Using Dynamic Work Design to Help Cure Cancer (and other diseases)

July 2017

Don Kieffer
MIT Sloan School of Management
don@shiftgear.work

Nelson Repenning
MIT Sloan School of Management
Milwaukee Consulting Group
nelson@mit.edu

Sheila Dodge
sdodge@broadinstitute.org
Broad Institute

Timothy De Smet
James Meldrim
Niall Lennon
Danielle Perrin
Steve Ferriera
Zachary Leber
Dennis Friedrich
Stacey Gabriel
Eric S. Lander
Broad Institute

MIT MANAGEMENT
SLOAN SCHOOL

ABSTRACT

Few words in the English language strike more fear than "cancer." Despite substantial progress in both its treatment and management, cancer remains a leading cause of death in the developed world, and the elusive search for a cure continues to attract both top scientific talent and vast sums of federal and private research dollars. Work design focuses on configuring productive human activity and is normally applied to settings like factories and information processing where the work is highly repetitive. As such, work design would seem just about as far from the cancer research lab as you could get. In this paper, we report the results of an extended effort to improve the productivity of the genomic sequence operation at the Broad Institute using an emerging framework that we call Dynamic Work Design. Improving the design of the Broad's operation and better matching that configuration to the humans doing that work has led to dramatic improvements in both the speed and cost of genomic sequencing. The gains in both the cost and cycle time of sequencing are enabling researchers focused on a variety of disease mechanisms (including cancer) to both run more experiments and get the results back sooner, thereby significantly speeding their progress towards better therapies and, ultimately, a cure. We outline three work design interventions that contributed to the Broad's success and then discuss the broader implications for designing work.

INTRODUCTION

Few words strike more fear in people's hearts than cancer. And, with good reason: Cancer is the second leading cause of death in the US (after heart disease) and is responsible for almost 600,000 deaths annually.[1] Not surprisingly, understanding the causes of cancer is an active area of research. The National Cancer Institute (a division of the National Institutes of Health) alone spends almost $5 billion annually. Among the most promising of these research avenues is the ongoing effort to understand cancer's genetic origins. Determining the genes disrupted by cancer offers the possibility of both identifying at-risk individuals sooner and developing targeted therapies that might offer a better benefit/harm ratio than current approaches.

In contrast to cancer, the words "work design" or "process design" probably elicit almost no response at all. If you studied management at some point in your career, it might conjure a vague memory of a long-forgotten class or case study in which you studied a factory making widgets or a logistics operation moving those widgets from one location to another. Alternatively, you may be picturing your firm's resident industrial engineer or "black belt" who tries to make the processes of widget making and widget moving cheaper and more productive. In either case, work or process design, while perhaps necessary for repetitive tasks, seems about as far away as you could get from the critical work of understanding the biological basis of cancer; after all, the cure for cancer will be found in the lab or in the clinician's office, not in a factory.

While there may have been some validity to this view a decade or more ago, an emerging framework known as Dynamic Work Design now plays an increasingly central role in understanding the genetic basis of a variety of diseases, ranging from cancer to Type II diabetes to schizophrenia to the Ebola virus. In the last five years, the Broad Institute, a leader in genomic research, has made dramatic reductions in cycle time and cost using dynamic work design principles to reconfigure the process of DNA sequencing. These gains allow genetic researchers to both run more experiments and get the results of those experiments back more quickly. Redesigning the Broad's work has played a critical role in identifying the roots of the Ebola virus and is quite literally speeding the search for a cure to cancer and many other diseases. Critically, in contrast to the conventional wisdom, the substantial gains in speed and efficiency have not come at the expense of adaptability and innovation. To the contrary, the Broad's efforts to redesign its work have increased its ability to keep pace with extremely rapid change in the sequencing technology and evolving scientific needs of the research community.

The changes at the Broad mirror large trends in the nature of innovation. Whereas we often cling to the image of the lone scientist working in her lab, as the world grows increasingly connected and computation becomes ever cheaper and faster, further gains in science and technology will increasingly come from collaboration among multiple disciplines—genetic research draws on biology, chemistry, mechanical and electrical engineering, computer science, and applied mathematics. Innovation thus becomes an *organizational* problem: How do we coordinate the activities of multiple individuals working from multiple disciplines to solve the critical problems facing our customers and society? Dynamic Work Design offers a set of principles for creating the structures necessary to facilitate the next generation of innovations.

[1] http://www.cdc.gov/nchs/fastats/leading-causes-of-death.htm.

In what follows, we begin with a brief description of DNA sequencing and the role the Broad Institute has played in developing and using that technology. Following that, we describe the challenge they faced in satisfying the growing demand on the part of genetic researchers for sequencing and how their existing organizational structure did not support those targets. We then outline the basic principles underlying the design of effective dynamic work and show how they were used to transform the Broad's operations. Finally, we discuss the implications of this experience for work that extends beyond genomic sequencing.

Genome Sequencing and the Broad Institute

DNA is quite literally the "source code" of life, and each of the more than 1 trillion cells in our body contains a full set. Only a few percent of that DNA differs across individuals, and those differences are encoded in a string of more than 3 billion DNA bases that is called a genome. Some stretches of these bases are called genes which, when transcribed and translated by the cell, encode the proteins required for the body to function. Elucidating the sequence of DNA in an individual's genome and studying the base changes (also known as variants) between individuals can identify the genes responsible for a wide range of diseases, including Type II diabetes, schizophrenia, early-onset heart attack, and many types of cancer. In some cases, these sequence variants help identify the most effective treatment; in others, they suggest which metabolic pathway might be a potential target for drug screening.

The first draft of the human genome was completed in 2001 and required more than ten years of effort and $2.7 billion.[2] Since then, the technology has evolved rapidly; the cost has dropped over 100,000 fold, and a sequence can now be completed in a matter of weeks. The continued decline in the cost of genome sequencing is critical to ongoing efforts to find a cure for cancer and several other equally devastating diseases. By allowing researchers, who typically have a fixed amount of funding, to increase the number of samples they analyze, cheaper sequencing enables the identification of increasingly rare genetic anomalies, ones that might only be seen in 1 out of 1,000 patients.

The mission of the Broad Institute is to transform medicine through systematically understanding the genetic underpinnings of disease. Following the completion of the first draft of the human genome sequence, the Broad Institute was founded to formalize the relationship that emerged among MIT, Harvard, and local hospitals. The unique organizational experiment catalyzed rapid progress in genomics, and the Broad continues to play a central role in unraveling the mysteries held within each cell, including its circuitry, genetic makeup, and response to perturbations. Completing a genetic sequence requires hundreds of physical and intellectual manipulations. It typically starts with a blood or tissue sample that then needs to be prepared before its DNA sequence is "read" by sequencing machines. Over the past 15 years, the Broad has built a significant operations infrastructure and can now process over 1 million patient samples annually.

The Broad's Sequencing Operations in 2012

In 2012 the Broad reached a crossroads with its sequencing operations. The operations function struggled to keep up with the growing demand for their services, and rapid change in the underlying technology made it increasingly difficult to focus on anything for an extended period of time. Congestion in the samples "pipeline" that fed the sequencing machines was growing. Lengthy turnaround times often meant that the team had to reprioritize work in process to meet specific customer requests. The average time required to analyze a sample had grown to 120 days,

[2] International Human Genome Consortium, Lander et al. (2001).

and researchers, who required timely analysis of their samples, began to complain and/or send their samples to other labs. By the middle of 2012, though industry demand for sequencing was growing rapidly, demand for the Broad's services was *declining*; the lab appeared to be incapable of handling the increasing volume of samples. By housing all facets of genetic research, from basic science to the sequencing operations, the Broad had been able to advance genomic science on several fronts. But, if its operations were not cost competitive, they would be difficult to fund and a key part of their "integrated" solution would be lost. By 2012, there were internal discussions concerning their ongoing viability.

The operational challenges were further exacerbated by the organizational structure. Operations consisted of three functions: samples acquisition, sequencing, and microarrays. The first two, samples acquisition and sequencing, represented key steps in the overall process of analyzing a sample. Microarrays was an increasingly outdated and redundant technology, but it still played a key role in quality control and early discovery for the other two groups. The sequencing operation was the largest and most functional and had already implemented many well-known improvement techniques (Six Sigma, 5S, Kanban); however, its utilization was severely limited by the other two functions. Though sequencing is capital intensive, each of the 50 sequencing machines costs about $1 million, those machines were operating at roughly 50 percent utilization.

The samples group was increasingly descending into chaos. Samples that needed processing were distributed throughout the lab, and there was no means of organizing them. When a technician wanted to start preparing a sample for sequencing, the first thing she had to do was find it, which could often take a week or more. Despite working increasingly longer hours, the technicians in the lab were falling ever further behind. Due to the lack of communication and reliance on middlemen to relay information, team managers felt disassociated from the final product and had little visibility into the operation's projects and goals. Not surprisingly, morale was low and arguments erupted daily as team leaders tried to figure out why yet another sample was about to miss its promised delivery date.

The Broad's challenges were similar to the growing pains that many organizations experience when their technologies move from the research to the commercial domains. Sequencing the first genome was a herculean task, occupying researchers around the world for more than a decade and costing almost $3 billion. Today, sequencing takes a few weeks, is done tens of thousands of times a year, and costs a few thousand dollars. Moreover, sequencing is no longer just an area of research; it now plays a critical role supporting other research programs and is increasingly used to treat individual patients. Much like microchips and innumerable other technologies, what started as a piece of research in the lab now plays a central role in contemporary life. The Broad's existing approach to organizing its work was thus rendered obsolete by its own success. The distributed, research-focused organization that helped sequence the first genome was increasingly ill suited to analyzing thousands of samples a month in a timely fashion.

In the summer of 2012, the Broad re-organized its operations. The three separate areas—samples, sequencing, and microarrays—were combined into a single organization, the Genomics Platform, reporting to a single operations manager (the first author). The new operations manager was charged with both improving communication and eliminating artificial disciplinary and functional boundaries. The organization chart was simplified and aligned around delivering timely analyses to the scientists whom the Broad supported. But, whereas people often equate changing the organization chart with designing effective work, this was only the beginning.

Dynamic Work Design

Process, Organization, and Work Design

Historically, scholars and managers have approached the challenge of constructing effective organizations either by starting at the top or the bottom. *Process design*, largely the domain of industrial engineers, typically starts on the shop floor or in the research and development lab and focuses on the mechanics of designing multi-step production processes. Process designers answer questions like: Which machines should we buy; in what sequence should we operate them; and where should we put them on the factory floor? The work of process design typically manifests as a process map, showing the major steps required to transform raw materials into a finished product or new ideas into a working design.[3]

At the other end of the spectrum, organization design focuses on the big "building blocks." Should we organize by function, business, or "matrix" ourselves? Who should report to whom? Who has the right to start projects, to kill projects, and, above all, to spend money? Organization design typically manifests in the organization chart, the familiar boxes and solid and dotted lines that show who reports to whom, and who gets to weigh in, even if they are not in charge.[4]

A critical limitation of both approaches is that neither takes particular account of the fallible human beings who will be doing the jobs that a given design specifies. While process designers do attend to ergonomic issues, they focus largely on human physical limitations and give little attention to the cognitive or emotional dimensions. A process designer might design work that is safe, but rarely considers whether that work fully capitalizes on the strengths and offsets the limitations of those who do it. Similarly, those focused on organization design typically presume both that people will execute their role in the organization chart in the way that the designer intended and that work will typically go as planned.

