MIT MANAGEMENT SLOAN SCHOOL

MIT Sloan School of Management

MIT Sloan School Working Paper 5175-16

What Problem Are You Trying To Solve: An Introduction to Structured Problem Solving

Todd Astor, Michael Morales, Don Kieffer, and Nelson Repenning

Last Revised: October 10, 2016

What Problem Are You Trying To Solve:

An Introduction to Structured Problem Solving

Todd Astor
Medical Director, Lung and Heart-Lung Transplant Program
Massachusetts General Hospital
Harvard Medical School
tastor@mgh.harvard.edu

Michael Morales
President
Corporacion Industrial, S.A.
mmorales@cisapanama.com

Don Kieffer
Senior Lecturer
MIT Sloan School of Management
Managing Partner
ShiftGear Work Design

don@shiftgear.work

Nelson Repenning

Distinguished Professor
MIT Sloan School of Management
Chief Social Scientist
ShiftGear Work Design

nelson@mit.edu

August 2016

Version 1.0

Introduction

In the late 1990s, Don Kieffer hired a seasoned consultant, Hajime Oba, to help implement the famed Toyota production system. At the time Don was the manager of Capitol Drive, one of Harley-Davidson's two engine manufacturing and assembly plants. Though Harley had a long history with the "lean" methods inspired by Toyota,[1] Don felt that he was still missing something. If anybody could help, it would be Mr. Oba, who had deep experience with the Toyota system and had been sent to the US to help its suppliers build the necessary capabilities. It took several months to schedule, but Mr. Oba finally showed up at Don's plant to help implement Toyota's system at Capitol Drive.

Don started the meeting in the conference room, showing multiple slides describing the improvement efforts already. After a few minutes, Mr. Oba asked politely if they might leave the office and go visit the factory. Once on the shop floor, Mr. Oba proceeded to give himself a tour of the facility, with Don (much to the amusement of those working on the line that day) chasing after him. After about 45 minutes, he asked Don to return to the conference room, where he drew a simple diagram on the white board that gave an overview of the plant. "This is what I observed, is this correct?" he asked.

"Yes," Don replied, "You have it about right."

Mr. Oba then circled a small part of the diagram, made a note, and said "You have a problem in this area, please make this change. When you have finished, call me, and I will come back and tell you the next step." Don stopped him. Having started his career on the shop floor as a piece worker and then working his way through every job in the factory to the corner office, Don

[1] Reid, *Well Made in America*.

wasn't used to being told what to do, especially in his own plant. So, he started asking Mr. Oba all the questions that you would normally pose to a prospective consultant: How much is it going to cost? When do we start? What kind of results should I expect? How often will you be here?

And, again Mr. Oba interrupted him. "But, Mr. Kieffer, what problem are you trying to solve?"

Don was starting to get upset--why was this guy being so evasive? "Mr. Oba," he said, the problem I am trying to solve is that I want to implement the Toyota production system and I don't know how."

Now it seemed like maybe Mr. Oba was finally starting to get it. He replied, "So, Mr. Kieffer, you want to implement the Toyota production system?"

"Yes!" Don replied.

"When do you want to start?" Mr. Oba asked.

"Right now!" Don replied.

"So you want to start implementing the Toyota production system right now, Mr. Kieffer?"

"Yes…" Don said, feeling the frustration starting to return, "…right now!"

"Ok, Mr. Kieffer…" Mr. Oba said, "…if you want to implement the Toyota production right now, then begin by telling me, what problem are you trying to solve?"

The day did not end well. Don grew increasingly frustrated with what seemed like an endless word game and Mr. Oba, tired of not getting an answer to his question, left the plant. But, despite the frustration on both sides, that day proved to be pivotal in our understanding of how to make organizations run more effectively. Although it was not immediately apparent on that fateful day, Mr. Oba was trying to teach Don one of the foundational skills in Toyota's system, formulating a clear problem statement. In the fifteen plus years since Mr. Oba's visit to Capitol Drive, we have studied and worked with dozens of organizations and taught almost a thousand students. We have helped organizations with everything from managing beds in a cardiac surgery unit to sequencing the human genome.[2] During this time we have developed an ever-deepening appreciation for the lesson that Mr. Oba was trying to give Don. Problem formulation is, in our view, the single most underrated skill in all of management practice, and there are few questions more powerful than "what problem are you trying to solve?" Leaders who can formulate clear problem statements both get more done with less effort and move more rapidly than their less-focused counterparts. Even better, clear problem statements are often the key to unlocking the energy and innovation that lies within those who do the core work of your organization, whether it be manufacturing, product development or service.

As valuable as a good problem formulation can be, it is rarely practiced. As we have dug into the underlying psychology it has become clear that the reason we often don't formulate clear problem statements is because doing so doesn't come naturally. The more psychologists and cognitive scientists dig into the inner workings of the human animal, the more we see that our brains are prone to leaping straight from a situation to a solution without pausing to define the problem clearly. Such "jumping to conclusions" can be remarkably effective, particularly when done by experts facing extreme time pressure like fighting a fire or performing emergency

---

[2] Mangi and Repenning, "Dynamic Work Design Decreases Post-Procedural Length of Stay and Enhances Bed Availability."; Dodge, S. et al., "Speeding the Search for a Cure: Using Dynamic Work Design to Improve Genetic Sequencing."

surgery. But, in daily work, neglecting to formulate a clear problem often prevents innovation and leads to wasted time and money. Good problem formulation and structured problem solving also offer a sustainable alternative to the endless stream of painful reorganizations and overblown initiatives that rarely deliver on their promise.

In this paper, we hope to both build your problem formulation skills and introduce a simple method for tackling those well-formulated problems. We begin with a brief summary of the relevant psychology and explain why problem formulation often does not come naturally and thus requires disciplined and conscious attention. Following that, we describe the basic composition of a good problem statement and summarize the most common failure modes. Then we introduce a simple structured problem-solving method, based on the "A3" developed by Toyota. Finally, we finish with two examples from recent projects that show the power of Mr. Oba's insight.

The Associative Machine

For many years management scholars assumed that the human brain worked like the computer on your desk or in your phone. Just like our computers, we assumed that the human brain takes a single approach to tackling any given problem. If, for example, you asked your spreadsheet to add a column of numbers or your word processor to check the spelling of a word, your computer does the task once and does it the same way each time, regardless of what else might be going on at that moment. Research done over the last few decades indicates that the human brain works differently. It often has at least two different methods for tackling a particular problem, and which method dominates (and thus determines the final answer) depends a lot on both the current state of the individual— are you feeling happy or sad, fresh or tired—and the surrounding context—is it noisy or quiet, do you feel time pressure or do you have time to think?

Following a large and growing collection of research, it is useful to distinguish between two basic modes of thinking, what psychologists and cognitive scientists sometimes call *automatic* *processing* and *conscious processing* (also sometimes known as system 1 and system 2).[3][4][5] The distinction between these two modes is important both because they tackle problems differently and do so at different speeds.

Conscious Processing

As the name suggests, conscious processing represents the part of our brain that you control. When you say you are thinking about something (meaning that you are aware of your mental effort) you are using conscious processing. Conscious cognition can be both powerful and precise. Most notably is it the only part of the brain capable of what psychologists call *cognitive decoupling* and *mental simulation*, the ability to form a mental picture of the situation at hand (acknowledging that the mental picture is not the same thing as the real situation) and then play out different possible scenarios, even if those scenarios have never happened before.[6] With this ability, humans can innovate and learn in ways simply not available to other species. As an example, an aspiring technology CEO can look at a cell-phone, GPS technology, and taxi services and imagine combining them into a new model for providing transportation. Critically, that same aspiring CEO can also look at the combination of a cell-phone, GPS technology, and air-travel and realize that it enables a different and probably more limited set of possibilities, *without ever having to experience either possibility directly*. Conscious processing is thus the domain of logic in the sense that it uses knowledge about the world to construct possibilities that extend beyond our past experience.

---

[3] For recent and very readable summaries, see: Kahneman, *Thinking, Fast and Slow*; Haidt, *The Happiness Hypothesis*.
[4] For recent overviews of scholarly work, see: Evans and Stanovich, "Dual-Process Theories of Higher Cognition Advancing the Debate"; Sloman, "Two Systems of Reasoning, an Update."
[5] For a collections of reviews, see: Sherman, Gawronski, and Trope, *Dual-Process Theories of the Social Mind*.
[6] Stanovich, *Rationality and the Reflective Mind*.

