Thought paper May 2012

Steven Spear

Reinventing healthcare delivery

In this thought paper, Professor Steven Spear argues that the way to achieve great operational performance and value has often been misunderstood. In order to ensure that good people and good science are facilitated, rather than overwhelmed, by systems, leaders have to expand their attention from 'what individuals do' to 'how the pieces come together'. High quality care and great performance involves leaders making problem solving, improvement, and innovation part of the regular routine of daily practice.

At the Health Foundation, the drive to find new ways of doing things that will improve the delivery of healthcare is central to our work. Our improvement programmes develop, test and spread new ideas. Our research also strives to identify and share best practice. For a number of years, we have been running leadership development programmes, and working with healthcare leaders on the front line, to strengthen and improve leadership within healthcare.

Health Foundation thought papers are the author's own views. We would like to thank Professor Spear for his work, which we hope will stimulate ideas, reflection and discussion.

About the author

Steven Spear
Professor Steven Spear is author of *The High Velocity Edge*, which won the Philip Crosby Medal. He is a senior lecturer in the MIT Sloan School of Management, a senior fellow at the Institute for Healthcare Improvement, and is a founder of See to Solve Corporation, a business process software company. Active in research and practice, he received *Harvard Business Review*'s McKinsey Award, for 'Fixing Healthcare from the Inside', and has been awarded the Shingo Prize five times. Professor Spear works with industrial, high tech, software, service, healthcare and military organisations to improve performance by accelerating problem solving and innovation in complex operating environments.

Summary

How the efforts of myriad individuals are harnessed towards a common purpose has a profound effect on meaningful measures of performance, including quality, cost, turnaround time, workplace safety, product and service security, and responsiveness to changes in market needs. These differences are observable in high-tech and heavy industry, in services and manufacturing, and across the spectrum of development, design, and delivery. They are even observable in healthcare.

However, the way that superiority in performance is achieved is often misunderstood. The conventional wisdom is that leadership is won by applying superior tools, techniques, and technologies ('technologies' for short), thereby implying that the job of an organisation's leaders is to make the correct decisions about what technologies to purchase and use. In fact, leadership is won through speed of discovery of new things and how to apply what is already understood. Therefore, an organisation's leaders must cultivate the behaviours and practices by which new insights and ideas can be developed and applied.

The differences in rewards that accrue to one approach (decisions about technology) versus the other (behaviour leading to discovery) have been historically strong. They will increase as the systems for which we are responsible and on which we depend become more complex in their structure and faster reconfiguring in their dynamics. Those rewards can be captured – measured in terms of cost, capacity, and quality – by leaders engaging their workforces in relentlessly identifying and addressing problems that impede the delivery of care and solving those problems as part of their everyday work. The penalties for not doing so will be equally great.

Introduction

Healthcare delivery presents terrible contradictions. On the one hand, exceptionally well-trained and well-motivated individuals, with the best science and technology at their disposal, endeavour to alleviate the pain and suffering of others. On the other hand, their collective efforts are marked by high cost, low quality, and relentless frustration for patients and staff alike. Fortunately, there is promise of a better situation: exceptional value, delivered far more broadly, with far less investment of time and effort.

This optimism is well grounded. First, if healthcare cannot break out of the current conundrum of high cost and low quality, it would be the unusual, if not the unique, industrial sector that is not characterised by year-in, year-out increases in functionality, reliability, and responsiveness, coupled with regular decreases in per unit costs. The experience of ever more value for ever less cost is characteristic of most other advanced sectors – aviation, high-tech manufacturing, services, software and other forms of engineering, etc. In fact, healthcare aside, it is nearly impossible to find examples of industries caught in the spiral of escalating costs and compromised value,

making it all the more difficult to imagine the compelling, intrinsic arguments as to why healthcare should resist subscribing to exactly the same trends.

Furthermore, even if one could construct an ambitiously creative argument as to why healthcare should be the outlier industry, that hypothetical of high cost being coupled inevitably with low quality is contradicted by inarguable data; a handful of healthcare providers have achieved breakout performance: great care, delivered broadly, at reduced cost relative to their counterparts. Complications like central line-associated blood stream infections? Proven avoidable. Patient falls? Avoidable. Surgical site infections? Avoidable. Readmissions? Avoidable. Mis-medication? Avoidable too. The list goes on, illustrating that the 'death spiral' can be broken when appropriately mature and sophisticated approaches for managing the delivery of care are adopted.

