Introduction

I. What happened after 1983

I wrote this book in 1983. It was the result of twenty years of managerial work during which I learned a variety of ways to make things take place more effectively. What I learned were the basics of managerial work, particularly as they pertained to middle managers. More than a decade has passed since, but I find that most of the things that were useful then are still useful now; the basics of management remain largely unaffected.

However, two critical events took place in the 1980s that altered the environment in which we managers do our work—and this made me realize that an updated Introduction to this book was necessary. Those events were the Japanese memory onslaught and e-mail.

Let me explain their implications.

By the mid-eighties, the Japanese producers of Dynamic Random Access Memories, or DRAMs for short—the most popular computer memory devices, used in computers of all kinds—had perfected their technological capability and honed their manufacturing prowess to the extent that they could take on the American producers (who had pioneered the market and totally dominated it for the first fifteen years of its existence). The mid-eighties were also when the personal computer revolution took place. And because personal computers require a lot of memory, the Japanese DRAM juggernaut had a ready market for its products centered in the United States. Everything was primed for an attack.

Intel, where I work, was one of the companies that got caught up in this assault. In fact, Intel was one of the early producers of DRAMs. More than that, in its earliest years, we had practically the

whole market to ourselves. However, by the mid-eighties, competition both from the United States and, increasingly, from Japanese manufacturers whittled down our share of the market. Under the ferocious attack of aggressively priced, high-quality Japanese DRAMs, we were forced to retreat and cut prices to a level where being in the DRAM business brought us major losses. Ultimately, the losses forced us to do something extraordinarily difficult: to back out of the business that the company was founded upon, and to focus on another business that we thought we were best at—the microprocessor business.

While this adjustment sounds quite logical and straightforward in theory, in reality its implementation required us to move and redeploy a lot of our employees, let some of them go, and shutter a number of factories. We did all this because under this strong attack, we learned that we must lead with our strength. Being second best in a tough environment is just not good enough.

Ultimately, we—Intel and the U.S. semiconductor industry—prevailed over the onslaught of the Japanese manufacturers. Intel grew to become the largest semiconductor manufacturer in the world, and U.S. manufacturers recently surpassed their Japanese counterparts overall. Nevertheless, in retrospect it's clear that this assault was just one wave of a much larger tide—the tide of globalization.

Globalization simply means that business knows no national boundaries. Capital and work—your work and your counterparts' work—can go anywhere on earth and do a job.

Some of us are fortunate to be residents of a country, the United States, that enjoys one of the highest standards of living. The U.S. market for goods and services is the largest in the world. And until recently, it has been easier to supply that market from inside the United States than from abroad.

Today, many markets outside the United States are growing faster than markets inside the U.S. And the domestic market can be supplied from anywhere in the world. For example, I recently bought a Gore-tex jacket from Patagonia (the clothing manufacturer, not the

region in South America), and I saw that it was made in China: American brand, American technology (the high-tech fabric was invented and made in the United States), and assembled to the specifications of the reseller (Patagonia) in a foreign country.

The consequence of all this is very simple. If the world operates as one big market, every employee will compete with every person anywhere in the world who is capable of doing the same job. There are a lot of them, and many of them are very hungry.

Another consequence also follows. When products and services become largely indistinguishable from each other, all there is by the way of competitive advantage is *time*. And that's where the second critical development of the eighties comes in—e-mail.

Just as the Japanese DRAM attack was the first wave of a much greater tide, e-mail is also the first manifestation of a *revolution in how information flows and how it is managed*.

The informed use of e-mail—short for computer-to-computer electronic messaging—results in two fundamentally simple but startling implications. It turns days into minutes, and the originator of a message can reach dozens or more of his or her co-workers with the same effort it takes to reach just one. As a result, if your organization uses e-mail, a lot more people know what's going on in your business than did before, and they know it a lot faster than they used to.

Let me interject a bit of irony. Back in the eighties, when the Japanese seemed invincible, one explanation advanced for their ability to act quickly and decisively was the way Japanese offices were set up. In a Japanese office, a manager and his subordinates all sit around a big long table. People work on their own assignments but when they need to exchange information, everybody they work with sits within speaking distance, right around the same table. So information is exchanged in minutes and everybody can be reached with the same effort. As a result, because of the ease with which Japanese office workers communicate, they have, in fact, been slow to embrace electronic mail.

But now the pendulum is swinging in the other direction. As businesses become more widely spread out around the globe and as time becomes the key competitive weapon, American organizations are often better positioned than their Japanese counterparts. Why? Because the same ease of communication that prevailed by natural means in the Japanese office now effectively travels around the world through electronic means.

And e-mail is only the first wave. Everything today is going to a digital format: sound, photos, movies, books, financial services. And everything that's digital can be shipped around the world just as fast as it can be shipped down the hall at your workplace.

Here is an interesting illustration of the consequence of such a capability. I am told that the post office sorts 90 percent of all letters automatically. For the 10 percent the machines can't decipher, a human reader types the addresses into a machine. Recently, to lower the cost of this work, the postal service tried a new system. A machine takes a digital photograph of the illegible envelopes, instantly ships the digital image to a lower-labor-cost region where someone reads and keys in the address from the digital image, then electronically ships the address back to the regional postal center. This is the beginning of a trend that will become all-encompassing in the next twenty-five years.

Simply put, the information revolution does away with hiding places anywhere, in any line of work. So the questions are: What are businesses to do and what are managers themselves to do?

II. Operating in the new environment

Let's back off for a moment and consider whom this book is aimed at. I am especially eager to reach the middle manager, the usually forgotten man or woman of any organization. The first-line supervisor on the shop floor and the chief executive officer of a company are both well appreciated. You'll find many courses designed to teach the former the fundamentals of his work, while practically all of our leading business schools are set up to turn out

the latter. But between the two is a large group of people—the middle managers, who supervise the shop-floor foremen, or who work as engineers, accountants, and sales representatives. Middle managers are the muscle and bone of every sizable organization, no matter how loose or "flattened" the hierarchy, but they are largely ignored despite their immense importance to our society and economy.

Middle managers are not confined to big corporations. In fact, they can be found in almost any business operation. If you run a small tax department at a law firm, you are a middle manager. The same is true if you are a school principal, an owner of a distributorship, or a small-town insurance agent. When people from each of these enterprises read the original manuscript of this book, their reactions confirmed what I suspected: the managerial ideas that were developed at Intel as it grew from a very small to a very large organization were broadly applicable.

Another group should also be included among middle managers—people who may not supervise anyone directly but who even without strict organizational authority affect and influence the work of others. These *know-how managers* are sources of knowledge, skills, and understanding to people around them in an organization. They are specialists and experts of some sort who act as consultants to other members of the organization; they are, in effect, nodes in a loosely defined network of information. Teachers, market researchers, computer mavens, and traffic engineers shape the work of others through their know-how just as much as or more than the traditional manager using supervisory authority. Thus a know-how manager can legitimately be called a middle manager. In fact, as our world becomes ever more information- and service-oriented, know-how managers will acquire greater importance as members of middle management. In short, know-how managers should also read on.

Whether you are a know-how manager or a traditional manager, your company has no choice but to operate in an environment shaped by the forces of globalization and the information revolution. Companies today basically have two choices: Adapt or die. Some have died in front of our eyes; others are struggling with the

adaptation. As they struggle, the methods of doing business that worked very well for them for decades are becoming history. Companies that have had generations of employees growing up under a no-layoff policy are now dumping ten thousand people at a time onto the street. Unfortunately, that's all part of the process of adaptation.

All managers in such companies need to adapt to the new environment. What are the rules of the new environment? First, everything happens faster. Second, anything that can be done will be done, if not by you, then by someone else. Let there be no misunderstanding: These changes lead to a less kind, less gentle, and less predictable workplace.

Again, as a manager in such a workplace, you need to develop *a higher tolerance for disorder*. Now, you should still not accept disorder. In fact, you should do your best to drive what's around you *to* order. The breakfast factory metaphor of this book—the idea that you should run your managerial processes like a well-oiled factory—is every bit as much the ideal now as it was when I wrote this book. But you as a manager need to be mentally and emotionally ready to be tossed into the turbulence generated by a mega-merger that takes place in your industry—perhaps in this country, perhaps on the other side of the globe. You should be prepared for the shockwaves engendered by a brand-new technique pioneered by someone you had never even heard of before.

You need to try to do the impossible, to anticipate the unexpected. And when the unexpected happens, you should double your efforts to make order from the disorder it creates in your life. The motto I'm advocating is *"Let chaos reign, then rein in chaos."*

Now, I'm sure that at various times you will take exception to what you read in this book. "This may be fine at Intel," you will say, "but it would never fly at PDQ, where I work. Nothing does until the Old Man himself decrees it. Short of a palace revolution, I can't use anything you recommend." Let me assure you that you *will* be able to use most of what I say. As a middle manager, of any sort, you are in effect a chief executive of an organization yourself. Don't wait for the

principles and practices you find appealing to be imposed from the top. As a *micro CEO,* you can improve your own and your group's performance and productivity, whether or not the rest of the company follows suit.

This book contains three basic ideas. The first is an output-oriented approach to management. That is to say, we apply some of the principles and the discipline of the most output-oriented of endeavors—*manufacturing*—to other forms of business enterprise, including most emphatically the work of managers. Consider Intel, which is a true manufacturing and production company, making highly complex silicon chips as well as computer-like products built from them. Our company now has over thirty thousand employees. Of these, about 25 percent actually work to make the products. Another 25 percent help them as they supervise the personnel, maintain the machines, and engineer and improve the manufacturing process. Another 25 percent work in administration, where they schedule production, keep personnel records, send bills to our customers, and pay our suppliers. Finally, the remaining 25 percent design new products, take them to the marketplace, sell them, and service them after the sale.

As we founded, organized, and managed Intel, we found that *all* our employees "produce" in some sense—some make chips, others prepare bills, while still others create software designs or advertising copy. We also found that when we approached any work done at Intel with this basic understanding in mind, the principles and discipline of production gave us a systematic way of managing it, much as the language and concepts of finance created a common approach to evaluating and managing investments of any sort.

The second idea is that the work of a business, of a government bureacracy, of most forms of human activity, is something pursued not by individuals but by teams. This idea is summed up in what I regard as the single most important sentence of this book: *The output of a manager is the output of the organizational units under his or her supervision or influence*. The question then becomes, what can managers do to increase the output of their teams? Put another

way, what specifically should they be doing during the day when a virtually limitless number of possible tasks calls for their attention? To give you a way to answer the question, I introduce the concept of *managerial leverage,* which measures the impact of what managers do to increase the output of their teams. High managerial productivity, I argue, depends largely on choosing to perform tasks that possess high leverage.

A team will perform well only if peak performance is elicited from the individuals in it. This is the third idea of the book. Can business use whatever it is that motivates an athlete to put out his "personal best" consistently? I think business can, which is why I examine the *sports analogy* and the role of something called task-relevant feedback to get and to sustain a high level of performance from the members of a business team.

We must recognize that no amount of formal planning can anticipate changes such as globalization and the information revolution we've referred to above. Does that mean that you shouldn't plan? Not at all. You need to plan the way a fire department plans. It cannot anticipate where the next fire will be, so it has to shape an energetic and efficient team that is capable of responding to the unanticipated as well as to any ordinary event.

Second, a responsive company should have fewer levels of managers. This concept is easier to apply today as electronic mail can carry information to anyone in the organization. One basic role of management—the role of disseminating information—is no longer as important a managerial function as it was in the past.

With fewer levels in today's organization, each manager will have larger numbers of employees reporting to him than was the case ten years ago. One of the fundamental tenets of Intel's managerial philosophy is the one-on-one meeting between a supervisor and a subordinate. Its main purposes are mutual education and the exchange of information. By talking about specific problems and situations, the supervisor teaches the subordinate his skills and know-how, and suggests ways to approach things. At the same time, the subordinate provides the supervisor with detailed information

about what he is doing and what he is concerned about. Obviously, one-on-ones take time, both in preparing for them and in actually holding them—time that today's busier manager may not have.

Are one-on-one meetings still needed? Absolutely. Can you have them as often with ten direct reports as with five? No. Do you need to? No again, because for the most part, these employees are more aware of what's going on in their business through their computer network than their counterparts were a decade ago; they no longer rely on you to bring them up to date. Nor do you need to rely on one-on-ones to catch up with what your subordinates have found in their lab, factory, or sales region; you've already read about those developments on your computer screen, minutes after they chose to inform you.

Now consider the proverbial Japanese employees sitting around the table in a Japanese office. They don't need to get together with their supervisors to be brought up to date. They may still need to leave the table and have a tête-a-tête with them to discuss their concerns or to bring up issues they're uncomfortable with, but a lot of the purposes of the one-on-ones are taken care of minute by minute. It's no different when you and your subordinates are working around the electronic equivalent of this table. So, yes, you still need one-on-ones. But you need them for fewer of the purposes I envisioned when I originally wrote this book. Therefore, you can deal with more employees less frequently and in meetings of shorter duration.

III. Managing your own career

But what about managers who are, after all, employees themselves?

I recently read an article saying that middle-aged men are twice as likely to lose their jobs today than they were in 1980, fifteen years ago. This trend is going to increase in the years ahead.

As a general rule, you have to accept that no matter where you work, you are not an employee—you are in a business with one employee: yourself. You are in competition with millions of similar

businesses. There are millions of others all over the world, picking up the pace, capable of doing the same work that you can do and perhaps more eager to do it. Now, you may be tempted to look around your workplace and point to your fellow workers as rivals, but they are not. They are outnumbered—a thousand to one, one hundred thousand to one, a million to one—by people who work for organizations that compete with your firm. So if you want to work and continue to work, you must continually dedicate yourself to retaining your *individual* competitive advantage.

In a slow or no-growth environment, there is another factor that you have to contend with as well: ambitious junior employees who desire to move upward in the organization. They may very well be ready to do so but can't because you're in the way. Sooner or later, your boss will inevitably have to make a choice: whether to hold on to you, who is doing a good job but is in the way of another person. The responsibility to avoid such situations is yours.

The recipe for success for the generation of managers who worked in the sixties, seventies, and much of the eighties was to join stable and enlightened companies and help them do well; these companies in turn would reward such managers with a career. Obviously, that is no longer the case.

The point is, the clichés of globalization and the information revolution have real meaning—potentially deadly meaning—for your career. The sad news is, nobody owes you a career. You own it as a sole proprietor. You must compete with millions of individuals every day, and every day you must enhance your value, hone your competitive advantage, learn, adapt, get out of the way, move from job to job, even from industry to industry if you must and retrench if you need to do so in order to start again. The key task is to manage your career so that you do not become a casualty.

I can offer you no surefire formula. But here are a few questions to ponder:

1. Are you adding real value or merely passing information along? How do you add more value? By continually looking for ways to make things truly better in your department. You are a manager. The

central thought of my book is that the output of a manager is the output of his organization. In principle, every hour of your day should be spent increasing the output or the value of the output of the people whom you're responsible for.

2. Are you plugged into what's happening around you? And that includes what's happening inside your company as well as inside your industry as a whole. Or do you wait for a supervisor or others to interpret whatever is happening? Are you a node connected to a network of plugged-in people or are you floating by yourself?

3. Are you trying new ideas, new techniques, and new technologies, and I mean personally trying them, not just reading about them? Or are you waiting for others to figure out how they can re-engineer your workplace—and you out of that workplace?

I am an engineer by training and a manager of a high-technology company by profession. As a manager, I am also a member of the group of individuals—many millions strong in the United States alone—that holds the key to increased productivity: generating more and better goods and services to meet people's needs. I am an optimist and believe our potential to increase our wealth has hardly been tapped.

But I also think that people do not always face up to the changes they have to deal with, so at times I feel I have to be a realist, too. You can't be optimistic about the future until you have survived the crucible of change. The key to survival is to learn to add more value—and that ultimately is what this book is about.

From my own experience at Intel, I strongly believe that applying the methods of *production,* exercising *managerial leverage,* and eliciting an athlete's desire for *peak performance* can help nearly everyone—lawyers, teachers, engineers, supervisors, even book editors; in short, middle managers of all kinds—to work more productively.

So, let us proceed by taking a field trip to a factory….

Andrew S. Grove
April 1995

Foreword to the Vintage Books Edition

I first read *High Output Management* in 1995. In those days, there were no blogs or TED Talks teaching us about entrepreneurship. In fact, there was almost nothing of use written for people like me who aspired to build and run a company.