Despite its power, conscious processing is also "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 space to tackle another one. If you have ever been so deeply engaged in a task that a spouse or child had to ask the same question several times before it registered, then you have experienced the limits on your conscious processing capacity. Third, conscious processing burns scarce energy and clearly declines as people grow tired and hungry or are distracted. Because of these costs, the human brain system has evolved to "save" conscious processing for when it is really needed. Consequently, when possible, the brain relies on the far "cheaper" automatic processing mode.

Automatic Processing

Automatic processing works quite differently from its conscious counterpart. Most notably, 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 someone stops suddenly in front of you. You cannot directly instruct your automatic processing functions to do something; instead, they constitute a kind of "back office" for your brain. If a piece of long sought after information has ever just "popped" into your head, hours or days after it was needed, you have experienced the workings of your automatic processing functions. Because we lack direct access or control, the workings of the automatic processing function sometimes feel magical and we often use words like gut, instinct, and intuition to describe them.

Importantly, the current science suggests that the automatic processing functions tackle problems very differently than their conscious counterparts. 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*. Pattern matching or association simply means that 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 the automatic processing's pattern matching to determine our choice of action.

In many circumstances, our "associative machine" is 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 enable feats of physical dexterity like returning a serve in tennis or hitting 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. He calls this response *recognition primed decision-making*.[7] Successful people in these environments rely on deep experience to almost immediately link the current situation (recognition) to the appropriate intervention. In one study he showed how a British Navy radar technician made a split second decision to fire on and destroy an object that appeared on his screen, an object that was later revealed to be a hostile missile. The technician's split second decision probably saved multiple lives. But, for months afterward, nobody could figure out how the technician had known that particular little blip on his screen was a Silkworm missile and not friendly aircraft. Only later did Klein's researchers determine that it was the timing with which the blip appeared on the radar that revealed its identity to the technician; the Silkworm missile flew at a lower altitude than aircraft and thus first showed up on the radar in a position farther from the coastline than the surrounding aircraft.[8] This is exactly the kind of setting in which the automatic processor excels, using past experience to identify subtle patterns in the environment, whether it be the way that an opponent sets up for a serve in tennis or the signature of an incoming threat.

---

[7] Klein, *Sources of Power*.
[8] For the complete story, see: Ibid., 35–39.

But, as powerful as automatic processing can be, it is not perfect. Because it relies on patterns identified from past experience and cannot engage in cognitive decoupling or mental simulation, automatic processing can bias us towards the status quo and away from new innovations that might constitute truly "breakthrough" or "disruptive" solutions. As a simple example, read the following question and jot down the first answer that jumps into your head:

How many animals of each kind did Moses bring on the ark?

If you are like the vast majority of people who have read this question, you answered "two." We all know that in the bible, Noah was responsible for building an ark and filling it with two of each kind of animal. However, if you return to the sentence and read a bit more carefully you will notice the Noah is not the object of the query; you were asked about Moses, who is typically not associated with arks and animals. So many people fail to answer this question correctly that psychologists have dubbed it the Moses illusion.[9] And, while its exact cause remains debated, scientists agree that the brain's propensity to match patterns is at the root. As we read the question we get at least two key pieces of information, animals and arks, indicating that the question concerns the famous parable of Noah and our pattern matching apparatus quickly returns the answer that we all think we know so well. The automatic processor can thus trap us in existing ways of doing things and work against innovation. Thus, 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 results from combining previously disparate perspectives and experiences.[10] Even worse, the propensity to rely on past experience can lead to major industrial accidents like Three Mile Island when a

---

[9] Park and Reder, "Moses Illusion: Implication for Human Cognition."

[10] For a review of literature and supporting evidence, see: Singh and Fleming, "Lone Inventors as Sources of Breakthroughs."

novel situation is misread as an established pattern and therefore leads to the wrong intervention.[11]

All of 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 that we can't solve with conscious effort, you need to step away for a while and "think about something else." There is some evidence for the existence of such "incubation" effects. Unconscious mental processes may also be better able to combine divergent ideas to create new innovations (such as cell phones, GPS technology, and taxi services).[12] But it also appears that such innovations can't happen without the assistance of the conscious machinery. Prior to the "a ha" moment, conscious effort is required to direct attention to the problem at hand and to immerse oneself in the relevant data. After the "flash" of insight, conscious attention is again needed to evaluate the resulting combinations.

The human brain's associative machine thus plays a complex role in the search for organizational performance. It enables highly efficient action and rapid response to patterns so subtle that they are not even perceived at the conscious level. It is also equally capable of trapping us in existing ways of thinking and acting. The challenge facing managers is to construct organizations and processes that capitalize on both thinking modes, leveraging our propensity to respond automatically *and* our ability to think things through. Our research suggests that structured problem solving is one of the four critical approaches for developing an effective Dynamic Work Design.[13] Organizations that use structured problem solving effectively both learn faster and

---

[11] Perrow, *Normal Accidents*.
[12] Dijksterhuis and Nordgren, "A Theory of Unconscious Thought"; Dijksterhuis, "Automaticity and the Unconscious."
[13] . Dodge, S. et al., "Speeding the Search for a Cure: Using Dynamic Work Design to Improve Genetic Sequencing."

make more efficient and targeted interventions. Effective structured problem solving begins with good problem formulation.

Problem Formulation

When the brain's associative machine is confronted with a problem, it jumps to a solution based on past 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 a diagnosis (i.e. root causes) and, eventually, to a solution. Developing this logical path increases the chance that you will leverage the strengths of conscious processing (the conscious function is the only one capable of constructing such an argument) and may also create the conditions for generating and then evaluating an unconscious breakthrough. Creating an effective logical chain starts with a clear question and in our experience, this is where most efforts fall short and predetermine failure.

Anatomy of a Good Problem Statement

A good problem statement has five basic elements:

- it references something that the organization cares about and connects that element to a clear and specific goal or target;
- 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, if not immediately measurable;
- it is neutral as possible concerning possible diagnoses or solutions;
- it is sufficiently small in scope that you can tackle it quickly.

Is your problem important?

As you may soon experience, structured problem solving is hard work and the conscious cognition that it requires might be the ultimate scarce organizational resource.[14] The first rule of

[14] Roy Baumeister and colleagues have done a fascinating series of studies showing that the ability to self-regulate is a finite resource that can easily get depleted. For example, having to resist a tempting

structured problem solving is simply to focus its considerable power on issues that really matter. As a first test, you should be able to draw a direct a clear path from the problem statement to your organization's overall mission and targets. All too many efforts to introduce new tools, like those embodied in TQM, Six Sigma and "Lean," have failed because their considerable power was directed at irrelevant problems (most famously, moving water coolers).[15] MIT great 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."[16] 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 need to do. Thus the first test of a good problem statement is: Does it reference an issue that is important to your organization?

Mind the Gap

Decades of research suggest that people work harder and are more focused when they face clear, easy-to-understand goals.[17] More recently, psychologists have shown that mentally comparing the desired state to the current one, a process known as *mental contrasting*, is most likely to lead people to change.[18] In contrast, focusing only on the future (fantasizing) or current challenges (dwelling) is less productive. Recent work also suggests that people draw considerable motivation from the feeling of *progress,* the sense that their efforts are moving them towards the goal in question.[19] Similarly, perhaps the principal insight of W. Edwards Deming, the father of the quality movement, was simply that structured problem solving could be applied repeatedly

---

dessert like chocolate cookies reduced persistence on a hard math problem compared to subjects who hadn't been confronted with temptation. Studies also show that problem solving capability declines as people become depleted. The ability to engage in structured problem solving appears to be a finite resource. For a summary, see: Bauer and Baumeister, "Self-Regulatory Strength."

15 Wiesendanger, "Deming's Luster Dims at Florida Power & Light."

16 Forrester, *Industrial Dynamics*, 449.

17 Locke and Latham, "Building a Practically Useful Theory of Goal Setting and Task Motivation."

18 Oettingen, Hönig, and Gollwitzer, "Effective Self-Regulation of Goal Attainment."

19. see Amabile and Kramer (2011). Amabile and Kramer, *The Progress Principle*.

(thus creating continuous improvement) and that the organization would develop new capabilities in the process. Making sure that people work hard, feel good about their efforts, and learn something in the process requires clarity around both the target and the current state. A good problem statement contains a clear articulation of the *gap* that you are trying to close. A clearly articulated gap helps you plan and focus your effort and also supports closing out the project to determine how much was learned.

