

Nelson P. Repenning
Don Kieffer
Todd Astor

The Most Underrated Skill in Management

There are few management skills more powerful than the discipline of clearly articulating the problem you seek to solve before jumping into action.

The Most Underrated Skill in Management

There are few management skills more powerful than the discipline of clearly articulating the problem you seek to solve before jumping into action.

BY NELSON P. REPENNING, DON KIEFFER, AND TODD ASTOR

IT'S HARD TO pick up a current business publication without reading about the imperative to change. The world, this line of argument suggests, is evolving at an ever-faster rate, and organizations that do not adapt will be left behind. Left silent in these arguments is which organizations will drive that change and how they will do it. Academic research suggests that the ability to incorporate new ideas and technologies into existing ways of doing things plays a big role in separating leaders from the rest of the pack,¹ and studies clearly show that it is easier to manage a sequence of bite-sized changes than one huge reorganization or change initiative.² But, while many organizations strive for continuous change and learning, few actually achieve those goals on a regular basis.³ Two of the authors have studied and tried to make change for more than two decades, but it was a frustrating meeting that opened our eyes to one of the keys to leading the pack rather than constantly trying to catch up.

In the late 1990s, one of the authors, Don Kieffer, was ready to launch a big change initiative: implementing the Toyota production system in one of Harley-Davidson Inc.'s engine plants. He hired a seasoned consultant, Hajime Oba, to help. On the appointed day, Mr. Oba arrived, took a tour of the plant, and then returned to Don's office, where Don started asking questions: When do we start? What kind of results should I expect? How much is it going to cost me? But, Mr. Oba wouldn't answer those questions. Instead he responded repeatedly with one of his own: "Mr. Kieffer, what problem are you trying to solve?" Don was perplexed. He was ready to spend money and he had one of the world's experts on the Toyota production system in his office, but the expert (Mr. Oba) wouldn't tell Don how to get started.



THE LEADING QUESTION

How can executives lead organizational change more effectively?

FINDINGS

- ▶ Articulate a clear statement of the problem you are trying to solve before initiating changes.
- ▶ Break big problems into a series of smaller ones that can each be tackled quickly.
- ▶ Follow a structured approach to problem-solving using the A3 form originally developed by Toyota Motor Corp.

The day did not end well. Don grew exasperated with what seemed like a word game, and Mr. Oba, tired of not getting an answer to his question, eventually walked out of Don's office. But, despite the frustration on both sides, we later realized that Mr. Oba was trying to teach Don one of the foundational skills in leading effective change: formulating a clear problem statement. Since Mr. Oba's visit, two of the authors have studied and worked with dozens of organizations and taught over 1,000 executives. We have helped organizations with everything from managing beds in a cardiac surgery unit to sequencing the human genome.⁴ Based on this experience, we have come to believe that problem formulation is the single most underrated skill in all of management practice.

There are few questions in business more powerful than "What problem are you trying to solve?" In our experience, leaders who can formulate clear problem statements get more done with less effort and move more rapidly than their less-focused counterparts. Clear problem statements can unlock the energy and innovation that lies within those who do the core work of your organization.

As valuable as good problem formulation can be, it is rarely practiced. Psychologists and cognitive scientists have suggested that the brain is prone to leaping straight from a situation to a solution without pausing to define the problem clearly. Such "jumping to conclusions" can be effective, particularly when done by experts facing extreme time pressure, like fighting a fire or performing emergency surgery. But, when making change, neglecting to formulate a clear problem statement often prevents innovation and leads to wasted time and money. In this article, we hope to both improve your problem formulation skills and introduce a simple method for solving those problems.

How Our Minds Solve Problems

Research done over the last few decades indicates that the human brain has at least two different methods for tackling problems, and which method dominates depends on both the individual's current situation and the surrounding context. A large and growing collection of research indicates that it is useful to distinguish between two modes of thinking, which psychologists and cognitive scientists sometimes call automatic processing and

conscious processing (also sometimes known as system 1 and system 2).⁵ These two modes tackle problems differently and do so at different speeds.

Conscious Processing Conscious processing represents the part of your brain that you control. When you are aware that you are thinking about something, you are using conscious processing. Conscious cognition can be both powerful and precise. It is the only process in the brain capable of forming a mental picture of a situation at hand and then playing out different possible scenarios, even if those scenarios have never happened before.⁶ With this ability, humans can innovate and learn in ways not available to other species.