Competitiveness advantage through operations

That exceptional performance can be won through operations – how the work of myriad individuals is integrated towards a common purpose – is not in doubt. There is overwhelming evidence that certain organisations generate far more value with far less effort, while consuming fewer resources and in less time than their counterparts. For example, people took notice of the huge disparities in performance in the automobile industry – in particular during the 1980s and 1990s – noting that companies like Toyota

and Honda were able to sell more reliable cars, more affordably than their rivals, particularly the 'Big Three' of General Motors, Ford, and Chrysler. Originally a 'paradox' of doing much more with much less on the production side,[1] it was evident that there was also a paradox of doing much more with much less on the design side.[2] Less widely discussed than the performance gaps in the auto industry, but no less pronounced, are differences between device density and processor speed of state-of-the-art microchips produced by Intel and those made by other semiconductor manufacturers.[3] Alcoa achieved levels of workplace safety two orders of magnitude better than other US manufacturers,[4] and there are similar orders of magnitude differences in workplace safety, patient safety and delivery cost when comparisons are made for healthcare providers.[5]

However, how these feats were achieved is subject to some misunderstanding.

Conventional wisdom: decisions about technology

The conventional wisdom is that differences in operational performance are attributable to differences in tools, technologies, and techniques. Have the right robot, just-in-time pull system, or software, and great results will follow. Without those assets, greatness cannot be accomplished. We see just such a bias manifest itself in the operational space, with the lean manufacturing movement, for example, exhorting its followers to 'eliminate waste' by the programmatic creation of continuous flow, paced by just-in-time, stabilised

1 Krafcik 1988; Womack *et al* 1990.

2 Ward *et al* 1995.

3 Based on in person interview.

4 Spear 2010 (chapter 4).

5 Spear 2005.

by standardised work performed within production cells that are well organised, with 5S discipline.[6]

This conventional wisdom likely exists and persists for several reasons. One is the way differences in performance have been explored and explained; another is how the role of management more generally is taught.

Differences in performance are typically identified and investigated in a cross-sectional fashion – comparing similar organisations by performance and approach at similar points in time. Such was surely the case in the automotive industry, the seed research of which involved comparing performance across all the world's 186 final assembly plants, leading to the realisation that a small number were able to generate far more product with fewer people, in less space, using less equipment, and requiring less material on hand than was typical. The term 'lean production' grew out of the observation that these few did so much more with so much less as a contrast to the 'mass production' of the rest of the industry.[7]

The follow-up on that realisation of differences in performance led to observations about differences in the operating environments between the 'lean' and the 'mass' producers. The explanation that emerged was that the lean producers were using techniques for managing production that were different from those employed by mass producers.

They used value stream maps to create continuous flows rather than a batch production approach of job shops; they controlled material flow with pull

Figure 1
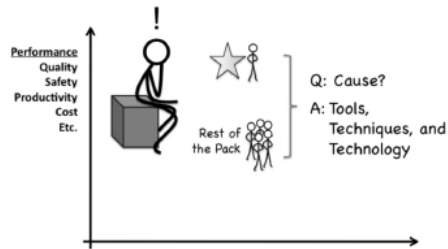
systems rather than relying on schedules from material requirement planning (MRP) systems; and they had far more work standardisation than mass producers did.[8]

Such a diagnosis – that superiority comes to those making the right decisions about using the right 'technology' (or techniques) – certainly fell on fertile soil, as it was wholly consistent with the premise of modern management education: that the job of an organisation's leader is to make decisions, and the better the decisions, the better the organisation's prospects. Consider, for example, the core courses of any Master of Business Administration (MBA) curriculum: finance to quantify value on decisions, and accounting to track them. Strategy has historically been taught as a decision making challenge – the right combination of differentiations to change an organisation's position from one of weakness to one of strength. Even operations management has had a heavy dose of decision making emphasis: economic order quantity, lot sizing, queuing theory, constrained optimisation and the like.[9]

6 See, for example, Womack and Jones 1996.

7 Krafcik 1988.

8 Womack and Jones 1996.

9 See, for instance, *Factory Physics* by Hopp and Spearman.

Reality: technology is not a source of sustained differentiation

For all the emphasis on informed decision making, the compelling evidence is that leadership is not gained in a step-like fashion (lack a critical technology and lose; decide to acquire a critical technology and win). In fact, quite the contrary; there is ample evidence, dating back over decades, that (a) there are enormous disparities in performance among those who use the same tools and technologies; and (b) these disparities in performance accrue to those who learn the most, the fastest, for the longest stretches of time.

