

Preface: On the Need for Resilience in Health Care

Erik Hollnagel, Jeffrey Braithwaite and Robert L. Wears

This book provides the first comprehensive description of *resilient health care*, henceforth referred to as RHC. Since there are probably only a few who, at the time of publication, have a clear idea about what this means, some words of introduction are appropriate. The simple explanation is that RHC is the application of the concepts and methods of resilience engineering to the field of health care, and in particular to the problems of patient safety. A more detailed explanation and exemplification are provided in the chapters that follow.

The motivation for RHC is threefold, as explained in this preface. The first is the sorry state of affairs in health care. The second is that attempts to improve this so far have had limited success. The third is the potential offered by resilience engineering as an alternative approach to safety and safety management.

The Sorry State of Affairs

The ‘sorry state of affairs’ of patient safety care is an expression of the fact that the general situation is not acceptable. Care is simply not as safe as it ought to be. To paraphrase Shakespeare, we may say that ‘there is something rotten in the state of health care.’ Justification for this view is not hard to find, but a few examples will suffice for the present. In a paper appropriately entitled ‘Is health care getting safer?’ Vincent et al. (2008) noted that 10 per cent of patients admitted to hospitals in the United Kingdom were subject to iatrogenic harm. This level of harm has been found wherever studies of care, via medial record review or other methods, have been conducted. And depending on how iatrogenia is measured, adverse events may occur in every third or fourth admission (Landrigan et al., 2010). In addition, a RAND study by McGlynn et al. (2003) showed poor adherence to many recommended clinical practices, with only 55 per cent of patients receiving care deemed to be appropriate. Almost 10 years after RAND, and in Australia rather than the United States, Runciman and colleagues (2012a; 2012b) found the proportion of patients receiving appropriate levels of care remained at a similar level, this time 57 per cent. There is clearly room for systemic improvement, to put it mildly.

This ‘sorry state of affairs’ has been recognised for some time, since several large studies in the USA, UK, and Australia as far back as 1955 provided ‘clear evidence that medical error is a common and sometimes preventable fact of the delivery of care in several highly developed and well-funded healthcare systems’ (Baker and Norton, 2001). We agree with this statement, except that we would substitute ‘often’ for ‘sometimes’. The clincher was the landmark report ‘To Err is Human’ from the Institute of Medicine in the USA (Kohn, Corrigan and Donaldson, 2000), which made clear that a majority of adverse events were due to systemic faults rather than to individual incompetence (cf. also Chapter 2).

That these circumstances have arisen and that they persist in sophisticated, well-resourced health systems can be attributed to a number of overarching factors. One is that the demand for care is rising because of population ageing. It is also rising because the provision of care is increasingly intense and complicated due largely to inter-linked factors such as technological, diagnostic and therapeutic advances. Another is that opportunities to provide the right care to the right patient at the right time are diminishing because of work pressures and associated demands on clinicians, exacerbated by workforce shortages and ageing workers. A third is the rising costs that are spiralling in many countries, perhaps uncontrollably so. According to estimates based on data from OECD Health Data and the World Health Organization, 2010 health care expenditure ranged from 6.28 per cent of GDP in Mexico to 17.6 per cent of GDP in the USA, while the OECD average was 9.5 per cent – with an estimated growth rate of around 4 per cent annually. This is a vast commitment of the world’s resources, especially in times of strained budgets and austerity – whether endogenous or exogenous.

The health system thus has to meet several irreconcilable goals – customer demands, performance pressures, work and workforce stresses, and cost challenges – and to meet them simultaneously. Not surprisingly, this creates stressed circumstances and working conditions (for the system and for the people in it) that not only are far from optimal but are in many respects detrimental to providers and patients. It is under these conditions of pressure, uncertainty, and risk that people are expected to deliver care which is both safe for patients, cost effective, and of high quality. The diminishing capacity to do so is an urgent problem that has triggered many studies and led to many attempts to solve it. Yet progress has been painfully slow – perhaps because we have tried to solve the problem based on the symptoms but without a proper diagnosis.