Despite its power, conscious processing is "expensive" in at least three senses. First, it is much slower than its automatic counterpart. Second, our capacity to do it is quite finite, so a decision to confront one problem means that you don't have the capacity to tackle another one at the same time. Third, conscious processing burns scarce energy and declines when people are tired, hungry, or distracted. Because of these costs, the human brain system has evolved to "save" conscious processing for when it is really needed and, when possible, relies on the "cheaper" automatic processing mode.

Automatic Processing Automatic processing works differently from its conscious counterpart. We don't have control over it or even feel it happening. Instead, we are only aware of the results, such as a thought that simply pops into your head or a physical response like hitting the brake when the car in front of you stops suddenly. You cannot directly instruct your automatic processing functions to do something; instead, they constitute a kind of "back office" for your brain. When a piece of long-sought-after information just pops into your head, hours or days after it was needed, you are experiencing the workings of your automatic processing functions.

When we tackle a problem consciously, we proceed logically, trying to construct a consistent path from the problem to the solution. In contrast, the automatic system works based on what is known as association or pattern matching. When confronted with a problem, the automatic processor tries to

match that current challenge to a previous situation and then uses that past experience as a guide for how to act. Every time we instinctively react to a stop sign or wait for people to exit an elevator before entering, we rely on automatic processing's pattern matching to determine our choice of action.

Our "associative machine" can be amazingly adept at identifying subtle patterns in the environment. For example, the automatic processing functions are the only parts of the brain capable of processing information quickly enough to return a serve in tennis or hit a baseball. Psychologist Gary Klein has documented how experienced professionals who work under intense time pressure, like surgeons and firefighters, use their past experience to make split-second decisions.⁷ Successful people in these environments rely on deep experience to almost immediately link the current situation to the appropriate action.

However, because it relies on patterns identified from experience, automatic processing can bias us toward the status quo and away from innovative solutions. It should come as little surprise that breakthrough ideas and technologies sometimes come from relative newcomers who weren't experienced enough to "know better." Research suggests that innovations often result from combining previously disparate perspectives and experiences.⁸ Furthermore, the propensity to rely on previous experiences can lead to major industrial accidents like Three Mile Island if a novel situation is misread as an established pattern and therefore receives the wrong intervention.⁹

That said, unconscious processing can also play a critical and positive role in innovation. As we have all experienced, sometimes when confronting a hard problem, you need to step away from it for a while and think about something else. There is some evidence for the existence of such "incubation" effects. Unconscious mental processes may be better able to combine divergent ideas to create new innovations.¹⁰ But it also appears that such innovations can't happen without the assistance of the conscious machinery. Prior to the "aha" moment, conscious effort is required to direct attention to the problem at hand and to immerse oneself in relevant data. After the flash of insight, conscious attention is again needed to evaluate the resulting combinations.

The Discipline of Problem Formulation

When the brain's associative machine is confronted with a problem, it jumps to a solution based on experience. To complement that fast thinking with a more deliberate approach, structured problem-solving entails developing a logical argument that links the observed data to root causes and, eventually, to a solution. Developing this logical path increases the chance that you will leverage the strengths of conscious processing and may also create the conditions for generating and then evaluating an unconscious breakthrough. Creating an effective logical chain starts with a clear description of the problem and, in our experience, this is where most efforts fall short.

A good problem statement has five basic elements:

- It references something the organization cares about and connects that element to a clear and specific goal;
- it contains a clear articulation of the gap between the current state and the goal;
- the key variables — the target, the current state, and the gap — are quantifiable;
- it is as neutral as possible concerning possible diagnoses or solutions; and
- it is sufficiently small in scope that you can tackle it quickly.

Is your problem important? The first rule of structured problem-solving is to focus its considerable power on issues that really matter. You should be able to draw a direct path from the problem statement to your organization's overall mission and targets. The late MIT Sloan School professor Jay Forrester, one of the fathers of modern digital computing, once wrote that "very often the most important problems are but little more difficult to handle than the unimportant."¹¹ If you fall into the trap of initially focusing your attention on peripheral issues for "practice," chances are you will never get around to the work you really need to do.

Mind the gap. Decades of research suggest that people work harder and are more focused when they face clear, easy-to-understand goals.¹² More recently, psychologists have shown that mentally comparing a desired state with the current one, a process known as mental contrasting, is more likely to lead people to change than focusing only on the future or on

current challenges.¹³ Recent work also suggests that people draw considerable motivation from the feeling of progress, the sense that their efforts are moving them toward the goal in question.¹⁴ A good problem statement accordingly contains a clear articulation of the gap that you are trying to close.