Against this backdrop, *High Output Management* had an almost legendary status. All the best managers knew about it. The top venture capitalists gave copies of it to their entrepreneurs, and aspiring leaders in Silicon Valley devoured its contents. It amazed all of us that the CEO of Intel had taken the time to teach us the essential skill of entrepreneurship: how to manage.

This was no small thing because Intel was known as the best company in the technology industry. It had pulled off the greatest transformation in the history of the business: moving from the memory business to microprocessors more than a decade after its founding. Beyond that, Intel ran with legendary precision, which gave it the ability to make multibillion dollar investments with high confidence. If you wanted to hire a great operational manager, then Intel was the place to go—but good luck getting one to leave the best-managed company in Silicon Valley.

Andy himself was a legendary figure. He had grown up Jewish in Hungary during a time when the country was occupied by the Nazis and, later, by the Soviet Communists. Arriving in New York, he spoke no English and had almost no money. He enrolled himself at the City College of New York, overcame his language deficiency, and went on to get a PhD from UC Berkeley. This nonnative English speaker would then write an important textbook on semiconductors *in English* while working at Fairchild Semiconductor. As a result, he

was considered a scientific pioneer even before helping to launch Intel in 1968, building it into the seminal technology company of the era. Later, in 1997, *Time* magazine would recognize his nearly impossible accomplishments and name him Man of the Year.

This is in part what made *High Output Management* so extraordinary. Andy Grove, who built himself from nothing to run Intel, stopped what he was doing to teach us his magic. And not through some ghostwriter either—Andy wrote this book himself. What an incredible gift.

When I finally got my hands on the book, the paperback cover took me aback. The 1995 version featured a picture of Andy Grove standing next to the Intel sign. Unlike every other CEO photo that I had ever seen, Andy was not wearing a designer suit. He did not have perfectly combed hair, and he did not strike an arms-folded power pose. No, Andy Grove was dressed for work right down to his key card hanging from his belt. I did a double take. "Was that a key card? He didn't remove his key card for the book's cover photo?"

In retrospect, the cover was perfect. As you will see when you read this book, Andy Grove was all substance. He did not have time for pretty photo shoots or self-promotion. He wrote the book for us, but if we had to be sold on it by how he looked in the photo, then that would be our loss. The time that he did not spend styling fancy photos, he put into writing the book. He did not just give us the lessons; he articulated them in a way that connected both logically and emotionally. We would come to understand him and feel what he meant in our core.

I immediately got a jolt of this style with the title of the very first chapter: "The Basics of Production: Delivering a Breakfast (or a College Graduate, or a Compiler, or a Convicted Criminal…)." Okay, I am interested. What does making a soft-boiled egg have to do with how many prisons we build? It turns out quite a bit. *High Output Management* opens by teaching us the importance of proper system design even when we are dealing with a system of human beings—*especially* when we are dealing with a system of human beings.

Andy then shows us how you can use these same principles to understand how society should operate. It doesn't accomplish anything to declare that we need more kids going to college than to jail and demand that we build more schools than jails. In fact, it's counterproductive. Identifying complex system problems is one thing. Solving them is something else entirely, and Andy lays out the tools to do just that.

Over the years, I have come to consider *High Output Management* a true masterpiece, and there are at least three core aspects to its genius. First, in as little as one sentence, it lucidly explains concepts that require entire books from lesser writers. Second, it consistently uncovers brand-new management ideas or finds new insights into old standards. Finally, while most management books attempt to teach basic competency, *High Output Management* teaches the reader how to be great.

Andy introduces management with this classic equation:

> A manager's output = the output of his organization + the output of the neighboring organizations under his influence.

On the surface it may seem simple, but he clarifies the essential difference between a manager and an individual contributor. A manager's skills and knowledge are only valuable if she uses them to get more leverage from her people. So, Ms. Manager, you know more about our product's viral loop than anyone in the company? That's worth exactly nothing unless you can effectively transfer that knowledge to the rest of the organization. That's what being a manager is about. It's not about how smart you are or how well you know your business; it's about how that translates to the team's performance and output.

As a means to obtain this leverage, a manager must understand, as Andy writes: "When a person is not doing his job, there can only be two reasons for it. The person either can't do it or won't do it; he is either not capable or not motivated." This insight enables a manager

to dramatically focus her efforts. All you can do to improve the output of an employee is motivate and train. There is nothing else.

As he describes the planning process, Andy sums up his essential point with this eloquent nugget of wisdom: “I have seen far too many people who upon recognizing today’s gap try very hard to determine what decision has to be made to close it. But today’s gap represents a failure of planning sometime in the past.” Hopefully, the value of this short insight is not lost on the young reader. If you only understand one thing about building products, you must understand that energy put in early in the process pays off tenfold and energy put in at the end of the program pays off negative tenfold.

The book has an entire section dedicated to an often neglected, but critically important management tool: meetings. Andy makes us see this oldest of business principles in a new light. He teaches meetings from first principles, beginning with how to conduct a one-on-one. It seems incredible that the CEO of Intel would take the time to explain how to have a one-on-one.

Why is he doing this? It turns out that the one-on-one is not only a fundamental element in the manager/employee relationship, but perhaps the best source for organizational knowledge that a manager can get. In my experience, managers who don’t have one-on-ones understand very little about what’s happening in their organizations.

It is by understanding the simple things that Andy goes deep. For example, when people visit today’s technology companies they often remark about how casual the environments are, but with very little explanation about why they are that way. In fact, many CEOs do not understand why as they simply follow the trend, but Andy explains it perfectly:

> A journalist puzzled by our management style once asked me, “Mr. Grove, isn’t your company’s emphasis on visible signs of egalitarianism such as informal dress, partitions instead of offices…just so much affectation?” My answer was that this is not affectation, but a matter of survival. In our business we have to mix knowledge-power people with position-power people daily,

> and together they make decisions that could affect us for years to come.

In this fashion, the book quickly gets to the heart of complex issues. It raises and deals with the stickiest management issues. Andy asks the question of whether you should be friends with the people you manage:

> Everyone must decide for himself what is professional and appropriate here. A test might be to imagine yourself delivering a tough performance review to your friend. Do you cringe at the thought? If so, don't make friends at work. If your stomach remains unaffected, you are likely to be someone whose personal relationships will strengthen work relationships.

By breaking down the process, he makes hard things manageable.

Ultimately, the power of *High Output Management* is that it creates expert rather than merely competent managers.

A great example of this is the section on task-relevant maturity. This part of the book became very personal for me as it taught me how to formulate the most useful management question that I use in interviews: *"Is it better to be a hands-on or hands-off manager?"*

It seems like a simple enough question, but it sorts out the 95 percent of managers who never think deeply about their craft from the 5 percent who do. The answer, as Andy explains, is that it depends. Specifically, it depends on the employee. If the employee is immature in the task, then hands-on training is essential. If the employee is more mature, then a delegate approach is warranted. Andy presents a great example of this: "The subordinate did poor work. My associate's reaction: 'He has to make his own mistakes. That's how he learns!' The problem with this is that the subordinate's tuition is paid by his customers. And that is absolutely wrong."

Perhaps the chapter that best reflects Andy Grove is the last, "Why Training Is the Boss's Job." Often, people who manage in the so-called knowledge economy believe their employees are so smart that

they need no training at all. Andy brilliantly corrects this notion by explaining why as customers we are flabbergasted when we encounter employees who are insufficiently trained at relatively simple tasks such as taking restaurant reservations. He then challenges us to imagine how furious customers of complex jobs will be if an employee isn't properly trained. Finally, he reiterates his thesis that there are only two ways in which a manager can impact an employee's output: motivation and training. If you are not training, then you are basically neglecting half the job.

Throughout the chapter, the reader feels Andy's intense passion toward training and teaching, because in the end—more than anything else—he is a teacher…in the very best sense of the word.

Many years after reading *High Output Management,* I met Andy for the first time. Upon seeing him, I was so excited that I immediately blurted out how much I loved the book. In classic Andy Grove style, he shot back: "Why?" I did not expect that. I thought that he would say, "Thank you" or "I appreciate that," but not "Why?" But that was Andy. He was always teaching and always expecting more from every student.

Caught completely off guard, I scrambled for the reason and came across a good one: "Every other management book that I've read explains the trivial, but yours gets to the real issues." Upon hearing that, the master teacher softened and replied with a priceless story:

> It's funny that you say that about management books. I recently ran out of space on my bookshelf at home, so I was faced with a choice. I either had to throw away some books or buy a bigger house. Well, that was an easy decision, but which books to throw out? Then I thought, the management books! But I had a problem. Nearly every management book that I'd received was sent to me by the author and was autographed with a kind inscription. I felt badly about throwing away all those nice notes. So, I went through each book and tore out the inscription page then threw away the book. So now I have a large stack of pages of nice notes to me and plenty of space for good books.

I have never met anyone other than Andy Grove who would have a story like that. He uniquely balances the highest standards for clear thinking and performance with an undying belief in the underlying person. Who else would require so high a bar for writing that you had to be good enough to fit on his one bookshelf and still be so touched by the fact that you wanted him to read your work that he would save the page that you inscribed?

Later, in 2001, I met with Andy again and I asked him about a recent run of CEOs missing their numbers despite having told investors that their businesses were strong. The bubble had burst for the first wave of Internet companies nearly a year prior, so it surprised me that so many many of them had not seen this coming. Andy replied with an answer that I did not expect: "CEOs always act on leading indicators of good news, but only act on lagging indicators of bad news."

"Why?" I asked him. He answered in the style resonant of his entire book: "In order to build anything great, you have to be an optimist, because by definition you are trying to do something that most people would consider impossible. Optimists most certainly do not listen to leading indicators of bad news."

But this insight won't be in any book. When I suggested he write something on the topic, his response was: "Why would I do that? It would be a waste of time to write about how to not follow human nature. It would be like trying to stop the Peter Principle.[*] CEOs must be optimists and all in all that's a good thing." This is classic Andy Grove. He is amazingly perceptive and can see every flaw in every person, yet despite that he believes in human potential more than anyone. Maybe that's why he has spent so much time teaching us to be better.

It has been an honor for me to learn from Andy Grove through the years and I am excited for everyone who is new to *High Output Management* to join me in this experience. I know that you will enjoy this marvelous book written by the best teacher that I have ever known.

---

[*] The Peter Principle is a concept in management theory in which the selection of a candidate for a position is based on the candidate's performance in their current role, rather than on abilities related to the intended role. Thus, "managers rise to the level of their incompetence."

I

The Breakfast Factory

1

The Basics of Production: Delivering a Breakfast

(or a College Graduate, or a Compiler, or a Convicted Criminal…)

The Three-Minute Egg

To understand the principles of production, imagine that you're a waiter, which I was while I went to college, and that your task is to serve a breakfast consisting of a three-minute soft-boiled egg, buttered toast, and coffee. Your job is to prepare and deliver the three items simultaneously, each of them fresh and hot.

The task here encompasses the basic requirements of production. These are to build and deliver products in response to the demands of the customer at a *scheduled* delivery time, at an *acceptable* quality level, and at the *lowest* possible cost. Production's charter cannot be to deliver whatever the customer wants whenever he wants it, for this would require an infinite production capacity or the equivalent—very large, ready-to-deliver inventories. In our example, the customer may want to have a perfect three-minute egg with hot buttered toast and steaming coffee waiting for him the moment he sits down. To fulfill such an expectation, you would either have to have your kitchen idle and poised to serve the customer whenever he drops in, or have a ready-to-consume inventory of perfectly boiled eggs, hot buttered toast, and coffee. Neither is practical.

Instead, a manufacturer should accept the responsibility of delivering a product at the time committed to—in this case, by implication, about five to ten minutes after the customer arrives at

our breakfast establishment. And we must make our breakfast at a cost that enables us to sell it at a competitive price and still make an acceptable profit. How are we going to do this in the most intelligent way? We start by looking at our production flow.

The first thing we must do is to pin down the step in the flow that will determine the overall shape of our operation, which we'll call the *limiting step*. The issue here is simple: which of the breakfast components takes the longest to prepare? Because the coffee is already steaming in the kitchen and the toast takes only about a minute, the answer is obviously the egg, so we should plan the entire job around the time needed to boil it. Not only does that component take the longest to prepare, the egg is also for most customers the most important feature of the breakfast.

What must happen is illustrated opposite. To work back from the time of delivery, you'll need to calculate the time required to prepare the three components to ensure that they are all ready simultaneously. First you must allow time to assemble the items on a tray. Next you must get the toast from the toaster and the coffee from the pot, as well as the egg out of the boiling water. Adding the required time to do this to the time needed to get and cook the egg defines the length of the entire process—called, in production jargon, the total throughput time.

Now you come to the toast. Using the egg time as your base, you must allow yourself time to get and toast the slices of bread. Finally, using the toast time as your base, you can determine when you need to pour the coffee. The key idea is that we construct our production flow by starting with the longest (or most difficult, or most sensitive, or most expensive) step and work our way back. Notice when each of the three steps began and ended. We planned our flow around the most critical step—the time required to boil the egg—and we staggered each of the other steps according to individual throughput times; again in production jargon, we *offset* them from each other.

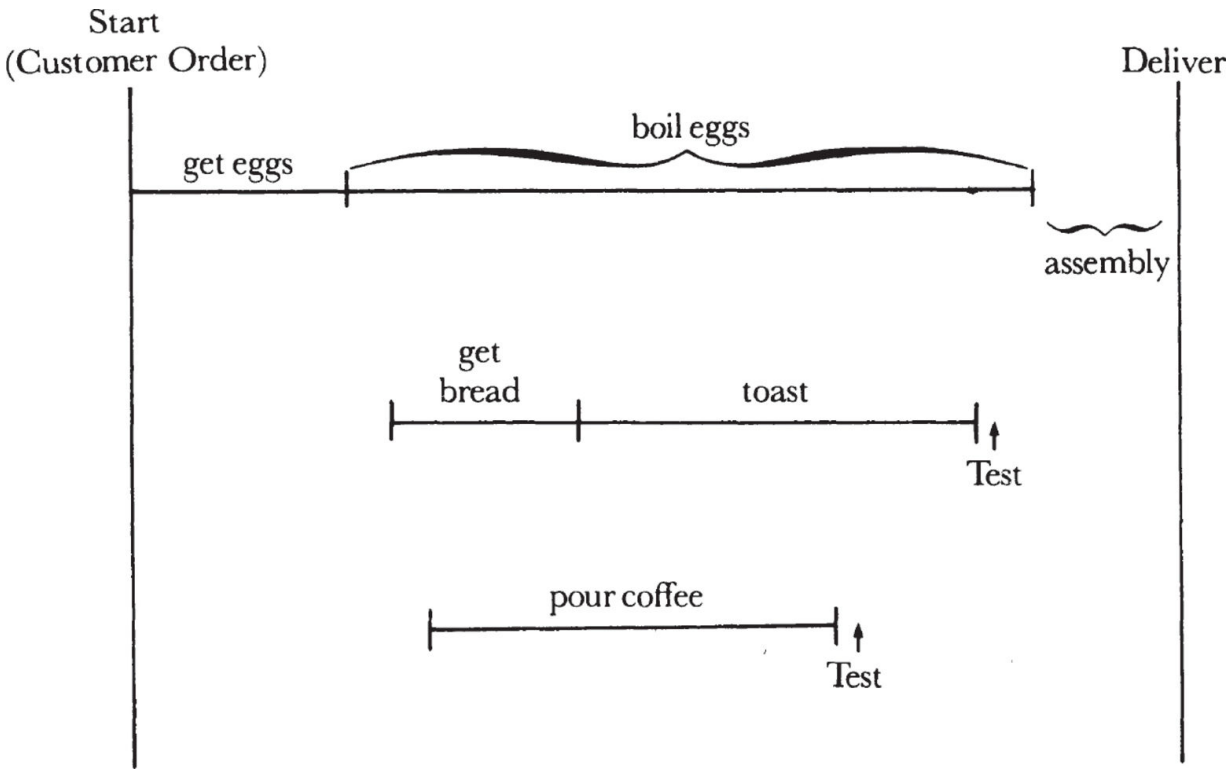

Making the eggs is the limiting step.