As early as the mid-1980s, there was compelling evidence of gross differences in performance using exactly the same tools – for instance, in how flexible manufacturing systems were performing in US and Japanese firms. The Japanese flexible systems used less equipment, were developed in half the time, but nevertheless supported 50% more production volume spread over 900% greater variety. Uptime was greater, employee productivity was far better, and time to market with new products was a fraction of the US experience.[10]

More recently, research showed large disparities in the impact of healthcare providers using electronic medical records. Some recorded gains as predicted by advocates: greater data accuracy, benefits in ease and cost of storing and transmitting data, accuracy in billing, etc. However, others demonstrated limited gain relative to the investment made.[11]

In the knowledge-intensive clinical setting – similar cases being treated with similar science and technology by people with similar training – there is only a loose relationship between cost of care and quality, with the highest-quality providers often the lowest in cost, and the most expensive providers often offering the most risky service.[12]

In short, buying the technology or using the technique does not generate the anticipated gain. It is how the technology is used that is critical.

Reality: learning behaviour is a source of sustained differentiation

Additional evidence suggests that the 'how' to use equipment to good effect is not due to some serendipitous realisation. Rather, great performance is achieved by repeatedly accumulating insights, improvements, and innovations, and putting them to good use. This leads to the conclusion that *the sources of competitive advantage are not technological, but behavioral.*

In the healthcare sector, for instance, cardiology teams, attempting to master a new, minimally invasive approach to performing surgery, had wide disparities in the rate at which they achieved mastery (if they achieved it at all). The rate depended on how well the team was fostering high-velocity learning.[13]

There is a certain irony that Toyota – whose success motivated major investment in 'lean' tools and techniques – has much of its success attributable to the relentlessness with which it generated new discoveries. Though the company garnered much attention for its market-leading performance in the 1980s and 1990s, in fact,

10 Jaikumar 1986.

11 Karsh *et al* 2010. Emphasising the behavioural side of success with new technology, IT in particular, Karsh *et al* write: 'Foundational cognitive and human factors engineering research and development are essential to better inform HIT development, deployment, and use.'

12 Institute of Medicine, 2001.

13 Edmondson *et al* 2001.

Toyota was not consistently a strong competitor. Its first entry into the US market, in 1958, was a flop. The Toyopet had the unusual distinction of getting to the top of hills with greater certainty in reverse gear rather than forward gear – a particular problem as the car was marketed in hilly California. But it wasn't just on quality that Toyota grossly under-performed; it was lousy on productivity and (consequently) on cost. Productivity of labour and capital (not just one or the other) was one-eighth that of the American automakers. However, by 1962, Toyota was equally productive, and by 1968, doubly so.[14]

When Toyota re-entered the US market in response to the 1970s oil price spikes, it was both an affordability leader (thanks to the steady gains in productivity it had generated) and a reliability leader – a far cry from the Toyopet experience. From that positioning of 'affordable reliability', Toyota broadened its product portfolio under the Toyota brand, introduced new brands (Lexus and Scion), and new technology (the Prius in 1996) – all while making itself increasingly self-sufficient for design, supply, and production in its major markets. During that period, its profitability dwarfed that of its competitors and its market capitalisation equalled the total of the other 'majors' combined. Which begs the question, *How did Toyota get so much better, so much faster, and for a longer period than its rivals?*

The answer is that Toyota did a far superior job of managing work so that doing work and learning to do work much better were exceptionally well coupled in contrast to other firms, which were far less capable of generating and incorporating new insights, regardless of whether they took a hierarchical, command and control posture[15] or a (faux) empowerment approach.[16] More specifically, Toyota had learned how to build validation into all of its work – that is, whether what was planned was actually occurring. When the reality contradicted expectations and predictions (*with very high granularity, very high frequency, and nearly immediate feedback*),[17] it triggered rapid swarming and solving of problems as a precursor for sharing/incorporating what was being regularly learned.[18] This was happening in production and in design.[19]