Quantify even if you can't measure. Being able to measure the gap between the current state and your target precisely will support an effective project. However, structured problem-solving can be successfully applied to settings that do not yield immediate and precise measurements, because many attributes can be subjectively quantified even if they cannot be objectively measured. Quantification of an attribute simply means that it has a clear direction — more of that attribute is better or worse — and that you can differentiate situations in which that attribute is low or high. For example, many organizations struggle with so-called “soft” variables like customer satisfaction and employee trust. Though these can be hard to measure, they can be quantified; in both cases, we know that more is better. Moreover, once you start digging into an issue, you often discover ways to measure things that weren't obvious at the outset. For example, a recent project by a student in our executive MBA program tackled an unproductive

weekly staff meeting. The student began his project by creating a simple web-based survey to capture the staff's perceptions of the meeting, thus quickly generating quantitative data.

Remain as neutral as possible. A good problem formulation presupposes as little as practically possible concerning why the problem exists or what might be the appropriate solution. That said, few problem statements are perfectly neutral. If you say that your “sales revenue is 22% behind its target,” that formulation presupposes that problem is important to your organization. The trick is to formulate statements that are actionable and for which you can draw a clear path to the organization's overarching goals.

Is your scope down? Finally, a good problem statement is “scoped down” to a specific manifestation of the larger issue that you care about. Our brains like to match new patterns, but we can only do so effectively when there is a short time delay between taking an action and experiencing the outcome.¹⁵ Well-structured problem-solving capitalizes on the natural desire for rapid feedback by breaking big problems into little ones that can be tackled quickly. You will learn more and make faster progress if you do 12 one-month projects instead of one 12-month project.

To appropriately scope projects, we often use the “scope-down tree,” a tool we learned from our colleague John Carrier, who is a senior lecturer of system dynamics at MIT. The scope-down tree allows the user to plot a clear path between a big problem and a specific manifestation that can be tackled quickly. (See “Narrowing a Problem's Scope.”)

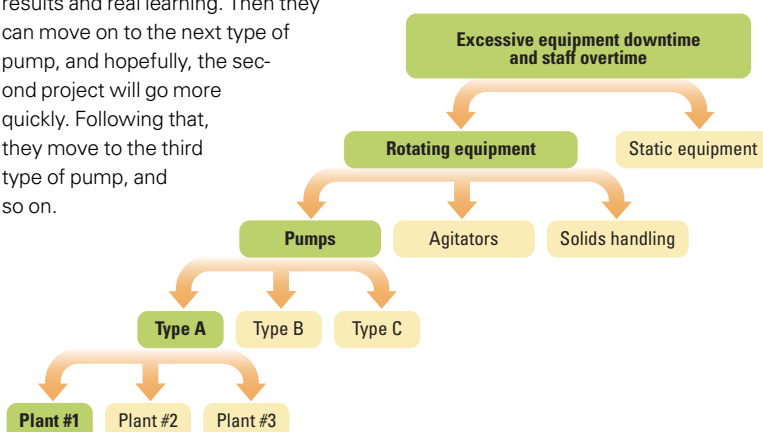
Managers we work with often generate great results when they have the discipline to scope down their projects to an area where they can, say, make a 30% improvement in 60 days. The short time horizon focuses them on a set of concrete interventions that they can execute quickly. This kind of “small wins” strategy has been discussed by a variety of organizational scholars, but it remains rarely practiced.¹⁶

Four Common Mistakes

Having taught this material extensively, we have observed four common failure modes. Avoiding these mistakes is critical to formulating effective problem statements and focusing your attention on the issues that really matter to you and your organization.

NARROWING A PROBLEM'S SCOPE

Good structured problem-solving involves breaking big problems into smaller ones that can be tackled quickly. In this “scope-down tree,” developed by John Carrier of MIT, the overall problem of excessive equipment downtime at a company's plants is broken down first into two types of equipment (rotating and nonrotating), and then further into different subcategories of equipment, ultimately focused on a specific type of pump in one plant. The benefit of reducing the problem's scope is that instead of a big two-year maintenance initiative, a team can do a 60-day project to improve the performance of the selected pumps and generate quick results and real learning. Then they can move on to the next type of pump, and hopefully, the second project will go more quickly. Following that, they move to the third type of pump, and so on.



1. Failing to Formulate the Problem The most common mistake is skipping problem formulation altogether. People often assume that they all already agree on the problem and should just get busy solving it. Unfortunately, such clarity and commonality rarely exist.