Quantify Even if You Can't Measure

Not surprisingly, being able to measure the gap between the current state and your target precisely will support an effective project. That said, not everything that matters can be measured with three digit accuracy. Contrary to popular belief, 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—is more of that attribute 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. We also can probably rank different situations (e.g., this work is more satisfying than that work, or this organization has more trusting employees than that one, etc.). While projects are easier to execute when you can measure the outcomes, subjective quantification can guide new action. Moreover, in our experience, once you start digging into an issue, you often discover new ways to measure things that weren't obvious at the outset. For example, a recent student project tackled a weekly staff meeting that "sucked." The student began his project by creating a simple web-based survey to capture the staff's perceptions of the meeting, thus quickly translating a subjective sense of frustration into quantitative data.

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. Remember that the automatic processor doesn't do problems statements, it simply jumps from a situation to a solution. While this is very valuable when facing extreme time pressure, in other settings it means that we are not engaging the full power of conscious processing. As mentioned above, engaging conscious processing requires constructing a logical path from problem to solution and doing so requires separating the problem from its causes and potential remedies. Following Mr. Oba's challenge to articulate "the problem you are trying to solve" is often central to breaking existing frames and moving people towards more focused, innovative solutions. That said, few problem statements are perfectly neutral. If, for example, you say that your "sales revenue is 22% behind its target," that formulation presupposes that sales revenue is important to your organization and that a more significant problem doesn't reside on the cost side of the income statement. The trick is to formulate statements that are both actionable and for which, as mentioned earlier, 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. In the last two decades many organizations have become overly enamored with large "change initiatives," usually labeled with acronyms that are subsequently enshrined on t-shirts and coffee mugs. But such an approach to change is often a bad match to our natural propensity for pattern matching. Our brains like to match new patterns, it quite literally feels physically good, but we can only do so effectively when there is a short time delay between taking an action and experiencing the outcome.[20] It shouldn't come as a surprise then that most organizations struggle to maintain commitment to their two-year change programs. In

[20] For a summary, see Sterman, *Business Dynamics.*

contrast, well-structured problem solving capitalizes on the natural desire for rapid feedback by focusing on decomposing big problems into little ones that can be tackled quickly. You will learn more and make faster progress if you do twelve one-month projects instead of one twelve-month project.

To appropriately scope projects we often use the "scope down tree". A tool we learned from our MIT colleague John Carrier. The scope down tree is just a simple map that allows the user to plot a clear path between a big problem and a specific manifestation that can be tackled quickly. John's favorite example is shown in figure one:

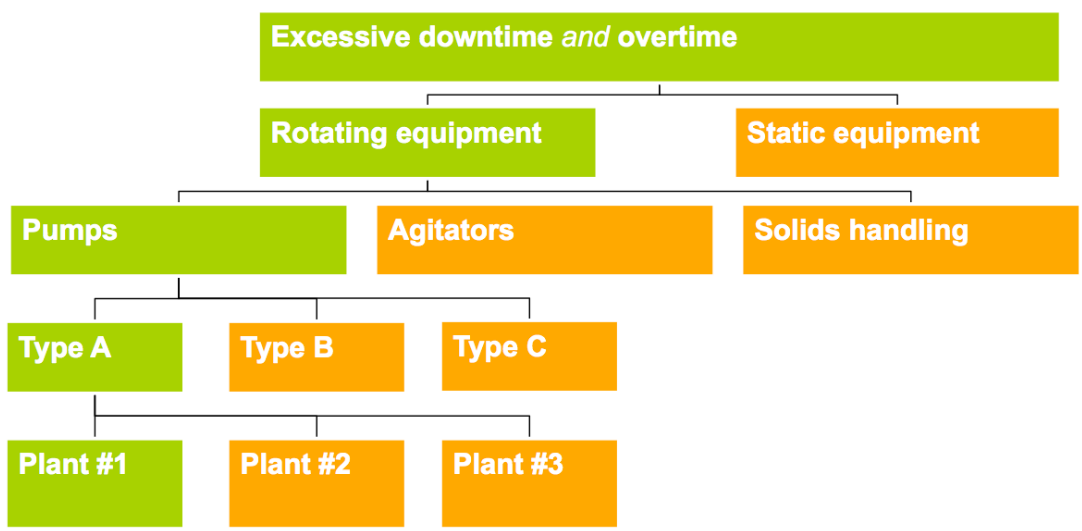

Figure 1: Example of Scope Down a Problem

The overall problem of excessive equipment downtime is first broken down into the two types of equipment (rotating and non-rotating), and then further decomposed into different sub-categories of equipment, ultimately focused on a specific type of pump in one plant. The benefit of the scope down is that instead of a big two-year maintenance initiative, now we can do a sixty-day project to improve the performance of the selected pumps that will generate quick results and real learning. After completion, we can move on to the next type of pump, and hopefully, the second project will go more quickly. Following that, we can move to the third type and so on. This "small wins" strategy has been discussed extensively by a variety of organizational scholars, but it remains all too rarely practiced.[21] Our students and intervention partners have typically generated great results when they have had the discipline to continue to scope down their projects until they identify an area where they can "make a 30% improvement in 60 days." The short time horizon focuses them on a set of concrete interventions that they can execute quickly and generate some real learning rather than endless planning sessions that never generate any action.

Common Mistakes

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

We All Know the Problem Already…

The first and probably most common mistake is skipping problem formulation altogether. Sometimes people assume that because they all already "agree" on the problem and they should just get busy trying to achieve them. Unfortunately, such clarity and commonality rarely exist. Without explicit attention to and discussion of problem statements, we all rely on our individual past experiences to guide our actions and there is no guarantee that those differing past experiences will generate focused action. This failure mode is often most evident in meetings. If you are in a meeting that seems to wander without a lot of focus, chances are that the lack of a clear problem statement is at the root. Nothing brings aimless conversation to a halt faster than our favorite question, "What problem are you trying to solve?"

---

21 Weick, "Small Wins"; Kotter, *Leading Change*; Amabile and Kramer, "The Power of Small Wins."

Problem Statement as Diagnosis or Solution in Disguise

The next most common failure mode 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 marketing capability", or the "problem is the people in manufacturing are poor communicators." Both of these statements could easily be true and represent major problems for the organization, but neither is an effective problem statement because they don't reference goals or targets that the organization really cares about. The overall target is implicit and the person formulating the statement has jumped straight to a diagnosis for why that target is not being met. In the first case, the true problem is probably something related to inadequate sales or revenue, the lack of marketing capability is a diagnosis. Similarly, in the second case, the true problem is probably too many product defects that irritate customers and the lack of communication skill is, again, a diagnosis. Allowing diagnoses to creep into problem statements means that that you have skipped a step 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. If you find yourself constantly adjudicating disputes by R&D and manufacturing or between sales and finance, look for two competing problem formulations, both of which are really diagnoses in disguise.

Confusing problem statements with solutions produces a similar, but often even more costly outcome. You will often hear problem statements like "The problem is that we don't have enough sales people" or "The problem is that we haven't spent the money to upgrade our IT system." In both cases, again thanks to automatic processing, the problem formulation *and* diagnosis steps have been skipped and the person has jumped straight to the solution, hiring more people or upgrading the IT infrastructure. These might be the right interventions, but, because of the overreliance on automatic processing and pattern matching, they are likely to reinforce existing ways of operating and limit the chance of finding innovative solutions. This mistake also tends to reinforce existing disputes, particularly around funding. Those in favor of their preferred solution will constantly feel like they don't have resources to "…really do it right." If you find yourself constantly fending off requests for more money, look for a problem statement that is really a solution in disguise.

Lack of a Clear Gap

The next failure mode is simply the lack of a clearly articulated gap. In some cases, these problem statements will sound like "We need to improve our brand" or "Sales have to go up". This mistake often shows up with so-called soft variables like morale; "we need to change our culture" is a common one. The lack of clear gap means that people are not engaging in clear mental contrasting and creates two related problems. First, people don't know when they are done and it is thus difficult for them to feel good about their efforts since they haven't achieved an identifiable goal. Second, when people address such poorly formulated problems, they often do so with large, one size fits all solutions such as team-building or culture change initiatives that, while expensive, rarely produce the desired results. If you find your organization doing lots of big change efforts that don't seem to work, look for problem statements that have no sense of a clear gap and revise them so that they identify specific shortfalls. If, for example, you feel that you need to change your organization's culture, identify the specific situations where the problem manifests. "People don't speak up in meetings" is a better choice than "we need to change our culture" and "Engineers don't surface failed test data in product review meetings" is even better. And, "at least 50% percent of our defects are revealed to the team four or more weeks after they were discovered by an engineer, falling far short of our one week target" would certainly earn an "A" in our class.