In the auto industry, Toyota was hardly alone in achieving exceptional performance through the quality of its learning processes. MacDuffie (1997) compared problem solving at three auto plants dealing with exactly the same issue – water leaks around window seals. In one, the problem was largely ignored – left for rework at best, and for the customer to deal with at worst. In the second, the problem was addressed, but only after data were collected, aggregated, sorted, and analysed, far from where and when the problem had actually occurred. Not surprisingly, the ratio of resolution to effort made was rather poor, as subtle information about the problem's presentations and causes 'spoiled' before the problem could be adequately investigated and addressed. Only in the third plant, a Honda one, were individual problems swarmed to be solved as soon as they were seen.[20]

14 See Cusumano 1989; Lieberman and Dhawan 2005.

15 Adler 1993.

16 Adler and Cole 1993.

17 Spear and Bowen 1999; Spear 2002; Spear 2010.

18 Spear and Bowen 1999; Spear 2010.

19 Ward *et al* 1995; Durward *et al* 1999.

20 MacDuffie 1997.

The auto industry was not alone in rewarding those select organisations that adopted the behaviours of managing work to see problems, solving them when seen, and aggressively sharing what was learned by incorporating local discoveries systemically and systematically.

Figure 2
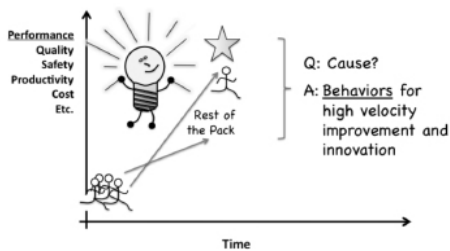
In the mid-1980s, Alcoa offered a relatively safe work environment compared with other US employers, considering the inherent threat and danger of the heavy industrial processes it used. Nevertheless, a 2% annual rate of workplace injury meant that it was more likely than not that an employee would be hurt in the course of their career, and nearly certain that they would know someone injured – given the communal nature of Alcoa hiring – likely a friend, neighbour, or family member. Over some 15 years, Alcoa went from being better than average to remarkably exceptional, driving its injury risk down by 97% to 0.07% per year. During this time, Alcoa's corporate numbers accounted for the fact that the company was in the midst of acquiring (and counting the results of) mills and other facilities in former Communist countries, hardly paradigmatic in their concerns for environment, health, and safety. Yet, in the

same period, the risk in US manufacturing declined by half, from 4% to 2%.

Alcoa achieved this remarkable performance relative to its counterparts even though they all had access to exactly the same tools, techniques, and technologies that all other commodity manufacturers used, and had access to the same tools, techniques, and technologies the rest could have used to drive down workplace risk. The difference was behavioural, rooted in the admonition that safety problems – injuries certainly, but also near misses – be called out when and where they occurred, and, once seen, be swarmed and solved before contextual information about their cause would perish. Then, once problem causes and solutions were discovered, they were incorporated as broadly as possible. Though a very different workplace setting than Toyota, the behavioural combination of see problems, solve them when seen, and share/incorporate what has been discovered was the key to Alcoa's success. Significantly, near-perfect workplace safety was not achieved at the expense of other 'hard-nosed' metrics like cost, yield, quality, and delivery timeliness. Rather, the close attention to problems, the energetic resolution of them, and the disciplined systematisation of newfound knowledge gave Alcoa sophistication with process management that led to across the board improvement.[21]

There are other examples of this same storyline in markedly different settings. For example, the US Navy developed nuclear propulsion for warships faster than the Soviets and then deployed it without any

21 For a detailed case about workplace safety at Alcoa, see Spear 2010 (chapter 4).

incident for nearly 70 years, whereas the Soviet/Russian experience has been closer to bi-annual catastrophes.[22] There are also the clinical settings able to learn better and faster, with greater results than typical in their fields.[23]

The long and short of it is: operational perational excellence – in development, design, and delivery – is a potential source of competitive advantage. Having certain tools and technologies may be a necessary ante to compete, but there are such wide disparities in their effectiveness that having them alone is not sufficient. Rather, it is how they are used that matters, and the sophistication with which they are used depends on an organisation's capacity to generate and sustain high-velocity, broad-based, non-stop discoveries. Tools are needed, but without the right learning behaviours, they are not sufficient.

Rising complexity demands even better discovery behaviour

There is a simple reason why superior performance is the reward for superior rates of broad-based, sustained learning. Anything that is complex, and designed by people, can possibly be perfect. It is only through iterative problem solving that we can discover the right answer,[24] so the faster the iterations and the more impact each has, the faster we can discover what needs to be known. And, even if a perfect system is designed or discovered, the perfect state is only temporary, as conditions – both external to the organisation and internally – continually change, meaning that even an optimised system must rapidly adjust as conditions change.