2. Problem Statement as Diagnosis or Solution

Another frequent mistake is formulating a problem statement that presupposes either the diagnosis or the solution. A problem statement that presumes the diagnosis will often sound like “The problem is we lack the right IT capabilities,” and one that presumes a solution will sound like “The problem is that we haven’t spent the money to upgrade our IT system.” Neither is an effective problem statement because neither references goals or targets that the organization really cares about. The overall target is implicit, and the person formulating the statement has jumped straight to either a diagnosis or a solution. Allowing diagnoses or proposed solutions to creep into problem statements means that you have skipped one or more steps in the logical chain and therefore missed an opportunity to engage in conscious cognitive processing. In our experience, this mistake tends to reinforce existing disputes and often worsens functional turf wars.

3. Lack of a Clear Gap A third common mistake is failing to articulate a clear gap. These problem statements sound like “We need to improve our brand” or “Sales have to go up.” The lack of a clear gap means that people are not engaging in clear mental contrasting and creates two related problems. First, people don’t know when they have achieved the goal, making it difficult for them to feel good about their efforts. Second, when people address poorly formulated problems, they tend to do so with large, one-size-fits-all solutions that rarely produce the desired results.

4. The Problem Is Too Big Many problem statements are too big. Broadly scoped problem formulations lead to large, costly, and slow initiatives; problem statements focused on an acute and specific manifestation lead to quick results, increasing both learning and confidence. Use John Carrier’s scope-down tree and find a specific manifestation of your problem that creates the biggest headaches. If you can solve that instance of the

problem, you will be well on your way to changing your organization for the better.

Formulating good problem statements is a skill anybody can learn, but it takes practice. If you leverage input from your colleagues to build your skills, you will get to better formulations more quickly. While it is often difficult to formulate a clear statement of the challenges you face, it is much easier to critique other people’s efforts, because you don’t have the same experiences and are less invested in a particular outcome. When we ask our students to coach each other, their problem formulations often improve dramatically in as little as 30 minutes.

Structured Problem-Solving

As you tackle more complex problems, you will need to complement good problem formulation with a structured approach to problem-solving. Structured problem-solving is nothing more than the essential elements of the scientific method — an iterative cycle of formulating hypotheses and testing them through controlled experimentation repackaged for the complexity of the world outside the laboratory. W. Edwards Deming and his mentor Walter Shewhart, the grandfathers of total quality management, were perhaps the first to realize that this discipline could be applied on the factory floor. Deming’s PDCA cycle, or Plan-Do-Check-Act, was a charge to articulate a clear hypothesis (a Plan), run an experiment (Do the Plan), evaluate the results (Check), and then identify how the results inform future plans (Act). Since Deming’s work, several variants of structured problem-solving have been proposed, all highlighting the basic value of iterating between articulating a hypothesis, testing it, and then developing the next hypothesis. In our experience, making sure that you use a structured problem-solving method is far more important than which particular flavor you choose.

In the last two decades, we have done projects using all of the popular methods and supervised and coached over 1,000 student projects using them. Our work has led to a hybrid approach to guiding and reporting on structured problem-solving that is both simple and effective. We capture our approach in a version of Toyota’s famous A3 form that we have modified to enable its use for work in settings other than manufacturing.¹⁷ (See “Tracking Projects Using an A3 Form,” p. 44.)

TRACKING PROJECTS USING AN A3 FORM

To track problem-solving projects, we have modified the A3, a famous form developed by Toyota, to better enable its use for tracking problem-solving in settings other than manufacturing. The A3 form divides the structured problem-solving process into four main steps, represented by the big quadrants, and each big step has smaller subphases, captured by the portions below the dotted lines. To view a completed A3 form, visit the online version of this article at <http://sloanreview.mit.edu/x/58330>.

PROBLEM STATEMENT	TARGET DESIGN		
Background	Improvement Goal Leadership Guidelines		
CURRENT DESIGN (based on seeing the work)	EXECUTION PLAN	Track Results	
		Date	Target
Root Causes	What Did We Learn and What's Next?		

The original A3 form was developed by Toyota Motor Corp. to support knowledge sharing in its factories by summarizing a structured problem-solving effort in a single page. Though the form may often have supporting documentation, restricting the project summary to a single page forces the user to be very clear in his or her thinking. The A3 divides the structured problem-solving process into four main steps, represented by the big quadrants, and each big step has smaller subphases, captured by the portions below the dotted lines. The first step (represented by the box at the upper left) is to formulate a clear problem statement. In the Background section (in the bottom part of the Problem Statement box), you should provide enough information to clearly link the problem statement to the organization's larger mission and objectives. The Background section gives you the opportunity to articulate the why for your problem-solving effort.

Observing the Current Design The next step in the A3 process is to document the current design of the process by observing the work directly. Due to

automatic processing, most people, particularly those who do repetitive tasks, cannot accurately describe how they actually execute their work. Through pattern matching, they have developed a set of habitual actions and routine responses of which they may not be entirely aware.