Unfortunately, people often do not find their work particularly engaging and rarely execute it in the way that the designer had hoped. Worse, operators and technicians, who are either bored with tedious, repetitive work or frustrated with a work design that does not account for all the problems they encounter, are often remarkably adept at finding ways to "work around" the process that the designer worked so hard to configure.[5] Such changes can include holding secret stashes of inventory, skipping steps, and making physical modifications to the production process. Similarly, on the macro side, even the best organization design often falls victim to local self-interest and the associated politics.[6]

The academic literature offers a complement to process and organization design known as "work design.[7]" Work design departs from its process and organization-focused brethren by focusing on creating work that matches the skills and capabilities of those who do it. Rather than presuming an "ideal" employee, work design takes the human condition as given and attempts to create work that is consistent with our inherent strengths and limitations. This line of research identifies three features that comprise engaging work. First, to be engaging a particular task needs to be clearly connected to an important outcome, and the person doing it needs to be able to use a variety of skills in

[3] See, for example, Stevenson (2015) for manufacturing and Ulrich and Eppinger (2011) for new product development.

[4] See Combe (2014) and Ancona et al. (2005) for a general overview and Galbraith (2002) and Kates and Galbraith (2007) for more detail.

[5] Repenning and Sterman (2001, 2002).

[6] Pfeffer (1992).

[7] Contemporary work design theory was first articulated by Hackman and Oldham (1979). See Grant et al. (2010) for summaries of recent updates. See Deci (1996) for a summary of research on human motivation.

executing that task. Second, humans find a task more engaging when they exert a measure of control over how it is executed. And third, we draw more satisfaction from tasks where we get regular and rapid feedback concerning our performance. To bring this point home, Hackman and Oldham use golf as the archetypal task that is likely to engage the human brain: Golf requires a variety of skills (driving, chipping, putting, etc.); the player is in total control of the outcome (only I determine whether I hit a good shot or a bad shot); and golf provides immediate feedback on performance (I know immediately whether I hit a good shot or a bad one).

Creative work, such as doing basic science, satisfies many of the work design criteria—it's focused on something important, the investigator is in control of her own actions, and you can get regular feedback on whether or not you are making progress. When the work design approach was first developed, several firms tried to make similar changes to more repetitive processes. For example, one automaker experimented with having teams focus on assembling a single car rather than working on an assembly line. Unfortunately, though operators may have found this work more engaging, the performance of this system was dominated by the efficiency of the more traditional assembly line.[8] Thus, in contrast to the skill variety and autonomy experienced in golf and other sports, contemporary work outside of basic science and research and development is often highly repetitive and offers few opportunities for exercising control or developing new skills (though as highlighted above technicians and operators often surreptitiously create these opportunities). Similarly, because the work is subdivided, those doing it rarely get immediate and accurate feedback concerning the contribution of their work to the organization's targets.

The existing approaches to process, work, and organization design thus leave today's leader with an apparent dilemma. On the one hand, the basic principles that have guided work since the beginning of the Industrial Revolution can be used to subdivide and routinize tasks, thereby creating work that is efficient. But, this approach also creates a system that both disenfranchises those doing the work and makes adapting to new innovations difficult. In contrast, managers can create work that is highly engaging (more like golf) and thus a workforce that is highly innovative, but this flexibility comes at the expense of the profitable efficiency that accompanies standardization and repetition.

In 2012, the Broad was on the cusp of moving from one regime to the other. As a research organization, its work had been loosely configured and its staff exercised considerable control over their day-to-day activities. The consequent flexibility and engagement produced dramatic gains in the science of genomics. But, due to their success in advancing the technology, they now faced an increasing demand for efficient, fast, and reliable sequencing that was at odds with their existing configuration; they needed to become more like a factory and less like a research lab. And, the reorganization discussed above (a classic piece of organization design), was intended to do just that: move genomic sequencing from a freewheeling collection of autonomous groups to a single focused entity that could deliver efficiently and reliably. Not surprisingly, many of the existing employees were concerned; they had joined the Broad because they wanted to work on the cutting edge of an important technology, not to become a cog in a tightly controlled machine.

The notion that organizations face a trade-off between innovation and engagement and grim efficiency has been at the heart of management theory and practice for decades and is a central tenet of organization design.[9] But, is it really true? The business literature certainly contains numerous examples of once successful and highly efficient firms that failed to adapt to a "disruptive" technology. In contrast, however, other large firms seem to be able, at least to a degree, to "have their cake and eat it too" in the sense that they are both efficient and innovative. Perhaps most

[8] Womack, Jones, and Roos (1990: 101-102) and Adler and Cole (1993).
[9] This argument was first advanced by Abernathy (1978) and then March (1991). See Adler et al. (2009) for a recent review.

notably, Toyota has remained highly successful in its core business of producing internal combustion engines while simultaneously pioneering hybrid drive and fuel cell vehicles, and was recently listed by *Fast Company* as one of the 25 most innovative companies in 2015.[10] Our ongoing research and intervention work reveals an alternative which, while perhaps not entirely eliminating the trade-off between efficiency and engagement, certainly weakens it considerably. We call our emerging approach Dynamic Work Design.

The Foundations of Dynamic Work Design

Dynamic work design builds on a simple but profound idea: Things don't always go as planned, and how an organization reacts to those inevitable hiccups plays a critical role in determining its long-term effectiveness. We have found it useful to think about a work design as having two components, one that is static and one that is dynamic. As a simple metaphor, imagine designing an intersection of multiple roads. The engineer doing the work is first faced with the task of choosing the type of intersection, should it be a simple four-way crossing, a roundabout, or an overpass with or without off and on ramps? These elements comprise the static design. But, as every driver knows, while necessary, these elements are far from sufficient; we also need "rules of the road." What happens, for example, when two people want to use the two crossing roads simultaneously, or somebody is trying to enter a rotary while someone else is trying to leave? These problems are all addressed by more dynamic elements of the design. Stop signs and stoplights help prevent collisions from two cars trying to cross simultaneously. Similarly, rules of the road dutifully learned in driver's education help resolve conflicts like who has the right-of-way in a rotary or who gets to go first when you want to make a left turn across traffic. The dynamic portion of the design tells us what to do when, for whatever reason, the static one is insufficient.

In a similar vein, organizations have both static and dynamic designs. The static design is typically captured in the aforementioned process diagrams, project plans, and organization charts and details how things are *supposed* to work. The dynamic design describes the way an organization responds when things don't go as planned. But whereas the designers of the traffic system have given a lot of attention to both elements, organization and process designers focus almost exclusively on the static components and, with a few important exceptions, give scant attention to designing the "rules of the road." Critically, the lack of attention *does not* mean that there is no dynamic design. To the contrary, all organizations have ways of responding to hiccups and conflicts, but if they have not been given explicit attention, then there is no guarantee that those rules will produce a good outcome.

A growing body of research suggests that many of these implicit "rules" can be deeply pathological. For example, almost every manager has been asked at one point or another to cut her budget mid-year because the organization is not meeting its annual earnings targets. Many (if not most) managers have an implicit rule of responding to such challenges by cutting investments with longer pay-off horizons (such as maintenance, training, and basic research) so as not to sacrifice current revenue or customers. But, while initially successful, such an approach can create a vicious cycle of declining capability and increasing short-term pressure. This basic dynamic has been implicated in a variety of organizational dysfunctions ranging from major product recalls such as the recent GM and Volkswagen challenges to industrial accidents like the explosion at BP's Texas City refinery.[11]

Dynamic Work Design offers managers a simple guide to creating a more effective set of "rules of the road" for their organization and builds on two foundational elements. First, following the contributions of the work design school, Dynamic Work Design recognizes that humans come with both strengths and limitations. For example, under certain

[10] See http://www.fastcompany.com/section/most-innovative-companies-2015.
[11] See Rahmandad, Henderson, and Repenning (2015).

conditions we are remarkably effective learners, but in other situations we are quite error-prone. Similarly, we find some tasks (like golf) deeply engaging, but others (like filing expense reports) painfully tedious. Moreover, research suggests that when we are engaged, we are likely to both work harder and be more creative.[12] Thus, the first foundation of dynamic work design is simply that work needs to "fit" the humans who do it, meaning that it needs to be designed in ways that match our cognitive and emotional processes. Research in cognitive and neuroscience has exploded in the last two decades and now provides an increasingly solid foundation on which to build work that is engaging, satisfying, and productive.[13] You wouldn't sit in a chair that was too small or ride a bike that was too big, so why should you work in an environment that is poorly matched to the way your brain works? Our work to date suggests that designing work that both capitalizes on our strengths and offsets our limitations offers a path to designing work that is both more satisfying and more productive.

Second, although they may not have recognized it, several of the major contributors to management theory tackled elements of the dynamic design problem.[14] Frederick Winslow Taylor's "one best way" created the foundation for recognizing when a design wasn't delivering the desired output. The founders of the quality movement, principally Deming and Shewhart, recognized the potential for learning that comes with those gaps and with the famed PDCA cycle (plan-do-check-act) outlined a way to realize that potential. Taiichi Ohno, through the Toyota production system, created a system to take control of that learning and make it happen on a regular basis. But each of these insights was both located in the specific context of manufacturing and was neither fully integrated with the other contributions nor the underpinning science. Dynamic work design unites these contributions into a set of general and cohesive principles that allows these critical insights to be extended beyond the original settings in which they were created.

Principles of Dynamic Work Design

The business community has become obsessed with the search for best practice. Academics and consultants spend hours trying to identify those tools and activities that separate high-performing organizations from the rest, and managers make a similar investment in trying to implement them. Unfortunately, efforts to implement someone else's best practice rarely produce significant gains.[15] There are many reasons for such failures, but one major contributor is the failure to distinguish between *practices* and *principles*. The contemporary organization is often a complex, multifaceted place that represents a melding of many different cultures and norms. The Broad, for example, is a US-based organization with close ties to two major educational institutions (Harvard and MIT) and several large hospitals. The operations team is young, largely US born, and the majority of its members have advanced degrees in biology, mathematics, or computer science. Given that it represents an amalgamation of national, local, occupational, and professional cultures, it is not surprising that practices developed on the factory floor in Japan could not be immediately transferred to the Broad. The key to transforming the Broad's operations was to identify the behavioral principles that underlie well-designed work—whether it be at Toyota or a local restaurant—and then to use those principles to create practices that were well adapted to the Broad's unique context. Consequently, before detailing the specific practices that proved so effective at the Broad, we briefly outline the basic principles of designing effective, dynamic work.