In adjusting, new problems arise:

- *Output*: are the right objectives set (and are they being met)?
- *Architecture*: are the right people and tools doing the right work, at the right time?
- *Interfaces*: are information and materials being exchanged with the right content, format, timing, and setting?
- *Component methods*: are the individual elements functioning as they should?[25]

As demanding as these questions were historically, they have only grown in significance. The idea that we have to 'discover our way to the right answer' rather than 'think our way to the right answer' has gained more currency as the systems on which organisations depend to deliver value to market have become more complex: a more complex architecture, more and more challenging interfaces, and tighter tolerances on both the system's overall functionality and that of its individual components.

22 See 'High velocity under the sea, in the air, and on the web', Spear 2010 (chapter 5), citing personal interviews, NASA benchmarking study, Rockwell Theodore and others.

23 See, for example, Edmondson *et al* 2001; Shannon et al 2006.
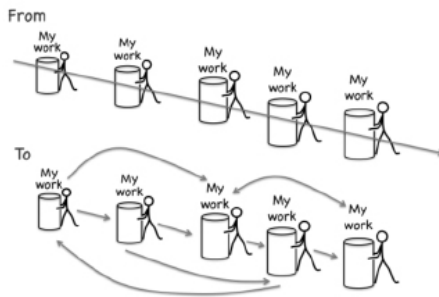
24 Nelson and Winter 1982.

25 See Spear 2010 for elaboration on this system design framework.

This increased complexity is expressed in at least three ways:
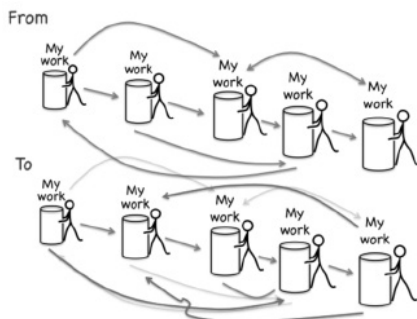
1. Increasing the number of actors/nodes

2. Intertwined rather than linear patterns of interdependency

3. Accelerating rates of required reconfiguration

These trends of increasing complexity are broad, not isolated.

Healthcare services have experienced (and continue to experience) this explosion in system complexity in each of the three dimensions described: number of actors involved, pattern of interdependency, and rate of required reconfiguration.

It used to be that medical care was 'simple' on all three dimensions:
1. Number of nodes: the doctor was the main actor, supported by staff.
2. Patterns of interaction were 'orbital', with a few staff lending support to the doctor, but with minimal interaction with each other.
3. The rate of system reconfiguration was minimal, with the same pattern of behaviour applied to each patient.

In a second phase, marked by an increase in specialisation and division of labour from more highly trained doctors to less highly trained professional staff:
1. The number of nodes increased, as doctors' work was parsed to other professions (e.g., Registered Nurse, MA, technician).
2. The patterns of interaction changed from orbital to sequential/linear.
3. Some (minimal) reconfiguration was required around what work, in what order, and by whom, depending on patient type.

The current, third phase is far more complex:
1. There is a vast increase in the number of professionals involved in providing care.

2. The patterns of interaction and interdependency have become highly intertwined and non-linear.
3. The system has to reconfigure diagnosis and treatment patient by patient.

Consider an example. Colleagues and I were asked to do a risk assessment at a major academic medical centre. The invitation's motivation was simple; the hospital had done severe harm to several patients – mis-medication, wrong-side surgery, and the like. They wanted to know how they compared with best practice so as to know how to improve their approaches.

The gap between the critique they thought we would deliver and the one we actually did is rooted in this transition from simple to complex and stable to dynamic – for both the external operating environment and internal operating systems. They probably expected that we would walk through their facility, compare what we saw to some benchmarked best practice seen elsewhere, and critique their techniques (and the professionals responsible for implementing those techniques) vis-à-vis testing, diagnosis, therapy, or recovery.

One can certainly understand their expectation. We examined their cancer centre. Thirty years ago, when their senior leaders were trained, medical science and technology was far less advanced than it is today. There were a limited number of tests that could be done, a limited number of diagnoses that could be generated, few therapies available (a small number of chemotherapies and a relatively small number of surgeries), and hence, a limited number of recovery procedures that had to be followed. In that environment, the imperative for achieving greatness (relative to the state-of-the-art at the time) was to make sure that the individuals responsible for each step were up to date on the latest approaches and well practised in them. If not, problems would arise based on individual shortcomings. Those problems were addressed by additional training and practice (or dismissal in egregious cases of shortcoming).

