way it is structured, any implemented technology becomes increasingly out of step with the place in which it is embedded [62]. The adaptations and workarounds that emerge over time to fix this mismatch will yield new, previously unpredicted unintended consequences.

Conclusions

The goal in studying unintended consequences is not to subvert the necessary progress in moving health care into the digital world, but rather to make sure that this journey is as fruitful as possible, and is not

sidetracked by clearly avoidable obstacles. Health care is a safety critical industry, just like the airline industry, and it is deep in the process of developing the safety culture and systems that we would expect of such an enterprise. Information technology must be added to the patient safety agenda, both because it is a new and still not well understood cause of patient harm, but also because it has the potential to avoid harm when it is well designed and effectively used.

Crucially, technologies do not stand still, nor do practice, organization, or the emergence of the unexpected. As long as we work in complex adaptive organizations, we will always have more to learn about unintended consequences if we are to manage them. HIT unintended consequences are not going to go away.

References

- Coiera E. Putting the technical back into socio-technical systems research. *Int J Med Inform* 2007;76):S98-S103.
- Ash JS, Berg M, Coiera E. Some unintended consequences of information technology in health care: the nature of patient care information system-related errors. *J Am Med Inform Assoc* 2004;11:104-12.
- Ash J. Organizational factors that influence information technology diffusion in academic health sciences centers. *J Am Med Inform Assoc* 1997;4(2):102-11.
- Berg M. Rationalizing medical work: decision-support techniques and medical practices. MIT press; 1997.
- Coiera E. When conversation is better than computation. *J Am Med Inform Assoc* 2000;7(3):277-86.
- Foundation NPS. Free from Harm: Accelerating Patient Safety Improvement Fifteen Years after To Err Is Human. Boston, MA: National Patient Safety Foundation; 2015.
- Coiera E, Aarts J, Kulikowski C. The dangerous decade. *J Am Med Inform Assoc* 2012;19:2-5.
- Magrabi F, Ong MS, Runciman W, Coiera E. Patient safety problems associated with healthcare information technology: an analysis of adverse events reported to the US Food and Drug Administration. *AMIA Annu Symp Proc* 2011:853-7.
- Ash JS, Sittig DF, Dykstra RH, Guappone K, Carpenter JD, Seshadri V. Categorizing the unintended sociotechnical consequences of computerized provider order entry. *Int J Med Inform* 2007 Jun;76 Suppl 1:S21-7.
- Yu P, Zhang Y, Gong Y, Zhang J. Unintended adverse consequences of introducing electronic health records in residential aged care homes. *Int J Med Inform* 2013;82(9):772-88.
- Novak LL, Holden RJ, Anders SH, Hong JY, Karsh BT. Using a sociotechnical framework to understand adaptations in health IT implementation. *Int J Med Inform* 2013 Dec;82(12):e331-44.
- Koppel R, Wetterneck T, Telles JL, Karsh BT. Workarounds to barcode medication administration systems: their occurrences, causes, and threats to patient safety. *J Am Med Inform Assoc* 2008;15(4):408-23.
- Kuperman G, McGowan J. Potential Unintended Consequences of Health Information Exchange. *J Gen Intern Med* 2013;28(12):1663-66.
- Richardson JE, Ash JS. The effects of hands-free communication device systems: communication changes in hospital organizations. *J Am Med Inform Assoc* 2010;17(1):91-8.
- Hodgson T, Coiera E. Risks and benefits of speech recognition for clinical documentation: a systematic review. *J Am Med Inform Assoc* 2016 Apr;23(e1):e169-79.
- Magrabi F, Ong MS, Runciman W, Coiera E. Using FDA reports to inform a classification for health information technology safety problems. *J Am Med Inform Assoc* 2012;19(1):45-53.
- Magrabi F, Liaw ST, Arachi D, Runciman W, Coiera E, Kidd MR. Identifying patient safety problems associated with information technology in general practice: an analysis of incident reports. *BMJ Qual Saf* 2015 Nov 5.
- Magrabi F, Ong M-S, Runciman W, Coiera E. An analysis of computer-related patient safety incidents to inform the development of a classification. *J Am Med Inform Assoc* 2010 Nov-Dec;17(6):663-70.
