Harvard Business Review

ANALYTICS

No-Nonsense Guide to Measuring Productivity

by W. Bruce Chew

FROM THE JANUARY 1988 ISSUE

A few years ago, a major manufacturing-based conglomerate asked a gifted mathematician to join its corporate staff. One of his first assignments was to design a system that senior managers could use to evaluate the operating efficiency of the company's various divisions. He devoted many months to the assignment and also tapped the knowledge of several academic experts. The result was a truly sophisticated model that combined historical performance data with economic forecasts to set target productivity levels for each business unit.

Much to management's dismay, however, when the results were in, the model suggested that not one unit was performing up to snuff. So headquarters asked the obvious question—Why? Why was an organization that was generating handsome profits and cash flows showing such disappointing productivity? The expert could not answer the question, nor was his model designed to do so. Not surprisingly, executives saw little value in the new system and scrapped it.

It's essential to measure productivity *appropriately*. Many companies that want to raise their competitiveness are investing a lot of money and faith in methods to track their plants' and offices' efficiency. Staff specialists or outside consultants —experts in cost accounting, statistics, and economics—usually play an important role in designing these systems. But specialists are often trained to focus on the technical elegance and statistical accuracy of productivity indexes. All too often, they introduce methods that are very precise but ignore the real challenges managers face.

While collecting information on productivity measurement systems and interviewing managers at plants across the United States during the last several years, I have seen many examples of effective productivity measurement—systems that have led to big strides in operating efficiency. But more frequently I have encountered frustration and confusion. A manager

will look at a productivity index developed by a specialist and say, "Whoever came up with this has no idea what my business is like."

Productivity measurement is simply too important to be delegated to productivity specialists. But managers don't have to become experts themselves to ensure that existing systems meet their needs—or that new systems are relevant. A set of practical guidelines can help them understand, evaluate, and apply productivity measurement techniques effectively.

Look Beyond Direct Labor

What is productivity? Remarkably, many people who make decisions every day about improving plant efficiency don't know how to answer this simple question. Let's begin with what productivity is *not*.

Productivity is not about wages. High wages can present a problem, not because workers are paid too much but because they produce too little. In deciding how best to measure productivity, managers should focus not on dollars per hour but on labor dollars per product. That is, on labor content, not labor cost. Workers who are very productive can be paid thousands of dollars more than employees elsewhere and the business can still prosper, as manufacturers like Lincoln Electric have demonstrated.

Productivity measurement should focus on overall capabilities, not on one set of costs. How good is your company at taking a pile of raw materials, a bunch of machines, stacks of paperwork, and groups of employees, and turning out useful goods or services? That's what a productivity index should address. It is, as much as possible, a relationship between physical inputs and outputs. The formula is disarmingly simple.

$$\text{Productivity} = \frac{\text{Units of output}}{\text{Units of inputs}}$$

Productivity Formula

administration?

The company producing more with a given set of inputs (capital, labor, and materials) or using fewer inputs to produce the same output has an advantage over the company producing less. Lower input costs create an added advantage—but not the principal advantage that productivity measures must identify. The central mission of a productivity index is to illuminate how a business can get more units of output per labor hour, per machine, or per pound of materials than its competitors.

Still, much of U.S. industry remains preoccupied with direct labor. At the national level, productivity figures *do* mean labor productivity. The Bureau of Labor Statistics, the primary source of productivity information, logically enough focuses on labor productivity. Cost accounting also reinforces this bias. The allocation of overhead, for example, is often based exclusively on labor hours. This approach may have been reasonable when labor hours represented a large percentage of total costs, but today, for many businesses, labor is a minor cost element. Or the bias may simply come from too many years when managing operations meant "kicking butts and taking names." If persuading people to work harder is all there is to efficiency, then stressing labor productivity makes sense.

But there is much more to productivity, and many companies miss opportunities to bolster efficiency in nonlabor areas. Consider one U.S. plant manager's experience at a company with extensive fabrication and assembly operations. For a long time, he was uncomfortable with the way his division was allocating its annual $2 million productivity improvement budget. On being promoted to run all the division's plants, he at once reviewed spending on productivity programs. His intuition proved correct, as Exhibit I illustrates. Although direct labor accounted for only 10% of manufacturing costs, nearly 40% of the productivity budget was allocated to upgrading direct-labor efficiency. His subordinates are now looking for ways to reduce overhead and make better use of technology.