Occasionally, of course, there were problems that had their root cause in the 'architecture' of the system (someone was doing the wrong task, at the wrong time or in the wrong order) or the 'connective tissue' of the organisation (critical information was poorly conveyed).

Those problems were so infrequent and typically of such low consequence that their diagnosis and resolution was left to informal processes – over the water cooler, at the end of the day, in the course of a serendipitous hallway meeting.

Figure 3: Cancer care – then…

Whatever they expected as a risk assessment based on how they had been trained, the reality was very different. At each stage of treatment, there has been a huge explosion in the available science and technology,

with a consequent explosion in the number of professional skills and capabilities required to employ them.

It wasn't just the increase in numbers of system 'elements' that mattered. The 'order' of the system – the number of interactions that had to be considered at any given time – had exploded too.

Figure 4: …and now
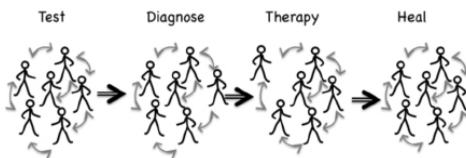
For example, what was once 'radiology' has become 'imaging', with X-rays complemented by CT scans (layered, multiple perspective X-rays), magnetic resonance imaging (MRI), sonograms, and so forth. Not only are there new skills required to operate the equipment and interpret the images, there is skill in making sense of what the images show by integrating the data generated from the various techniques. Similarly, therapy might once have been limited to chemotherapy or limited surgical alternatives. Now, there are many more drugs, which can be combined in tailored 'cocktail' treatment plans depending on the patient's condition. This deepens the required knowledge of oncologists, pharmacists, and nurses. But now, the type of chemotherapy regimen that is given depends on and affects the radiation oncology and surgical options.

The increased complexity of the system at the medical centre was not only reflected in the number of healthcare professionals and the 'relationships' that had to be managed, but in the types of failures that occurred. The greater the number of relationships and interdependencies, the more 'connective tissue' there is to be managed. As a result, each individual might do a good job, but the system as a whole can fail.

Consider one adverse event – fortunately one with minimal impact on the patient's wellbeing. A radiation oncology patient presented with what appeared to be severe sunburn on her forearm. Diagnosis of the system showed that the treatment plan developed by the radiation oncologist (an MD) was appropriate, the equipment was rightly calibrated (an important issue in its own right, given the overexposures experienced by patients at a New York City hospital when software updates were loaded without adequate notification to operators) and correctly operated, and the 'flight path' physicist (a PhD) had correctly calculated deflections and absorptions of the therapeutic beam as it travelled through the patient's body.

The problem, it turned out, was that the patient – too frail to hold her arms crossed on her chest – was told by a well-meaning technician to put her arms by her side on the treatment bed. He didn't know that the beam would enter her body precisely where her arm was. The flight path physicist hadn't known about this change in the patient's position, so hadn't accommodated with a change in flight path. The good news was that the harm done was minimal. The 'sunburn' was not terribly severe, and the problem was caught early enough and

re-irradiation done quickly enough so as not to upset the treatment plan. However, this was but one example where 'slips' were not the shortcomings of an individual in his or her role, but problems in the 'architecture' of the system – a misfire in terms of who was doing what, and when – or the connective tissue informational linkages of the complex system in which the work of individuals was embedded.

Of course, it is not just an increase in complexity that made work challenging in this modern cancer centre. It was the 'need for speed and agility' that raised the stakes. With the multiplication in available therapeutic combinations and huge advances in understanding the variety of disease states that exist, treatment plans can be better tailored and targeted, patient by patient by patient. So the care system is not only more complex, it has to adaptively reconfigure too – ideally, patient by patient. It would be one thing if the system were complex but, once found, the structure provided a predictable, steady state of challenge. However, with the need to reconfigure and the absence of a steady state, not only does the number of problems to be solved increase, but so too does the pace at which they have to be addressed. And, as with the case of the 'sunburned' patient, *as the system gets more complex, more problems – and their potential solutions – come from the shaping of the informational space in which people work, and not only from the physical environment or how they perform their individual professional tasks.*

Establishing discovery-based healthcare delivery

It is taken as given that superior healthcare results from superior research and superior training within professions. This is undoubtedly the case. However, the delivery of care is now so dependent on integrating the contributions of many individual professionals towards one common purpose (a purpose that is repeatedly re-targeted) that bad systems can overwhelm good people and good science any day of the week.