The idea of a limiting step has very broad applicability. Take, for example, the need to recruit college graduates to work for Intel. Certain of our managers visit the colleges, interview some of the seniors, and invite the more promising candidates to visit the company. We bear the expense of the candidates' trip, which can be considerable. During the trip, the students are closely interviewed by other managers and technical people. After due consideration, employment is offered to some of the students whose skills and capabilities match our needs best, and those who accept the offers eventually come to work for the company.

To apply the basic principle of production, you need to build the sequence here around its most expensive feature, which is the students' trip to the plant, thanks to the cost of travel and the time that Intel managers spend with the candidates. To minimize the use of this step per final college hire, we obviously have to increase the ratio of accepted offers to applicants invited to visit the plant, which we do by using phone interviews to screen people before issuing

invitations. The technique saves money, substantially increases the ratio of offers extended per plant visit, and reduces the need to use the expensive limiting step per hire.

The principle of time offsets is also present here. Working back from the time the students will graduate, the recruiter staggers the various steps involved to allow time for everything—on-campus interviews, phone screening, plant visits—to take place at the appropriate times during the months preceding graduation.

Production Operations

Other production principles underlie the preparation of our breakfast. In the making of it, we find present the three fundamental types of production operations: *process* manufacturing, an activity that physically or chemically changes material just as boiling changes an egg; *assembly,* in which components are put together to constitute a new entity just as the egg, the toast, and the coffee together make a breakfast; and *test,* which subjects the components or the total to an examination of its characteristics. There are, for example, visual tests made at points in the breakfast production process: you can see that the coffee is steaming and that the toast is brown.

Process, assembly, and test operations can be readily applied to other very different kinds of productive work. Take, for instance, the task of training a sales force to sell a new product. The three types of production operations can be easily identified. The conversion of large amounts of raw data about the product into meaningful selling strategies comprehensible to the sales personnel is a process step, which transforms data into strategies. The combination of the various sales strategies into a coherent program can be compared to an assembly step. Here the appropriate product-selling strategies and pertinent market data (such as competitive pricing and availability) are made to flow into one presentation, along with such things as brochures, handouts, and flip charts. The test operation comes in the form of a “dry run” presentation with a selected group of field sales personnel and field sales management. If the dry run

fails the test, the material must be “reworked” (another well-established manufacturing concept) to meet the concerns and objections of the test audience.

The development of a “compiler,” a major piece of computer software, also demonstrates process, assembly, and test. A computer understands and uses human instruction only if it receives such instruction in its own language. A compiler is an interpreter, enabling the computer to translate into its language material written in terms and phrases resembling English. With a compiler, a programmer can think more or less like a human being rather than having to adapt himself to the way the computer processes information. The task of getting a machine to interpret and translate in this fashion is obviously formidable; thus the development of a compiler takes strenuous effort on the part of skilled and gifted software engineers. The effort, however, is justified by the simplification it brings to computer use.

In any case, the development of the individual pieces out of which a compiler is built represents a series of processing steps. Actual working pieces of software are generated out of specifications and basic design know-how. Each piece then undergoes an individual operation called a “unit test.” When one fails, the defective portion of the software is returned to the process phase for “rework.” After all the pieces pass their respective unit tests, they are assembled to form the compiler. Then, of course, a “system test” is performed on the complete product before it is shipped to the customer. Time offsets are used extensively in the task. Because throughput times for the various engineering steps are well established, the timing of the releases of various bodies of software from one stage to another can all be calculated and staged in advance.

Breakfast preparation, college recruiting, sales training, and compiler design are very much unlike one another, but all of them possess a basically similar flow of activity to produce a specific output.

A Few Complications

Real life, as you know, is full of thickets and underbrush. In a schematic flow chart, our breakfast operation assumed infinite capacity, meaning that nobody had to wait for an available toaster or for a pot to boil an egg in. But no such ideal world exists. What would happen if you had to stand in a line of waiters, waiting for your turn to use the toaster? If you didn't adjust your production flow to account for the queue, your three-minute egg could easily become a six-minute egg. So limited toaster capacity means you have to redo your flow around the new limiting step. The egg still determines the overall quality of the breakfast, but your time offsets must be altered.

How would our model reflect the change in manufacturing flow? Working back from the time of breakfast delivery, let's see how the production is affected, as illustrated opposite. The egg cycle remains the same, as does the one for coffee. But limited toaster capacity makes for quite a difference. Now you must account for the delivery time of the toast and the wait for a free toaster. This means the whole production process has to be conceived differently. Toaster capacity has become the limiting step, and what you do has to be reworked around it.

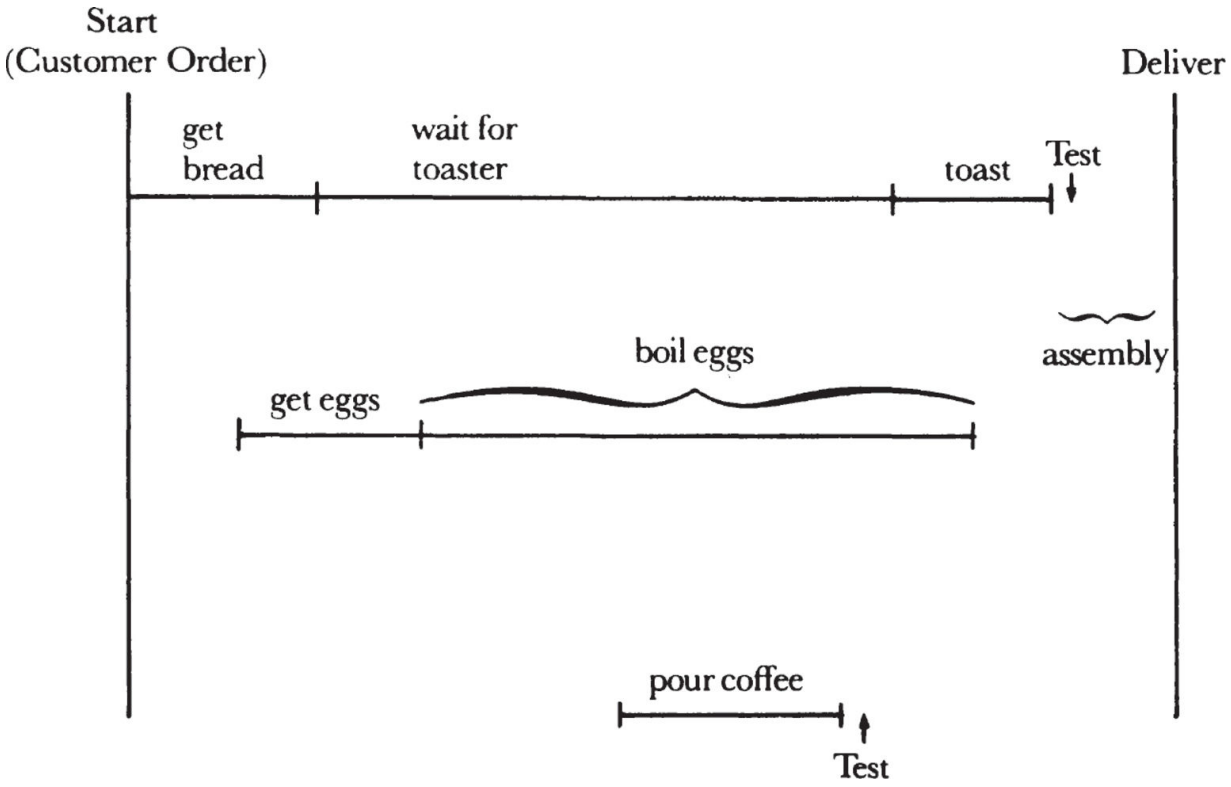

With limited toaster capacity, making the toast becomes the limiting step.

Now let's complicate things a little further. What happens if you are stuck in line waiting for a toaster when it's time to start boiling your egg? Your conflict is seemingly irreconcilable, but it really isn't. If you were managing the restaurant, you could turn your personnel into *specialists* by hiring one egg-cooker, one toast-maker, one coffee-pourer, and one person to supervise the operation. But that, of course, creates an immense amount of *overhead,* probably making it too expensive to consider.

If you were a waiter, you could ask the waiter in line next to you to help out—to put your toast in while you ran off to start your egg. But when you have to depend on someone else, the results are likely to be less predictable. As the manager, you could add another toaster, but this becomes an expensive addition of *capital equipment*. You could run the toaster continuously and build up an *inventory* of hot toast, throwing away what you can't use but always having immediate access to product. That means waste, which can also become too

expensive for the operation. But at least you know that alternatives do exist: equipment capacity, manpower, and inventory can be traded off against each other and then balanced against delivery time.

Because each alternative costs money, your task is to find the *most cost-effective* way to deploy your resources—the key to optimizing all types of productive work. Bear in mind that in this and in other such situations there is a right answer, the one that can give you the best delivery time and product quality at the lowest possible cost. To find that right answer, you must develop a clear understanding of the trade-offs between the various factors—manpower, capacity, and inventory—and you must reduce the understanding to a quantifiable set of relationships. You probably won't use a stopwatch to conduct a time-and-motion study of the person behind a toaster; nor will you calculate the precise trade-off between the cost of toast inventory and the added toaster capacity in mathematical terms. What is important is the thinking you force yourself to go through to understand the relationship between the various aspects of your production process.

Let's take our manufacturing example a step further and turn our business into a high-volume breakfast factory operation. First, you buy a *continuous egg-boiler* that will produce a constant supply of perfectly boiled three-minute eggs. It will look something like what's drawn in the figure opposite. Note that our business now assumes a high and predictable demand for three-minute eggs; it cannot now readily provide a four-minute egg, because automated equipment is not very flexible. Second, you match the output of the continuous egg-boiler with the output of a continuous toaster, as specialized personnel load each piece of equipment and deliver the product. We have now turned things into a *continuous operation* at the expense of flexibility, and we can no longer prepare each customer's order exactly when and how he requests it. So our customers have to adjust their expectations if they want to enjoy the benefits of our new mode: lower cost and more predictable product quality.

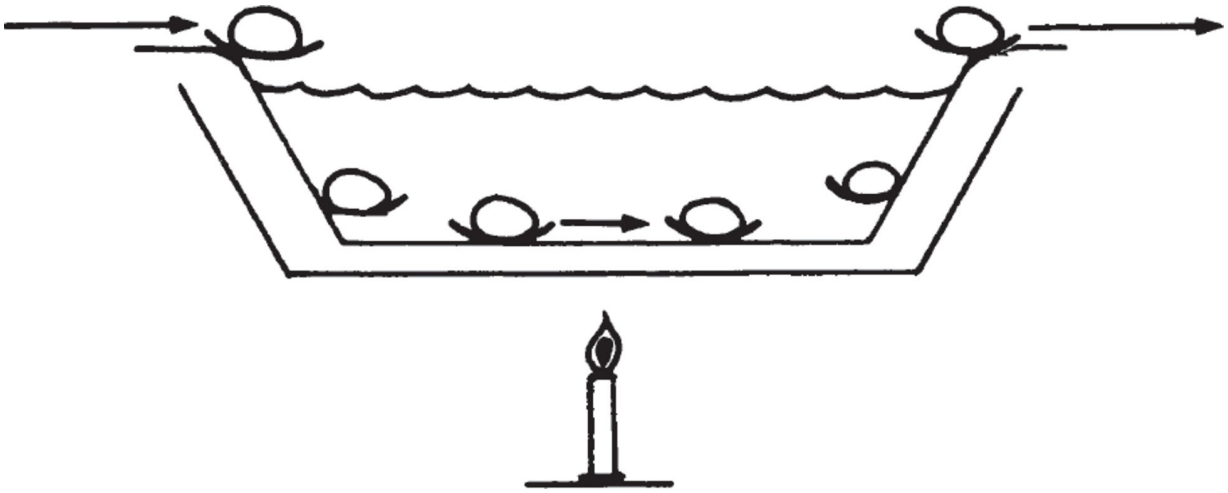

The continuous egg-boiler: a constant supply of three-minute eggs.

But continuous operation does not automatically mean lower cost and better quality. What would happen if the water temperature in the continuous egg-boiler quietly went out of specification? The entire work-in-process—all the eggs in the boiler—and the output of the machine from the time the temperature climbed or dropped to the time the malfunction was discovered becomes unusable. All the toast is also wasted because you don't have any eggs to serve with it. How do you minimize the risk of a breakdown of this sort? Performing a *functional test* is one way. From time to time you open an egg as it comes out of the machine and check its quality. But you will have to throw away the egg tested. A second way involves *in-process inspection,* which can take many forms. You could, for example, simply insert a thermometer into the water so that the temperature could be easily and frequently checked. To avoid having to pay someone to read the thermometer, you could connect an electronic gadget to it that would set off bells anytime the temperature varied by a degree or two. The point is that whenever possible, you should choose in-process tests over those that destroy product.

What else could go wrong with our continuous egg-machine? The eggs going into it could be cracked or rotten, or they could be over- or undersized, which would affect how fast they cook. To avoid such problems, you will want to look at the eggs at the time of receipt,

something called *incoming* or *receiving inspection*. If the eggs are unacceptable in some way, you are going to have to send them back, leaving you with none. Now you have to shut down. To avoid that, you need a *raw material inventory*. But how large should it be? The principle to be applied here is that you should have enough to cover your consumption rate for the length of time it takes to replace your raw material. That means if your egg man comes by and delivers once a day, you want to keep a day's worth of inventory on hand to protect yourself. But remember, inventory costs money, so you have to weigh the advantage of carrying a day's supply against the cost of carrying it. Besides the cost of the raw material and the cost of money, you should also try to gauge the *opportunity at risk:* what would it cost if you had to shut your egg machine down for a day? How many customers would you lose? How much would it cost to lure them back? Such questions define the opportunity at risk.

Adding Value

All production flows have a basic characteristic: the material becomes more valuable as it moves through the process. A boiled egg is more valuable than a raw one, a fully assembled breakfast is more valuable than its constituent parts, and finally, the breakfast placed in front of the customer is more valuable still. The last carries the perceived value the customer associates with the establishment when he drives into the parking lot after seeing the sign "Andy's Better Breakfasts." Similarly, a finished compiler is more valuable than the constituent parts of semantic analysis, code generation, and run time, and a college graduate to whom we are ready to extend an employment offer is more valuable to us than the college student we meet on campus for the first time.

A common rule we should always try to heed is to detect and fix any problem in a production process at the *lowest-value* stage possible. Thus, we should find and reject the rotten egg as it's being delivered from our supplier rather than permitting the customer to find it. Likewise, if we can decide that we don't want a college candidate at the time of the campus interview rather than during the

course of a plant visit, we save the cost of the trip and the time of both the candidate and the interviewers. And we should also try to find any performance problem at the time of the unit test of the pieces that make up a compiler rather than in the course of the test of the final product itself.

Finally, at the risk of being considered hard-hearted, let's examine the criminal justice system as if it were a production process aimed at finding criminals and putting them into jail. The production begins when a crime is reported to the police and the police respond. In many instances, after some questions are asked, no further action can be taken. For those crimes which the police can pursue, the second step is more investigation. But the case often ends here for lack of evidence, complaints being dropped, and so on. If things move to the next stage, a suspect is arrested, and the police try to find witnesses and build a case, hoping to get an indictment. Once again, an indictment is often not returned because of insufficient evidence. For the cases that actually do go ahead, the next stage is trial. Sometimes the suspect is found not guilty; sometimes the case is dismissed. But when a conviction is secured, the process moves to the sentencing and appeals round. At times a person found guilty of a crime will be given a suspended sentence and probation, and at others the conviction will be overturned on appeal. For the small fraction that remains, the final stage is jail.

If we make some reasoned assumptions about the percentages that move forward at each stage and the costs associated with each, we arrive at some striking conclusions. If we compile the cost of the effort that goes into securing a conviction and assign it only to those criminals who actually end up in jail, we find that the cost of a single conviction works out to be well over a million dollars—an absolutely staggering sum. The number is so high, of course, because only a very small percentage of the flow of accused persons makes it all the way through the process. Everyone knows that prisons are overcrowded, and that many criminals end up serving shorter jail terms or no jail terms at all because cells are in such short supply. So a terribly expensive trade-off results, violating the most important

production principles. The limiting step here should clearly be obtaining a conviction. The construction cost of a jail cell even today is only some $80,000. This, plus the $10–20,000 it costs to keep a person in jail for a year, is a small amount compared to the million dollars required to secure a conviction. Not to jail a criminal in whom society has invested over a million dollars for lack of an $80,000 jail cell clearly misuses society's total investment in the criminal justice system. And this happens because we permit the wrong step (the availability of jail cells) to limit the overall process.