Exhibit I How One Company Didn't Match Productivity Improvement Spending to Production Costs

PRODUCTION COSTS

PRODUCTIVITY IMPROVEMENT SPENDING

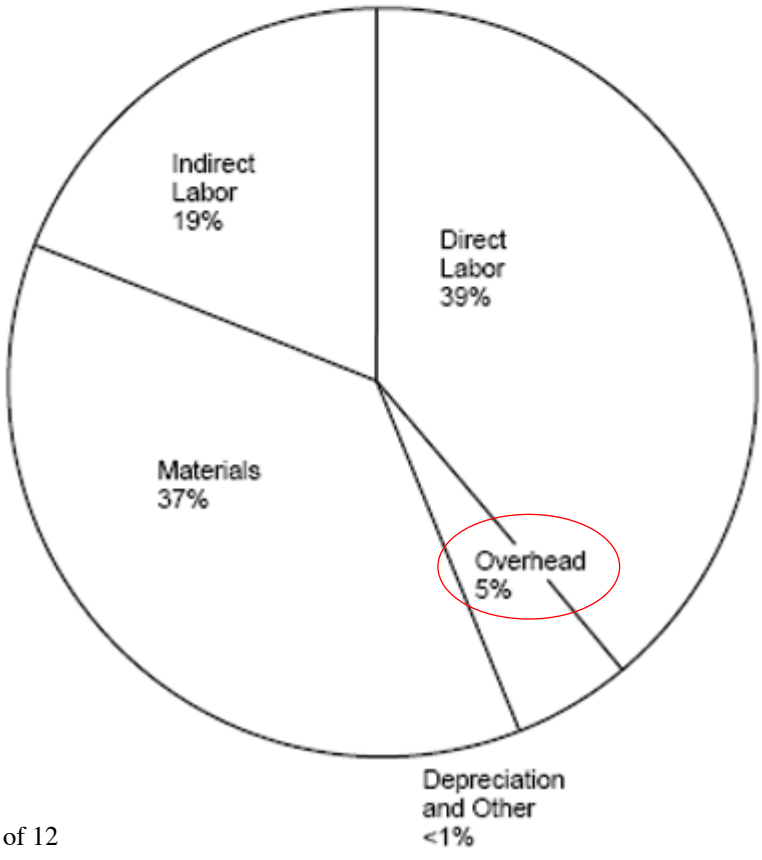

Exhibit I How One Company Didn't Match Productivity Improvement Spending to Production Costs

Single-minded attention to direct labor can produce unexpected consequences. Several years ago, a big New York bank concerned about labor costs in its back office implemented a department-by-department system to measure productivity, defined as transactions per employee. Senior management gave high visibility to the new system and even used it to calculate a large portion of the bonuses it paid to line managers. So the line managers computerized everything in sight. The result was increased productivity in every department but one—data processing. While staff was shrinking in the rest of the bank, data processing came under incredible pressure. It boosted its staff as well as its spending on hardware and software.

If that expansion in overhead was best for the bank, executives could never say for sure; their measurement system focused only on the productivity of direct labor. It didn't analyze the trade-offs between the productivity of departments (shorter turnaround times on letters of credit vs. growth in the data processing head count) or between cuts in variable costs (labor) and higher capital spending (new computers).

Take a Multifactor Perspective

The trouble with single-factor productivity measures (whether output per labor hour, output per machine, or output per ton of material) is that it is easy to increase the productivity of one factor by replacing it with another. Labor, capital, and materials are all potential substitutes for each other. Effective productivity measurement requires the development of an index that identifies the contribution of each factor of production and then tracks and combines them.

Take a hypothetical plant that machines purchased castings as one step in its production of motors. Now the company decides to purchase this component premachined. The premachined parts cost 20% more than standard castings, but buying them allows the company to lay off skilled workers and sell its machine tools. What happens to productivity? Output has remained constant, but the number of workers has fallen, so labor productivity is up. So too is capital productivity, by virtue of the lower asset base. But materials productivity has declined by 20% since output is unchanged, while the value of purchased materials has gone up.

In such a situation, a productivity index that focuses only on labor or capital would create strong incentives to reduce the value added by operations, which may not at all be management's goal. In theory, a company's make vs. buy decision process should prevent such unintended results. But with top management pushing hard for identifiable productivity increases, there is a real risk that defining productivity too narrowly will lead to unsound decisions by subordinates.

A multifactor view of productivity is important, therefore, but it is difficult for one index to encompass all inputs. Using several different single-factor measures can also yield a multifactor perspective. Indeed, even if a plant uses one aggregate measure, it still makes sense to use single-factor measures because they help identify the sources of aggregate productivity trends. A big change in a multifactor productivity measurement raises obvious questions: Is the change due to simultaneous shifts in the productivity of labor, capital, and materials, or has only one dimension changed?

Don't Sacrifice Function for Form

A multifactor index to track productivity gives managers a convenient scorecard to answer the question, "How are we doing?" But an index can play this role only if managers and workers understand it, which may require certain compromises in mathematical elegance and accuracy. Economists and productivity specialists like to use sophisticated functional forms when they combine labor, materials, and capital into one index. Rather than simply adding everything up or averaging inputs, they prefer logarithmic and multiplicative techniques. When the chief goal is to study productivity behavior, as in statistical research, these approaches have theoretical advantages. But when the primary goal is to influence behavior, the simpler the better must be the rule If the people who use an index can't understand it at a gut level, it probably will not affect their decisions and priorities.