Because those who do the work often cannot fully describe what they do, you as a manager must get as close to the locus of the problem as you can and watch the work being done. Taiichi Ohno, one of the founding fathers of the Toyota production system, developed the Gemba walk (Gemba is a Japanese word that roughly translates to "the real place") as a means for executives to find out what really happens on a day-to-day basis. The goal is to understand how the work is really done. This could mean watching a nurse and a doctor perform a medical procedure, engineers in a design meeting, or salespeople interacting with a customer.

Senior executives are often quite removed from the day-to-day work of the organizations that they lead. Consequently, observing and thoroughly understanding the current state of the work often suggests easy opportunities for improvement. We

give our students the following rule of thumb to guide their efforts: When you go see the work, if you aren't embarrassed by what you find, you probably aren't looking closely enough. Recently, we helped a team tackle the problem of reducing the time to process invoices. In walking through the process, the team observed that each invoice spent several days waiting for the proper general ledger code to be added. The investigation, however, revealed that for this type of invoice, the code was always the same; each invoice spent several days waiting for a piece of information that could have been printed on the form in advance!

Root Causes Observing the work closely often shakes loose a variety of preconceptions. The next step in filling out the A3 is to analyze root causes and engage your conscious processing by explicitly linking your observations to the problem statement.

There are a variety of techniques and frameworks to guide a root cause analysis. Perhaps most famously, Sakichi Toyoda, founder of Toyota Industries, suggested asking the “5 whys,” meaning that for each observed problem, the investigator should ask “why” five times in the hope that five levels of inquiry will reveal a problem's true cause. Later, Kaoru Ishikawa developed the “fishbone” diagram to provide a visual representation of the multiple chains of inquiry that might be required to dig into the fundamental cause of a problem.¹⁸ Since then, just about all structured problem-solving methods have offered one or more variants of the same basic method for digging into a problem's source.¹⁹

The purpose of all root-cause approaches is to help the user understand how the observed problem is rooted in the existing design of the work system. Unfortunately, this type of systems thinking does not come naturally. When we see a problem (again, thanks to pattern matching) we have a strong tendency to attribute it to an easily identifiable, proximate cause. This might be the person closest to the problem or the most obvious technical cause, such as a broken bracket. Our brains are far less likely to see that there is an underlying system that generated that poorly trained individual or the broken bracket. Solving the immediate problem will do nothing to prevent future manifestations unless we address the system-level cause.

A good root-cause analysis should build on your investigation to show how the work system you are analyzing generates the problem you are studying as a part of normal operations. If the root-cause analysis identifies a series of special events that are unlikely to happen again, you haven't dug deeply enough. For example, customer service hiccups often differ from instance to instance and are easily attributed to things that “are once in a lifetime and could never happen again.” Digging deeper, however, might reveal a flawed training process for those in customer-facing jobs or an inconsistent customer on-boarding process. A good root-cause analysis links the data obtained in your investigation to the problem statement to explain how the current system generates the observed challenges not as a special case but as a part of routine conduct.

Target Design One you have linked features of the work system to the problem you are trying to solve, use the Target Design section of the A3 form to propose an updated system to address the problem. Often the necessary changes will be simple.²⁰ In the Target Design section, you should map out the structure of an updated work system that will function more effectively. This might be as simple as saying that from now on we will print the general ledger code on the invoice form or something more complicated, such as changes to training and on-boarding programs. The needed changes will rarely be an entirely new program or initiative. Instead, they should be specific, targeted modifications emerging from the root-cause analysis. Don't try to solve everything at once; propose the minimum set of changes that will help you make rapid progress toward your goal.

Goals and Leadership Guidelines Completing the Target Design section requires two additional components. First, create an improvement goal — a prediction about how much improvement your proposed changes will generate. A good goal statement builds directly from the problem statement by predicting both how much of the gap you are going to close and how long it will take you to do it. If your problem is “24% of our service interactions do not generate a positive response from our customers, greatly exceeding our target of 5% or less,” then an improvement goal might be “reduce the number of

negative service interactions by 50% in 60 days.” Clear goals are highly motivating, and articulating a prediction facilitates effective learning.

Finally, set the leadership guidelines. Guidelines are the “guardrails” for executing the project; they represent boundaries or constraints that cannot be violated. For example, the leadership guidelines for a project focused on cost reduction might specify that the project should identify an innovation that reduces cost without making trade-offs in quality.