2

Managing the Breakfast Factory

Indicators as a Key Tool

A hungry public has loved the breakfast you've been serving, and thanks to the help of your many customers and a friendly banker, you've created a *breakfast factory,* which among other things uses specialized production lines for toast, coffee, and eggs. As manager of the factory, you have a substantial staff and a lot of automated equipment. But to run your operation well, you will need a set of good *indicators,* or *measurements*. Your output, of course, is no longer the breakfasts you deliver personally but rather all the breakfasts your factory delivers, profits generated, and the satisfaction of your customers. Just to get a fix on your output, you need a number of indicators; to get efficiency and high output, you need even more of them. The number of possible indicators you can choose is virtually limitless, but for any set of them to be useful, you have to *focus* each indicator on a specific operational goal.

Let's say that as manager of the breakfast factory, you will work with five indicators to meet your production goals on a daily basis. Which five would they be? Put another way, which five pieces of information would you want to look at each day, immediately upon arriving at your office?

Here are my candidates. First, you'll want to know your *sales forecast* for the day. How many breakfasts should you plan to deliver? To assess how much confidence you should place in your forecast, you would want to know how many you delivered yesterday compared to how many you planned on delivering—in other words,

the *variance* between your plan and the actual delivery of breakfasts for the preceding day.

Your next key indicator is *raw material inventory*. Do you have enough eggs, bread, and coffee on hand to keep your factory running today? If you find you have too little inventory, you can still order more. If you find you have too much, you may want to cancel today's egg delivery.

Another important piece of information is the condition of your *equipment*. If anything broke down yesterday, you will want to get it repaired or rearrange your production line to meet your forecast for the day.

You also must get a fix on your *manpower*. If two waiters are out sick, you will have to come up with something if you are still going to meet the demand forecasted. Should you call in temporary help? Should you take someone off the toaster line and make him a waiter?

Finally, you want to have some kind of *quality* indicator. It is not enough to monitor the number of breakfasts each waiter delivers, because the waiters could have been rude to the customers even as they served a record number of breakfasts. Because your business depends on people wanting what you sell, you must be concerned with the public's opinion of your service. Perhaps you should set up a "customer complaint log" maintained by the cashier. If one of your waiters elicited more than the usual number of complaints yesterday, you will want to speak to him first thing today.

All these indicators measure factors essential to running your factory. If you look at them early every day, you will often be able to do something to correct a potential problem before it becomes a real one during the course of the day.

Indicators tend to direct your attention toward what they are monitoring. It is like riding a bicycle: you will probably steer it where you are looking. If, for example, you start measuring your inventory levels carefully, you are likely to take action to drive your inventory levels down, which is good up to a point. But your inventories could become so lean that you can't react to changes in demand without creating shortages. So because indicators direct one's activities, you

should guard against overreacting. This you can do by *pairing* indicators, so that together both effect and counter-effect are measured. Thus, in the inventory example, you need to monitor both inventory levels and the incidence of shortages. A rise in the latter will obviously lead you to do things to keep inventories from becoming too low.

The principle here was evident many times in the development of a compiler. Measuring the completion date of each software unit against its capability is one example. Watching this pair of indicators should help us to avoid working on the perfect compiler that will never be ready, and also to avoid rushing to finish one that is inadequate. In sum, joint monitoring is likely to keep things in the optimum middle ground.

Nowhere can indicators—and paired indicators—be of more help than in administrative work. Having come to this realization, our company has been using measurements as a key tool to improve the productivity of administrative work for several years. The first rule is that a measurement—any measurement—is better than none. But a genuinely effective indicator will cover the *output* of the work unit and not simply the *activity* involved. Obviously, you measure a salesman by the orders he gets (output), not by the calls he makes (activity).

The second criterion for a good indicator is that what you measure should be a *physical, countable* thing. Examples of effective measures of administrative output are shown below. Because those listed here are all quantity or output indicators, their paired counterparts should stress the *quality* of work. Thus, in accounts payable, the number of vouchers processed should be paired with the number of errors found either by auditing or by our suppliers. For another example, the number of square feet cleaned by a custodial group should be paired with a partially objective/partially subjective rating of the quality of work as assessed by a senior manager with an office in that building.

ADMINISTRATIVE FUNCTION	WORK OUTPUT INDICATOR

Accounts payable	# Vouchers processed
Custodial	# Square feet cleaned
Customer service	# Sales orders entered
Data entry	# Transactions processed
Employment	# People hired (by type of hire)
Inventory control	# Items managed in inventory

Examples of administrative work output indicators.

Such indicators have many uses. First, they spell out very clearly what the objectives of an individual or group are. Second, they provide a degree of objectivity when measuring an administrative function. Third, and as important as any, they give us a measure by which various administrative groups performing the same function in different organizations can be compared with each other. The performance of a custodial group in one major building can now be compared with that of another group in a second building. In fact, if indicators are put in place, the competitive spirit engendered frequently has an electrifying effect on the motivation each group brings to its work, along with a parallel improvement in performance. More about this later when we examine the "sports analogy."

The Black Box

We can think of our breakfast factory as if it were a "black box": input (the raw materials) and the labor of waiters, helpers, and you, the manager, flowing into the box, and the output (the breakfast) flowing out of it as illustrated below. In general, we can represent any activity that resembles a production process in a simple fashion as a black box. Thus, we can draw a black box to represent college recruiting, where the input is the applicants on campus and the output is college graduates who have accepted our employment

offers. The labor is the work of our on-campus interviewers and the managers and technical people who interview back at the plant. Similarly, the process of field sales training can be seen as a black box with the input being the raw product specifications and the output being trained sales personnel. The labor here is the work of the marketing and merchandising people who turn raw information into usable sales tools and train the field sales personnel to exploit them. In fact, we can represent most, if not all, administrative work by our magical black box. A group whose job is to bill customers has as its input the information about the customer—what he has purchased, the pricing data, and the shipment records; and output is the final bill sent to the customer through which payment is collected. The labor is the work of all personnel involved.

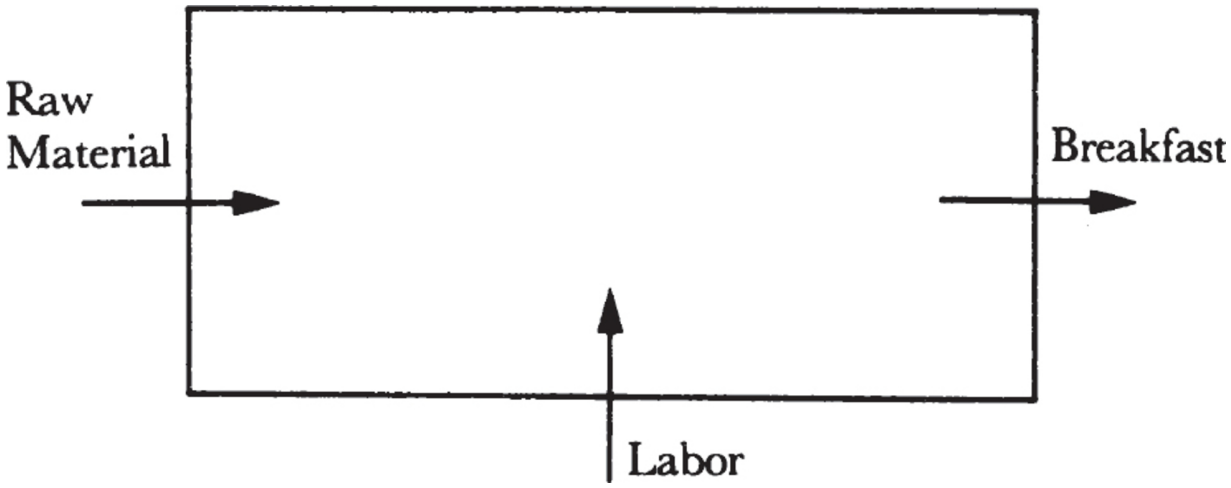

The breakfast factory—as a "black box."

The black box sorts out what the inputs, the output, and the labor are in the production process. We can improve our ability to run that process by cutting some *windows* in our box so that we can see some of what goes on within it. By looking through the openings, as illustrated below, we can better understand the internal workings of any production process and assess what the future output is likely to be.

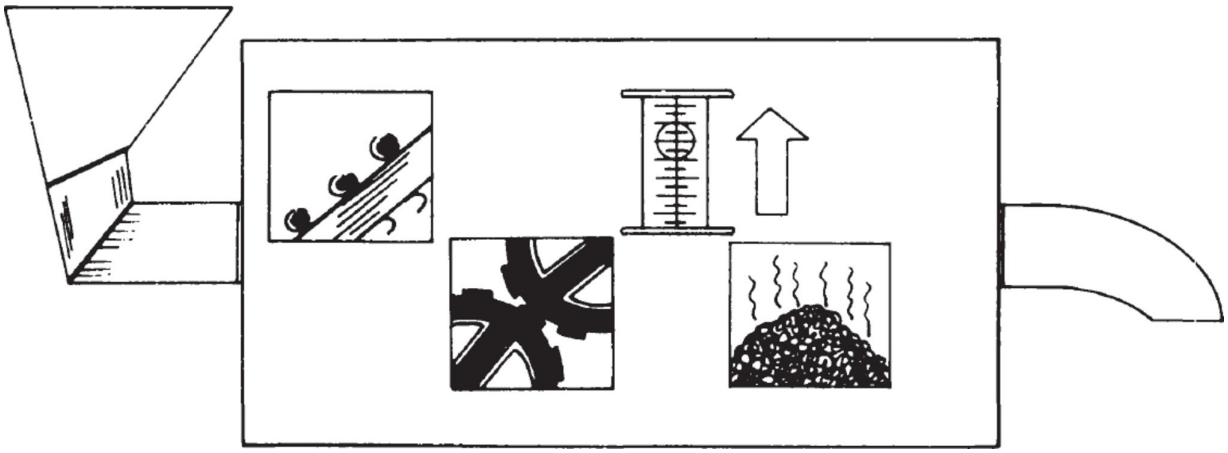

By peering through the windows in the black box, we can get an idea of what the future output is likely to be.

Leading indicators give you one way to look inside the black box by showing you in advance what the future might look like. And because they give you time to take corrective action, they make it possible for you to avoid problems. Of course, for leading indicators to do you any good, *you must believe in their validity*. While this may seem obvious, in practice, confidence is not as easy to come by as it sounds. To take big, costly, or worrisome steps when you are not yet sure you have a problem is hard. But unless you are prepared to act on what your leading indicators are telling you, all you will get from monitoring them is anxiety. Thus, the indicators you choose should be credible, so that you will, in fact, act whenever they flash warning signals.

Leading indicators might include the daily monitors we use to run our breakfast factory, from machine downtime records to an index of customer satisfaction—both of which can tell us if problems lie down the road. A generally applicable example of a "window" cut into the black box is the *linearity indicator*. In the figure below, we provide one for the college recruiting process. Plotted here is the number of college graduates who have accepted our offers versus the month of the year. If all went ideally, we would move along the straight line that would yield our hiring target for the semester by the month of June. If by April the actual progress is as shown here, we will find ourselves far below the ideal straight line. So from reading the

indicator, we know that the only way we can hit our target is by getting acceptance at a much higher rate in the remaining two months than we had gotten in the preceding four. Thus, the linearity indicator flashes an early warning, allowing us time to take corrective action. Without it, we would discover that we had missed our target in June, when nothing can be done about it.

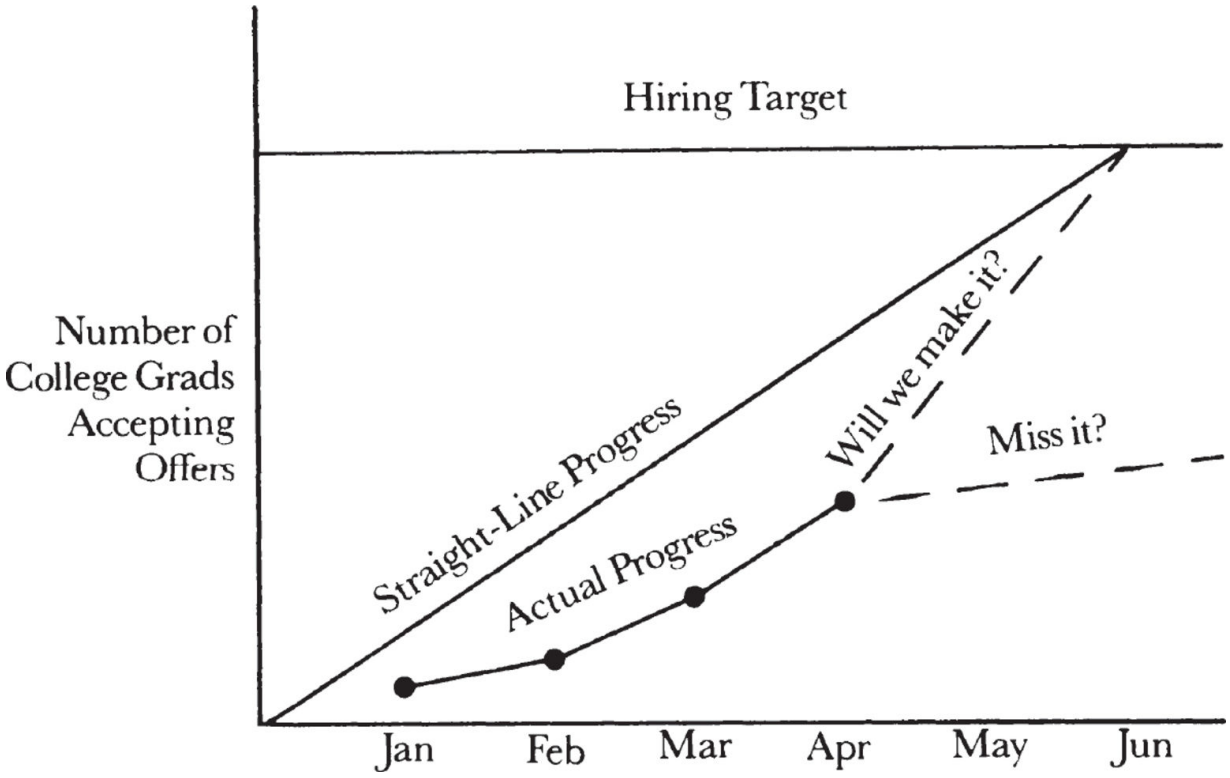

The linearity indicator can give us an early warning that we are likely to miss our target.

If we consider a manufacturing unit in this fashion, we may assume that because it makes monthly goals regularly, all is well. But we can cut a window into the black box here, measure production output against time as the month proceeds, and compare that with the ideal linear output. We may learn that output performance is spread evenly throughout the course of the month or that it is concentrated in the last week of the month. If the latter is the case, the manager of the unit is probably not using manpower and equipment efficiently. And if the situation is not remedied, one

minor breakdown toward month's end could cause the unit to miss its monthly output goal entirely. The linearity indicator will help you anticipate such a problem and is therefore quite valuable.

Also valuable are *trend indicators*. These show output (breakfasts delivered, software modules completed, vouchers processed) measured against time (performance this month versus performance over a series of previous months), and also against some standard or expected level. A display of trends forces you to look at the future as you are led to extrapolate almost automatically from the past. This extrapolation gives us another window in our black box. Also, measurement against a standard makes you think through *why* the results were what they were, and not what the standard said they would be.

Another sound way to anticipate the future is through the use of the *stagger chart,* which forecasts an output over the next several months. The chart is updated monthly, so that each month you will have an updated version of the then-current forecast information as compared to several prior forecasts. You can readily see the variation of one forecast from the next, which can help you anticipate future trends better than if you used a simple trend chart.

In my experience, nowhere has the stagger chart been more productive than in forecasting economic trends. The way it works is shown in the figure below, which gives us forecasted rates of incoming orders for an Intel division. The stagger chart then provides the same forecast prepared in the following month, in the month after that, and so on. Such a chart shows not only your outlook for business month by month but also how your outlook varied from one month to the next. This way of looking at incoming business, of course, makes whoever does the forecasting take his task very seriously, because he knows that his forecast for any given month will be routinely compared with future forecasts and eventually with the actual result. But even more important, the improvement or deterioration of the forecasted outlook from one month to the next provides the most valuable indicator of business trends that I have ever seen. I would go as far as to say that it's too

bad that all economists and investment advisers aren't obliged to display their forecasts in a stagger chart form. Then we could really have a way to evaluate whatever any one of them chooses to say.