Within Northern Telecom, some divisions make sure managers and workers understand multifactor productivity by including them in the design of department-specific indexes and by keeping the indexes simple. A department develops several performance ratios (no fewer than three, no more than seven) that it believes capture the essence of its mission. For example, one design engineering team proposes six ratios, among which are: reworked drawings as a percentage of total drawings, overdue drawings as a percentage of total drawings, and overtime hours as a percentage of total hours. Next the department identifies current performance, long-term goals, and interim goals for each ratio. Finally, managers assign weights to the ratios to reflect their relative importance, with the sum totaling 100. Thus the index produces a single productivity "score" (a weighted average of the ratios) that measures progress towards agreed-on goals in a way that everyone can understand.

This approach is not analytically perfect; there is no statistical reason to limit the number of ratios to seven, for example, and the weighting scheme is undeniably subjective. But Northern Telecom follows a basic principle that many other companies fail to appreciate: when deciding whether you need greater measurement precision, ask first whether greater precision will make a real difference in subsequent actions to improve productivity. Executives should seek the measure that promises the greatest impact, not the measure boasting the greatest accuracy or technical elegance.

The same principle applies to data collection. There are real costs associated with developing and implementing elaborate

productivity systems. My research suggests that the point comes—sometimes very early on—where increased accuracy is not worth additional cost. For example, the mismatch between the information provided by some accounting systems and what is needed for productivity analysis often means that bypassing the accounting data and developing data specifically for the productivity index will raise accuracy. But it is seldom worth the cost.

The costs can go even higher if you consider another factor: the time it takes to develop and implement a productivity measurement system. Lost time can translate directly into lost opportunities in to-day's turbulent business climate. As one manager complained to her staff, "Explain to me how taking an extra six months to get the measure more accurate is going to raise my productivity during those six months."

Measure the Unmeasurable

It's a real challenge to develop a productivity index that captures the roles of the direct factors of production in a way that workers and managers can understand. But the challenge goes further. Conventional systems to measure productivity often overlook two aspects of the production process that are becoming very important in determining international competitiveness: production time and the role of employees other than shop-floor workers. Since neither lends itself well to direct measurement, productivity technicians often prefer to look the other way. Managers do so at their peril.

The first oversight, time, is not purchased, so it is usually ignored. But the production process certainly consumes time, and the fact that it is not purchased doesn't mean it's free. If two businesses use identical machines, the same number of people, and equivalent materials to produce identical products, most productivity indexes would produce identical scores. But suppose one business ships orders within three days of receiving them and the other takes three weeks. Is their productivity the same? Obviously not.

This is not an exaggerated example. Increasingly, companies are discovering the competitive power of shortening their production cycles—or the dangers of not doing so. But unless a productivity index assigns some value to the amount of time consumed, it is unrealistic to expect managers to focus on shortening turnaround times. Assigning inventory carrying costs is a step in the right direction, although most companies record carrying charges far below their true competitive costs. Carrying costs should not only be realistic but they should also reflect where the inventory sits (in terms of value added) as well as how long it sits. An additional time-based charge that captures how long it takes to complete an order can focus attention even more directly on possible gains in turnaround time.

The new head of a sheet metal plant owned by a major electronics company learned this lesson soon after he took over. His operation's primary role was to process prototypes of new products, but he discovered that several company divisions

were routinely sourcing prototype work to outside suppliers because of his plant's unacceptably long lead times. Productivity measures of labor, capital, and materials were quite satisfactory; the plant's work flow had been structured to get the most out of its people and machines. But that represented a misunderstanding of its mission. So the new manager introduced a productivity index that focused on turnaround time, and he posted the results prominently. Eventually, the plant cut prototype production time from 20 weeks to three days. The reconfigured operation makes less "efficient" use of labor and materials, but can anyone argue that it is less productive?

The second crucial but often overlooked aspect of many productivity measurement systems regards whose performance is being measured. Most systems target inputs on the shop floor, but manufacturing efficiency is not only a function of who and what are located there. Engineers, supervisors, and other white-collar employees make significant contributions to manufacturing productivity, but few systems measure their roles. (The Northern Telecom system cited earlier is a notable exception.)

To a large extent, the absence of such measures reflects two principal difficulties of quantifying productivity in any service setting: measuring output and connecting employee actions to outputs. For the line worker in an auto plant, output is basically the number of cars or components produced. The connection between worker activity and output is also straightforward—the person tightens three bolts on every car, and this action helps complete the car. Measuring the productivity of product designers is a much more subtle problem. Defining output as simply the number of models or prototypes completed does not begin to capture these workers' true productivity. Designing an item to make production smoother will improve the efficiency of the entire plant, for example. If such a design takes twice as long to complete as a simpler approach, it certainly does not mean that the engineer is less productive.