In order to ensure that good people and good science are, instead, facilitated and enabled by good systems, leaders have to expand their attention from 'what individuals do' to 'how the pieces come together'. This means building the daily routines of seeing problems that impede effectiveness, solving them so they don't recur, and sharing what's been discovered locally for systemic impact.

There is no magic potion or 'silver bullet' solution to close the gap between what is achieved today and what is theoretically possible. However, much of the hard work and discipline that is required involves making problem solving, improvement, and innovation part of the regular routine of daily practice. This may be the only way in which healthcare becomes like virtually all other sectors, delivering ever greater value at ever lower cost, day in and day out.

To share your thoughts about this paper, please visit **www.health.org.uk/SpearTP.** *You can also follow the Health Foundation on Twitter at* **www.twitter.com/HealthFdn**

References

Adler P. Time and motion regained. *Harvard Business Review*, January–February 1993.

Adler P, Cole R. Designed for learning: a tale of two auto plants. *Sloan Management Review* spring 1993;87.

Cusumano, Michael *The Japanese Automobile Industry: Technology and Management at Nissan and Toyota*, Cambridge, MA: Harvard University Press; 1989.

Durward K, Sobek II, Ward P, Liker J. Toyota's principles of set-based concurrent engineering. *Sloan Management Review* 15 January 1999.

Edmondson A, Bohmer R, Pisano G. Disrupted routines: team learning and new technology implementation in hospitals. *Adm Sci* Q December 2001.

Hopp W, Spearman M. *Factory physics*. Illinois: Waveland Press; 2011.

Institute of Medicine, *Crossing the Quality Chasm*, Washington DC: National Academy Press, 2001.

Jaikumar R. Postindustrial manufacturing. *Harvard Business Review* November 1986:64(6):69–76.

Karsh B-T, Weinger MB, Abbott PA, Wears RL. Health information technology: fallacies and sober realities. *J Am Med Inform Assoc* 2010;17:617e623.

Krafcik JF. Triumph of the lean production system. *Sloan Management Review* 1988;30(1):41–52.

Lieberman MB, Dhawan R. Assessing the resource base of Japanese and US auto producers: a stochastic frontier production function approach. *Management Science* 2005;51(7):1060–75.

MacDuffie JP. The road to "root cause": shop-floor problem-solving at three auto assembly plants. *Management Science* 1997;43(4):479–502.

Nelson R, Winter S. *An evolutionary theory of economic change*. Cambridge MA: Belknap Press; 1982.

Shannon R, Frndak D, Grunden N, Lloyd JC, Herbert C, Patel B et al. Using real-time problem solving to eliminate central line infections. *Jt Comm J Qual Patient Saf* 2006;32(9):479–87.

Spear SJ, Kent Bowen H. Decoding the DNA of the Toyota production system. *Harvard Business Review* September 1999.

Spear SJ. The essence of just-in-time: embedding diagnostic tests in work-systems to achieve operational excellence. *Production Planning & Control* 2002;13(8):754–67.

Spear SJ. Fixing healthcare from the inside, today. *Harvard Business Review* September 2005.

Spear SJ. *The high-velocity edge: how market leaders leverage operational excellence to beat the competition*. McGraw Hill; 2010.

Ward A, Liker JK, Cristiano JJ, Sobek II. The second Toyota paradox: how delaying decisions can make better cars faster. *Sloan Management Review* 1995;36(3):43–61.

Womack J, Jones D, Roos D. *The machine that changed the world*. New York: Simon & Schuster; 1990.

Womack J, Jones D. *Lean thinking: banish waste and create wealth in your corporation*. Second edition. New York: Free Press; 1996 and 2003.

The Health Foundation is an independent
charity working to continuously improve
the quality of healthcare in the UK.

We want the UK to have a healthcare
system of the highest possible quality
– safe, effective, person centred, timely,
efficient and equitable.

We believe that in order to achieve this,
health services need to continually improve
the way they work. We are here to inspire
and create the space for people to make
lasting improvements to health services.

Working at every level of the system,
we aim to develop the technical skills,
leadership, capacity and knowledge,
and build the will for change, to secure
lasting improvements to healthcare.