Too Big

The final failure mode is pretty obvious but nonetheless important: most 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 particular instance, you will be well on your way to actually changing your organization for the better.

A Final Tip for Better Problem Formulation

In our experience, formulating good problem statements is a skill that anybody can learn, but it does take continued practice and discipline. More than fifteen years after Mr. Oba's visit, we still find ourselves devolving to the instincts emerging from our own automatic processing. Learning to put the instinctual response produced by your deep experience temporarily aside and engage your conscious thought processes is not easy; you are quite literally fighting the hardwiring of your brain. That said, in the years that we have taught this material and coached managers, we have noticed something very interesting: while it is often difficult to formulate a clear statement of the challenges *you* face, it is much easier to critique other peoples' efforts. In retrospect, it is easy to understand why. Our problem formulations are informed by our past experience and the resulting patterns that our automatic processes have encoded, thus making it difficult to separate problem, diagnosis, and solution. Remember Gary Klein's firefighters: They went straight from situation to intervention. When you look at someone else's problem formulation, you don't have the same experiences or set of patterns and are often less invested in a particular diagnosis or solution, thus making it much easier to see the failure modes outlined above. When we use the standard business school trick of having our students coach each other, their problem formulations often improve dramatically in as little as thirty minutes. As you build your skills, leverage the input of your colleagues. You will get to better formulations more quickly.

Structured Problem Solving and the A3

A well-formulated problem statement will help you move quickly to a more focused, efficient solution. As you begin to tackle more complex problems, however, you will need to complement good problem formulation with a more structured approach to problem solving. Structured problem is nothing more than essential elements of the scientific method re-packaged for the complexity of the world outside the laboratory. The basics of the scientific method—an iterative cycle of formulating hypotheses and testing them through controlled experimentation—were developed to offset our natural and well-documented tendency to select data from the environment consistent with our pre-existing beliefs and ignore those things that were inconsistent. Only through careful hypothesis articulation and controlled experimentation did, for example, Louis Pasteur disprove the theory of spontaneous generation and thereby develop the germ theory of disease that would eventually save countless human lives. W. Edwards Deming and his mentor Walter Shewhart, the grandfathers of total quality management, were perhaps the first to realize that this basic discipline could be applied on the factory floor with similar results. Deming's famous 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 hypotheses, testing them, and then developing the next hypothesis. In our experience, making sure that you actually *use* a structured problem solving method is far more important than which particular flavor you choose.

In the last two-plus decades, we have done projects using all of the popular methods and supervised and coached almost a thousand student projects using them. We have also watched them being used by others and tried to understand the strengths and weaknesses of each approach based on the underlying psychological and organization theory. Our work has led to a hybrid approach to guiding and reporting on structured problem solving that is both simple and effective. The format is essentially a simplified version of Toyota's famous A3 and is shown in figure 2.[22] The A3 was developed by Toyota to support problem solving and knowledge sharing in its factories. The label "A3" refers to the size of the paper that was initially used and it was chosen because, at the time, an A3 sheet was the largest piece of paper that could be faxed (thus promoting easy knowledge sharing).

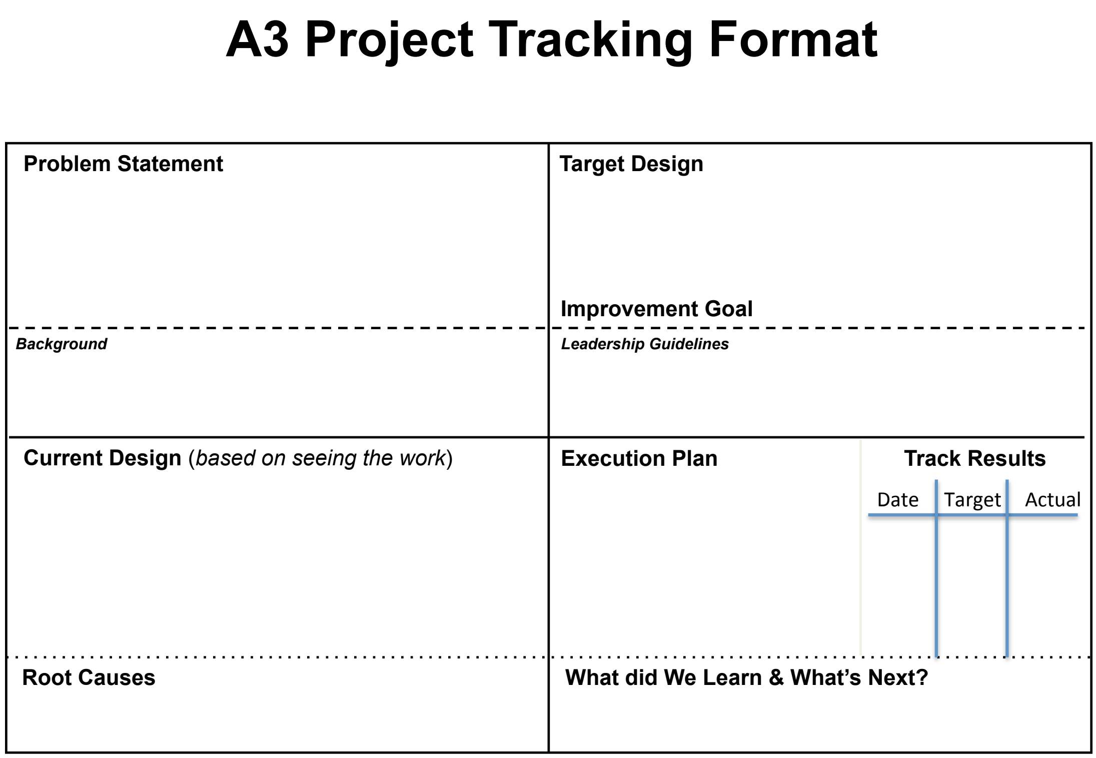

Figure 2: A3 Problem Solving Format

Problem Statement

The A3 divides the structured problem solving process into four main steps, represented by the big quadrants, and each big step has smaller sub-phases, captured by the portions below the dotted lines. The first step (represented by the box in the upper left) is (not surprisingly at this point) to formulate a clear problem statement following the guidelines above. In the *Background*

[22] Shook, "Toyota's Secret."

section (in the bottom third of the Problem Statement box), you should provide enough information to clearly link the problem statement to the organization's larger mission and objectives. Much of structured problem solving is simply stating assumptions that would otherwise be implicit (thereby hopefully engaging more conscious processing) and the background section gives you the opportunity to articulate the *why* for your problem solving effort.

Current Design (*based on seeing the work*)

The next step in the A3 process is to document the current design of the process *by observing the work directly*. The manager's job is a bit different from that of the bench scientist. When tackling a problem in the field, whether it is on the shop floor, in the R&D lab, or at the customer service desk, all of the available knowledge is not codified in easily searchable scientific papers. Instead, the relevant knowledge resides in the heads and hands of the people who actually do the work. Even more challenging, 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. As a simple example, recall the last time you tried to teach somebody a skill, such as driving a car, for which you are deeply experienced and competent.[23] Only when you are forced to articulate everything you do to somebody else—change lanes when you see a merge approaching—do you realize how much of your daily activity is unconscious.

Because those who do the work often cannot fully describe what they really do, when you begin digging into a problem you cannot rely on self-reports. Instead, you 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

[23] Special thanks to Rena Repenning, whose efforts to learn to drive suggested this example.

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. Going and seeing is a simple but powerful method for cutting through all of the psychological and organizational distortions that happen when people try to describe what they do. The goal of the Current Design phase is to understand how the work is *actually* done (not what is written in process manuals and standard procedures) and observe the problem as it happens. This could mean watching an operator on the shop floor doing her job, a nurse and/or a doctor performing a medical procedure, engineers having a design meeting, or sales people interacting with the customer.

Due to their busy schedules, senior executives are often quite removed from the work of the organizations that they lead and can easily lose touch with what is actually happening. Consequently, thoroughly understanding the current state of the work will often suggest multiple, easy to exercise, opportunities for improvement. Put differently, as we often tell our students, if, when you go see the work, 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 (GL) code to be added. The investigation also, however, revealed that for this particular type of invoice the GL 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! Get of out of your office and watch the work.