Crushed Expectations: The Bane of Conventional Solutions

Dissatisfaction with system performance levels is not unique to health care, but is common across industries – and indeed in practically all types of organised human activity. Other industries have throughout their development either experienced calamitous events, such as spectacular accidents and major disasters, or extended periods of intolerable problems, but have in most cases eventually managed to find workable solutions. When dissatisfaction with the performance of the health care system became common in the 1990s, the obvious reaction was therefore to look to other industries that appeared to do better, in the hope that simply imitating what they did would be a panacea – or at least keep the wolves from the door a little longer.

A pivotal point we want to make is that people who championed learning from other industries (aviation, chip manufacturing, etc.) lacked an adequate appreciation either of what these other industries were doing, or why things worked there, or both. Although few will openly admit that they believe in silver bullets, there was in the beginning the hope among many that it would be possible to find quick remedies that could be ‘rolled out’, which would make all the problems dissolve. (This is affectionately known as ‘picking the low-hanging fruit’.)

Dating from the 1970s, health care has, with varying degrees of optimism and conviction, tried putative solutions such as intensifying bureaucracy, quality circles, quality assurance, root cause analysis, ‘lean manufacturing’, standardised therapeutic programmes via clinical guidelines, teamwork, use of check-lists, accreditation, and above all, information technology (IT) in various forms. Most times solutions have been introduced in local settings or systems by convinced champions with much enthusiasm but with little or no thought about the principles and values underlying their efforts, or how the initiatives would fit together strategically or, indeed, how they would affect the existing equilibrium. An under-recognised common feature of such initiatives is that they practically represent a highly rationalised, Taylorist approach that presumes predictability and an inherent linearity and proportionality of causes and effects. This is regrettably nowhere to be found in the real world of care delivery. Health systems are not simplistic production lines for car assembly or mining operations for mineral extraction, and can therefore not be precisely described, specified, codified, mechanised, and controlled. Even staunch health care supporters have gradually realised that real progress will require abandoning the Taylorist

approach. Indeed, Berwick (2003) has indicated that: ‘... prevailing strategies rely largely on outmoded theories of control and standardization of work.’ It seems to be a cornerstone of the human condition that people believe – or want to believe – that they will be able to solve today’s problems, improve things, reduce errors, and ameliorate harm – all with just a few more resources, a bit more effort, another set of recommendations from a wise enquiry, a little more knowledge of the amount and rate of harm being delivered, increasingly precise measurements of system features, tightening up practices or a new whizz-bang IT system that is just around the corner.

It would, of course, be encouraging, even gratifying, if one could see that the accumulated experience had gradually led to a changing perspective – or even better, demonstrable systems improvements. However, there seems to be little system-wide improvement experience, perhaps because most solutions have been applied opportunistically and piecemeal, based on a trial-and-error ‘philosophy’ (Hollnagel, 1993). With opportunistic control, solutions are chosen because they look good but without much purposeful exploration. If they succeed, they have limited effect, leading to intermittent and localised gains. They can cause more harm than good, suffer unanticipated consequences and burn a lot of money in the process. When they fail, they are simply abandoned; and because they were not chosen for a clear reason, there would be little to learn from the failure. Indeed, project-itis and faddism are rife in health care.)

Solutions based on the use of IT have, since the turn of the century, tended to be looked at with unbridled – although generally unfounded – optimism. We can thus find statements like: ‘It is widely believed that, when designed and used appropriately, health IT can help create an ecosystem of safer care ...’ (Institute of Medicine, 2012). The recently published document ‘Health Information Technology Patient Safety Action & Surveillance Plan for Public Comment’, published by the Department of Health and Human Services in the USA, defines two health IT patient safety objectives, namely *Use health IT to make care safer*, and *Continuously improve the safety of health IT*. While few would disagree with these, the unspoken assumption is that health IT is a (or the?) solution to the problems, even though health IT is mostly a label. A closer look at actual experiences from other industries might be advised.

One of the lessons that could have been learned, had anyone bothered to make an effort, is that most presumed solutions – with IT being no exception – tend to make the health system more complex and less tractable. Regardless of whether a solution is added

to the system as something new or substitutes for something that is already present, it will affect what already goes on in ways that can be difficult to anticipate or even imagine. The assumption that solutions are neutral in their effects and that their introduction into a system therefore only has intended and no unintended consequences is palpably false. Several of the chapters in this book address this issue. The growing complexity shows itself in systems, in compliance costs, in quality and safety programs, in bureaucracy, in clinical technology, in IT, and the like. At the very least, we must make a concerted effort to tease out the nature of this growing complexity, and take steps to understand health care as a complex adaptive system (cf. Chapter 6).