Execution Plan The next step is running the experiment. In the upper portion of the Execution Plan box of the A3 form, lay out a plan for implementing your proposed design. Be sure that the plan is broken into a set of clear and distinct activities (for example, have the invoice form reprinted with the general ledger code or hold a daily meeting to review quality issues) and that each activity has both an owner and a delivery date.

Now execute your plan and meet your target. But, even as you start executing, you are not done engaging in conscious learning. Instead, you want to make sure that you are not only solving the problem but also absorbing all the associated lessons. Track each activity relative to its due date and note those activities that fall behind. These gaps can also be the subject of structured problem-solving. During this phase, interim project reports should be simple: The owner of the action should report whether that element is ahead of or behind schedule, what has been learned in the latest set of activities, and what help he or she may need.

In the Track Results section of the form, measure progress toward your goal. For example, if the overall target is to reduce the number of poor service interactions by 50% in 60 days, then set intermediate goals, perhaps weekly, based on your intervention plan. Put these intermediate targets in the first column of the Track Results section and then measure your progress against them. Also, make sure that you continue to track the results for an extended period after you have met your target. You want results that stick.

Once the project is complete, document what you learned in the What Did We Learn and What’s Next section. Here you should both outline the main lessons from the project and articulate the new opportunities that your project revealed. If you

exceeded your predictions, what does that tell you about future possibilities? In contrast, falling short of your target may reveal parts of the work system that you don’t understand as well as you thought. Finally, and perhaps most importantly, what problem are you going to tackle next? A well-functioning process, whether in manufacturing, customer service, or new product development, is the product of numerous small changes, and fixing one real problem often reveals many additional pressing issues. Close out your A3 by outlining the next problem you and your organization need to solve.

A Case Study in a Hospital

How does this process work in practice? To illustrate, we describe a recent case where one of the authors, a hospital executive who had been introduced to the basics of problem formulation and structured problem-solving, used the techniques to improve organizational performance.

Todd Astor and his team transplant human lungs at Massachusetts General Hospital in Boston, Massachusetts. Although the lung transplant procedure is highly complex, its complexity pales in comparison to managing the recipient’s health after the transplant. The human body often responds to the transplanted organs in dangerous ways. A big part of Todd’s job is staying in close contact with his patients and carefully managing the complicated suite of medicines needed to suppress the body’s natural immune response.

Several times a week Todd’s lung transplant unit has a clinic in which transplant recipients come to be evaluated and receive any necessary adjustments in their treatment. Each clinic session lasts for three hours and utilizes three dedicated exam rooms. Based on the evaluation criteria of Todd’s hospital, that should allow him to see 27 patients (three per hour in each room). But at the outset of the project, the team was able to see an average of seven patients per clinic session. Running the clinic at less than 30% of its ideal capacity potentially compromised care — patients might have to wait longer to be evaluated — and had significant revenue implications for the hospital. With a few iterations, Todd’s challenge led to the following problem statement and supporting background:

The post-lung transplant outpatient clinic session has an average volume of 7 patients, even

though the clinic has the recommended space capacity for up to 27 patients (20 minutes per patient) per session.

The “gap” between the actual and ideal utilization of clinic space (26% of ideal utilization) has resulted in a delay in timely access to care for many lung transplant patients and a loss of potential revenue/profit for the outpatient clinic and the hospital.

After adding some additional background information about the problem to the A3 form, Todd went to understand the work. (To see Todd’s completed A3 form, visit the online version of this article at <http://sloanreview.mit.edu/x/58330>. See “Additional Resources.”) He tracked 71 patients over nine sessions as they flowed through the clinic day. Todd discovered huge variability in both the patient arrival rates and the time that patients spent in the various stages of a clinic visit. A little digging into the root causes revealed numerous ambiguities and departures from the way the system was supposed to work. Patient arrival times were highly variable, due both to a lack of clarity on appointment details and to traffic patterns around the hospital; lab testing times varied depending on the time of day; different versions of the pulmonary function test (PFT) were conducted; there was often little coordination between the doctors and the nurse practitioners; and large amounts of time were spent checking each patient’s medication list.

Todd made two key decisions in analyzing the root causes and proposing changes. First, despite variability at all stages of the visit, he scoped down the problem to focus only on processes occurring in the clinic area. He and his team had more direct control over these processes (compared with those occurring in the laboratory, radiology area, etc.), and were more able to make changes. Second, Todd included every member of the team, from the administrative staff to the physicians, in analyzing the root causes and proposing changes. Widespread inclusion allowed every individual to think about specific ways to address the problem in his or her own assigned area.