- Magrabi F, Baker M, Sinha I, Ong MS, Harrison S, Kidd MR, et al. Clinical safety of England's national programme for IT: A retrospective analysis of all reported safety events 2005 to 2011. *Int J Med Inform* 2015 Mar;84(3):198-206.
- Schiff G, Amato M, Eguale T, Boehne JJ, Wright A, Koppel R, et al. Computerised physician order entry-related medication errors: analysis of reported errors and vulnerability testing of current systems. *BMJ Qual Saf* 2015 Apr;24(4):264-71.
- Turchin A, Shubina M, Goldberg S. Unexpected Effects of Unintended Consequences: EMR Prescription Discrepancies and Hemorrhage in Patients on Warfarin. *AMIA Annu Symp Proc* 2011;2011:1412-17.
- Reckmann MH, Westbrook JI, Koh Y, Lo C, Day RO. Does computerized provider order entry reduce prescribing errors for hospital inpatients? A systematic review. *J Am Med Inform Assoc* 2009 Sep-Oct;16(5):613-23.
- Slight SP, Eguale T, Amato MG, Seger AC, Whitney DL, Bates DW, et al. The vulnerabilities of computerized physician order entry systems: a qualitative study. *J Am Med Inform Assoc* 2016 Mar;23(2):311-6.
- Strom BL, Schinnar R, Abera F, Bilker W, Hennessey S, Leonard CE, et al. Unintended effects of a computerized physician order entry nearly hard-stop alert to prevent a drug interaction: a randomized controlled trial. *Arch Intern Med* 2010;170(17):1578-83.
- Campbell EM, Sittig DF, Ash JS, Guappone KP, Dykstra RH. Types of Unintended Consequences Related to Computerized Provider Order Entry. *J Am Med Inform Assoc* 2006;13(5):547-56.
- Weiner JP, Kfuri T, Chan K, Fowles JB. "e-Iatrogenesis": The Most Critical Unintended Consequence of CPOE and other HIT. *J Am Med Inform Assoc* 2007;14(3):387-8.
- Institute of Medicine. Health IT and Patient Safety: Building Safer Systems for Better Care. The National Academies Press; 2012.
- Miller M, Clark J, Lehmann C. Computer based medication error reporting: insights and implications. *Qual Saf Health Care* 2006;15(3):208-13.
- Shojania KG. The frustrating case of incident-reporting systems. *Qual Saf Health Care* 2008;17(6):400-02.
- Hewett R, Kulkarni A, Seker R, Stringfellow C. On effective use of reliability models and defect data in software development. Region 5 Conference, 7-9 April 2006 IEEE, San Antonio, TX, USA. 2006. p. 67-71.
- Stringfellow C, Andrews A, Wohlin C, Petersson H. Estimating the number of components with defects post-release that showed no defects in testing. *Software Testing, Verification and Reliability* 2002;12(2):93-122.
- Walsh KE, Adams WG, Bauchner H, Vinci RJ, Chessare JB, Cooper MR, et al. Medication Errors Related to Computerized Order Entry for Children. *Pediatrics* 2006;118(5):1872-9.
- Westbrook JI, Reckmann M, Li L, Runciman WB, Burke R, Lo C, et al. Effects of two commercial electronic prescribing systems on prescribing error rates in hospital in-patients: a before and after study. *PLoS Med* 2012;9(1):e1001164.
- Westbrook JI, Baysari MT, Li L, Burke R, Richardson KL, Day RO. The safety of electronic prescribing: manifestations, mechanisms, and rates of system-related errors associated with two commercial systems in hospitals. *J Am Med Inform Assoc* 2013;20(6):1159-67.
- Ong MS, Magrabi F, Coiera E. Syndromic surveillance for health information system failures: a feasibility study. *J Am Med Inform Assoc* 2012;20(3):506-12.
- Adelman JS, Kalkut GE, Schechter CB, Weiss JM, Berger MA, Reissman SH, et al. Understanding and preventing wrong-patient electronic orders: a randomized controlled trial. *J Am Med Inform Assoc* 2013;20(2):305-10.
- Han YY, Carcillo JA, Venkataraman ST, Clark RS, Watson RS, Nguyen TC, et al. Unexpected increased mortality after implementation of a commercially sold computerized physician order entry system. *Pediatrics* 2005;116(6):1506-12.
- Del Beccaro MA, Jeffries HE, Eisenberg MA, Harry ED. Computerized provider order entry implementation: no association with increased mortality rates in an intensive care unit. *Pediatrics* 2006;118(1):290-95.
- Longhurst CA, Parast L, Sandborg CI, Widen E, Sullivan J, Hahn JS, et al. Decrease in hospital-wide mortality rate after implementation of a commercially sold computerized physician order entry system. *Pediatrics* 2010;126(1):14-21.
- Keene A, Ashton L, Shure D, Napoleone D, Katyal C, Bellin E. Mortality before and after initiation of a computerized physician order entry system in a critically ill pediatric population. *Pediatr Crit Care*