[12] See Meyer and Allen (1997) and Amabile (1997, 1996).

[13] There are several popular accounts describing the growing knowledge of human decision-making. Kahneman (2011) and Haidt (2006) are excellent and readable summaries.

[14] Taylor (1911), Ford (1926), Deming (1982), and Ohno (1988).

[15] See Pfeffer and Sutton (2000).

Principle #1. The first principle underlying effective work is the *constant reconciliation of activity and intent*. Frederick Winslow Taylor, the father of process design (often known as "Taylorism"), made a key contribution to work design by recognizing the value of systematically studying work and clearly articulating the state of current knowledge about how to do a given task, what he called the "one best way," and a clear set of targets to be achieved by those activities. Unfortunately, Taylor also believed strongly in separating those in the organization who think from those in the organization who do. Only managers were, in Taylor's view, smart enough to understand the "why" of work; Taylor once wrote that if a man chose iron handling as his profession then "…he was too phlegmatic and stupid" to understand the design of that work.[16] Consequently, the responsibility of identifying the one best way fell entirely to managers, and those doing the work were to be given very specific instructions in how to execute their specific tasks with little if any explanation as to how that work might fit into the larger system or why it might be important. Not surprisingly, Taylor's ideas were (and continue to be) widely unpopular with organized labor. Nonetheless, even today, Taylor's legacy remains strong as most work is constructed in ways that give those who do that work very little idea for why they are doing it.

Despite its popularity, research suggests that the practice of separating the intent of an activity from its actual execution is highly counterproductive because it neither motivates people nor leverages their inherent capabilities and experience. Humans draw considerable motivation from connecting to something larger than themselves (note the popularity of connecting sports and hobbies to charitable causes), but separating activity from intent makes this difficult. Being fully engaged in a job typically requires understanding *why* you are doing it.[17] In addition, standard processes (the modern version of the one best way), no matter how carefully designed, can never cover all the variations faced in day-to-day work. When those doing the work understand why they are doing it, they are in a better position to make the appropriate adjustments. Thus, the first feature of an effective work system is simply that everybody knows *why* they are doing what they are doing, and there is an ongoing method for everyone, from the leader down to the technician in the lab, to compare the activities that they are doing with the intended outcome. Put more simply, in dynamic work everybody knows when the static design (the one best way) has not produced the desired outcome.

The more regularly and rapidly this comparison happens, the more useful it will be. Humans are remarkably adept at identifying patterns in their environment and learning from that experience. It is exactly this pattern matching that enables feats of athletic prowess like hitting a fastball in baseball or returning a serve in tennis. Moreover, we draw considerable motivation from identifying new patterns in our environment and capitalizing on them (this is why sports like golf can be so engaging—we are always trying to improve—and why gambling remains popular). However, our cognitive pattern-matching apparatus only works effectively when we get rapid feedback on our efforts. When there is a long delay between taking an action and getting an outcome, we are both less likely to learn to do that task better and find doing that task less satisfying. Imagine how your experience would change if when you went to practice golf or tennis, instead of seeing the result immediately, you would get a report three months later detailing where each shot had landed. In well-designed, dynamic work, those doing a task get rapid feedback showing them whether or not that activity met the intent. Work is more dynamic in this setting because, with rapid feedback, people will quite naturally learn to do that task more effectively and will draw considerable motivation from the process of improving.

[16] Taylor (1911), p.28.
[17] See Deci (1996), chapter 5.

Principle #2. The second principle defining an effective dynamic design is the *regular use of structured problem solving.* Gaps between activity and intent are inevitable and humans come equipped with a natural drive to improve (remember sports again). If, following the first principle, work is designed to clearly reveal those gaps, then those doing the work are likely to become engaged in solving them. However, while our natural pattern-matching instincts are very effective in some situations, we are not always effective learners, particularly as the complexity of the system grows. Most notably, a variety of psychological biases reinforce existing ways of thinking and prevent the identification of novel solutions. For example, we are prone to a phenomenon known as confirmation bias, which simply means that once we believe something, we selectively focus on the data consistent with those beliefs and conveniently ignore that which is contrary.

One of the central contributions of the quality movement (Shewhart and Deming) was to recognize that a simple application of the scientific method can complement our natural developmental instincts and greatly enhance the ability of those doing the work to find innovative ways to close the gap between activity and intent. By forcing people to articulate the problem they are trying to solve, to be systematic in collecting data, to formulate clear hypotheses and carefully test them, structured problem solving allows people to take a metaphorical "step back" from their day-to-day work and, thereby, leverage their accumulated experience far more effectively. Since Shewhart and Deming's work, several variants of structured problem solving have become popular, including PDCA (plan, do, check, act) and DMAIC (define, measure, analyze, improve, control).[18] In our experience, which specific method you use is far less important than making sure *some* kind of structured method is used to address any significant gap between activity and intent.

Principle #3. The third principle underlying dynamic work is the use of *optimal challenge*. Research suggests that humans have a strong developmental drive—we like to improve—and we respond well to targets that we perceive to be just beyond our current capability.[19] If people can see gaps between activity and intent (Principle #1) and respond effectively (Principle #2), they will both be more productive and more engaged. However, if the gap between the results produced by the current activity level and the intent (i.e., the targets) grows too large, what was once a functional response can quickly become counterproductive.[20] As people start to feel that they are too far behind, they can become more stressed, less attentive to details, more prone to making mistakes, and less creative. Case studies suggest that overstretched organizations produce fewer innovations, have higher rates of employee turnover, and experience more accidents and product recalls. Worse, as discussed above, if the pressure continues to mount, people begin to take shortcuts and sacrifice capability-building and maintenance activities that are critical to the long-run health of the organization. Not surprisingly, overly aggressive targets have been implicated in several major industrial accidents, including both space shuttle disasters and the explosion at BP's Texas City refinery.

Creating an engaging and effective work environment requires constantly managing the gap between activity and intent. If the gap is too small, people become complacent and work becomes boring (though we have rarely observed this situation in knowledge work). If the gap is too large, people begin to take shortcuts, and quality and innovation begin to suffer. Carefully regulating the degree of challenge is central to creating an environment that is both innovative and efficient. Taiichi Ohno, one of the fathers of the Toyota production system, understood the value of optimal challenge in manufacturing and assembly work. To this day, Toyota uses the speed of the assembly line to regulate the level of challenge. If the line is moving too slowly, there are never any gaps between activity and intent. If the line is

[18] See Deming (1982) and http://asq.org/learn-about-quality/six-sigma/overview/dmaic.html.
[19] See Deci (1996) and Csikszentmihalyi (2008). Pink (2011) provides a recent popular summary.
[20] See Rudolph and Repenning (2002) and Repenning, Goncalves, and Black (2001).

moving too quickly, then much like the famous *I Love Lucy* episode in which Lucy works in a chocolate factory,[21] the operation quickly descends into chaos and firefighting. Toyota managers will continually adjust that speed to get the optimal amount of challenge, enough gap to create opportunities for learning, but not so much that people never have the time to exercise those opportunities.

Principle #4: The fourth principle is to *connect the human chain*. Golf and other sports are engaging in part because we are in total control of the outcome and it is easy to reconcile and activity and intent—I know immediately whether I have hit a good shot or a bad one, and nobody hit that shot but me. Similarly, the first three principles can be applied to any solitary activity. Moving from the practice tee (or field) to the contemporary organization complicates the design problem considerably since work is usually a team sport; most work begins with something received from somebody else (specifications, information, a partially completed idea) and ends when the output of that effort is transferred to somebody else. In this case, the person or people doing the task neither control the final outcome nor necessarily receive rapid feedback on the quality of their contribution. Worse, the "hand-off" from one person to the next introduces multiple opportunities for misunderstanding and miscommunication; medical errors, for example, often originate when shift changes are not properly executed.

Fortunately, evolution appears to have equipped us with the ability to at least partially compensate for this additional complexity. Under the right circumstances, humans can be adept and effective communicators, and direct human interaction can often resolve ambiguity far faster than more technology-mediated (e.g., email) or more automated approaches (compare your experience with automated "help" lines versus talking to an actual person).[22] Moreover, recent research shows that humans are strongly "pro-social," meaning we draw considerable energy and satisfaction from interacting with and helping others. For example, a study by Adam Grant showed that call center fundraisers raised twice as much money when they had a better understanding of those whom they would be helping. Other research shows that people will often work harder when they are doing something as a favor than when they are being paid a small amount.[23] Thus, though we will rarely experience the autonomy and direct connection at work that we experience with sports and hobbies, work inside complex organizations can still be satisfying and engaging due to the human connection.

Unfortunately, organizations often fail to capitalize on the benefits that can accrue from direct human connection. Taylor's system removed much of the human connection in work, and it's not clear that the situation has improved much since then. Consider, for example, how your organization uses email, staff meetings, and hallway conversations. In our experience, staff meetings are typically used for dry and extensive updates (essentially information sharing), decisions are made afterwards in the hallway, and most problem solving is done via endless, repetitive chains of email. This is a highly counterproductive design: One-way "update" meetings provide few opportunities to genuinely help (you are just supposed to listen); hallway decisions by a subset of the meeting attendees both disenfranchise those who are left out and leave them wondering why they were supposed to come to the meeting, thereby further undermining their sense of purpose; and email, because it lacks nuance and context, can be an extremely inefficient method for ambiguity resolution and problem solving. At the Broad, simple operational problems could often generate dozens of emails as the different functions worked to understand one another.

Connecting the human chain typically requires two changes to the way work is configured. First, processes (chains of activity) should be constructed so that each individual knows which *person* (not function or department) supplies the inputs she receives and which *person* gets the output she produces. Connecting people is often central to rein-

[21] See https://youtu.be/8NPzLBSBzPI.
[22] See Calefato et al. (2012) for a summary.
[23] See Grant (2014) and Ariely (2010).

troducing the human element and helps capitalize on our natural pro-social motives. Though we know of no direct experimental study, in our experience knowing *who* receives the output of one's work and why she needs that work (remember activity and intent) often re-energizes people and significantly improves both productivity and engagement. Knowing that Chris needs to receive my work to accomplish hers creates a very different level of motivation than "do it because finance said so." In a similar vein, compare your own experience of customer service when you call a designated representative that you know by name with calling an anonymous help line.

Second, the face-to-face interactions required to resolve ambiguity efficiently need to be explicitly designed into the human chain. Organizations are composed of standard ways of doing things (standard processes)—I submit my budget this way; I present a new strategy that way. But, standard approaches, no matter how detailed, cannot cover every contingency faced in doing real work. Even on assembly lines where each task is precisely specified, those doing the work are constantly making adjustments in response to variations in the materials and the way other work might have been done on a particular day. On assembly lines and in restaurants, because of the close proximity, these interactions often happen naturally—"make sure that bolt is tight" or "go easy on the seasoning for that order," but in many other types of work, these interactions need to be designed into the process. In some cases, creating such interaction requires nothing more than giving somebody a name and a phone number—if there is a problem call Chris. However, we find that short, structured, face-to-face meetings are often the most effective. At the Broad, a 15-minute, stand-up meeting focused on solving a problem could often replace dozens or even hundreds of emails. Connecting the human chain means clearly identifying the points in the work flow that benefit from face-to-face collaboration and then carefully designing those interactions to meet that need.