	Forecasted incoming orders for:											
Forecast made in:	JUL	AUG	SEP	OCT	NOV	DEC	1982 JAN	FEB	MAR	APR	MAY	JUN
JUL '81	22	28	34	29								
AUG	*23	27	33	31	29							
SEP		*21	30	30	35	33						
OCT			*29	32	32	32	29					
NOV				*27	32	31	32	31				
DEC					*27	27	31	30	40			
JAN '82						*26	28	29	39	30		
FEB							*24	30	36	32	34	
MAR												

(* means the actual number for that month)

I have found the "stagger chart" the best means of getting a feel for future business trends.

Finally, indicators can be a big help in solving all types of problems. If something goes wrong, you will have a bank of information that readily shows all the parameters of your operation, allowing you to scan them for unhealthy departures from the norm. If you do not systematically collect and maintain an archive of indicators, you will have to do an awful lot of quick research to get the information you need, and by the time you have it, the problem is likely to have gotten worse.

Controlling Future Output

There are two ways to control the output of any factory. Some industries *build to order*. For example, when you go shopping for a sofa, you are going to have to wait a long time to get what you bought, unless you buy it right off the floor. A furniture factory builds to order. When it learns what you want, the factory looks for a hole in its manufacturing schedule and makes the item for you. If you order a new car rather than buying one right off the lot, the same thing happens: the plant will paint the car in the color you want and provide the options you want, but you will have to wait for it. And our breakfast factory, of course, builds breakfasts to order.

But if your competition in the sofa business makes the same product but has it ready in four weeks while you need four months, you are not going to have many customers. So even though you would much rather build to order, you will have to use another way to control the output of your factory. In short, you will have to *build to forecast,* which is a *contemplation of future orders*. To do this, the manufacturer sets up his activities around a reasoned speculation that orders will materialize for specific products within a certain time.

An obvious disadvantage here is that the manufacturer takes an inventory risk. Since the forecast is an assessment of future requirements, which the manufacturer commits resources to satisfy, the factory could be in an immense amount of trouble if the orders do not materialize or if they materialize for a product other than the one anticipated. In either case, unwanted inventory is the result. To build to forecast, you risk capital to respond to anticipated future demand in good order.

At Intel, we build to forecast because our customers demand that we respond to their needs in a timely fashion, even though our manufacturing throughput times are quite long. Our breakfast factory makes its product to customer order, but buys from its suppliers—like the egg man—on the basis of forecasted demand. Similarly, most companies recruit new college graduates to fill anticipated needs—rather than recruiting only when a need develops, which would be foolish because college graduates are turned out in a

highly seasonal fashion. Computer software products, such as compilers, are also typically developed in response to an anticipated market need rather than to specific customer order. So “building” to forecast is a very common business practice.

Delivering a product that was built to forecast to a customer consists of two simultaneous processes, each with a separate time cycle. A manufacturing flow must occur in which the raw material moves through various production steps and finally enters the finished goods warehouse, as illustrated below. Simultaneously, a salesman finds a prospect and sells to that prospect, who eventually places an order with the manufacturer. Ideally, the order for the product and the product itself should arrive on the shipping dock at the same time.

Because the art and science of forecasting is so complex, you might be tempted to give all forecasting responsibility to a single manager who can be made accountable for it. But this usually does not work very well. What works better is to ask both the manufacturing and the sales departments to prepare a forecast, so that people are responsible for performing against their own predictions.

At Intel we try to match the two parallel flows with as much precision as possible. If there’s no match, we end up with a customer order that we can’t satisfy or with a finished product for which we have no customer. Either way we have problems. Obviously, if the match does come off, with a forecasted order becoming a real order, the customer’s requirements can be nicely satisfied with the factory’s product delivery.

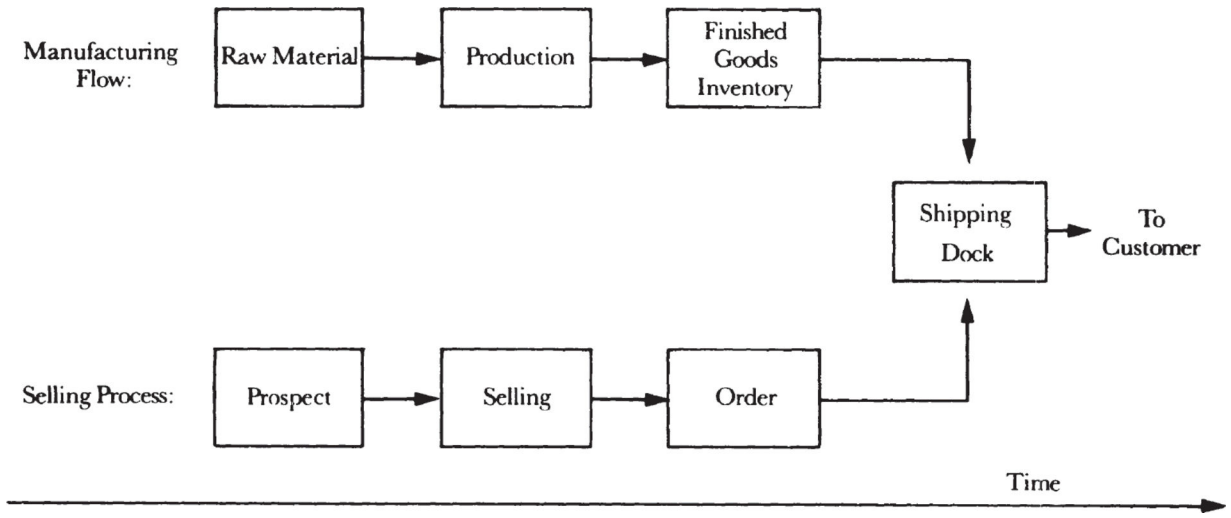

The order for the product and the product itself should arrive at the shipping dock at the same time.

The ideal is rarely found in the real world. More often, customer orders don't develop in time or the customer changes his mind. As for the other flow, manufacturing could miss deadlines, or make mistakes, or encounter unforeseen problems. Because neither the sales flow nor the manufacturing flow is completely predictable, we should deliberately build a reasonable amount of *"slack"* into the system. And inventory is the most obvious place for it. Clearly, the more inventory we have, the more change we can cope with and still satisfy orders. But inventory costs money to build and keep, and therefore should be controlled carefully. Ideally, inventory should be kept at the *lowest-value stage,* as we've learned before, like raw eggs kept at the breakfast factory. Also, the lower the value, the more production flexibility we obtain for a given inventory cost.

It is a good idea to use stagger charts in both the manufacturing and sales forecasts. As noted, they will show the trend of change from one forecast to another, as well as the actual results. By repeatedly observing the variance of one forecast from another, you will continually pin down the causes of inaccuracy and improve your ability to forecast both orders and the availability of product.

Forecasting future work demands and then adjusting the output of an "administrative factory" represents a very important way in which

its productivity can be increased. Though an old and honored way of operating "widget factories," the application of forecasting techniques is hardly common as a way to control administrative work. Such work has up to now been considered qualitatively different from work in a widget factory, and has also lacked objective performance standards needed to size or scale the work unit.

But if we have carefully chosen indicators that characterize an administrative unit and watch them closely, we are ready to apply the methods of factory control to administrative work. We can use de facto standards, inferred from the trend data, to forecast the number of people needed to accomplish various anticipated tasks. By rigorous application of the principles of forecasting, manpower can be reassigned from one area to another, and the headcount made to match the forecasted growth or decline in administrative activity. Without rigor, the staffing of administrative units would always be left at its highest level and, given Parkinson's famous law, people would find ways to let whatever they're doing fill the time available for its completion. There is no question that having standards and believing in them and staffing an administrative unit objectively using forecasted workloads will help you to maintain and enhance productivity.

Assuring Quality

As we have said, manufacturing's charter is to deliver product at a quality level acceptable to the customer at minimum cost. To assure that the quality of our product will in fact be acceptable, all production flows, whether they "make" breakfasts, college graduates, or software modules, must possess inspection points. To get acceptable quality at the lowest cost, it is vitally important to reject defective material at a stage where its accumulated value is at the lowest possible level. Thus, as noted, we are better off catching a bad raw egg than a cooked one, and screening out our college applicant before he visits Intel. In short, reject before investing further value.

In the language of production, the lowest-value-point inspection where we inspect raw material is called *incoming material*

inspection or *receiving inspection*. If we again use a black box to represent our production process, inspections that occur at intervening points within it are called, logically enough, *in-process* inspections. Finally, the last possible point of inspection, when the product is ready to be shipped to the customer, is called *final inspection* or *outgoing quality inspection*. The three types are depicted below.

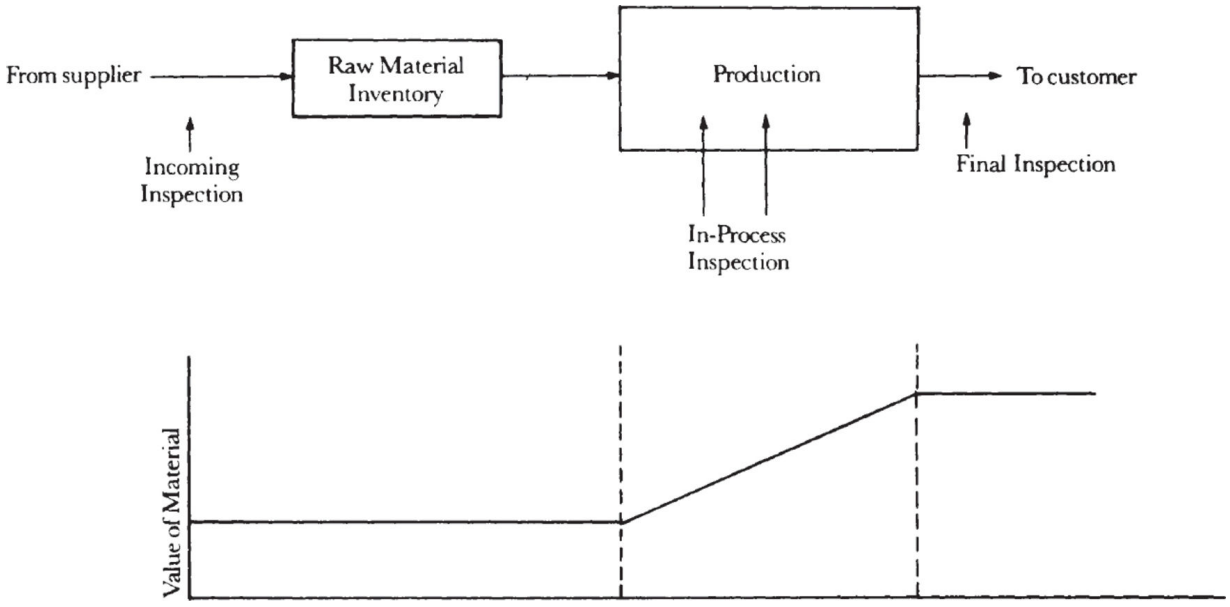

The key principle is to reject the defective "material" at its lowest-value stage.

When material is rejected at incoming inspection, a couple of choices present themselves. We can send it back to the vendor as unacceptable, or we can waive our specifications and use the substandard material anyway. The latter would result in a higher reject rate in our production process than if we had used thoroughly acceptable material, but that might be less expensive than shutting down the factory altogether until our vendor provides better material. Such decisions can only be made properly by a balanced group of managers, which typically consists of representatives from the quality assurance, manufacturing, and design engineering departments. This group can weigh all the consequences of rejecting or accepting substandard raw material.

While in most instances the decision to accept or reject defective material at a given inspection point is an economic one, one should *never* let substandard material proceed when its defects could cause a complete failure—a *reliability problem*—for our customer. Simply put, because we can never assess the consequences of an unreliable product, we can't make compromises when it comes to reliability. Think of a component going into the making of a cardiac pacemaker. If some of the components don't work upon receipt by the manufacturer, he can replace them while the unit is still in the factory. This will probably increase costs. But if the component fails later, after the pacemaker has been implanted, the cost of the failure is much more than a financial one.

Inspections, of course, cost money to perform and further add to expense by interfering with the manufacturing flow and making it more complicated. Some material has to be recycled through steps already performed, upsetting the smoothness with which the rest of the material moves. Accordingly, one should approach the need to inspect recognizing that a balance exists between the desired result of the inspection, improved quality, and minimum disturbance to the production process itself.

Let's consider a few techniques commonly used to balance the two needs. There is a *gate-like* inspection and a *monitoring* step. In the former, all material is held at the "gate" until the inspection tests are completed. If the material passes, it is moved on to the next stage in the production process; if the material fails, it will be returned to an earlier stage, where it will be reworked or scrapped. In the latter, a sample of the material is taken, and if it fails, a notation is made from which a failure rate is calculated. The bulk of the material is not held as the sample is taken but continues to move through the manufacturing process. The smoothness of the flow is maintained, but if, for example, three successive samples fail the monitoring test, we can stop the line. What is the trade-off here? If we hold all the material, we add to throughput time and slow down the manufacturing process. A monitor produces no comparable slowdown but might let some bad material escape before we can act

on the monitor's results and shut things down, which means that we might have to reject material later at a higher-value stage. Clearly, for the same money we can do a lot more monitoring than gate-type inspections; if we do the former, we may well contribute more to the overall quality of the product than if we choose less frequent gate-like inspections. The trade-off here is not obvious, and any choice has to be made with a specific case in mind. As a rule of thumb, we should lean toward monitoring when experience shows we are not likely to encounter big problems.

Another way to lower the cost of quality assurance is to use *variable inspections*. Because quality levels vary over time, it is only common sense to vary how often we inspect. For instance, if for weeks we don't find problems, it would seem logical to check less often. But if problems begin to develop, we can test ever more frequently until quality again returns to the previous high levels. The advantage here is still lower costs and even less interference with the production flow. Yet this approach is not used very often, even in widget manufacturing. Why not? Probably because we are creatures of habit and keep doing things the way we always have, whether it be from week to week or year to year.

Suitably thought through, intelligent inspection schemes can actually increase the efficiency and productivity of any manufacturing or administrative process. Let's take an example very different from the making of widgets or breakfasts.

I recently read a story in a news magazine that said that the American Embassy in London could not deal with a deluge of visa applications. Some one million Britons apply for visas each year, of which about 98 percent are approved. The embassy employs sixty people, who process as many as 6,000 applications a day. Most applications are received by mail, and at any time, from 60,000 to 80,000 British passports are in the embassy's hands. Meanwhile, lines of one hundred or more British and other nationals stand in front of the building, looking for an opportunity to walk their passports through. The embassy has tried a number of ways to handle matters more efficiently, including newspaper advertisements

asking tourists to apply early and to expect a three-week turnaround. The embassy also installed boxes where applicants could drop off their passports and visa applications if they really needed same-day service. Even so, the lines at the embassy remained long.

In fact, the embassy's expediting schemes only made the problem worse, because nothing was done to address the basic issue: to speed the processing of visas overall. Time and money were spent to classify various kinds of applications slated for different processing times, but this only created more logistical overhead with no effect on output.

If our government wants British tourists to visit the United States, our government should not irritate these would-be visitors. And if the embassy can't get the money to increase its staff, a simple solution can be borrowed from basic production techniques. We need, in short, to replace their present scheme with a quality assurance test.

For that, the bureaucratic minds at the embassy would need to accept that a 100 percent check of the visa applicants is unnecessary. Some 98 percent of those applying are approved without any question. So if the embassy were to institute a sampling test of visas (a quality assurance test), and a thorough one at that, the logjam of applications could be broken without materially increasing the chance that the undesirable will enter our country. Moreover, the embassy could select the sample to be checked according to predetermined criteria. The visa processing could then work rather like the Internal Revenue Service. Through the checks and audits that the IRS performs, that government agency induces compliance among most taxpayers without having an agent look at every single return.