Root Causes

A good investigation will often shake loose a variety of preconceptions and therefore lead to significant gains. That said, as you analyze the results of your investigation, your automatic processing functions will still be mapping those observations into past experiences in ways that will be consistent with your existing beliefs, thus making it difficult to find new solutions. To help offset this tendency, the next step in filling out the A3 is to analyze root causes. The idea here is simply to engage conscious processing by constructing the next step in the logical chain and 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, famed loom inventor and founder of Toyota industries, suggested asking the "5 whys", meaning that for each observed problem, the investigator should ask "why" five times—why won't the car start, the battery is dead; why is the battery dead? The alternator isn't charging it; why is the alternator not working?[24] etc. — in the hope that five levels of inquiry will reveal the fundamental 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.[25] Since then just about all structured problem solving methods offer one or more variants of the same basic method for digging into a problem's source.

The various tools for doing a root cause analysis are important, but the emphasis on a specific approach often gets in the way of understanding *why* a tool is needed. 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*. In other words, a proper root cause analysis shows how the problem is generated by normal operations rather than special, one-time events. Unfortunately, this type of 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. Retraining the person or

[24]. Ohno, *Toyota Production System*.
[25] "Fishbone Diagram (Ishikawa) - Cause & Effect Diagram | ASQ."

replacing the bracket will solve the immediate problem but will do nothing to prevent future manifestations unless we address the system-level cause.

The upshot of this is that a good root cause analysis should build on your investigation to show how the work system you are analyzing generates that 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, then you probably 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 and/or inconsistent training process for those in customer facing jobs or an inconsistent customer on-boarding process, either of which will continue to generate a series of seemingly unrelated service problems. All too often, we see root cause analyses that end with special cases that will likely never happen again. If the problem is important enough to be worth your attention as a senior manager, then chances are it is not a special case. 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

With a solid analysis in hand, structured problem solving moves to action. In some sense, the *Target Design* section is just a mirror image of the previous root cause analysis. Having linked features of the work system to the problem you are trying to solve, now propose an updated system that won't generate (or will at least generate less of) the problem. Often the necessary changes will be simple.[26] 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

[26] In other work we have proposed four principles for effective work that may be helpful in more complex situations. See Dodge, S. et al., "Speeding the Search for a Cure: Using Dynamic Work Design to Improve Genetic Sequencing."

from now on we will print the General Ledger code on the invoice form or something more complicated such as changes to the employee training or supplier qualification programs. Note that in these cases the changes will rarely be an entirely new program or initiative. Instead, they should be specific, targeted modifications to the existing system that are built on the root cause analysis. One of the biggest benefits of good problem formulation and structured problem solving is that they enable focused and effective interventions. And, as we mentioned earlier, you will usually learn more and make more progress by doing multiple, fast improvement cycles. Don't try to solve everything at once. Instead, propose the minimum set of changes that can help you make rapid progress towards your goal.

Goals and Leadership Guidelines

Completing the Target Design section requires two additional components. First, make a prediction for how much improvement your proposed changes will make in the form of an *Improvement Goal*. 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. Thus, if your problem was "24% of our service interactions do not generate a positive response from our customers, greatly exceeding our target of 5% or less," then a target statement might be, "reduce the number of negatively-related service interactions by 50% in 60 days." The benefit of the clear target statement is twofold, clear goals are highly motivating, and a good goal statement constitutes a clear prediction and thereby enables effective learning.

Finally, set the *Leadership Guidelines*. Guidelines are the "rules of the road" or "guard rails" for executing the project, boundaries or constraints that cannot be violated. For example, a project focused on cost reduction could easily achieve its target by sacrificing quality. Such an intervention, however, doesn't represent the generation of any new knowledge or competitive advantage. Instead, a structured inquiry should identify an innovation that reduces cost without making other trade-offs. The guideline section exists to rule out those obvious trade-offs and focus people on more fundamental innovation.

Execution Plan

The final step is running the experiment and seeing what you learned. In the upper portion of the box, you should lay out a plan for implementing your proposed design. This is probably the most familiar part of the A3 since it's just a standard project plan, something you probably have done dozens of times. Be sure that the plan is broken into a set of clear and distinct activities (e.g., have invoice form reprinted with the GL code or hold a daily meeting to review quality issues) and that each activity has both an owner and a clear delivery date.

Now execute your plan and meet your target! But, even as you start executing, you are not done with structured problem solving. 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 very simple: the owner of the action should report whether that element is ahead 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 track your progress to your target. 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. How much progress do you expect to make in the first week, in the second, and so on? Put these intermediate targets in the first column of the Track Results section and then track 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, not ones that leak away as soon as the project is over.

Once the project is complete, document what you learned in the *What Did We Learn and What Do We Do 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 (not uncommon in early projects), 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 might have thought. Finally, and perhaps most importantly, what are you going to do next? Experienced leaders will tell you that a well-functioning process, whether it be manufacturing, service or new product development, is the product of numerous small changes. Toyota has been at this for decades and still manages to find new improvements every year. Perhaps not surprisingly fixing one real problem typically reveals three more pressing issues (often because this is the first time you have really looked at how to work in your organization is actually done). Thus, close out your A3 by outlining your next steps. What is the next problem you and your organization need to solve?

Two Case Studies: Making Boxes and Treating Patients

Transforming organizations is a contact sport and you can't do it sitting in your office. Though they are busy like everybody else, effective leaders find ways to connect with the fundamental work processes in their organizations. And, as with any real activity, there are often gaps between theory and reality and filling these gaps takes experience born out of practice. In what follows we describe two recent cases of using good problem formulation and structured problem solving to improve organizational performance. As a testament to the power of the framework, these results were not generated by an expensive consulting project staffed by experts with extensive training or Ninja-like process improvement skill. Instead, both projects were done by practicing managers with active day jobs who were introduced to the basic of problem formulation and structured problem solving in a course here at MIT.

Making Boxes

Michael Morales is the President of Corporacion Industrial, S.A., a company in Panama that makes corrugated boxes. Conceptually, making corrugated boxes is relatively straightforward: wood is first turned into pulp, which is then treated, pressed and dried to form large sheets of paper; those sheets are then glued together to create flat pieces of corrugated cardboard; and, finally, the cardboard is cut to size, folded and glued to make a finished box. In practice, however, it entails moving heavy loads among large, complex machines and executing a series of precision operations. Large rolls of paper must first be moved from the loading dock to a staging area and then fed into a cutting press. As with all complex industrial processes, it was designed carefully at the outset. Nonetheless, it was never as efficient as its owners had hoped. Prior to the project, Mike and his team had concluded the biggest shortfall was due to unplanned paper losses, paper that was damaged during movement between operations or lost due to incorrect cutting. Paper represented over 80% of the production costs and over 60% of the total cost, so any losses could have a significant effect on the bottom line. Prior to 2012, paper losses had been as high as 18% of the total production. In 2012 the company did a significant change initiative focused on reducing paper losses that included spending $3.5 million on new equipment and another $.5 million in training. Despite the effort and expense, paper losses hit a record high of 21.1% in 2012.

In the summer of 2014, Mike chose this area as the focus for the project he was required to do in our course (the completed A3 is shown in figure three). The first step was to define a clear problem statement. Mike and his team knew that the paper losses were too high, but they had never quite taken the time develop a clear and agreed upon formulation. After several iterations, the problem statement became:

> **Our paper losses are 5% higher than industry average, this negatively affects our financial position and ability to be competitive.**

Supported by the following background statement:

> **Paper represents nearly 88% of our production cost and 65% of our total cost. In the years leading up to 2011, our paper losses were as high as 18%. In 2012 we**

invested more than $3.5 million in PPE and an additional $.5 million in training but even so, our losses increased to a record of 21.1% in 2014.

The problem statement clearly highlights the gap between the desired and actual state (though we might prefer those be explicitly stated) and the background does an excellent job of indicating why the problem is important.

The A3 document shows the results of Mike's visit to the factory floor to watch the work understand the current design. Numerous departures from the desired state were quickly revealed: The paper was often too wide, resulting in extra losses from cutting; paper rolls were often damaged by the forklifts that moved them, and various machines were not properly calibrated. Perhaps most notably, while Mike was observing, the main corrugator machine stopped at 11:30am. Assuming it was an unplanned outage, Mike rushed to the machine only to learn that it was stopped *every day* at lunch, both reducing productivity and leading to thermal cycling that increased wear and tear on the machine. Interestingly, the lunch break turned out to be a legacy response to instability in the power provided by the local utility, a problem that had long since been fixed.

Such seemingly obvious problems might lead one to think that the plant was poorly run. To the contrary, this is often exactly what happens when a piece of work is regularly executed by a dedicated and experienced staff. With time and experience, our automatic processing functions take control of our day-to-day work. The benefit is increased efficiency, often significantly so, but those gains comes at the cost of not noticing all of the little work arounds and accommodations that we all make to get things done. We have discovered similarly "obvious" issues in almost every piece of work we have ever studied and it is the rare organization that has used structured problem solving with sufficient intensity to exercise all the easy improvement opportunities.