Dynamic Work Design in Genomic Sequencing

The four principles of dynamic work design—constant reconciliation of activity and intent, structured problem solving, optimal challenge, and connecting the human chain—are abstract and their application to specific work systems is not immediately obvious. This remains an active area of ongoing research. Nonetheless, our work to date suggests that applying these principles can both produce significant gains in quality, throughput, and engagement and can foster a climate that supports ongoing innovation. Moreover, many of the practices associated with different work-focused techniques (e.g., Six Sigma, Lean) can now be understood as specific manifestations of these principles. To show how these principles can be applied, we describe three interventions focused on improving the Broad's sequencing operations.

Exome Express: Creating cancer vaccines to impact patient care

The exome is the portion of the human genome that encodes proteins. It represents less than one percent of the total genome, but mutations in the exome are thought to be responsible for a significant fraction of diseases with a genetic origin. Given its disproportionate impact, researchers often prefer to analyze just the exome, thus getting the data they need with far less sequencing effort. Researchers at the Broad invented the method now widely used to generate exome sequences and have produced more than 160,000 such analyses to date. However, as the popularity of exome sequencing grew, the Broad's historical focus on quality and process control increasingly came at the cost of turnaround time. In 2012, the time that elapsed between when a sample entered the exome process and when the associated data was returned to the researcher had grown to over 120 days. Meanwhile, as progress was being made in understanding specific disease pathways, researchers required increasingly rapid feedback on their samples and, not surprisingly, long turnaround times became the number-one complaint. Worse, as genomics moved from a research technology to a therapeutic one, long turnaround times not only slowed basic research, but also had the potential to impact the treatment of individual patients. The situation had grown so acute that the Broad leadership was considering building a separate lab, requiring millions of dollars of capital, that would be dedicated to processing samples coming from clinical research. Something needed to change.

Convoluted and cumbersome processes are a common consequence of losing the connection between activity and intent. New steps are often added to the process to solve a local problem or to support a particular group's or function's objectives, but those additions are not reconciled against the overall intent, resulting in a process that is inefficient and frustrating. In many cases at the Broad, the original designer was no longer with the organization and thus the rationale for a given step was sometimes difficult to determine; it was just "how things were done." The exome process had grown unwieldy and those working within the process had lost a sense of connection between the work they were doing and its overall purpose. The first intervention was thus targeted at helping participants reconnect to the importance of what they were doing (to establish the "why" of what they were doing) and then reconciling their daily activities with that intent.

To begin, the leadership team contacted several leading clinical researchers and asked them to tell their stories about how DNA sequencing could impact patient lives. The stories provided a human connection to the work that many of the team members had not previously experienced. Following this exercise, the team worked on selecting three words that described their identity as an operations organization—eventually choosing, Pioneering, Premium, and Responsive—and then set several annual goals underlying each word (see Figure 1).

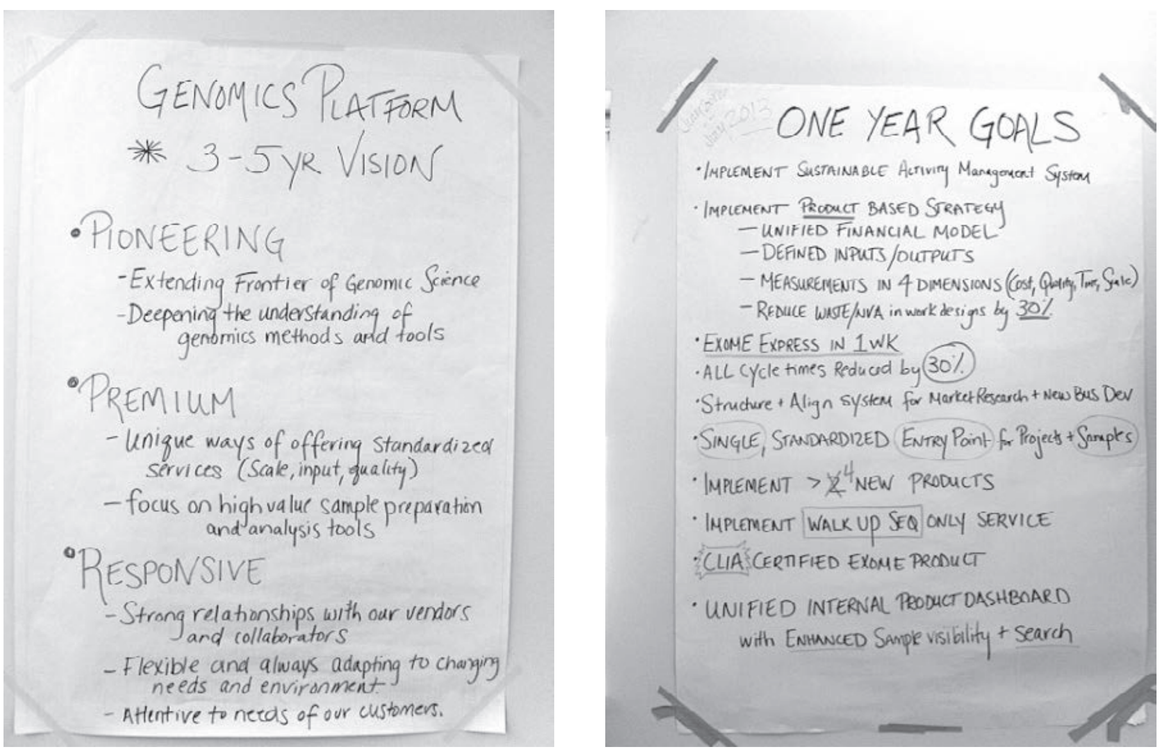

FIGURE 1: GENOMICS PLATFORM VISION AND GOALS

In attempting to reconcile their existing activities with their new articulated intent, the team quickly realized that the largest impediment to their desire to be viewed as "Premium" was an exome process that was slow and unreliable. With this gap in mind, the team began to re-evaluate the activities that comprised the exome process.

At the outset, the exome sequencing process comprised 48 process steps (see Figure 2). To facilitate this re-evaluation, the team used a simple process map. Mapping can be used to analyze processes from a variety of perspectives. In this case, it provided a means to systematically evaluate whether each activity in the process contributed to the intent of delivering a sequenced exome to the customer in a timely fashion. This flavor of mapping is similar to one of the fundamental practices in lean manufacturing, the value stream map, whereby each step in a process is scrutinized to determine whether or not it adds value for the final customer. In the case of the exome process, the process mapping exercise yielded two key benefits.

First, the map showed that many of the steps in the current process were either redundant or simply not needed. Second, the mapping exercise provided the first opportunity for the technicians to connect with the person who was processing the samples before or after them. Connecting the "human chain" of people who actually did the work almost immediately dissolved the boundaries between the samples and sequencing labs. Whereas under the old system Sarah in sequencing might complain about slow turnaround in the samples lab, now Sarah knew that Mike was handling her samples; going directly to Mike and asking where her samples were was far more effective than complaining to her manager.

FIGURE 2: INITIAL EXOME SEQUENCING PROCESS

The exercise also highlighted the need for more regular feedback to evaluate and improve the exome process. The turnaround time metric was redefined to focus on the customer's perspective, changing from a focus on local areas and specific steps to the overall process. The new measurement scheme helped technicians connect their individual activities to the overall intent of the process and catalyzed several local improvements. Further, measuring turnaround time more regularly allowed the leadership team to focus their energy on the most troubled parts of the process.

With the process mapping analysis, a more connected human chain, and a new metric in hand, the team proposed a new exome process composed of 17 process steps (see Figure 3). Theoretically, these steps could be completed within just a few days, but such a target did not account for mistakes, rework, or issues with equipment, etc. To make sure that those working in the process were not overwhelmed (and thus maintain something resembling an optimal challenge), the leadership set an initial target of completing samples within 21 days. Twenty-one days represented a big improvement over 120 but still provided some slack for the team to learn new routines and solve the inevitable problems associated with the new design.

Within a month of the initial mapping exercise, the team implemented "Exome Express," a new process that was immediately able to deliver a complete analysis within 21 days. Since the initial intervention, the cycle time had been further reduced to an average of 12 days. The team believes that they can soon get it to three. "Exome Express" has become a popular service and is now the foundation of the Broad's clinical platform, regularly returning results to physicians for patient care and clinical trials within 15 business days of receipt. The process generates over $1 million of revenue annually. Since the initial intervention, the team has made ongoing efforts to ensure that those working within the process remain connected to its overall intent and receive regular and rapid feedback on their performance. As an example, the Exome Express process is being used as part of an effort to create a cancer vaccine that is currently in Stage I clinical trials. The team regularly stays in contact with requesting physicians to see how their work is impacting the patients in their studies.