The root-cause analysis led to several proposed changes. The administrative assistant would call

patients both a week and a day in advance to remind them about their appointments and provide advice on managing traffic and parking. The PFT test was standardized with a clear rule for when a more detailed test was needed. When possible, the medication list reconciliation would happen the day before the clinic via the telephone. And, finally, the nurse practitioner and the doctor would coordinate their exams to eliminate asking the patient for the same information twice. With these changes, Todd set a target of adding two patients per clinic session until the clinic reached a throughput of 18 patients. Todd further outlined a clear set of guidelines, the most important being that quality of patient care could in no way be sacrificed during the project.

The results were impressive. In seven weeks, the throughput moved from the average of seven to a high of 17 in week seven, not quite meeting Todd’s target of 18, but more than doubling the existing patient flow. After the initial project was completed, the lung transplant clinic subsequently did reach a maximum flow of 18 patients per session.

The increased throughput had several positive benefits. The clinic was able to provide better, more timely care to its patients. Surveys suggested that despite the higher volume, patient satisfaction improved, due to shorter wait times and the perception that they were getting better, more consistent care. Revenue also improved significantly. Less obvious but equally important, improved throughput created space for more patients, thereby matching the growth in the transplant program. Finally, Todd’s team got to control their work and improve it, generating clear gains in motivation and engagement.

From Reorganization to Real Learning

We always ask executives in our MIT Sloan classes: “How many of your companies reorganize every 18 to 24 months?” Typically, more than half of the people in the class raise their hands. Change has become a big business, and any number of consultants will be more than happy to assist your company in your next reorganization. But be careful. Changing everything at once takes a lot of time and resources, and big initiatives often collapse under their own weight as senior executives, tired of waiting for the results, move on to the next big idea. By focusing

ADDITIONAL RESOURCES

To view a completed A3 form for Todd Astor’s patient flow project as well as read an additional case study about structured problem-solving in another setting, visit the online version of this article at <http://sloanreview.mit.edu/x/58330>.

your scarce resources on those issues that really matter and enabling rapid learning cycles, good problem formulation and structured problem-solving offer a sustainable alternative to the endless stream of painful reorganizations and overblown change initiatives that rarely deliver on their promises.

Nelson P. Repenning is the School of Management Distinguished Professor of System Dynamics and Organization Studies at the MIT Sloan School of Management in Cambridge, Massachusetts, as well as chief social scientist at the consulting firm ShiftGear Work Design LLC. **Don Kieffer** is a senior lecturer in operations management at the MIT Sloan School and managing partner of ShiftGear Work Design. **Todd Astor** is the medical director of the lung and heart-lung transplant program at Massachusetts General Hospital and an assistant professor of medicine at Harvard Medical School in Boston, Massachusetts. Comment on this article at <http://sloanreview.mit.edu/x/58330>, or contact the authors at smrfeedback@mit.edu.