With the Current Design in hand, it was straightforward to identify several root causes, mostly due to the plant operators and technicians not understanding the importance and cost associated with paper losses and how their actions influenced them. Mike and his team proposed several improvements that followed directly from the observations and the root cause analysis: the forklift was retrofitted and holes in the floor were repaired to prevent damaging the paper in transit; several standard procedures were either created or updated to improve the functioning of the equipment; a camera was installed to give the operators earlier feedback; and, perhaps, most importantly, the lunch schedule was staggered so the corrugator machine could run for the entire shift. Based on these interventions, Mike and his team thought they could reduce paper losses by 2 percentage points in 60 days (note the clear target statement and the short time frame), which would save the company about $120,000 on an annual basis. Based on the proposal, an execution plan was quickly developed with dates and assignments.

The results exceeded the target significantly. Losses fell from 21.4% to 15.45% in the first month and to 14.7% in the second month, generating over $50,000 savings in *less than 60 days*. And the project produced several collateral benefits, including reduced overtime, fuel use, and power consumption. When we asked Mike to explain the success of the project (and why it took a class project to generate it) he replied with the following email:

> *Believe it or not Nelson, our return to profitability as a company lies in our paper losses. It's just that simple. If we can work through some of these issues (others have come out as well), we'll make it through the storm.*
>
> *I've been in the industry all of my professional career and consider myself a hands on type of manager. That said, it took this process to get off my ass and actually go see and talk with our operators to understand what was going on. Funny thing is, they already knew what the problem was, we just weren't listening.*

We couldn't provide a better summary of the benefits of good problem formulation and structured problem solving. Formulating a clear gap between the target and actual state focuses attention on the things that matter. Tackling those gaps in a structured manner provides the means to engage both of our cognitive strengths, instincts, and more rational thought and the knowledge of others.

Define the Problem:

Our paper losses are at least 5% higher than industry average and this is negatively affecting our financial position and our ability to be competitive.

Background: Paper represents nearly 88% of our production cost and 65% of our total cost. In the years leading up to 2011, our total paper losses were as high as 18%. In 2012 we invested more than USD 3.5M in PPE and an additional 0.5M in training but our losses increased to a record high of 21.1% in 2014.

Current Design:

- Avg. yearly paper consumption is 8535 tons w/ an avg. cost of $658 / Ton (2014)
- Our controllable losses have increased by 42.3% with respect to 2013.
- Corrugator section of the plant represented 54.4% of total paper losses, represents cash loss of $270K.
- Longitudinal trim can be as much as 2 inches greater than required due to suboptimal paper roll width.
- Paper rolls are damaged while being moved within the plant which result in 4.4% of total losses.
- Corrugator line shuts down due to air pressure drop when demand peaks 2 to 4 times per week.
- Due to manual calibration of the preheaters, single facers are not setup properly which generates up to 5% of paper losses.
- No "best way" to regulate steam return to boiler
- Operator has limited visibility to detect paper misalign and cannot quickly detect problems.
- Corrugator is stopped between 11:30AM and 12:00PM (lunch) which results in paper waste in drier / plate section. Restarting line after cooling down for 30 min increases the probability of mechanical problems and breakage of the paper due to increased tension.

Root Cause:

- High Longitudinal trim due to production runs based on traditional width paper rolls.
- Damage to paper due to inadequate spacing between rolls. Forklift operator lack proper training and do not understand the consequences and effect of paper losses in this area.
- New dry-end equipment shuts down automatically when pressure is below 88psi.
- Valve assemblies in steam section are not properly labeled and no clear procedure exist.
- Visibility is reduced due to distance and because other machinery / equipment in the way

Target Design:

- Change planning and rearrange personnel to eliminate stoppage of corrugator during lunch
- Reorganize paper inv.. **Fix forklift clamps and patch holes in floor to avoid damage to rolls**
- Increase line pressure to compensate for pressure drop and relocated air buffer tank closer to high use areas. **Not necessary at this point. Higher line pressure seems to be working.**
- Label steam valve assembly and produce "best way" procedure, reduce errors. Improve TQC
- Install closed circuit cameras and screen at operator station to detect problems sooner.
- **Improvement Goal: Reduce losses by 2% within 60 days (Annual savings >$120K**

Guidelines:

- Limited CAPEX to 10K.
- Progress measured daily with managers and weekly with all involved (include operators)
- Involve managers, supervisors and operators

Execution Plan (Follow up – 03/11/14):

Task	Responsible	Date	Target Red.	Status
Redesign production runs (update)	E. Dorcy (Planning)	15/10/14	0.2%	No longer an action item
Eleminate line stoppage during lunch	R. Farugia (OPS)	7/10/2014	0.2%	Completed
Increase line pressure (+15psi)	G. De Gracia (Eng.)	30/09/14	0.1%	Completed
Move air buffer tank	G. De Gracia (Eng.)	10/10/2014	0.1%	Completed
Label steam valves (still pending)	G. De Gracia (Eng.)	24/10/14	0.2%	Delayed
Rework steam valve oper. procedure	**R. Farugia / G. DG**	**23/10/2014**	**0.1%**	**Reschedule completion date**
Install (6) cameras and screens (2)	G. De Gracia (Eng.)	7/10/2014	0.2%	Completed
Repair Forklift clamps and floor	**R. Farugia (OPS)**	**18/10/2014**	**1.0%**	Completed
			2.1%	

Track Results:

Production results for October and November showed extraordinary progress. Our overall paper losses decreased by 6%, from an average of 21.41% (Jan–Sept 14) to 15.45% in October and 14.7% in November which provided savings during a 2 month period of $56,274. The reduction of controllable losses by 50% during this period was a major driver. We are aligned with our objectives and have seen improvements in other areas. Bunker fuel and power consumption has decreased by 30% and 10% respectfully and overtime spending is down 70%.

What Did We Learn & What's Next?

Figure 3. Completed A3 for the Paper Loss Reduction Project.

Seeing More Patients

Todd Astor transplants human lungs. A lung transplant is a complicated undertaking on many dimensions: a candidate with end-stage lung disease is selected after a comprehensive evaluation; the candidate and potential donor need to be matched; the donor's lungs need to be explanted and then quickly moved to the hospital in which the recipient resides; and then there is the complicated surgical procedure required to remove the recipient's existing lungs and replace them with the new ones. All of this complexity, however, pales in comparison to managing the recipient's health after the transplant. The human body was not made to operate with someone else's organ(s) and often responds to the transplanted organs in dangerous ways. Consequently, a big part of Todd's job is staying in close contact with his patients and carefully managing the complicated suite of medicines they take to suppress the body's natural immune response. Staying in touch with patients requires physically seeing and examining them, and that's where Todd's project started.

Several times a week Todd has a "clinic" in which transplant recipients come to be evaluated and receive any necessary adjustments in their treatment. Each clinic session lasts for 3 hours and utilizes three dedicated examination rooms. Based on the evaluation criteria of Todd's hospital, that should allow him to see 27 patients (three per hour/room). At the outset of Todd's project, the average number was 7. Being less than 30% utilized both 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:

> **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 min/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.**

An added the additional background:

> **The post-lung transplant outpatient clinic serves as the venue for monitoring lung transplant recipients, and managing their lung allograft and non-allograft related complications. The clinic visit includes a comprehensive evaluation that includes laboratory testing, x-rays, pulmonary function testing, a nurse practitioner encounter, and a transplant physician examination. This clinic is a vital component to the successful care of the lung transplant recipients, and is relied upon to generate a significant component of the overall revenue for the lung transplant program.**

Todd's problem formulation is exemplary: both the gap between the desired and the actual state and the importance of the problem are crystal clear.

Problem statement in hand, Todd went to understand the work. He tracked 71 patients over nine sessions as they flowed through the clinic day (the A3 in figure 4 shows a basic process map). Much like Mike's trip to his factory floor, Todd discovered huge variability in both the patient arrival rates and the time that they spent in the various stages of a clinic visit. Just 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 to both a lack of clarity on appointment details and 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 sometimes spent on checking the patient's medication list.

Define the Problem

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 min/patient) per session. This "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 hospital.

Background

The post-lung transplant outpatient clinic serves as the venue for monitoring lung transplant recipients, and managing their lung allograft and non-allograft related complications. The clinic visit includes a comprehensive evaluation that includes laboratory testing, x-rays, pulmonary function testing, a nurse practitioner encounter, and a transplant physician examination. This clinic is a vital component to the successful care of the lung transplant recipients, and is relied upon to generate a significant component of the overall revenue for the lung transplant program.