- Med 2007;8(3):268-71.
41. Rosenbloom ST, Harrell FE, Lehmann CU, Schneider JH, Spooner SA, Johnson KB. Perceived increase in mortality after process and policy changes implemented with computerized physician order entry. *Pediatrics* 2006;117(4):1452-55.
 42. Coiera E. *Guide to Health Informatics* (3rd Edition). 3rd ed. London: CRC Press; 2015.
 43. Braithwaite J, Coiera E. Beyond patient safety Flatland. *J R Soc Med* 2010;103(6):219-25.
 44. Hollnagel E, Braithwaite J, Wears RL. *Resilient health care*. Ashgate Publishing, Ltd.; 2013.
 45. Novak LL, Holden RJ, Anders SH, Hong JY, Karsh BT. Using a sociotechnical framework to understand adaptations in health IT implementation. *Int J Med Inform* 2013;82(12):e331-e44.
 46. Sittig DF, Singh H. A new sociotechnical model for studying health information technology in complex adaptive healthcare systems. *Qual Saf Health Care* 2010;19(Suppl 3):i68-i74.
 47. Sittig DF, Ash JS, Singh H. *ONC issues guides for SAFER EHRs*. *JAHIMA* 2014;85:50-2.
 48. Harrison MI, Koppel R, Bar-Lev S. Unintended consequences of information technologies in health care—an interactive sociotechnical analysis. *J Am Med Inform Assoc* 2007;14(5):542-9.
 49. Koppel R. The health information technology safety framework: building great structures on vast voids. *BMJ Qual Saf* 2016 Apr;25(4):218-20.
 50. Lau AYS, Coiera EW. Do People Experience Cognitive Biases while Searching for Information? *J Am Med Inform Assoc* 2007;14(5):599-608.
 51. Brown GD, Matthews WJ. Decision by Sampling and Memory Distinctiveness: Range Effects from Rank-Based Models of Judgment and Choice. *Front Psychol* 2011;2:299.
 52. Coiera E, Westbrook JI, Wyatt JC. The Safety and Quality of Decision Support Systems. *Methods Inf Med* 2006;45(1):S20-5.
 53. Goddard K, Roudsari A, Wyatt JC. Automation bias: a systematic review of frequency, effect mediators, and mitigators. *J Am Med Inform Assoc* 2012;19(1):121-7.
 54. Coiera E. Technology, cognition and error. *BMJ Qual Saf* 2015;24(7):417-22.
 55. Wickens C, Hollands J, Parasuraman R. *Engineering Psychology and Human Performance*. 4th ed. New Jersey: Pearson; 2012.
 56. Skitka L, Mosier K, Burdick M, Rosenblatt B. Automation bias and errors: are crews better than individuals. *Int J Aviat Psychol* 2000;10(1):85-97.
 57. Lau AYS, Coiera EW. Can Cognitive Biases during Consumer Health Information Searches Be Reduced to Improve Decision Making? *J Am Med Inform Assoc* 2009;16(1):54-65.
 58. Koppel R, Lehmann CU. Implications of an emerging EHR monoculture for hospitals and healthcare systems. *J Am Med Inform Assoc* 2015;22(2):465-71.
 59. Coiera E. The science of interruption. *BMJ Qual Saf* 2012;21(5):357-60.
 60. Sweller J, Ayres P, Kalyuga S. *Cognitive load theory*. Springer; 2011.
 61. Sintchenko VS, Coiera E. Decision complexity affects the extent and type of decision support use. *AMIA Annu Symp Proc*; 2006:724-8.
 62. Coiera E. Communication spaces. *J Am Med Inform Assoc* 2014;21(3):414-22.
 63. Eccles M, Mittman B. *Welcome to Implementation Science*. *Implement Sci* 2006;1(1):1.

Correspondence to:

Enrico Coiera
 Australian Institute of Health Innovation
 Macquarie University
 Australia
 E-mail: enrico.coiera@mq.edu.au