Resilience Engineering

In the rush to find ‘implementation-amenable’, readily-packaged solutions to the undeniable raft of problems in health care, the focus turned to what other industries had done under similar conditions in the past. Little attention was paid to the fact that a number of these other industries, such as nuclear, aviation, and offshore activities, had gradually started to revise their own approach to safety, prompted by the realisation that the tried and trusted methods of the past, which were always limited anyway, were no longer adequate for the present, and would be even less so in the future. One important insight was that *adverse events increasingly needed to be explained as unfortunate combinations of a number of conditions*, rather than as failures of single functions or components – including ‘human error’. Another was that *failures should be seen as the flip side of successes*, in the sense that there was no need to evoke special failure mechanisms to explain the former. Both failures and successes have their origin in performance variability on individual and systemic levels. It is just as wrong to attribute successes to careful planning and diligence as it is to attribute failures to incompetence or error. Instead, both owe their occurrence to a mostly unpredictable, but not unimaginable, combination of a number of system characteristics.

The term ‘resilience engineering’ was put forward to represent this new way of thinking about safety (Hollnagel, Woods and Leveson, 2006), and quickly became recognised as a valuable complement to the established approaches to industrial safety – and soon also to patient safety. Enlightened thinkers in both industry and academia began to appreciate that resilience engineering provided an articulated basis for confronting the puzzles of such phenomena as complexity, interconnectedness, ‘system of systems’, and

ultra-high reliability. The concepts and principles of resilience engineering have since the beginning been continuously refined by applications in such fields as air traffic management, nuclear power generation, offshore production, commercial fishing, and accident investigation. Over time there has been a dawning realisation that resilience is neither limited to handling threats and disturbances, nor confined to situations where something can go wrong. Today, resilience is understood more broadly as the intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions. This definition emphasises the ability to continue functioning, rather than simply to react and recover from disturbances, as well as the ability to exploit opportunities that arise, rather than simply survive threats.

Toward Resilient Health Care

In line with this way of thinking, resilient health care can be defined as the ability of the health care system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required performance under both expected and unexpected conditions. In order to strive for health care resilience, it is therefore necessary to study and understand how health systems work, and not just to perpetuate the predominating myopic focus on how they fail. This realisation has led to the juxtaposition of two views on safety, Safety-I and Safety-II, which permeates this book. (The two views are presented in detail in Chapter 1.) Briefly stated, Safety-I is defined by the (relative) absence of adverse events (accidents, incidents). Safety-I is reactive, and assumes that safety can be achieved by first finding, and then eliminating, or weakening, the causes of adverse events. As a contrast to this way of thinking, Safety-II is defined as the ability to succeed under varying conditions, so that the number of intended and acceptable outcomes (in other words, everyday activities) is as high as possible. Where Safety-I focuses on what goes wrong, Safety-II focuses on what goes right, and the purpose of safety management is to achieve and maintain that ability. The importance of making this distinction and its practical consequences are amply illustrated by the chapters that follow.

In agreement with this distinction, and in order to examine various facets of RHC, the book is structured in three sections. Section I, entitled ‘Health Care as a Multiple Stakeholder, Multiple Systems Enterprise’, articulates the scope and depth of RHC and the players involved with it. This section highlights the complexities, adaptive capacity and self-

organisational features of health care. It provides a context for knowledge and perspectives about the nature of the health system and how it can be conceptualised, particularly emphasising its complexity, patient safety and the quality of care, in the light of the Safety-I, Safety-II dichotomy. Section II, ‘The Locus of Resilience – Individuals, Groups, Systems’, sharpens the focus, providing a microscope to display the detailed conditions under which resilience manifests itself at differing levels and in characteristic ways. Building on the earlier chapters, Section III, ‘The Nature and Practice of Resilient Health Care’, then scrutinises ways of ‘being’ resilient, and ‘doing’ resilience, in everyday activities. Finally, in the Epilogue, the book synthesises key learning points and implications for various stakeholder groups interested in health care resilience.