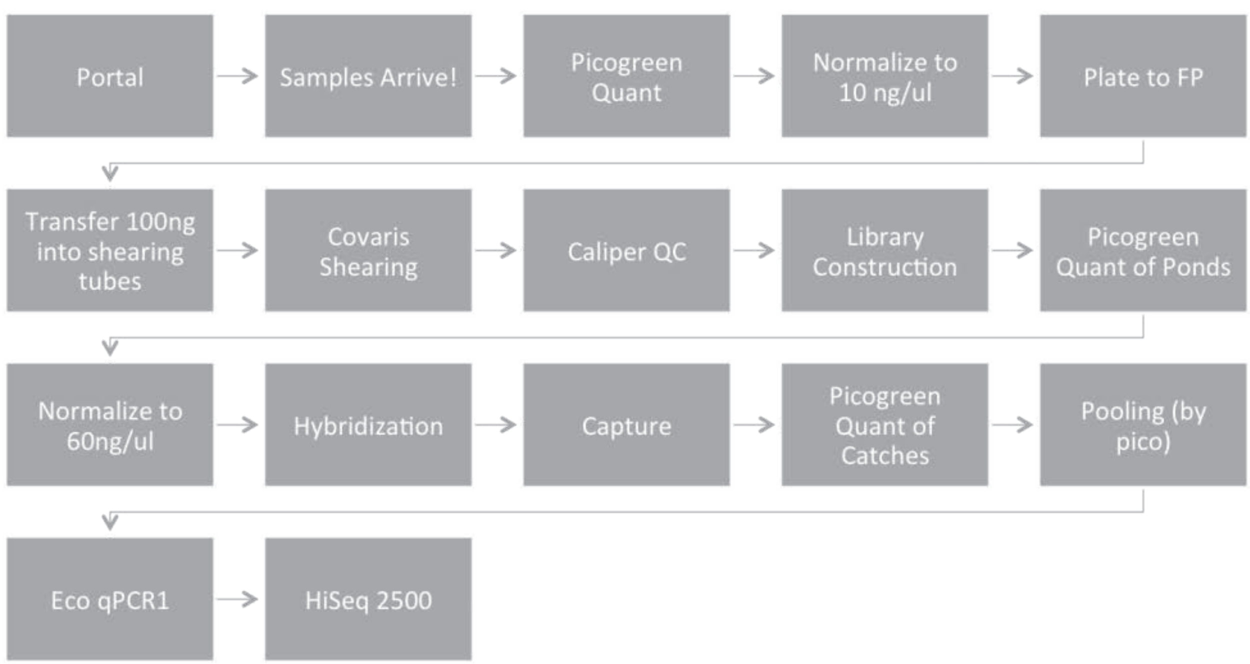

FIGURE 3. IMPROVED EXOME SEQUENCING PROCESS

Reducing Chaos = Faster Science

Having made a significant improvement in the exome process, the team turned its attention to the broader sequencing operations. Dealing with a diverse collection of biological samples and serving the needs of an active research community created the need for a complex set of offerings beyond the exome. While the Exome Express addressed an important segment of the Broad's customers, it represented only one percent of the total sequencing revenue, and the rest of the operations remained highly chaotic and unpredictable. The average turnaround time was still over 120 days, and there was substantial variation around that mean. In addition, the demand for products was highly variable due to fluctuations in sample availability and funding cycles. In-progress samples were distributed throughout the operation and when a customer inquired as to the state of a sample, it could take days worth of searching just to give her an answer. Worse, despite the large accumulation of in-process samples and an already overworked sequencing staff, utilization rates for the expensive sequencing equipment hovered around 50 percent. Customers increasingly complained about the slow turnaround times and new orders for sequencing were increasingly at risk as other labs, especially those with less complexity, could produce faster results.

Figure 4 provides a simplified view of the process. As soon as a sample arrived, it would be "pushed" into the process and would sit in the queue (in a refrigerator) awaiting "plating." Once a sample completed the plating step, it would then proceed to the queue in front of the next step where it would again wait its turn and so on until it was ready to be placed on a sequencing machine.

Analyzing the process yielded two key insights. First, there was simply too much work in the system. A large collection of in-process samples sat in front of each stage in the process, and the sequencing staff was completely overloaded. They would typically spend their entire day trying to complete samples and had no time for maintenance or to contribute to ongoing efforts to innovate. The degree of challenge far exceeded the optimal level.

Second, the samples were often re-prioritized at each step, meaning that the samples in a given pile were constantly being re-ordered due to a recent customer complaint. Constantly changing the schedule had several negative consequences. It made completion times highly variable as some samples repeatedly got shuffled to the bottom in favor

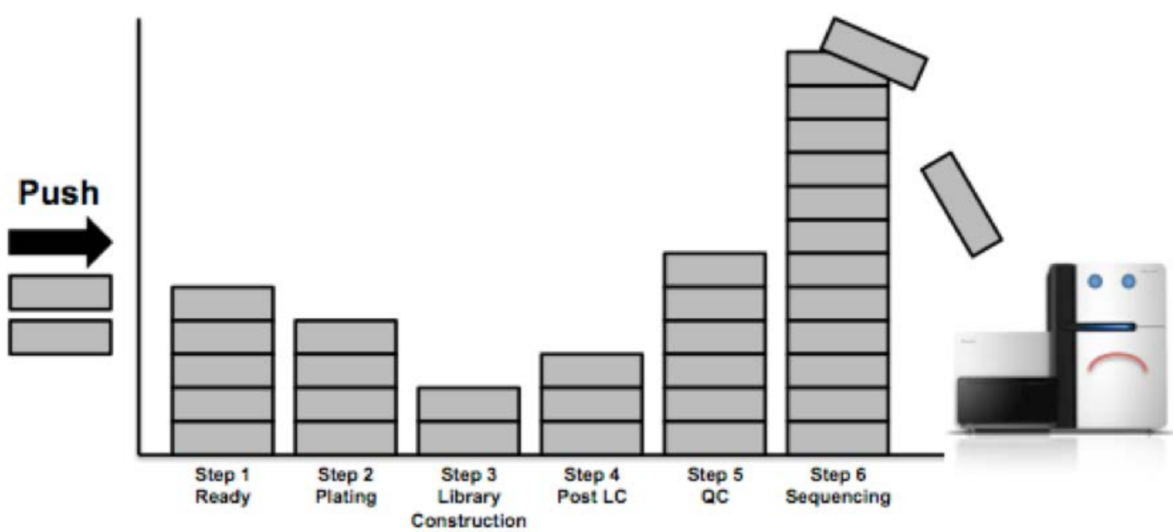

FIGURE 4. PUSH SYSTEM FOR FEEDING THE GENE SEQUENCING MACHINES

of “rush” jobs. It also dramatically increased the workload on the staff, as they were constantly looking at every pile in the operation trying to determine what to do next, while simultaneously reducing equipment utilization since the piles were rarely equally sized. Although it might not be immediately obvious, the local re-prioritization represents another divergence of activity and intent. While the intent of the sequencing process was to deliver timely and reliable analysis to all of its customers, the teams were responding to local pressures (get this done today!) with little view of the overall system. Local re-prioritization, while well intended, compromised that intent by creating a set of activities that were chaotic and inefficient.

To re-infuse the process with meaning and help the participants re-connect activity and intent, the leadership team started their intervention with an “all-hands” meeting focused on delivering three messages: (1) demand was declining due to increased turnaround times; (2) low utilization of expensive capital equipment was putting the entire operation at risk of being shut down; and (3) any solution to (1) and (2) required flexibility to meet rapidly changing scientific needs. Following the exome experience, they also introduced several new metrics that focused on the customer’s experience of the process. The teams, now having transparency on the current state of the operations, expressed a strong desire to be part of a working solution. They jointly agreed to a goal of reducing the turnaround times by 50 percent.

Using the basic tenets of structured problem solving,[24] the teams studied the existing process and collected data on both past demand and their ability to satisfy that demand. The analysis revealed several causes underlying the poor performance. Most importantly (as mentioned above), there was simply too much work in the system; in-progress samples were quite literally everywhere. In addition (and partially as a consequence of the excess loading), the teams got little regular feedback on the performance of the system. When they got a nasty call from an important researcher, they would rush her sample through the process, but they had no way to assess the overall health of the process. Finally, much like the exome process, there was no sense of a human connection. To the contrary, the different functional areas often bickered, and meetings frequently descended into counterproductive finger pointing.

To address these challenges, the team introduced a practice known as “pull” (a scheme widely used by proponents of lean manufacturing).[25] The idea behind pull is simply that in any work system there is an ideal amount of work (measured in a single common unit) and that to maintain that level a new activity should not be introduced until an old activity is completed. Further, in a pull scheme, tasks are prioritized once, before they enter the work system. To implement “pull,” the “waiting area” in front of each process step was reconfigured (see Figure 5). Under the old system, samples were allowed to accumulate as they arrived without restriction, often resulting in large, disorganized piles. Under the new system, the waiting area was configured so that only a small number of samples could be present and was further color coded to provide clear feedback to the technician at the upstream operation. If samples only filled the green area, then the technician could work as quickly as she wanted. Once samples began to occupy the yellow area, it was time to slow down. Once the red area was full, work should stop until the downstream operation had done enough work to open a slot.

[24] Deming (1982) and Shook (2009).
[25] Ohno (1988).

BEFORE

AFTER

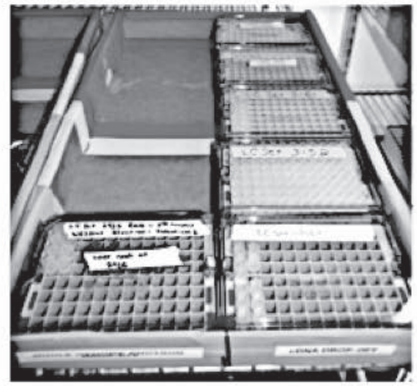

FIGURE 5. SAMPLE WAITING AREAS BEFORE AND AFTER IMPLEMENTING PULL

With these changes, the schematic for the new system is shown in Figure 6.

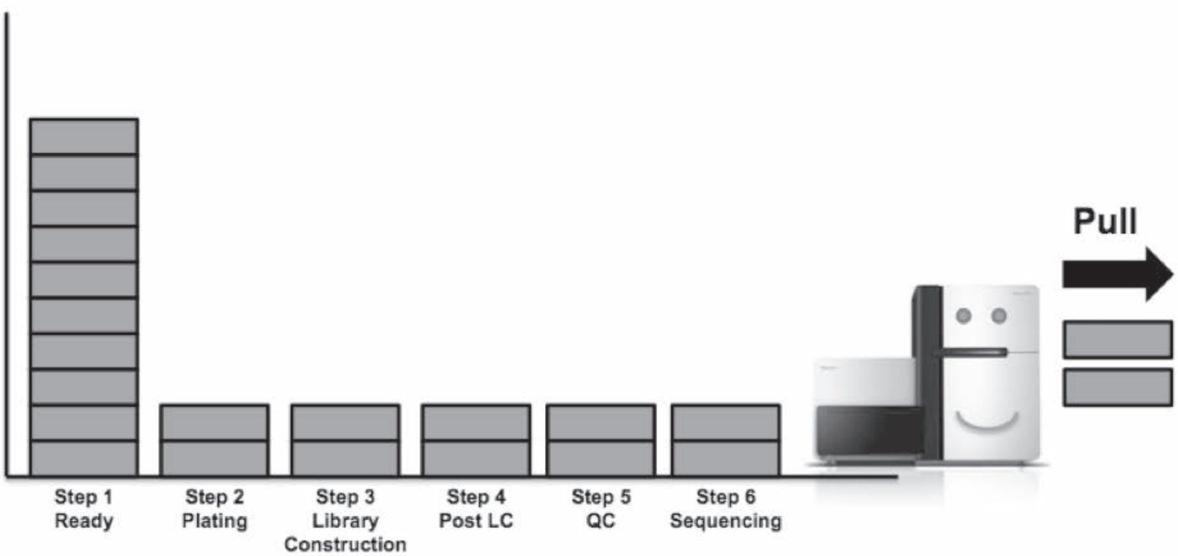

FIGURE 6. PULL SYSTEM FOR FEEDING THE GENE SEQUENCING MACHINES

The new configuration has three key properties. First, the amount of total work in the system (i.e., the level of challenge) is now directly controlled by the size of the staging areas. Work only enters the system when there is available capacity and the staging areas can be easily adjusted to make sure that the challenge level remains optimal (by adding or taking away slots). Second, prioritization is now only done once, before work is "pulled" into the system. A single prioritization stage ensures that all the subsequent activities meet the intent and that technicians don't waste endless hours reprioritizing samples. The pull system also produces highly regular and predictable process times, so re-prioritization isn't needed as often. Finally, the clearly labeled staging areas (see Figure 5) provide unambiguous, easy-to-interpret visual signals about the health of the process (rapid feedback). If, for example, a staging area is constantly full, then it is clear the process step that it feeds requires some attention. In the old system, because work in process accumulated everywhere, it was far less obvious when a particular machine was not operating at its specified capacity. Moreover, because a technician would stop once her downstream staging area was "in the red," she could go assist another technician who was falling behind (thus creating the human connection). The pull system, with its clear visual cues, caused "line balancing" (ensuring that all stages had similar capacity levels) to happen naturally and dynamically.