Later, when we examine managerial productivity, we'll see that when a manager digs deeply into a specific activity under his jurisdiction, he's applying the principle of variable inspection. If the manager examined everything his various subordinates did, he would be meddling, which for the most part would be a waste of his time. Even worse, his subordinates would become accustomed to not

being responsible for their own work, knowing full well that their supervisor will check everything out closely. The principle of variable inspection applied to managerial work nicely skirts both problems, and, as we shall see, gives us an important tool for improving managerial productivity.

Productivity

The workings of our black box can furnish us with the simplest and most useful definition of productivity. The productivity of any function occurring within it is the output divided by the labor required to generate the output. Thus, one way to increase productivity is to do whatever we are now doing, but *faster*. This could be done by reorganizing the work area or just by working harder. Here we've not changed what work we do, we've just instituted ways to do it faster—getting more *activities per employee-hour* to go on inside the black box. Because the output of the black box is proportional to the activity that occurs within it, we will get more output per hour.

There is a second way to improve productivity. We can change the *nature* of the work performed: what we do, not how fast we do it. We want to increase the ratio of output to activity, thereby increasing output even if the activity per employee-hour remains the same. As the slogan has it, we want to "work smarter, not harder."

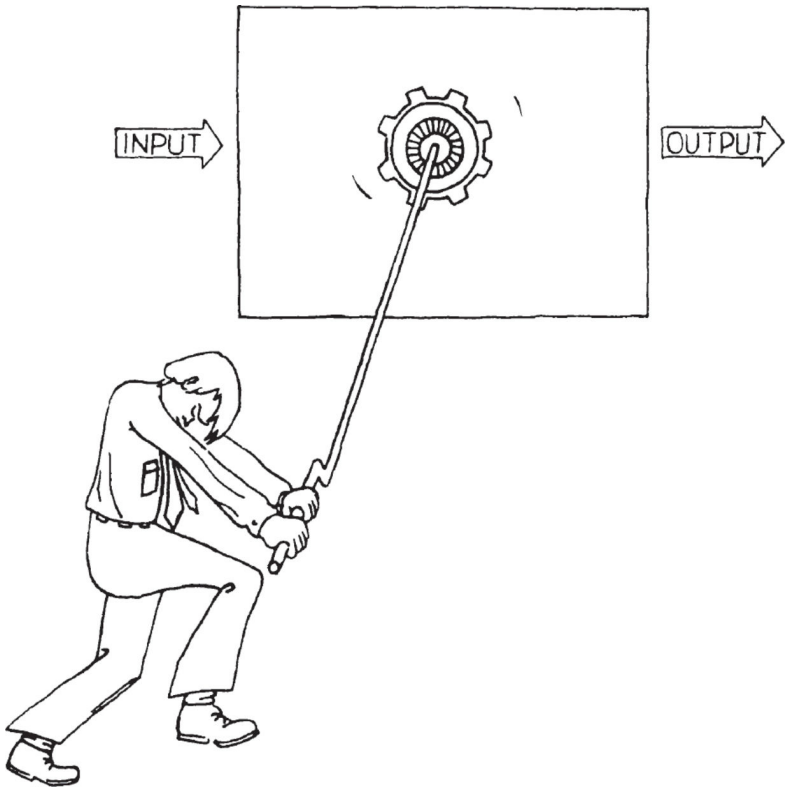

Productivity can be increased by performing the work activities at a higher rate…

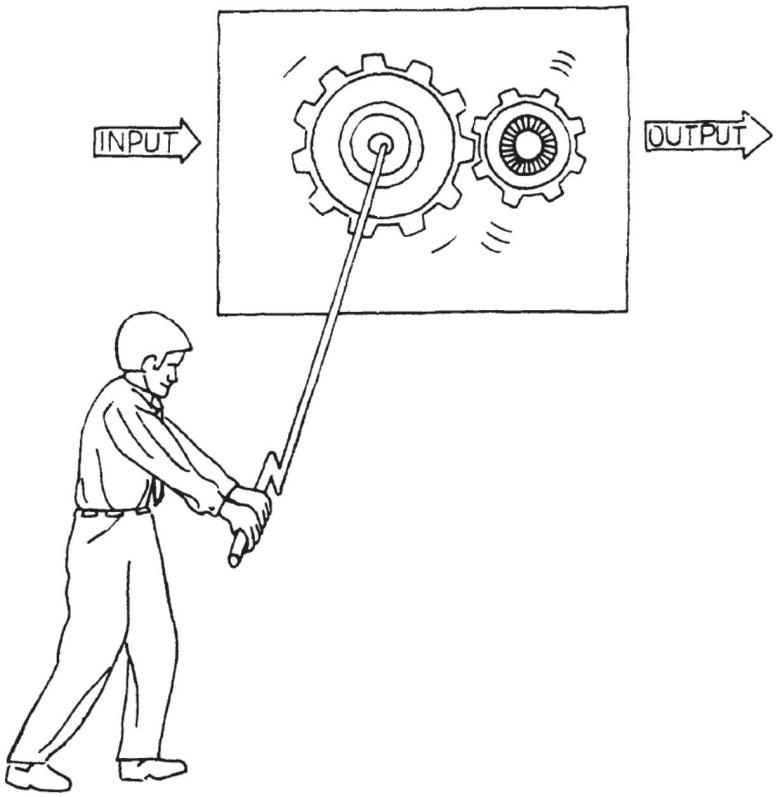

...or by increasing the leverage *of the activities.*

Here I'd like to introduce the concept of *leverage,* which is the output generated by a specific type of work activity. An activity with high leverage will generate a high level of output; an activity with low leverage, a low level of output. For example, a waiter able to boil two eggs and operate two toasters can deliver two breakfasts for almost the same amount of work as one. His output per activity, and therefore his leverage, is high. A waiter who can handle only one egg and one toaster at a time possesses lower output and leverage. The software engineer using a programming language rather like English, later to be translated by a compiler, can solve many problems per hour of programming. His output and leverage are high. A software engineer using a more cumbersome programming method of ones and zeros will require many more hours to solve the same number of problems. His output and leverage are low. Thus, a very important way to increase productivity is to arrange the work flow inside our black box so that it will be characterized by high output per activity, which is to say high-leverage activities.

Automation is certainly one way to improve the leverage of all types of work. Having machines to help them, human beings can create more output. But in both widget manufacturing and administrative work, something else can also increase the productivity of the black box. This is called *work simplification*. To get leverage this way, you first need to create a flow chart of the production process as it exists. Every single step must be shown on it; no step should be omitted in order to pretty things up on paper. Second, count the number of steps in the flow chart so that you know how many you started with. Third, set a rough target for reduction of the number of steps. In the first round of work simplification, our experience shows that you can reasonably expect a 30 to 50 percent reduction.

To implement the actual simplification, you must question *why* each step is performed. Typically, you will find that many steps exist in your work flow for no good reason. Often they are there by tradition or because formal procedure ordains it, and nothing practical requires their inclusion. Remember, the "visa factory" at our embassy in Britain didn't really have to process 100 percent of the applicants. So no matter what reason may be given for a step, you must critically question each and throw out those that common sense says you can do without. We found that in a wide range of administrative activities at Intel, substantial reduction—about 30 percent—could be achieved in the number of steps required to perform various tasks.

Of course, the principle of work simplification is hardly new in the widget manufacturing arts. In fact, this is one of the things industrial engineers have been doing for a hundred years. But the application of the principle to improve the productivity of the "soft professions"—the administrative, professional, and managerial workplace—is new and slow to take hold. The major problem to be overcome is defining what the output of such work is or should be. As we will see, in the work of the soft professions, it becomes very difficult to distinguish between output and activity. And as noted,

stressing output is the key to improving productivity, while looking to increase activity can result in just the opposite.

II

Management Is a Team Game

3

Managerial Leverage

What Is a Manager's Output?

I asked a group of middle managers just that question.
I got these responses:

judgments and opinions
direction
allocation of resources
mistakes detected
personnel trained and subordinates developed
courses taught
products planned
commitments negotiated

Do these things really constitute the output of a manager? I don't think so. They are instead activities, or descriptions of what managers do as they try to create a final result, or output. What, then, is a manager's output? At Intel, if she is in charge of a wafer fabrication plant, her output consists of completed, high-quality, fully processed silicon wafers. If he supervises a design group, his output consists of completed designs that work correctly and are ready to go into manufacturing. If a manager is the principal of a high school, her output will be trained and educated students who have either completed their schooling or are ready to move on to the next year of their studies. If a manager is a surgeon, his output is a fully recovered, healed patient. We can sum matters up with the proposition that:

A manager's output =	The output of his organization + The output of the neighboring organizations under his influence

Why? Because business and education and even surgery represent work done by *teams*.

A manager can do his "own" job, his individual work, and do it well, but that does not constitute his output. If the manager has a group of people reporting to him or a circle of people influenced by him, the manager's output must be measured by the output created by his subordinates and associates. If the manager is a knowledge specialist, a *know-how manager,* his potential for influencing "neighboring" organizations is enormous. The internal consultant who supplies needed insight to a group struggling with a problem will affect the work and the output of the entire group. Similarly, if a lawyer acquires a regulatory permit for a drug company, he will release the flow of the result of many years of research at that company to the public. Or a marketing analyst who reviews mountains of product, market, and competitive information, analyzes market research, and makes fact-finding visits can directly affect the output of many "neighboring" organizations. His interpretations of the data and his recommendations will perhaps guide the activities for the whole company. Thus, the definition of "manager" should be broadened: individual contributors who gather and disseminate know-how and information should also be seen as middle managers, because they exert great power within the organization.

But the key definition here is that the output of a manager is a result achieved by a group either under her supervision or under her influence. While the manager's own work is clearly very important, that in itself does not create output. Her organization does. By analogy, a coach or a quarterback alone does not score touchdowns and win games. Entire teams with their participation and guidance

and direction do. League standings are kept by team, not by individual. Business—and this means not just the business of commerce but the business of education, the business of government, the business of medicine—is a team activity. And, always, it takes a team to win.

It is important to understand that a manager will find himself engaging in an array of activities in order to affect output. As the middle managers I queried said, a manager must form opinions and make judgments, he must provide direction, he must allocate resources, he must detect mistakes, and so on. All these are necessary to achieve output. But output and activity are by no means the same thing.

Consider my own managerial role. As president of a company, I can affect output through my direct subordinates—group general managers and others like them—by performing supervisory activities. I can also influence groups not under my direct supervision by making observations and suggestions to those who manage them. Both types of activity will, I hope, contribute to my output as a manager by contributing to the output of the company as a whole. I was once asked by a middle manager at Intel how I could teach in-plant courses, visit manufacturing plants, concern myself with the problems of people several levels removed from me in the organization, and still have time to do my job. I asked him what he thought my job was. He thought for a moment, and then answered his own question, “I guess those things are your job too, aren’t they?” They are absolutely my job—not my entire job, but part of it, because they help add to the output of Intel.

Let me give another example. Cindy, an engineer at Intel, supervises an engineering group in a wafer fabrication plant. She also spends some of her time as a member of an advisory body that establishes standard procedures by which all the plants throughout the company perform a certain technical process. In both roles, Cindy contributes to the output of the wafer fabrication plants. As a supervising engineer, she performs activities that increase the output of the plant in which she works; as a member of the advisory body,

she provides specialized knowledge that will influence and increase the output of "neighboring organizations"—all the other Intel wafer fabrication plants.

Let's refer again to our black box. If the machinery within an organization can be compared to a series of gears, we can visualize how a middle manager affects output. In times of crisis, he provides power to the organization. When things aren't working as smoothly as they should, he applies a bit of oil. And, of course, he provides intelligence to the machine to direct its purpose.

"Daddy, What Do You Really Do?"

Most of us have had to struggle to answer that question. What we actually do is difficult to pin down and sum up. Much of it often seems so inconsequential that our position in the business hardly seems justified. Part of the problem here stems from the distinction between our activities, which is what we actually do, and our output, which is what we achieve. The latter seems important, significant, and worthwhile. The former often seems trivial, insignificant, and messy. But a surgeon whose output is a cured patient spends his time scrubbing and cutting and suturing, and this hardly sounds very respectable either.

To find out what we managers really do, let's take a look at one of my busier days, shown in the table below. Here I describe the activity in which I was engaged, explain it a bit, and categorize it into types we shall examine in the balance of the chapter.

A Day from My Life

Time and Activity	Explanation (Type of Activity)
8:00–8:30	
Met with a manager who had submitted his resignation to	I listened to his reasons *(information-gathering),* felt he could be turned around and saved for Intel. Encouraged him to

leave for another company.	talk to certain other managers about a career change *(nudge),* and decided to pursue this matter with them myself *(decision-making).*
Incoming telephone call from a competitor.	Call was ostensibly about a meeting of an industry-wide society, but in reality he was feeling out how I saw business conditions. I did the same. *(Information-gathering.)*
8:30–9:00	
Read mail from the previous afternoon.	I scribbled messages on about half of it, some of which were expressions of encouragement or disapproval, others exhortations to take certain types of action *(nudges).* One was the denial of a request to proceed with a particular small project *(decision-making).* (Of course, *information-gathering* took place in all of these cases, too.)
9:00–12:00	
Executive Staff Meeting (a regular weekly meeting of the company's senior management). Subjects covered at this particular one:	
—Review of the prior month's incoming order and shipment rates.	*(Information-gathering)*
—Discussion to set priorities for the upcoming annual planning process.	*(Decision-making)*

—Review of the status of a major marketing program (scheduled subject).	This came about through a prior decision that this program was faltering and required review. We found it was doing a little bit better than before *(information-gathering),* but the presentation still elicited a lot of comments and suggestions *(nudges)* from various members of the audience.
—Review of a program to reduce the manufacturing cycle time of a particular product line (scheduled subject).	The presentation indicated that the program was in good shape. (It represented only *information-gathering;* no further action was stimulated.)
12:00–1:00	
Lunch in the company cafeteria.	I happened to sit with members of our training organization, who complained about the difficulty they had in getting me and other senior managers to participate in training at our foreign locations *(information-gathering).* This was news to me. I made a note to follow up with my own schedule, as well as with my staff, and to *nudge* them into doing a better job supporting the foreign training program.
1:00–2:00	
Meeting regarding a specific product-quality problem.	The bulk of the meeting involved getting sufficient information on the status of the product and the corrective action that had been implemented (*information-gathering*). The meeting ended in a *decision* made by the division manager, with my concurrence, to resume shipment of the product.
2:00–4:00	

Lecture at our employee orientation program.	This is a program in which senior management gives all professional employees a presentation describing the objectives, history, management systems, etc., of the company and its major groups. I am the first lecturer in the series. This clearly represented *information-giving,* and I was a *role model* not only in communicating the importance we place on training, but also, by my handling of questions and comments, in representing, in living form, some of the values of the company. The nature of the questions, at the same time, gave me a feeling for the concerns and understanding level of a large number of employees to whom I would not otherwise have access. So this also represented *information-gathering,* characteristic of the "visit" type in its efficiency.
4:00–4:45	
In the office, returning phone calls.	I disapproved granting a compensation increase to a particular employee, which I thought was way outside of the norm (clearly a *decision*). I decided to conduct a meeting with a group of people to decide what organization would move to a new site we were opening in another state. (This was a *decision* to hold a decision-making meeting.)
4:45–5:00	
Met with my assistant.	Discussed a variety of requests for my time for a number of meetings in the upcoming week. Suggested alternatives where I *decided* not to attend.
5:00–6:15	
Read the day's mail,	As with the morning's mail reading, this

including progress reports.	was *information-gathering,* interspersed with *nudging* and *decision-making* through my annotations and messages scribbled on much of it.

When you look at what happened, you won't see any obvious patterns. I dealt with things in seemingly random fashion. My wife's reaction to my day was that it looked very much like one of her own. She was right in noting a similarity. My day always ends when I'm tired and ready to go home, not when I'm done. I am never done. Like a housewife's, a manager's work is never done. There is always more to be done, more that should be done, always more than can be done.

A manager must keep many balls in the air at the same time and shift his energy and attention to activities that will most increase the output of his organization. In other words, he should move to the point where his *leverage* will be the greatest.

As you can see, much of my day is spent acquiring information. And as you can also see, I use many ways to get it. I read standard reports and memos but also get information ad hoc. I talk to people inside and outside the company, managers at other firms or financial analysts or members of the press. Customer complaints, both external and internal, are also a very important source of information. For example, the Intel training organization, which I serve as an instructor, is an internal customer of mine. To cut myself off from the casual complaints of people in that group would be a mistake because I would miss getting an evaluation of my performance as an internal "supplier." People also tell us things because they want us to do something for them; to advance their case, they will sometimes shower us with useful information. This is something we should remember, apart from whether we do as they ask.