Current Design

I followed patient flow for 71 patients over 9 clinics sessions and recorded process times (Ave and STDEV values are in minutes):

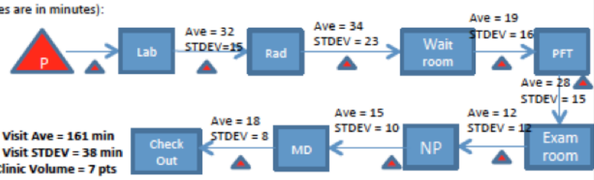

Total Visit Ave = 161 min
Total Visit STDEV = 38 min
Ave Clinic Volume = 7 pts

1)Patient arrival times were extremely variable and majority of patients did not arrive to hospital at scheduled time.
2)Laboratory wait time was shorter and less variable for the patients arriving earlier in the morning.
3)The earlier half of the clinic period runs much more efficiently than the later half.
4)There is a significant amount of time variability for the PFT (pulmonary function testing) process.
5)There is a significant amount of time variability for the NP (nurse practitioner) and MD (physician) evaluation; the average total time (NP+MD =45 min) exceeds the allotted 20 min exam room time.
6)There is a significant amount of variability in the total visit time (+/-38 minutes).

Identify Root Causes

1)Variability in patients arriving on-time due to lack of clarity regarding scheduled appointment time and unpredictable traffic patterns.
2)Variability in laboratory (and radiology) times due to early versus late morning arrival time for patients to the lab.
3)Improved efficiency in the earlier versus later half of the clinic due to patients arriving from lab and radiology on time (related to #2).
4)Variability in PFT times and PFT wait room times due to partial versus complete PFT testing (partial PFT = spirometry without lung volume testing, which in most cases is sufficient and takes less than half the time).
5)Variability in exam room and NP/MD evaluation times due to lack of a coordinated system of communication between clinic admin, NP, and MD.
6)NP time is largely spent on clarification and reconciliation of patient medication list.
7)Variability in total visit time due to combination of #1-#6.
8)"Gap" between average clinic volume (7 patients) and "ideal" number of potential clinic slots (27) due to #1-#6, but (27) 20min clinic slots may be unrealistic target.

Target Design

1)Administrative assistant could call patients 1 week and 1 day before clinic to remind patients of appointment, arrival time to hospital, and variability in traffic.
2)Patients could be scheduled in 2 shifts – a larger early shift that will arrive at the lab at 6:30 am, and a smaller, later shift that will arrive at lab at 8am.
3)All PFTs could be ordered as spirometry testing alone. If complete PFTs are required, they could be performed at a later time/date **(reduction of 15-20 min/patient)**.
4)A time limit for the exam room portion of the visit could be set, less than the current 45 min average, but longer than the current suggested 20min time slot - suggest 30 minutes **(reduction of 15 min/ patient)**.
5)NP could either directly or indirectly contact patients day before clinic visit to review medication list **(reduction of 5-10 min/patient)**.
6)NP could wait until MD in room to ask similar clinical questions to avoid redundancy **(reduction of 5 min/patient).**
7)Gradual addition of increased patient volume as changes are implemented, with a goal of adding 2 additional patients/clinic every 2 weeks for an end goal of an average of 18 patients/clinic session.

Guidelines

1)Listening to the patients is very important – they are an essential partner in this proposed plan, and their feedback is vital.
2)Quality of patient care must not be sacrificed in trying to improve efficiency – continue to ensure that enough time and attention is being devoted to each patient.
3)Communication between team members is essential for the clinic to run efficiently.

Execution Plan

	Responsibility	Due Date	Target
1)Patient phone reminder about appointment	M.H.	9/29/2014	100% patients
2)Scheduling patients in 2 timed shifts for lab	M.H./E.R.	10/6/2014	100% patients
3)Changing PFT ordering method to spirometry alone	M.H./E.R.	10/13/2014	90% patients
4)Imposing time limit of 30 min for exam room	E.R./J.P.	10/6/2014	> 75% patients
5)Contact patient prior to visit for medication review	E.R./J.P.	10/6/2014	> 75% patients
6)Combined NP/MD clinical questioning	E.R./J.P./T.A./I.N.	9/29/2014	100% patients
7)Increasing clinic patient volume by 2 every 2 weeks	M.H.	10/6/2014	13 patients
		10/20/2014	15
patients			
		11/3/2014	17
patients			
		11/17/2014	19
patients			

Validate & Sustain Results

Figure 4: A3 for Patient Flow Project

Todd made two important decisions in his approach to the root cause analysis and proposal of changes. First, despite variability being identified at all stages of the visit, he narrowed the focus to the processes occurring in the clinic area (i.e. scoping down the problem). He and his team had more direct control over these processes (compared to those occurring in the laboratory, radiology area, etc.), and would, therefore, be more likely to make rapid and impactful change. Second, Todd included every member of the team, from the administrative staff to the physicians, in constructing the root cause analysis and proposing changes. This approach afforded the opportunity for every individual to think about the problem personally and concentrate on specific ways that he/she could address the problem in his/her own assigned area.

As a result, every staff member became more invested in the successful outcome of the project.

The root cause analysis led to several proposed changes. The clinic administrative assistant would call patients both a week and a day in advance to remind them about their appointments and to provide advice on managing traffic and parking (a significant challenge in Boston). 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 redundant information. With all of these changes, Todd set a target of adding two patients per clinic session until they 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.

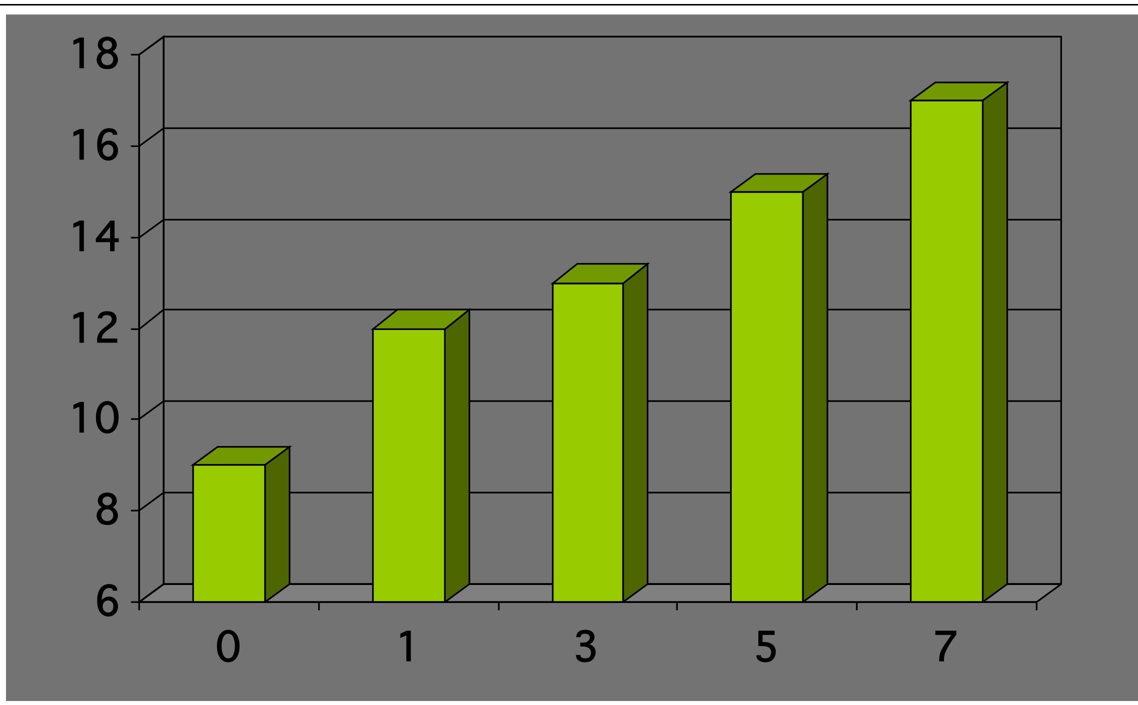

Figure 5. Weekly Patient Flow for Lung Transplant Clinic.

The results are shown in figure five. During the seven weeks that remained in the course, the throughput moved from the average of seven to a high of 17 in week seven, not quite meeting Todd's target to 18, but more than doubling existing the patient flow. The increased throughput had several positive benefits. Most importantly, 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, improving the throughput of the clinic created space to accommodate more transplant 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.