Implementing the pull system yielded significant gains. The system became far less chaotic, and the participants felt calmer and less stressed. Turnaround times fell by more than 50 percent and became far more predictable (see Figure 7).

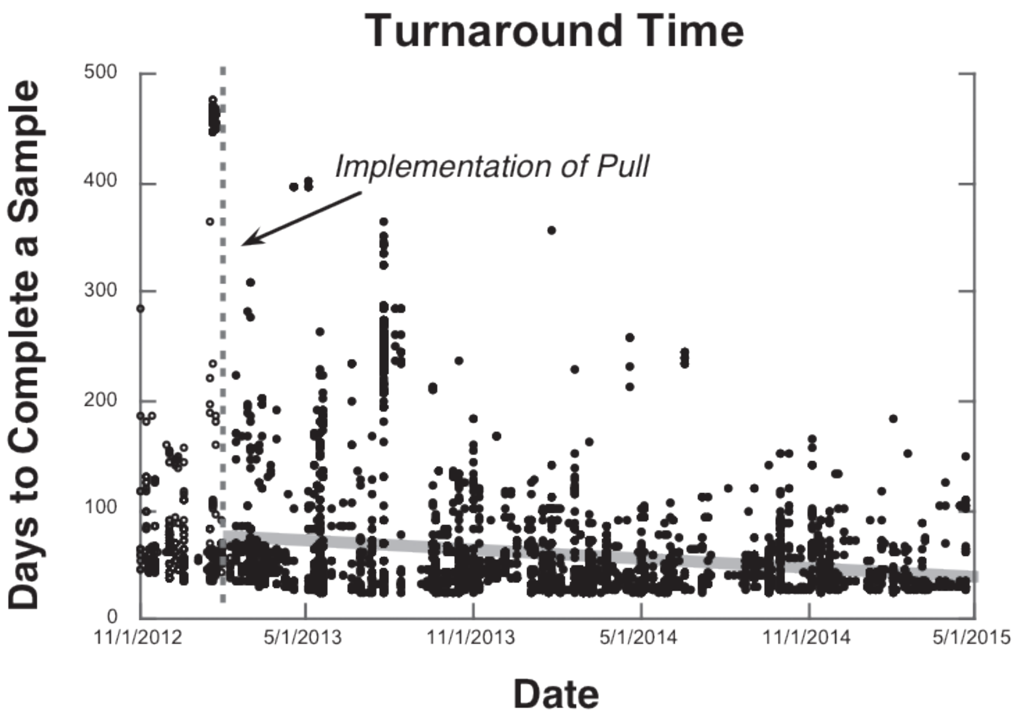

FIGURE 7. AVERAGE TURNAROUND TIME FOR GENE SEQUENCING

In addition, the machine utilization increased significantly without adding any additional staff, going from a low of less than 50 percent to over 85 percent. And, perhaps most importantly, the visual pull system activated the human chain. Each day, as depicted below in Figure 9, the line supervisors and technicians walk through each area reviewing daily targets and looking at the visual pull boxes. Issues are immediately visible and the team leverages their time together to resolve them, in most cases on the spot. More serious issues are escalated to more senior leaders within 24 hours so additional resources can be brought to bear. Creating an opportunity to solve problems in real time using direct human interaction (the "human chain") has reduced the chaos significantly, eliminating hundreds of daily emails in exchange for a daily, 30-minute conversation.

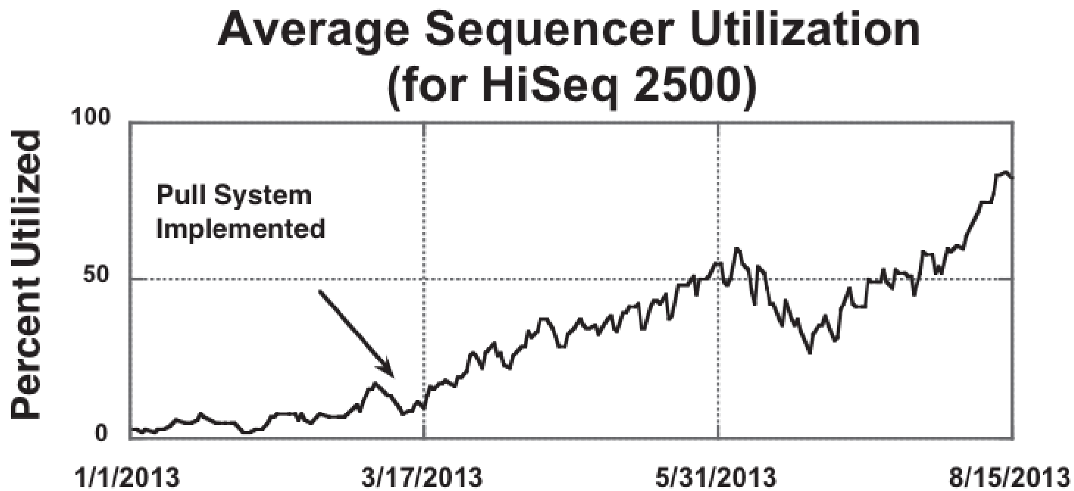

FIGURE 8. UTILIZATION OF GENE SEQUENCING MACHINES

FIGURE 9. MORNING PRODUCTION MEETING

Visual Management

If you are a student of process-focused change (TQM, Six Sigma, Lean, etc.), the first two interventions may strike you as fairly standard implementation of techniques that are well known in the world of manufacturing and assembly. And, at one level, your instincts are correct; these tools are well known and often used in manufacturing operations. Missing from the picture we have painted so far, however, is the environment in which all of this was taking place. As we mentioned at the outset, the science of genomics is changing rapidly and staying competitive requires that the Broad change its operations on an almost daily basis. The process mapping and pull techniques described above are normally associated with far more stable processes. More generally, organizational scholars have long argued that efficiency and adaptability are trade-offs and that an organization that is highly productive will not have the flexibility necessary to adapt to new innovations.[26] Academic theory notwithstanding, the Broad team did not have the luxury of choosing one or the other, so how could they "have their cake and eat it too" in the sense that they could benefit from effective production processes without losing the adaptability and innovation required by their industry? The third major intervention, visual management, tackled this challenge.

Incorporating innovation into existing ways of work takes time and energy. Managing a highly complex, highly dynamic set of biological processes requiring constant change can be overwhelming, especially in a highly competitive industry. The Broad faces a particular challenge in this regard because, though it aspires to lead the industry and build strong, collaborative relationships with its technology vendors, its research and development budget is often just a few percent of what its competitors spend. To maintain its leadership position in the genomics community, the Broad needs to deliver analyzed samples to its customers in a timely manner while simultaneously making very smart decisions about investments in the capabilities that will enable faster, more sophisticated analyses in the future. There are numerous examples of organizations that have failed to manage this balance effectively.[27] Some have focused so much on the innovation that they failed to deliver their core product or service in a timely manner and were forced to exit the market before they could realize the benefits that their innovations might have provided. On the flip side, lots of organizations get so focused on delivering efficiently within their existing technologies that they are eventually left behind when new innovations reach commercial scale.

To meet this challenge, the Broad team required an approach to managing innovation that embodied the principles of good dynamic work design. First, they needed a way to regularly reconcile the activities they undertook with their intent; would the innovations under consideration deliver their goals? Second, given the long cycle time often associated with new ideas, how could they maximize the feedback they were getting as their investigations progressed? It is currently popular to talk about "failing fast," but how do you collect and structure the emerging information to ensure that the right innovations are taken off the table?[28] Third, because people often get attached to ideas that they have helped develop, the basic elements of structured problem solving are needed to ensure that decisions on innovations were made based on the best available thinking. Finally, ensuring optimal challenge was critical. Introducing innovation can be extremely disruptive to ongoing operations and can put existing business at risk. In 2012, it became increasingly clear that the operations teams were taking on too much activity, making it difficult to deliver their work on time. To stay competitive, the Broad team needed a method for ensuring that they were tackling the right projects and introducing them into their existing operations in an efficient fashion.

Paradoxically, one of the key ingredients to meeting the innovation challenge is increasingly well known and seemingly obvious: Research suggests that, *when structured properly,* diverse groups make better decisions than individuals.[29] In the context of innovation, this finding means that success turns on having high-quality conversations that cross critical functional and disciplinary boundaries and ensuring that all the relevant issues get surfaced and feed a structured decision-making process. But, as easy as this is to say, putting it into practice is a significant challenge. As recently discussed by Cass Sunstein and Reid Hastie, it is often the structure of the collaborative process that makes all the difference. Absent a good structure, people are often not good at talking about their work and even less willing to discuss shortcomings and challenges. Worse, work is often implicit, residing only in the heads of those

[26] Abernathy (1978), March (1991), Adler et al. (2009).
[27] See, for example, Christensen (1997), Henderson and Clark (1990).
[28] Sastry and Penn (2014).
[29] For a recent summary see Sunstein and Hastie (2015).

who do it; and, in a chaotic, rapidly changing environment, the intent of implicit tasks is easily forgotten. Different disciplines often use different jargon and are not well acquainted with the challenges a particular innovation poses outside of their particular function. To meet these challenges and create the necessary discussions on a regular basis, the Broad turned to *visual management.*

Visual management is a technique for explicitly (and visually) linking activity and intent. It begins with clarifying the organization's goals and continues with the progressive articulation of each activity that the organization might undertake, typically written on a Post-it note, to support those goals. At the outset of the reorganization in 2012, the Broad team realized they had more activity underway than they could possibly complete given the available resources and required timeframe. To make good, collaborative decisions about which activities to prioritize, the team needed a method to see all of the work that they potentially faced, which, in turn, required surfacing the many activities that were done without the direct knowledge of leadership. To start, the team placed every activity they had underway on a Post-it and then categorized them according to both the three words that described their competitive position (Pioneering, Premium, and Responsive) and their one-year goals. The initial mapping exercise helped connect the "human chain" and forced the reconciliation of activity and intent as team members were forced to articulate how their activities supported the agreed-upon goals. With all of the work represented in a visual format, the team (with the help of more senior leadership) began eliminating work that did not directly support the one-year targets. Such a reduction would have been impossible under the prior system in which most of the work was hidden or implicit. Through this exercise, the target workload was reduced by more than 50 percent, allowing the teams to focus on completing the projects that really mattered.

Since that initial mapping exercise, the Broad's visual display of its work has evolved through a variety of different configurations, each one representing a better understanding of the essential activities required to meet their intent. The one that has proven most useful so far is shown in Figure 10.