I have to confess that the information most useful to me, and I suspect most useful to all managers, comes from quick, often casual verbal exchanges. This usually reaches a manager much faster than

anything written down. And usually the more timely the information, the more valuable it is.

So why are written reports necessary at all? They obviously can't provide timely information. What they do is constitute an archive of data, help to validate ad hoc inputs, and catch, in safety-net fashion, anything you may have missed. But reports also have another totally different function. As they are formulated and written, the author is forced to be more precise than he might be verbally. Hence their value stems from the discipline and the thinking the writer is forced to impose upon himself as he identifies and deals with trouble spots in his presentation. Reports are more a *medium* of *self-discipline* than a way to communicate information. *Writing* the report is important; reading it often is not.

There are many parallels to this. As we will see later, the *preparation* of an annual plan is in itself the end, not the resulting bound volume. Similarly, our capital authorization *process* itself is important, not the authorization itself. To prepare and justify a capital spending request, people go through a lot of soul-searching analysis and juggling, and it is this mental exercise that is valuable. The formal authorization is useful only because it enforces the discipline of the process.

To improve and maintain your capacity to get information, you have to understand the way it comes to you. There's a hierarchy involved. Verbal sources are the most valuable, but what they provide is also sketchy, incomplete, and sometimes inaccurate, like a newspaper headline that can give you only the general idea of a story. A headline can't give any of the details and might even give you a distorted idea of what the real story is. So you then read the newspaper article itself to find out who, what, where, why, and how. After this, you should have some reiteration and perspective, which can be compared to reading a news magazine or even a book. Each level in your information hierarchy is important, and you can rely on none alone. Though the most thorough information might come from the news magazine, you do not, of course, want to wait a full week after an event to find out about it. Your information sources

should complement one another, and also be redundant because that gives you a way to verify what you've learned.

There is an especially efficient way to get information, much neglected by most managers. That is to visit a particular place in the company and observe what's going on there. Why should you do this? Think of what happens when somebody comes to see a manager in his office. A certain stop-and-start dynamics occurs when the visitor sits down, something socially dictated. While a two-minute kernel of information is exchanged, the meeting often takes a half hour. But if a manager walks through an area and sees a person with whom he has a two-minute concern, he can simply stop, cover it, and be on his way. Ditto for the subordinate when he initiates conversation. Accordingly, such visits are an extremely effective and efficient way to transact managerial business.

Then why are they underutilized? Because of the awkwardness that managers feel about walking through an area without a specific task in mind. At Intel we combat this problem by using programmed visits meant to accomplish formal tasks, but which also set the stage for ad hoc mini-transactions. For example, we ask our managers to participate in "Mr. Clean" inspections, in which they go to a part of the company that they normally wouldn't visit. The managers examine the housekeeping, the arrangement of things, the labs, and the safety equipment, and in so doing spend an hour or so browsing around and getting acquainted with things firsthand.

As can be seen from my schedule, a manager not only gathers information but is also a source of it. He must convey his knowledge to members of his own organization and to other groups he influences. Beyond relaying facts, a manager must also communicate his objectives, priorities, and preferences as they bear on the way certain tasks are approached. This is extremely important, because only if the manager imparts these will his subordinates know how to make decisions themselves that will be acceptable to the manager, their supervisor. Thus, transmitting objectives and preferred approaches constitutes a key to successful delegation. As we will see later, a shared corporate culture becomes indispensable to a

business. Someone adhering to the values of a corporate culture—an intelligent corporate citizen—will behave in consistent fashion under similar conditions, which means that managers don't have to suffer the inefficiencies engendered by formal rules, procedures, and regulations that are sometimes used to get the same result.

The third major kind of managerial activity, of course, is decision-making. To be sure, once in a while we managers in fact *make* a decision. But for every time that happens, we *participate* in the making of many, many others, and we do that in a variety of ways. We provide factual inputs or just offer opinions, we debate the pros and cons of alternatives and thereby force a better decision to emerge, we review decisions made or about to be made by others, encourage or discourage them, ratify or veto them.

Just how decisions should be made, we'll talk about later. Meanwhile, let's say that decisions can be separated into two kinds. The forward-looking sort are made, for example, in the capital authorization process. Here we allocate the financial resources of the company among various future undertakings. The second type is made as we respond to a developing problem or a crisis, which can either be technical (a quality control problem, for example) or involve people (talking somebody out of quitting).

It's obvious that your decision-making depends finally on how well you comprehend the facts and issues facing your business. This is why information-gathering is so important in a manager's life. Other activities—conveying information, making decisions, and being a role model for your subordinates—are all governed by the *base of information* that you, the manager, have about the tasks, the issues, the needs, and the problems facing your organization. In short, information-gathering is the basis of all other managerial work, which is why I choose to spend so much of my day doing it.

You often do things at the office designed to influence events slightly, maybe making a phone call to an associate suggesting that a decision be made in a certain way, or sending a note or a memo that shows how you see a particular situation, or making a comment during an oral presentation. In such instances you may be

advocating a preferred course of action, but you are not issuing an instruction or a command. Yet you're doing something stronger than merely conveying information. Let's call it "nudging" because through it you nudge an individual or a meeting in the direction you would like. This is an immensely important managerial activity in which we engage all the time, and it should be carefully distinguished from decision-making that results in firm, clear directives. In reality, for every unambiguous decision we make, we probably nudge things a dozen times.

Finally, something more subtle pervades the day of all managers. While we move about, doing what we regard as our jobs, we are *role models* for people in our organization—our subordinates, our peers, and even our supervisors. Much has been said and written about a manager's need to be a leader. The fact is, no single managerial activity can be said to constitute leadership, and nothing leads as well as example. By this I mean something straightforward. Values and behavioral norms are simply not transmitted easily by talk or memo, but are conveyed very effectively by doing and doing *visibly*.

All managers need to act so that they can be seen exerting influence, but they should do so in their own way. Some of us feel comfortable dealing with large groups and talking about our feelings and values openly in that fashion. Others prefer working one-on-one with people in a quieter, more intellectual environment. These and other styles of leadership will work, but only if we recognize and consciously stress the need for us to be role models for people in our organization.

Don't think for a moment that the way I've described leadership applies only to large operations. An insurance agent in a small office who continually talks with personal friends on the phone imparts a set of values about permissible conduct to everyone working for him. A lawyer who returns to his office after lunch a little drunk does the same. On the other hand, a supervisor in a company, large or small, who takes his work seriously exemplifies to his associates the most important managerial value of all.

A great deal of a manager's work has to do with allocating resources: manpower, money, and capital. But the single most important resource that we allocate from one day to the next is our own time. In principle more money, more manpower, or more capital can always be made available, but our own time is the one absolutely finite resource we each have. Its allocation and use therefore deserve considerable attention. How you handle your own time is, in my view, the single most important aspect of being a role model and leader.

As you can see, in a typical day of mine one can count some twenty-five separate activities in which I participated, mostly information-gathering and -giving, but also decision-making and nudging. You can also see that some two thirds of my time was spent in a meeting of one kind or another. Before you are horrified by how much time I spend in meetings, answer a question: which of the activities—information-gathering, information-giving, decision-making, nudging, and being a role model—could I have performed outside a meeting? The answer is practically none. Meetings provide an occasion for managerial activities. Getting together with others is not, of course, an activity—it is a *medium*. You as a manager can do your work in a meeting, in a memo, or through a loudspeaker for that matter. But you must choose the most effective medium for what you want to accomplish, and that is the one that gives you the greatest leverage. More about meetings later.

Leverage of Managerial Activity

We've established that the output of a manager is the output of the various organizations under his control and his influence. What can a manager do to increase his output? To find out, let's look at the concept of *leverage*. Leverage is the measure of the output generated by any given managerial activity. Accordingly, managerial output can be linked to managerial activity by the equation:

Managerial Output	=	Output of organization

$$= L_1 \times A_1 + L_2 \times A_2 + \ldots$$

This equation says that for every activity a manager performs—A_1, A_2, and so on—the output of the organization should increase by some degree. The extent to which that output is thereby increased is determined by the leverage of that activity—L_1, L_2, and so on. A manager's output is thus the sum of the result of individual activities having varying degrees of leverage. Clearly the key to high output means being sensitive to the *leverage* of what you do during the day.

Managerial productivity—that is, the output of a manager per unit of time worked—can be increased in three ways:

1. Increasing the rate with which a manager performs his activities, speeding up his work.
2. Increasing the leverage associated with the various managerial activities.
3. Shifting the mix of a manager's activities from those with lower to those with higher leverage.

Let us consider first the leverage of various types of managerial work.

HIGH-LEVERAGE ACTIVITIES

These can be achieved in three basic ways:

- When many people are affected by one manager.
- When a person's activity or behavior over a long period of time is affected by a manager's brief, well-focused set of words or actions.
- When a large group's work is affected by an individual supplying a unique, key piece of knowledge or information.

The first is the most obvious example. Consider Robin, an Intel finance manager, responsible for setting up the annual financial planning process for the company. When Robin defines in advance

exactly what information needs to be gathered and presented at each stage of the planning process and lays out who is responsible for what, she directly affects the subsequent work of perhaps two hundred people who participate in the planning process. By spending a certain amount of time *in advance* of the planning activities, Robin will help to eliminate confusion and ambiguity for a large population of managers over an extended period of time. Consequently, her work contributes to the productivity of the entire organization and clearly has great leverage, leverage that depends, however, on *when* it is performed. Work done in advance of the planning meeting obviously has great leverage. If Robin has to scramble later to help a manager define guidelines and milestones, her work will clearly have much less leverage.

Another example of leverage that depends on timely action is what you do when you learn that a valued subordinate has decided to quit. In such a case, you must direct yourself to the situation immediately if you want to change the person's mind. If you put it off, all your chances are lost. Thus to maximize the leverage of his activities, a manager must keep *timeliness,* which is often critical, firmly in mind.

Leverage can also be negative. Some managerial activities can *reduce* the output of an organization. I mean something very simple. Suppose I am a key participant at a meeting and I arrive unprepared. Not only do I waste the time of the people attending the meeting because of my lack of preparation—a direct cost of my carelessness—but I deprive the other participants of the opportunity to use that time to do something else.

Each time a manager imparts his knowledge, skills, or values to a group, his leverage is high, as members of the group will carry what they learn to many others. But again, leverage can be positive or negative. An example of leverage that I hope is high and positive is my talk in the orientation course. During the two hours I have, I try to impart a great deal of information about Intel—its history, its objectives, its values, its style—to a group of two hundred new employees. Besides what I say specifically, my approach toward answering questions and my conduct in general communicate our

way of doing things to these employees when they are most impressionable.

Here is another example of this kind of leverage. To train a group of salesmen, Barbara, an Intel marketing engineer, sets out to teach them what the organization's products are. If she does her job well, the salespeople will be better equipped to sell the line. If she does it poorly, great and obvious damage is done.

A final, less formal, example here: Cindy, as you recall, is a member of a technical coordinating body in which she tries to disseminate her understanding of a specific technology to all of the company's manufacturing groups. In effect, she uses the coordinating body as an informal training vehicle to effect high leverage on her counterparts in neighboring Intel organizations.

A manager can also exert high leverage by engaging in an activity that takes him only a *short* time, but that affects another person's performance over a *long* time. A performance review represents a good example of this. With the few hours' work that a manager spends preparing and delivering the review, he can affect the work of its recipient enormously. Here too a manager can exert either positive or negative leverage. A subordinate can be motivated and even redirected in his efforts, or the review can discourage and demoralize him for who knows how long.

Another seemingly trivial piece of work—creating a tickler file—can improve daily work significantly for a long time. Setting up the simple mechanical aid is a one-time activity, yet it is likely to improve the productivity of the manager who uses it indefinitely. Thus the leverage here is very, very high.

Examples of high negative leverage abound. After going through the annual planning process, an Intel manager saw that, in spite of successful cost reduction efforts in the prior year, his division was still not going to make any money in the coming year. The manager became depressed. Though he didn't realize it, he almost immediately began to affect people around him and soon depression spread throughout his organization. He snapped out of it only when someone on his staff finally told him what he was doing to the people

under him. Another example is waffling, when a manager puts off a decision that will affect the work of other people. In effect, the lack of a decision is the same as a negative decision; no green light is a red light, and work can stop for a whole organization.

Both the depressed and the waffling manager can have virtually unlimited negative leverage. If people are badly affected by a poor sales training effort, the situation can be handled by retraining the group. But the negative leverage produced by depression and waffling is very hard to counter because their impact on an organization is both so pervasive and so elusive.

Managerial meddling is also an example of negative leverage. This occurs when a supervisor uses his superior knowledge and experience of a subordinate's responsibilities to assume command of a situation rather than letting the subordinate work things through himself. For example, if a senior manager sees an indicator showing an undesirable trend and dictates to the person responsible a detailed set of actions to be taken, that is managerial meddling. In general, meddling stems from a supervisor exploiting too much superior work knowledge (real or imagined). The negative leverage produced comes from the fact that after being exposed to many such instances, the subordinate will begin to take a much more restricted view of what is expected of him, showing less initiative in solving his own problems and referring them instead to his supervisor. Because the output of the organization will consequently be reduced in the long run, meddling is clearly an activity having negative managerial leverage.

The third kind of managerial activity with high leverage is exercised by a person with *unique skills* and *knowledge*. One such person is an Intel marketing engineer responsible for setting prices for the product line. Hundreds of salespeople in the field can be negatively affected if prices are set too high: no matter how hard they may try, they won't be able to get any business. Of course, if the prices were set too low, we would be giving money away.

Take another example. An Intel development engineer who has uniquely detailed knowledge of a particular manufacturing process

effectively controls how it is used. Since the process will eventually provide the foundation for the work of many product designers all over the company, the leverage the development engineer exerts is enormous. The same is true for a geologist in an oil company or an actuary in an insurance firm. All are specialists whose work is important for the work of their organization at large. The person who comprehends the critical facts or has the critical insights—the "knowledge specialist" or the "know-how manager"—has tremendous authority and influence on the work of others, and therefore very high leverage.

The *art* of management lies in the capacity to select from the many activities of seemingly comparable significance the one or two or three that provide leverage well beyond the others and concentrate on them. For me, paying close attention to customer complaints constitutes a high-leverage activity. Aside from making a customer happy, the pursuit tends to produce important insights into the workings of my own operation. Such complaints may be numerous, and though all of them need to be followed up by someone, they don't all require or wouldn't all benefit from my personal attention. Which one out of ten or twenty complaints to dig into, analyze, and follow up is where art comes into the work of a manager. The basis of that art is an intuition that behind this complaint and not the other lurk many deeper problems.

DELEGATION AS LEVERAGE

Because managerial time has a hierarchy of values, delegation is an essential aspect of management. The "delegator" and "delegatee" must share a common information base and a common set of operational ideas or notions on how to go about solving problems, a requirement that is frequently not met. Unless both parties share the relevant common base, the delegatee can become an effective proxy only with specific instructions. As in meddling, where specific activities are prescribed in detail, this produces low managerial leverage.

Picture this. I am your supervisor, and I walk over to you with pencil in hand and tell you to take it. You reach for the pencil, but I won't let go. So I say, "What is wrong with you? Why can't I delegate the pencil to you?" We all have some things that we don't really *want* to delegate simply because we like doing them and would rather not let go. For your managerial effectiveness, this is not too bad so long as it is based on a *conscious* decision that you will hold on to certain tasks that you enjoy performing, even though you could, if you chose, delegate them. But be sure to know exactly what you're doing, and avoid the charade of insincere delegation, which can produce immense negative managerial leverage.

Given a choice, should you delegate activities that are familiar to you or those that aren't? Before answering, consider the following principle: delegation without follow-through is *abdication*. You can never wash your hands of a task. Even after you delegate it, you are still responsible for its accomplishment, and monitoring the delegated task is the only practical way for you to ensure a result. Monitoring is not meddling, but means checking to make sure an activity is proceeding in line with expectations. Because it is easier to monitor something with which you are familiar, if you have a choice you should delegate those activities you know best. But recall the pencil experiment and understand before the fact that this will very likely go against your emotional grain.