Summary

As we said in the introduction, Mr. Oba's visit to Don's plant did not end well. Don was frustrated because nobody would tell him about how to implement the Toyota production system. Mr. Oba was frustrated with an American manager who wouldn't listen. But despite the difficult start, the story has a happy ending. It took a while for us to figure it out, but Mr. Oba was teaching Don the most basic principle of the Toyota production system, problem formulation. Since then, we have come to believe that good problem formulation is not just the foundation of running an efficient factory, it is the basis of being an effective leader in *any* domain. Following Mr. Oba's visit, Don eventually mastered the art of good problem formulation and used structured problem solving to great effect, creating millions of dollars of new value for Harley-Davidson. Since then, we have coached dozens of intervention projects in domains ranging from drilling oil wells[27] to sequencing the human genome[28], and significant gains are often revealed with just a good problem formulation and a thoroughly understanding the way the work is

---

[27] Mangi and Repenning, "Dynamic Work Design Decreases Post-Procedural Length of Stay and Enhances Bed Availability."
[28] Dodge, S. et al., "Speeding the Search for a Cure: Using Dynamic Work Design to Improve Genetic Sequencing."

actually done. We have also used this material to launch almost a thousand student projects with similar impact.

It's hard to pick up a current business publication without reading about the imperative that every organization has to change. The world, this line of argument suggests, is evolving at an ever-faster rate and those who do not adapt will be left behind. As one piece of anecdotal evidence of the extent to which people subscribe to this idea, we always ask our MIT Sloan classes "how many of your companies re-organize every eighteen to twenty-four months?" and more than half of the class always raises their hands. Left silent in these arguments though is the question of which organizations drive the changes to which everyone else must adapt. Many of these come from the much-vaunted start-ups that often "disrupt" markets with new, innovative technologies. But, such changes are also often driven by perennial industry leaders, many of who have a remarkable track record of combining existing performance with future innovation. Several studies suggest that even the most established industries often show dramatic performance differentials between the top players and those who make up the middle.[29]

Academic research suggests that the ability to manage change plays a big role in separating the leaders from the rest of the pack and our studies clearly show that it is easier to manage a sequence of bite-sized changes instead of one huge re-organization or "initiative." And it's not just the mighty Toyota that has turned those little changes into a strong competitive position, Southwest, Alcoa, CostCo and several others have made similar moves in their industries. Good problem formulation and structured problem solving offer a sustainable alternative to the endless stream of painful reorganizations and overblown initiatives that rarely deliver on their promise. And if this seems hard to believe, run your own experiment. We have all been to lots of bad meetings. Next time you find yourself facing a seemingly endless, one-way stream of

[29] For recent review, see: Gibbons and Henderson, *What Do Managers Do?*

information occasionally interspersed by a circuitous conversation that never seems to go anywhere, try Mr. Oba's question: Ask the presenter or the meeting leader, "What problem are we trying to solve?" In our experience, this question either leads to dramatic improvements for the whole team or being invited to fewer meetings, which at least makes our days more productive.

References

Amabile, Teresa, and Steven Kramer. *The Progress Principle : Using Small Wins to Ignite Joy, Engagement, and Creativity at Work*. Boston, Mass. : Harvard Business Review Press, c2011., 2011.

Amabile, Teresa M., and Steven J. Kramer. "The Power of Small Wins." *Harvard Business Review* 89, no. 5 (May 2011): 70–80.

Bauer, Isabelle M., and Roy F. Baumeister. "Self-Regulatory Strength." *Handbook of Self-Regulation: Research, Theory, and Applications*, 2011, 64–82.

Dijksterhuis, Ap. "Automaticity and the Unconscious." In *Handbook of Social Psychology*. John Wiley & Sons, Inc., 2010. http://onlinelibrary.wiley.com/doi/10.1002/9780470561119.socpsy001007/abstract.

Dijksterhuis, Ap, and Loran F. Nordgren. "A Theory of Unconscious Thought." *Perspectives on Psychological Science* 1, no. 2 (June 1, 2006): 95–109. doi:10.1111/j.1745-6916.2006.00007.x.

Dodge, S., T Desment, J Meldrim, N Lennon, D Perrin, S Ferriera, Z Leber, et al. "Speeding the Search for a Cure: Using Dynamic Work Design to Improve Genetic Sequencing." *MIT Sloan Working Papers* 5159–16 (2016). http://mitsloan.mit.edu/shared/ods/documents/Repenning_Cancer_full.pdf&PubID=15032.

Evans, Jonathan St B. T., and Keith E. Stanovich. "Dual-Process Theories of Higher Cognition Advancing the Debate." *Perspectives on Psychological Science* 8, no. 3 (May 1, 2013): 223–41. doi:10.1177/1745691612460685.

"Fishbone Diagram (Ishikawa) - Cause & Effect Diagram | ASQ." Accessed August 12, 2016. http://asq.org/learn-about-quality/cause-analysis-tools/overview/fishbone.html.

Forrester, Jay W. *Industrial Dynamics*. Cambridge, Mass.: M.I.T. Press, 1961.

Gibbons, Robert, and Rebecca Henderson. *What Do Managers Do?: Exploring Persistent Performance Differences Among Seemingly Similar Enterprises*. Harvard Business School, 2012. http://www.hbs.edu/faculty/Publication%20Files/13-020.pdf.

Haidt, Jonathan. *The Happiness Hypothesis : Finding Modern Truth in Ancient Wisdom*. New York : Basic Books, c2006., 2006.

Kahneman, Daniel. *Thinking, Fast and Slow*. New York : Farrar, Straus and Giroux, 2011., 2011.

Klein, Gary A. *Sources of Power : How People Make Decisions*. Cambridge, Mass. : MIT Press, 1999, c1998., 1999. http://libproxy.mit.edu/login?url=https://search.ebscohost.com/login.aspx?direct=true&AuthType=cookie,sso,ip,uid&db=cat00916a&AN=mit.001496552&site=eds-live.

Kotter, John P. *Leading Change*. Boston, Mass: Harvard Business School Press, 1996.

Locke, Edwin A., and Gary P. Latham. "Building a Practically Useful Theory of Goal Setting and Task Motivation." *American Psychologist* 57, no. 9 (September 2002): 705.

Mangi, A., and N. Repenning. "Dynamic Work Design Decreases Post-Procedural Length of Stay and Enhances Bed Availability.," Working paper available from authors, 2016.

Oettingen, Gabriele, Gaby Hönig, and Peter M. Gollwitzer. "Effective Self-Regulation of Goal Attainment." *International Journal of Educational Research* 33, no. 7–8 (2000): 705–32. doi:10.1016/S0883-0355(00)00046-X.

Ohno, Taiichi. *Toyota Production System: Beyond Large-Scale Production*. Cambridge, Mass: Productivity Press, 1988.

Park, H, and L Reder. “Moses Illusion: Implication for Human Cognition.” In *Cognitive Illusions : A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*. New York, NY : Psychology Press, 2004., 2004.
Perrow, Charles. *Normal Accidents: Living with High-Risk Technologies*. Princeton Paperbacks. Princeton, N.J: Princeton University Press, 1999.
Reid, Peter C. *Well Made in America: Lessons from Harley-Davidson on Being the Best*. New York: McGraw-Hill, 1990.
Sherman, Jeffrey W., Bertram Gawronski, and Yaacov Trope. *Dual-Process Theories of the Social Mind*. New York : The Guilford Press, [2014], 2014.
Shook, John. “Toyota’s Secret: The A3 Report.” *MIT Sloan Management Review* 50, no. 4 (Summer 2009): 30–33.
Singh, Jasjit, and Lee Fleming. “Lone Inventors as Sources of Breakthroughs: Myth or Reality?” *Management Science* 56, no. 1 (October 16, 2009): 41–56. doi:10.1287/mnsc.1090.1072.
Sloman, Steven A. “Two Systems of Reasoning, an Update.” *Dual-Process Theories of the Social Mind*, 2014, 107–120.
Stanovich, Keith E. *Rationality and the Reflective Mind*. New York: Oxford University Press, 2011.
Sterman, John. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston: Irwin/McGraw-Hill, 2000.
Weick, Karl E. “Small Wins: Redefining the Scale of Social Problems.” *American Psychologist* 39 (January 1984): 40–49. doi:10.1037/0003-066X.39.1.40.
Wiesendanger, Betsy. “Deming’s Luster Dims at Florida Power & Light.” *The Journal of Business Strategy* 14, no. 5 (October 1993): 60.