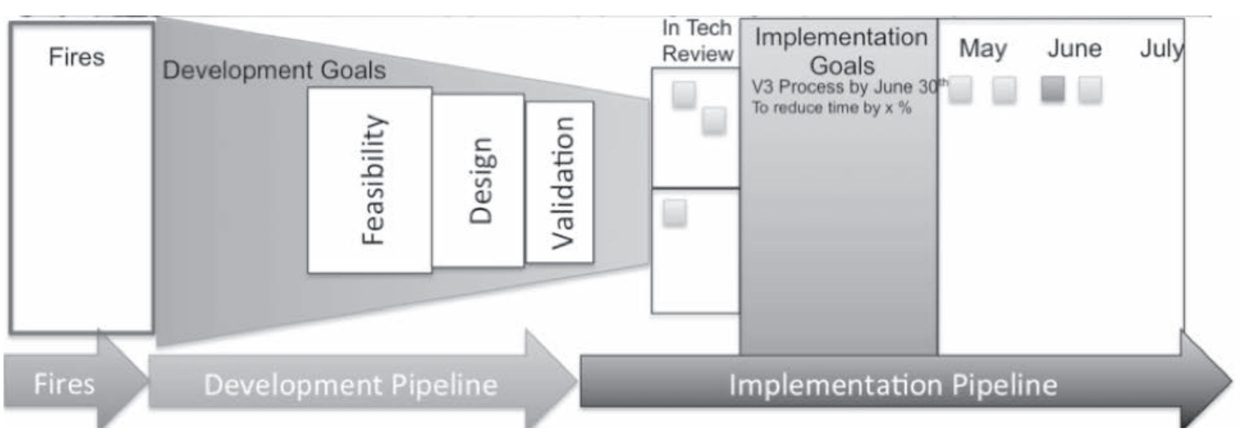

FIGURE 10. OVERVIEW OF VISUAL MANAGEMENT BOARD

The board is divided into three area—Fires, Development Pipeline, and Implementation Pipeline—and is read from left to right. The board serves as the agenda for both a strategic-level, senior leadership review that occurs every six to eight weeks and the weekly operations review meeting (which sometimes will be attended by as many as 50 people). The weekly meeting begins with the review of the “fire board.” Each team member reports whatever “fires” she might be facing—fires being issues that threaten meeting the team’s objectives. Each fire is captured on a Post-it and the board is divided into four columns, each representing a different phase of the structured problem-solving process—formulation, investigation and analysis, propose improvements, and implementation. Once surfaced, each fire is assigned an owner who is responsible for shepherding it through the problem-solving process. Fires that threaten daily production are highlighted with a big red stop sign, making them immediately obvious to anyone walking by the board. The owner of any such fire also updates the entire organization nightly on the progress of her investigation. Often, fires are surfaced and resolved during a single meeting. For example, a junior lab supervisor might surface an issue that is preventing work from happening, and someone from the relevant function (e.g., IT) will immediately look into the issue on his computer, ask a few questions, and resolve the issue in real time. The fire board “connects the human chain” and thereby creates highly productive and rapid face-to-face interactions. Using the fire board has eliminated hours of discussion that previously took place in lengthy email chains.

Following the fire board, the team reviews its development goals, focusing on what technologies are on the horizon that might contribute to the operations mission (as captured by Pioneering, Premium, and Responsive). If an idea is deemed at all promising, it goes on a Post-it and sits in the blue area to the left of the box marked “feasibility.” Simple visual symbols (usually a red line) are used to show when activity (represented by a Post-it) is being worked on. When the team deems that an idea is worth a serious investigation (meaning it’s worth allocating appreciable resources), its Post-it is moved into the feasibility box. As an idea progresses from feasibility to design to validation, the required level of resources typically grows. Being able to see where each project resides allows heavily burdened teams to allocate their scarce energy appropriately without much conversation. For example, the software engineering team can look at the board and focus their attention on the new laboratory protocols that are nearing implementation rather than relying on conversations from the development team members who often want to prioritize their particular protocol improvement, whether or not it is ready for implementation. Critically, ideas are only moved into feasibility (and thus receive appreciable time and attention) when the team deems that they have the collective capacity to work on it. Put differently, the board facilitates a pull system for knowledge work; work doesn’t enter the system until something else is completed. This ensures both that the group stays in the region of optimal challenge and that the activities are clearly aligned with the intent of the work. “Secret” or “pet” projects, which are common in other organizations and are often highly distracting, are heavily discouraged, and the board helps enforce that discipline. Once a project goes on the board, it has both a person and a projected launch date into production assigned to it. Thus, to meet your commitments to the rest of the team, you need to work on the stuff that is on the board.

Once something enters the development funnel (via feasibility), it then proceeds through the two subsequent phases, design and validation. When a project exits a phase, the relevant Post-it is either moved to the next phase or removed if the team determines that the idea in question is not worth pursuing further. As a project nears its launch date, a due date and the name of the person responsible are placed on each Post-it note. When a project slows or work stops, it is visible to the entire team since its Post-it is not moving (recall the clear feedback principle). Now the team can have a structured conversation about the source of the missed expectations, what was learned, and what they might do next. Often, these conversations result in the project being abandoned (we learned it didn't work) or offers of help from other functions (we didn't realize we needed software engineering). In either case, the visual tool gives the team rapid feedback on their progress and catalyzes the use of structured inquiry.

All the ideas that successfully exit validation (meaning that they have been proven to work) are then reviewed with stakeholders in and outside the immediate group. The technology review provides critical input to strategic decisions about how the new protocol might manifest as a new product and helps identify the set of activities required to implement it. Those ideas that have high anticipated impact but require little effort are prioritized first. Using the pull concept again, the best ideas are only moved into the implementation phase when the team deems that they have available capacity. Once an idea (and its Post-it) are moved into implementation, it is given a clear deadline and the required activities and associated deadlines are captured on the implementation board. Again, when a date is missed, it is easily identified and forces a structured conversation: Why has the activity not been completed? What was learned? What are we going to do?

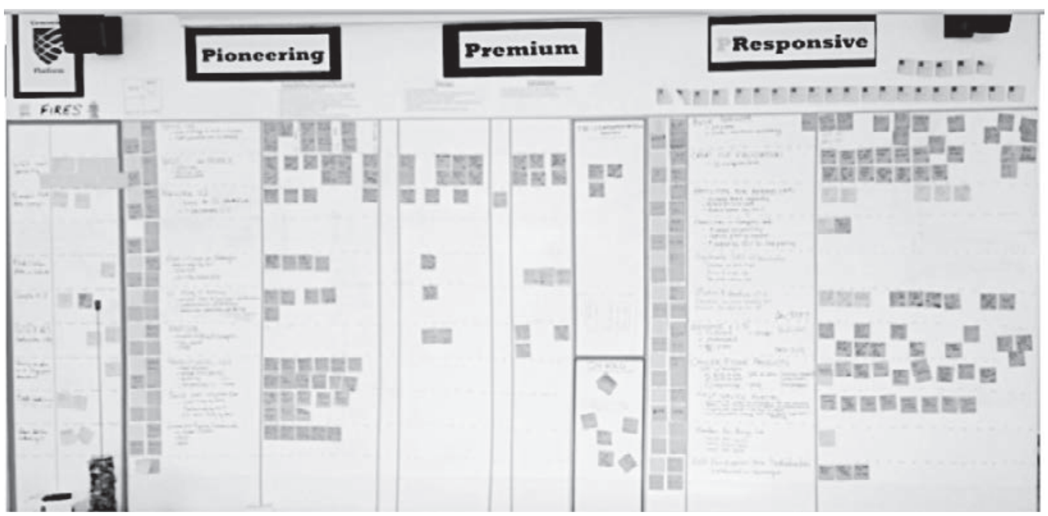

FIGURE 11. THE BROAD OPERATIONS VISUAL MANAGEMENT BOARD

By giving work a physical manifestation (in this case Post-its), visual management enables the team to apply the dynamic work design principles to their "knowledge work." With visual management, the team can regularly assess whether or not they believe that their current activity set will yield the desired outcomes. By attaching a due date to each Post-it and clearly marking those that miss those dates, the board gives the team clear and immediate feedback concerning whether or not the activities are being executed. The clear feedback catalyzes structured problem solving when an important activity falls short. Particularly vexing issues may be placed on the fire board, which further enforces good problem-solving discipline and clearly highlights the issues that most threaten the team's objectives. Finally, the board allows an easy visualization of the amount of work in the system and enables the use of a simple "pull" system in which new projects aren't allowed into a particular phase until there is available capacity. Using this scheme has allowed the team to move their workload far closer to the optimal challenge.

The impact of using visual management has been significant. All of the operational improvements discussed previously were made while the underlying technology continued to evolve rapidly. During the three years of the effort to date, the team has successfully implemented a variety of new technologies and practices while still making significant operations improvements. The team further reports deeper engagement in their work and far more success with cross-functional efforts. Whereas previously the team struggled with the turf battles that often ensue when a collection of different functions, all with their own goals and projects, try to coordinate, they are now far more integrated. And the quantitative results bear out this qualitative assessment. The Broad is now capable of analyzing many samples in less than a few days, overall cycle times have fallen by more than 50 percent and due to efficiency improvements they have dropped their prices by more than 30 percent since 2012. Shorter cycle times and lower prices mean that researchers studying everything from cancer to Ebola can run more experiments and get the results back faster, thus quite literally speeding the pace of science. And, all of these changes have taken place while the Broad continues to keep pace with, and often lead, the rapid changes in sequencing technology.

DISCUSSION

As we have begun to teach this material, by far the most common question goes something like this: "Sure, this works in your examples (we typically discuss manufacturing, genomic sequencing, and drilling oil wells), but will it work in *my* organization?" We always give the same reply: "The stuff we teach only works in organizations that have people in them." Toyota and other Asian manufacturers catalyzed a revolution in manufacturing, and Western companies spent the better part of two decades mastering quality management and Lean production (an effort that continues to this day). But, while the companies that have used lean have developed significant capability in the manufacturing and supply chain operations, they may have missed the larger message. The assumption underlying the understanding of lean methods is that they represent a better way to organize manufacturing activity. The message underlying our quip about the contexts in which dynamic work design might be useful is that many of the tools and practices associated with quality management and lean approaches work, not because they are better ways of organizing manufacturing activity, but because they are better ways of organizing *human* activity.

With a few important exceptions, both consultants and management scholars have historically analyzed work systems using a set of deeply flawed assumptions about the humans doing that work. Taylor believed that work was inherently distasteful and that people would only do it for a buck, a presumption that continues to this day. But a growing line of research convincingly demonstrates that not all work is created equally; we find some tasks far more engaging than others, and no amount of money can induce the kind of motivation that comes from deep engagement. Money clearly matters, but our work suggests that managers who try to induce effort and creativity exclusively through incentives, bonuses, and semiannual performance reviews are using a severely restricted set of tools.