Please turn back to the table of my day's activities on this page. During the executive staff meeting we heard two follow-up presentations, one on the status of an extremely important marketing program and the other on the progress of a program aimed at reducing manufacturing throughput times. Both reviews are examples of monitoring. Earlier, we had assigned each to a middle manager and made sure these managers and the senior staff agreed about what the programs were to be. The middle managers then went about their business expecting to report back to the executive staff, the body that delegated the programs to them.

Monitoring the results of delegation resembles the monitoring used in quality assurance. We should apply quality assurance

principles and monitor at the lowest-added-value stage of the process. For example, review *rough drafts* of reports that you have delegated; don't wait until your subordinates have spent time polishing them into final form before you find out that you have a basic problem with the contents. A second principle applies to the frequency with which you check your subordinates' work. A variable approach should be employed, using different sampling schemes with various subordinates; you should increase or decrease your frequency depending on whether your subordinate is performing a newly delegated task or one that he has experience handling. How often you monitor should not be based on what you believe your subordinate can do *in general,* but on his experience with a specific task and his prior performance with it—his task-relevant maturity, something I'll talk about in detail later. As the subordinate's work improves over time, you should respond with a corresponding reduction in the intensity of the monitoring.

To use quality assurance principles effectively, the manager should only go into details randomly, just enough to try to ensure that the subordinate is moving ahead satisfactorily. To check into *all* the details of a delegated task would be like quality assurance testing 100 percent of what manufacturing turned out.

Making certain *types* of decisions is something managers frequently delegate to subordinates. How is this best done? By monitoring their decision-making *process*. How do you do that? Let's examine what Intel goes through to approve a capital equipment purchase. We ask a subordinate to think through the entire matter carefully before presenting a request for approval. And to monitor how good his thinking is, we ask him quite specific questions about his request during a review meeting. If he answers them convincingly, we'll approve what he wants. This technique allows us to find out how good the thinking is without having to go through it ourselves.

Increasing Managerial Activity Rate: Speeding Up the Line

Of course, the most obvious way to increase managerial output is to increase the *rate,* or speed, of performing work. The relationship here is:

$$\frac{\text{Managerial Output}}{\text{Time}} = L \times \frac{\text{Activity Performed}}{\text{Time}}$$

where L is the leverage of the activity.

The most common approach to increasing a manager's productivity—his output over time—has been time-management techniques, which try to reduce the denominator on both sides of this equation. Any number of consultants will tell a manager that the way to higher productivity is to handle a piece of paper only once, to hold only stand-up meetings (which will presumably be short), and to turn his desk so that he presents his back to the door.

These time-management suggestions can be improved upon, I think, by applying our production principles. First, we must identify our *limiting step:* what is the "egg" in our work? In a manager's life some things really have to happen on a schedule that is absolute. For me, an example is the class I teach. I know when it is going to meet, and I know I must prepare for it. There is no "give" in the time here, because over two hundred students will be expecting me. Accordingly, I have to create offsets and schedule my other work around this limiting step. In short, if we determine what is immovable and manipulate the more yielding activities around it, we can work more efficiently.

A second production principle we can apply to managerial work is *batching* similar tasks. Any manufacturing operation requires a certain amount of set-up time. So for managerial work to proceed efficiently, we should use the same set-up effort to apply across a group of similar activities. Think about our continuous egg-boiler, which was installed to produce fine-quality, identical, three-minute eggs. Should we now decide to serve our customers four-minute eggs, we would have to slow down the conveyor belt moving them through the hot water. The adjustment takes time: not only do we

adjust nuts and bolts on the machine, we also have to inspect the quality of the four-minute eggs by sampling a few of them.

Set-up time has many parallels in managerial work. For example, once we have prepared a set of illustrations for a training class, we will obviously increase our productivity if we can use the same set over and over again with other classes or groups. Similarly, if a manager has a number of reports to read or a number of performance reviews to approve, he should set aside a block of time and do a batch of them together, one after the other, to maximize the use of the *mental* set-up time needed for the task.

What makes running a factory different from running a job shop? The latter is prepared to service any customer who drops in; the owner handles the job required and moves on to the next one. A factory, on the other hand, is usually run by *forecast* and not by individual order. From my experience a large portion of managerial work *can be* forecasted. Accordingly, forecasting those things you can and setting yourself up to do them is only common sense and an important way to minimize the feeling and the reality of fragmentation experienced in managerial work. Forecasting and planning your time around key events are literally like running an efficient factory.

What is the *medium* of a manager's forecast? It is something very simple: his *calendar*. Most people use their calendars as a repository of "orders" that come in. Someone throws an order to a manager for his time, and it automatically shows up on his calendar. This is mindless passivity. To gain better control of his time, the manager should use his calendar as a "production" planning tool, taking a firm initiative to schedule work that is not time-critical between those "limiting steps" in the day.

Another production principle can be applied here. Because manufacturing people trust their indicators, they won't allow material to begin its journey through the factory if they think it is already operating at capacity. If they did, material might go halfway through and back up behind a bottleneck. Instead, factory managers say "no" at the outset and keep the start level from overloading the

system. Other kinds of managers find this hard to apply because their indicators of capacity are not as well established or not as believable. How much time do you need to read your mail, to write your reports, to meet with a colleague? You may not know precisely, but you surely have a feel for the time required. And you should exploit that sense to schedule your work.

To use your calendar as a production-planning tool, you must accept responsibility for two things:

1. You should move toward the *active* use of your calendar, taking the initiative to fill the holes between the time-critical events with non-time-critical though necessary activities.
2. You should say "no" at the outset to work beyond your capacity to handle.

It is important to say "no" earlier rather than later because we've learned that to wait until something reaches a higher value stage and then abort due to lack of capacity means losing more money and time. You can obviously say "no" either explicitly or implicitly, because by not delivering you end up saying what amounts to "no." Remember too that your time is your one finite resource, and when you say "yes" to one thing you are inevitably saying "no" to another.

The next production principle you can apply is to allow *slack*—a bit of looseness in your scheduling. Highway planners, for example, know that a freeway can handle an optimum number of vehicles. Having fewer cars means that the road is not being used at capacity. But at that optimum point, if just a few more cars are allowed to enter the traffic flow, everything comes to a crunching halt. With the new metering devices that control access during the rush hour, planners can get a fix on the right number. The same thing can be done for managerial work. There is an optimum degree of loading, with enough slack built in so that one unanticipated phone call will not ruin your schedule for the rest of the day. You need some slack.

Another production principle is very nearly the opposite. A manager should carry a raw material *inventory* in terms of projects.

This is not to be confused with his work-in-process inventory, because that, like eggs in a continuous boiler, tends to spoil or become obsolete over time. Instead this inventory should consist of things you need to do but don't need to finish right away—discretionary projects, the kind the manager can work on to increase his group's productivity over the long term. Without such an inventory of projects, a manager will most probably use his free time *meddling* in his subordinates' work.

A final principle. Most production practices follow well-established procedures and, rather than reinventing the wheel repeatedly, use a specific method that has been shown to work before. But managers tend to be inconsistent and bring a welter of approaches to the same task. We should work to change that. As we become more consistent, we should also remember that the value of an administrative procedure is contained not in formal statements but in the real thinking that led to its establishment. This means that even as we try to standardize what we do, we should continue to think critically about what we do and the approaches we use.

Built-In Leverage: How Many Subordinates Should You Have…

An important component of managerial leverage is the number of subordinates a manager has. If he does not have enough, his leverage is obviously reduced. If he has too many, he gets bogged down—with the same result. As a rule of thumb, a manager whose work is largely supervisory should have six to eight subordinates; three or four are too few and ten are too many. This range comes from a guideline that a manager should allocate about a half day per week to each of his subordinates. (Two days a week per subordinate would probably lead to meddling; an hour a week does not provide enough opportunity for monitoring.)

The six to eight rule is right for the classically hierarchical manager whose primary work is the supervision of others. What about a know-how manager, the middle manager who mainly supplies expertise and information? Even if he works without a single subordinate, servicing a number of varied "customers" as an internal

consultant can in itself be a full-time job. In fact, anyone who spends about a half day per week as a member of a planning, advisory, or coordinating group has the equivalent of a subordinate. So as a rule of thumb, if a manager is both a hierarchical supervisor and a supplier of know-how, he should try to have a total of six to eight subordinates or their equivalent.

Sometimes a business is organized in a way that makes the ideal fan-out of six to eight subordinates hard to reach. A manufacturing plant, for example, may have an engineering section and a production section, in which case the plant manager would only have two people reporting directly to him. The manager might then choose to "act" as one of the two subordinates, choosing to be his own engineering manager, for instance. If he does that, the manufacturing manager will still report to him, and he will have added the people who would ordinarily report to the head of engineering. So the plant manager will actually have six direct reports: five engineers and the manufacturing manager. The arrangement, shown below, does not have the engineers *appearing to be* at the same organizational level as the manufacturing manager —something he would surely take exception to.

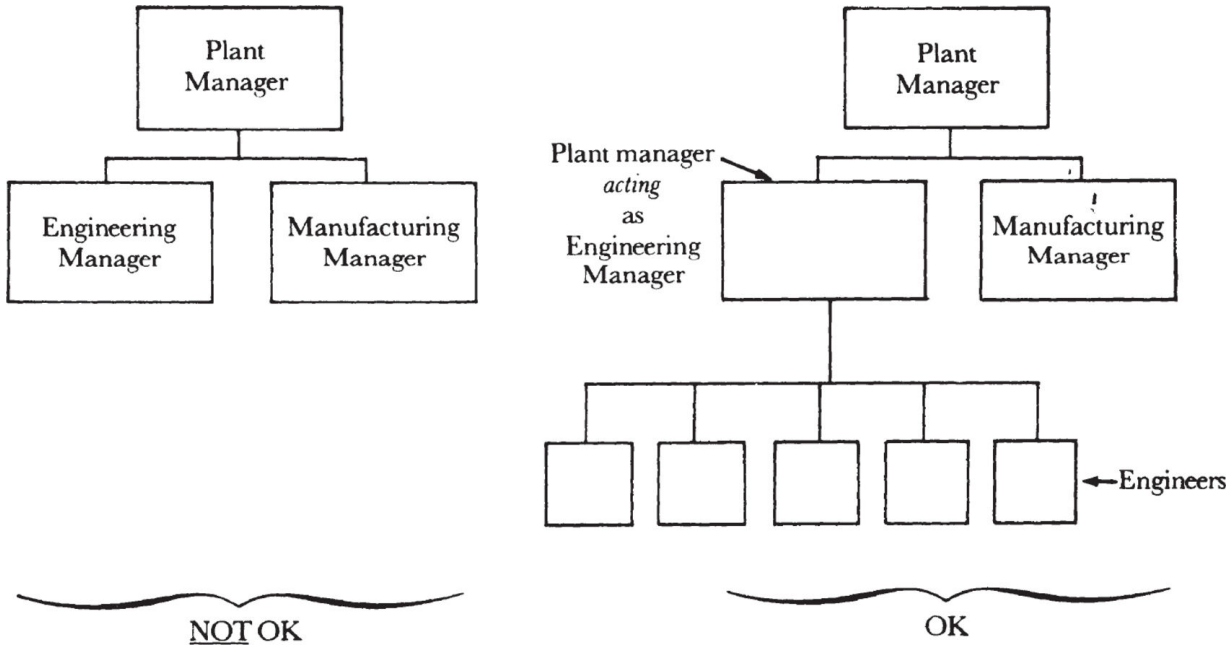

This arrangement will avoid forcing the plant manager either into on-the-job retirement or into meddling.

Interruptions—The Plague of Managerial Work

The next important production concept we can apply to managerial work is to strive toward *regularity*. We could obviously run our breakfast factory more efficiently if customers arrived in a steady and predictable stream rather than dropping in by ones and twos. Though we can't control our customers' habits, we should try to smooth out our workload as much as possible. As noted, we should try to make our managerial work take on the characteristics of a factory, not a job shop. Accordingly, we should do everything we can to prevent little stops and starts in our day as well as interruptions brought on by big emergencies. Even though some of the latter are unavoidable, we should always be looking for sources of future high-priority trouble by cutting windows into the black box of our organization. Recognizing you've got a time bomb on your hands means you can address a problem *when you want to,* not after the bomb has gone off.

But because you must coordinate your work with that of other managers, you can only move toward regularity if others do too. In

other words, the same blocks of time must be used for like activities. For example, at Intel Monday mornings have been set aside throughout the corporation as the time when planning groups meet. So anybody who belongs to one can count on Monday for that purpose and be free of scheduling conflicts.

About twenty middle managers at Intel were once asked to be part of an experiment. After pairing up, they tried some role-playing in which one manager was to define the problem most limiting his output and the other was to be a consultant who would analyze the problem and propose solutions.

The most common problem cited was *uncontrolled interruptions,* which in remarkably uniform fashion affected both supervisory and know-how managers. Everyone felt that the interruptions got in the way of his "own" work. Interruptions had a common source, most frequently coming from subordinates and from people outside the managers' immediate organization but whose work the managers influenced. For those in manufacturing, the interruptions most often came from production operators, and for marketing people, from outside customers: in short, from the consumers of the middle managers' authority and information.

The most frequently proposed solutions were not very practical. The idea mentioned most often was to create blocks of time for individual work by hiding physically. But this is a less than happy answer, because the interrupters obviously have legitimate problems, and if the manager responded by hiding, these would pile up. One "solution" was a suggestion that customers not call marketing managers at certain hours. No good.

There are better ways. Let's apply a production concept. Manufacturers turn out *standard products*. By analogy, if you can pin down what kind of interruptions you're getting, you can prepare standard responses for those that pop up most often. Customers don't come up with totally new questions and problems day in and day out, and because the same ones tend to surface repeatedly, a manager can reduce time spent handling interruptions using

standard responses. Having them available also means that a manager can delegate much of the job to less experienced personnel.

Also, if you use the production principle of *batching*—that is, handling a group of similar chores at one time—many interruptions that come from your subordinates can be accumulated and handled not randomly, but at staff and at one-on-one meetings, the subject of the next chapter. If such meetings are held regularly, people can't protest too much if they're asked to batch questions and problems for *scheduled* times, instead of interrupting you whenever they want.

The use of indicators, especially the bank of indicators kept over time, can also reduce the time a manager spends dealing with interruptions. How fast he can answer a question depends on how fast he can put his finger on the information he needs for a response. By maintaining an archive of information, a manager doesn't have to do ad hoc research every time the phone rings.

If the people who interrupt you knew how much they were disturbing you, they would probably police themselves more closely and cut down on the number of times they felt they had to talk to you right away. In any case, a manager should try to force his frequent interrupters to make an *active* decision about whether an issue can wait. So, instead of going into hiding, a manager can hang a sign on his door that says, "I am doing individual work. Please don't interrupt me unless it really can't wait until 2:00." Then hold an open office hour, and be completely receptive to anybody who wants to see you. The key is this: understand that interrupters have legitimate problems that need to be handled. That's why they're bringing them to you. But you can channel the time needed to deal with them into organized, scheduled form by providing an *alternative* to interruption—a scheduled meeting or an office hour.

The point is to impose a *pattern* on the way a manager copes with problems. To make something regular that was once irregular is a fundamental production principle, and that's how you should try to handle the interruptions that plague you.

4

Meetings—The Medium of Managerial Work

Meetings have a bad name. One school of management thought considers them the curse of the manager's existence. Someone who did a study found that managers spend up to 50 percent of their time in meetings, and implied that this was time wasted. Peter Drucker once said that spending more than 25 percent of his time in meetings is a sign of a manager's malorganization, and William H. Whyte, Jr., in his book *The Organization Man,* described meetings as "non-contributory labor" that managers must endure.

But there is another way to regard meetings. Earlier we said that a big part of a middle manager's work is to supply information and know-how, and to impart a sense of the preferred method of handling things to the groups under his control and influence. A manager also makes and helps to make decisions. Both kinds of basic managerial tasks can only occur during face-to-face encounters, and therefore only during meetings. Thus I will assert again that a meeting is nothing less than the *medium* through which managerial work is performed. That means we should not be fighting their very existence, but rather using the time spent in them as efficiently as possible.

The two basic managerial roles produce two basic kinds of meetings. In the first kind of meeting, called a *process-oriented* meeting, knowledge is shared and *information* is exchanged. Such meetings take place on a regularly scheduled basis. The purpose of the second kind of meeting is to solve a specific problem. Meetings of