REFERENCES

1. R. Gibbons and R. Henderson, "What Do Managers Do? Exploring Persistent Performance Differences Among Seemingly Similar Enterprises" in "The Handbook of Organizational Economics," ed. R. Gibbons and J. Roberts (Princeton, New Jersey: Princeton University Press, 2013), 680-731.
2. N.P. Repenning and J.D. Sterman, "Nobody Ever Gets Credit for Fixing Problems That Never Happened: Creating and Sustaining Process Improvement," *California Management Review* 43, no. 4 (summer 2001): 64-88.
3. A study by Towers Watson reported that only about one in four change efforts are effective in the long run. See Towers Watson, "How the Fundamentals Have Evolved and the Best Adapt: 2013 - 2014 Change and Communication ROI Study," (December 2013), www.towerswatson.com. Others have reached similar conclusions; for example, see J.P. Kotter, "Leading Change" (Boston, Massachusetts: Harvard Business School Press, 1996); and M. Beer, R.A. Eisenstat, and B. Spector, "Why Change Programs Don't Produce Change," *Harvard Business Review* 68, no. 6 (November-December 1990): 158-166.
4. A. Mangi and N.P. Repenning, "Dynamic Work Design Decreases Post-Procedural Length of Stay and Enhances Bed Availability," manuscript available from the author; S. Dodge et al., "Using Dynamic Work Design to Help Cure Cancer (And Other Diseases)," MIT Sloan School of Management working paper 5159-16, June 2016, www.mitsloan.mit.edu.
5. For very readable summaries, see D. Kahneman, "Thinking, Fast and Slow" (New York: Farrar, Straus, and Giroux, 2011); and J. Haidt, "The Happiness Hypothesis: Finding Modern Truth in Ancient Wisdom" (New York: Basic Books, 2006). For recent overviews of scholarly work, see J. St. B.T. Evans and K.E. Stanovich, "Dual-Process Theories of Higher Cognition: Advancing the Debate," *Perspectives on Psychological Science* 8, no. 3 (May 1, 2013): 223-241; and S.A. Sloman, "Two Systems of Reasoning, an Update" in J.W. Sherman, B. Gawronski, and Y. Trope, "Dual-Process Theories of the Social Mind" (New York: Guilford Press, 2014), 107-120. For a collection of reviews, see Sherman, Gawronski, and Trope, "Dual-Process Theories of the Social Mind."
6. K.E. Stanovich, "Rationality and the Reflective Mind" (New York: Oxford University Press, 2011).
7. G.A. Klein, "Sources of Power: How People Make Decisions" (Cambridge, Massachusetts: MIT Press, 1998).
8. J. Singh and L. Fleming, "Lone Inventors as Sources of Breakthroughs: Myth or Reality?" *Management Science* 56, no. 1 (January 2010): 41-56.
9. C. Perrow, "Normal Accidents: Living With High-Risk Technologies" (Princeton, New Jersey: Princeton University Press, 1999).
10. A. Dijksterhuis and L.F. Nordgren, "A Theory of Unconscious Thought," *Perspectives on Psychological Science* 1, no. 2 (June 1, 2006): 95-109; and A. Dijksterhuis, "Automaticity and the Unconscious," in "Handbook of Social Psychology," 5th ed., vol. 1, ed. S.T. Fiske, D.T. Gilbert, and G. Lindzey (Hoboken, N.J.: John Wiley & Sons, 2010), 228-267.
11. J. W. Forrester, "Industrial Dynamics" (Cambridge, Massachusetts: MIT Press, 1961), 449.
12. E.A. Locke and G.P. Latham, "Building a Practically Useful Theory of Goal Setting and Task Motivation," *American Psychologist* 57, no. 9 (September 2002): 705-717.
13. G. Oettingen, G. Hönig, and P.M. Gollwitzer, "Effective Self-Regulation of Goal Attainment," *International Journal of Educational Research* 33, no. 7-8 (2000): 705-732.
14. T.M. Amabile and S.J. Kramer, "The Power of Small Wins," *Harvard Business Review* 89, no. 5 (May 2011): 70-80; and T.M. Amabile and S.J. Kramer, "The Progress Principle: Using Small Wins to Ignite Joy, Engagement, and Creativity at Work" (Boston, Massachusetts: Harvard Business Review Press, 2011).
15. For a summary, see J. Sterman, "Business Dynamics: Systems Thinking and Modeling for a Complex World" (Boston, Massachusetts: Irwin/McGraw-Hill, 2000).
16. K.E. Weick, "Small Wins: Redefining the Scale of Social Problems," *American Psychologist* 39 (January 1984): 40-49; Kotter, "Leading Change"; and T.M. Amabile and S.J. Kramer, "The Power of Small Wins."
17. J. Shook, "Toyota's Secret: The A3 Report," *MIT Sloan Management Review* 50, no. 4 (summer 2009): 30-33.
18. "Fishbone Diagram (Ishikawa) — Cause & Effect Diagram | ASQ," <http://asq.org>.
19. For a summary of root-cause analysis techniques, see en.wikipedia.org/wiki/Root_cause_analysis.
20. In other work, we have proposed four principles for effective work that may be helpful in more complex situations. See Dodge et al., "Using Dynamic Work Design."

Reprint 58330.

Copyright © Massachusetts Institute of Technology, 2017.

All rights reserved.

**PDFs ■ Reprints ■ Permission to Copy ■ Back Issues**

Articles published in MIT Sloan Management Review are copyrighted by the Massachusetts Institute of Technology unless otherwise specified at the end of an article.

MIT Sloan Management Review articles, permissions, and back issues can be purchased on our Web site: sloanreview.mit.edu or you may order through our Business Service Center (9 a.m.-5 p.m. ET) at the phone numbers listed below. Paper reprints are available in quantities of 250 or more.

To reproduce or transmit one or more MIT Sloan Management Review articles by electronic or mechanical means (including photocopying or archiving in any information storage or retrieval system) **requires written permission.**

To request permission, use our Web site: sloanreview.mit.edu

or

E-mail: smr-help@mit.edu

Call (US and International): 617-253-7170 Fax: 617-258-9739

Posting of full-text SMR articles on publicly accessible Internet sites is prohibited. To obtain permission to post articles on secure and/or password-protected intranet sites, e-mail your request to smr-help@mit.edu.