Returning to a distinction we highlighted earlier, while significant effort has been invested in documenting the tools and practices used by high-performing organizations, less attention has been given to the question of why these practices work. If we analyze these practices under the assumption that humans are machines or the coolly rational beasts depicted in traditional economic theory, then we end up focusing on the details. But, if we instead assume both that humans come equipped with significant strengths and equally significant limitations and that strengths and limitations exert a strong influence on the types of tasks that we find engaging and the types that we find annoyingly repetitive, then we might learn lessons that extend beyond the factory floor.

The emerging principles of dynamic work design result from analyzing work systems in light of our growing knowledge of the human animal. It is unlikely that the creators of the Toyota system had a deep understanding of cognitive and behavioral science. But, through insight and trial and error, they discovered practices that created more engaging, more productive work. Through our work, we have discovered similar pockets of excellence—places where the work is well matched to the people doing it. But so far, these remain isolated pockets of success and we have yet to work with an organization that has fully capitalized on the learning that those pockets have to offer. Analyzing these practices has yielded four principles for designing work that works: regularly reconcile activity and intent, use structured thinking, ensure optimal challenge, and connect the human chain. As the experience at the Broad demonstrates, simple application of these principles, often embodied in well-known practices, can lead to significant improvements. Despite the gains, we believe we are just scratching the surface. The principles we have advanced are underpinned by decades of careful research; we expect that, while they will evolve incrementally and new ones may be added to the list, it is unlikely that they will be fundamentally overturned (though it is of course possible). In contrast, the connection between those principles and the practices that actually create better work represents mostly uncharted territory. The Broad team used two very familiar practices, process mapping and pull, and one much less well known and competently practiced concept, visual management, to improve their work. But, while they are clearly more productive, it is not clear that these are in fact the most effective manifestations of the dynamic work design principles. The visual management board, for example, went through several iterations before the team settled on the current configuration. With the design principles in hand, there is a huge opportunity to identify the tools and practices that are most useful in turning them into action. Exercising these opportunities, however, requires overturning a few deeply held assumptions about work.

Most importantly, all leaders need to revisit the question of why people come to work. As discussed above, historically management theory presumed that people's primary motivation was to maximize income and minimize effort. Money does matter, but as basic needs are increasingly satisfied, other dimensions of work have a growing influence on our experience. Research shows that people are more engaged when they have the opportunity to develop in new capabilities (and receive feedback on their progress) in the service of something larger than themselves.[30] Rejecting Taylor's outdated conception of motivation in favor of a more nuanced, accurate view is the first step towards work that is more engaging and more productive. A big part of the Broad's success was reconnecting the operations staff with their mission. In the flux and chaos created by poorly designed work, it is easy to lose sight of the "why."

Second, it is commonly believed that work can only be designed if it is repetitive. People who do so-called "knowledge" work, work that requires innovation and creativity, often argue that their work cannot be designed or even improved. Nothing could be further from the truth for several reasons. Even the most creative job has multiple routine elements and requires several inputs (customer specifications, for example). Properly structuring such tasks ensures both that the routine tasks take less time, thus leaving *more* time for creative work, and that the proper inputs are there when needed. In addition, excess time pressure and task switching have been shown to hamper creativity.[31] At Broad, visual management coupled with a simple pull system ensures a reasonable level of task loading and challenge. Returning to our quip mentioned above, if work involves people, whether it be on the assembly line, in the research lab, or in the executive suite, then there is a good chance it can be redesigned to be both more engaging and more productive.

[30] Deci (1996) and Grant (2014).
[31] Amabile (1996).

Finally, returning to a distinction that we have raised several times, managers and leaders need to be careful not to confuse principles and practices. Efforts to transfer ideas developed in the manufacturing area into other domains rarely succeed. A tool or practice is likely to be successful if it is well adapted to the specific situation in which it is used. But, it is exactly this context-specific customization that makes it difficult to transfer to other domains. Toyota makes cars, the Broad sequences genomes, and it's probably unreasonable to expect that the exact technique that facilitates car assembly would be equally applicable to analyzing blood and tissue samples. But, because humans do critical tasks in both processes, it is reasonable to expect that a common set of design principles might be used to craft effective work in each setting. In the past two decades, the business community has grown obsessed with documenting best practice, despite the fact that evidence on the utility of transferring such practices is decidedly mixed. In contrast, managers have shown little patience for digging into the principles underlying these practices, principles that, if identified, might enable successful transfer. Consequently, the corporate world now has a profound learning disability. Industries spend billions annually on documenting best practices, but it is the rare organization that reaps any significant value. It should come as little surprise that the same improvement and change techniques periodically reappear, having been relabeled to capture the attention of a new generation of managers. Dynamic work design only scratches the surface. Moving the conversation from best practices to best principles offers the possibility of creating work that is both more productive and more satisfying for the people who do it.

REFERENCES

Abernathy, W.J., 1978. *The Productivity Dilemma Roadblock to Innovation in the Automobile Industry*. Johns Hopkins University Press, Baltimore.

Adler, P. and R. Cole (1993). "Designed for Learning: A Tale of Two Auto Plants," *Sloan Management Review*, Spring: 85–94.

Adler, P., M. Benner, D. Brunner, J. MacDuffie, B. Staats, H. Takeuchi, M. Tushman, and S. Winter (2009). "Perspectives on the Productivity Dilemma," *Journal of Operations Management*, 27: 99–113.

Amabile, T. M. (1997). "Motivating Creativity in Organizations: On Doing What You Love and Loving What You Do," *California Management Review,* 40, 1: 39–58

Amabile, T. M. (1996). *Creativity in Context: An Update to The Social Psychology of Creativity.* Westview Press.

Ancona, D., T. Kochan, M. Scully, J. Van Maanen, and D. Westney (2005). *Managing for the Future: Organizational Behavior and Processes*. Thompson Southwestern.

Ariely, D. (2010). *Predictably Irrational*. Harper Perennial.

F. Calefato, D. Damian, and F. Lanubile (2012). "Computer-Mediated Communication to Support Distributed Requirements Elicitations and Negotiations Tasks," *Empirical Software Engineering Journal*, 17, 6: 640–674

Christensen C.M. (1997). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston, Harvard Business School Press.

Combe, C. (2014). *Introduction to Managemen*t. Oxford, Oxford University Press.

Deci, E. (1996). *Why We Do What We Do: Understanding Self-Motivation*. Penguin Books.

Deming, W. (1982). *Out of Crisis.* MIT Press.

Ford, H. (1926). *Today and Tomorrow*. Double Day, Page and Co.

Galbraith, J. (2002). *Design Organizations: An Executive Guide to Strategy, Structure, and Process*. San Francisco, Jossey-Bass.

Grant, A. (2014). *Given and Take*. Penguin Books.

Grant, A., Y. Fried, S. Parker, and M. Frese (2010). "Putting Job Design in Context: Introduction to the Special Issue," *Journal of Organizational Behavior*, 2–3:145–157.

Hackman, J. and G. Oldham (1980). *Work Redesign*. Reading, MA, Addison-Wesley.

Haidt, J. (2006). *The Happiness Hypothesis.* Basic Books.

Henderson, R. and K. Clark (1990). "Architectural Innovation: The Reconfiguration of Existing Product Technologies and The Failure of Established Firms," *Administrative Science Quarterly*; 35, 1.

International Human Genome Consortium and Lander et al. (2001). *Nature*. 409, 860–921 (15 February 2001).

Kahneman, D. (2011). *Thinking Fast and Slow*. Macmillan / Holtzbrinck.

Kates, A., and J. Galbraith (2007). *Designing Your Organization.* San Francisco, Jossey-Bass.

March, J.G., (1991). "Exploration and Exploitation in Organizational Learning." *Organization Science,* 2 (1), 71–87.

Ohno, T. (1988). *Work Place Management.* Productivity Press.

Pink, D. (2011). *Drive: The Surprising Truth About What Motivates Us*. Riverhead Books.

Pfeffer, J. (1992). *Managing with Power Politics and Influence in Organizations.* Boston, Harvard Business School Press.

Pfeffer, J. and R. Sutton (2000). *The Knowing Doing Gap.* Harvard Business School Press.

Sastry, A. and K. Penn (2014). *Fail Better: Design Smart Mistakes and Succeed Sooner.* Harvard Business School Press.

Repenning, N., P. Goncalves, and L. Black (2001). "Past the Tipping Point: The Persistence of Fire Fighting in Product Development." *California Management Review,* 43, 4: 44–63.

Repenning, N. and J. Sterman (2002). "Capability Traps and Self-Confirming Attribution Errors in the Dynamics of Process Improvement." *Administrative Science Quarterly*, 47, 265–295.

Repenning, N. and J. Sterman (2001). "Nobody Ever Gets Credit for Fixing Defects that Didn't Happen: Creating and Sustaining Process Improvement." *California Management Review*, 43, 4: 64–88.

Rudolph, J. and N. Repenning (2002). "Disaster Dynamics: Understanding the Role of Stress and Interruptions in Organizational Collapse." *Administrative Science Quarterly*, 47: 1–30.

Shook, J. (2009). "Toyota's Secret: The A3 Report." *Sloan Management Review*, 50, 5.

Spear, S. (2010). *The High-Velocity Edge: How Market Leaders Leverage Operational Excellence to Beat the Competition.* McGraw Hill.

Sterman, J. D. (2000). *Business Dynamics.* McGraw Hill.

Stevenson, W. (2015). *Operations Management.* McGraw Hill.

Sunstein, C. and R. Hastie (2015). *Wiser: Getting Beyond Groupthink to Make Groups Smarter.* Harvard Business Review Press.

Taylor, F. (1911). *The Principles of Scientific Management.* New York, Harper and Brothers Publishers.

Wheelwright, S. and K. Clark (1992). *Revolutionizing Product Development.* Simon and Schuster.

Womack, J., D. Jones, and D. Roos (1990). *The Machine that Changed the World.* New York: Harper Collins.

ABOUT THE AUTHORS

DON KIEFFER

Don started working in factories at the age of 17 with a high school education. Starting as a piece-worker on a metal lathe, he eventually got his electrical engineering degree on the road to becoming vice president of manufacturing excellence at Harley-Davidson. He later was senior vice president of operations for a multinational corporation, responsible for 5,000 people globally in manufacturing, engineering, supply chain, quality, and logistics. He is an expert in Toyota methods, especially at the leadership level. As founder of the consulting company ShiftGear Work Design, Don has been consulting around the world and has taught at business schools in Moscow, Copenhagen, Luxembourg, and Zagreb. He is a Senior Lecturer at MIT Sloan, where he runs a two-day Executive Education course and teaches in the Executive MBA program.

NELSON REPENNING

Nelson Repenning is the MIT Sloan School of Management Distinguished Professor of System Dynamics and Organization Studies and Associate Dean of Leadership and Special Projects. He also is a partner and chief social scientist at the consulting firm ShiftGear Work Design. Throughout his career, Nelson has focused on understanding the factors that contribute to the successful implementation, execution, and improvement of business processes. His award-winning research has been recognized by the *California Management Review* and the *Journal of Product Innovation Management*. He is also a winner of the prestigious Jay Forrester Award from the International System Dynamics Society.