It isn't possible to measure white-collar outputs or inputs fully, but this fact doesn't mean that only blue-collar productivity can or should be measured. It does mean, however, that managers must be creative and open to new ways of thinking about an operation.

A plant manager of an important supplier to the auto industry met with resistance at headquarters to his request to augment his engineering staff. He knew that the additional money would be well spent, but he had no measurement system for making his case. The engineering staff focused heavily on improving the plant's use of materials. So aiming for a surrogate index, the manager argued that changes in the materials productivity index could be one indicator of the engineering team's productivity. Indeed, over time, as the size and expertise of the engineering group increased, the ratio of total output to materials input showed dramatic increases. As a result, the division's entire perspective on the relationship between the engineering function and manufacturing productivity changed.

Do surrogate measures give a complete picture of a group's performance? No. Are they true gauges of productivity? For the economist or measurement specialist, no. Can they focus managers and employees on critical aspects of the production process and, therefore, lead to improved performance? Yes. Productivity measures in the white-collar world can be relevant and effective, even if they aren't perfect.

Compare Apples and Apples

Ultimately, any productivity measurement system is only helpful if it's used appropriately. Management carries the burden of usage almost entirely. Productivity indexes today are being used to compare the performance of companies in an industry, plants in a company, and departments in a plant. The results influence investment choices, judgments about factory closings, and decisions on management compensation, so managers must be careful to make fair comparisons.

What is fair is not always obvious. Consider some of the ambiguities on the output side of the productivity ratio. Exhibit II describes the output in 1986 and 1987 of a hypothetical plant making two related products. In 1987, the price of product A rose, so many customers switched to product B. From the set of facts presented, what conclusions can be drawn about the change in output—and, therefore, the change in productivity? Depending on your point of view, output went up, down, or stayed the same. If you look at nominal revenues, output rose dramatically. If you adjust for the price change by comparing revenues using 1987 prices, output went down. If you focus on physical units, output stayed the same. You might look to standard costs for guidance, but they may also present a confusing picture as well as concerns about accuracy.

Exhibit II What Is Output?

The Facts

	1986		1987	
Product	*Units Produced*	*Price*	*Units Produced*	*Price*
A	2,000	$10	1,000	$25
B	2,000	$20	3,000	$20

Change in Total Output Equals

	1986	1987	Percent Change
In Units	4,000	4,000	0%
Using Nominal Prices	$60,000	$85,000	+ 41.7%
Using 1986 Prices	$60,000	$70,000	+ 16.7%
Using 1987 Prices	$90,000	$85,000	– 5.6%

Exhibit II What Is Output?

What really happened at the company depends on what really happened at the plant and in the marketplace, not in the numbers. Was product A radically redesigned? Was the old price relationship between the two products somehow incorrect? Was there a dramatic change in input costs for product A? A manager must consider questions like these before evaluating productivity trends in such a case.

Price changes, of course, are not the only important factor affecting output. Quality has an impact on productivity measures. The most productive plant or company does not necessarily have the lowest cost per unit of output. It does have the lowest cost per *comparable* unit of output, though.

Suppose one company manufactures tires that last 15,000 miles and another company uses 10% more labor and materials to produce tires that last 30,000 miles. If the two businesses produce the same number of tires, it is not immediately clear which is more productive. Or suppose one enterprise, by virtue of a product development breakthrough, uses the same number of workers and machines to manufacture one million 15,000-mile tires in 1986 and one million 30,000-mile tires in 1987. Ignoring price for the moment, is output (and thus productivity) constant?

Comparing the performance of plants making different products requires a method to determine equivalencies. The three most common alternatives are standard costs, price, and technical parameters (like miles of tire life) that quantify product performance. No one focus is best; each has its strengths and weaknesses, and managers should select methods consistent with their company's strategy. If a company is and wants to remain a low-cost producer, it might focus on prices. If it wants to promote innovation, it might use technical parameters. Standard costs will focus attention on internal improvements independent of developments in the marketplace.

Managers also need to interpret trends, which can create further ambiguities. There is a fundamental distinction, for example, between levels of productivity and rates of productivity change. Looking at actual trends in a multifactor productivity index for two plants making identical products (Exhibit III), plant A's productivity exceeds plant B's in any given month and over the entire period. But whose performance should management worry about? As it turns out, plant B dramatically changed the run lengths in its production process, which improved productivity enormously. Corporate policies pushed plant A in the opposite direction, and its performance suffered. Here, focusing on absolute performance could mask important trends in relative performance.

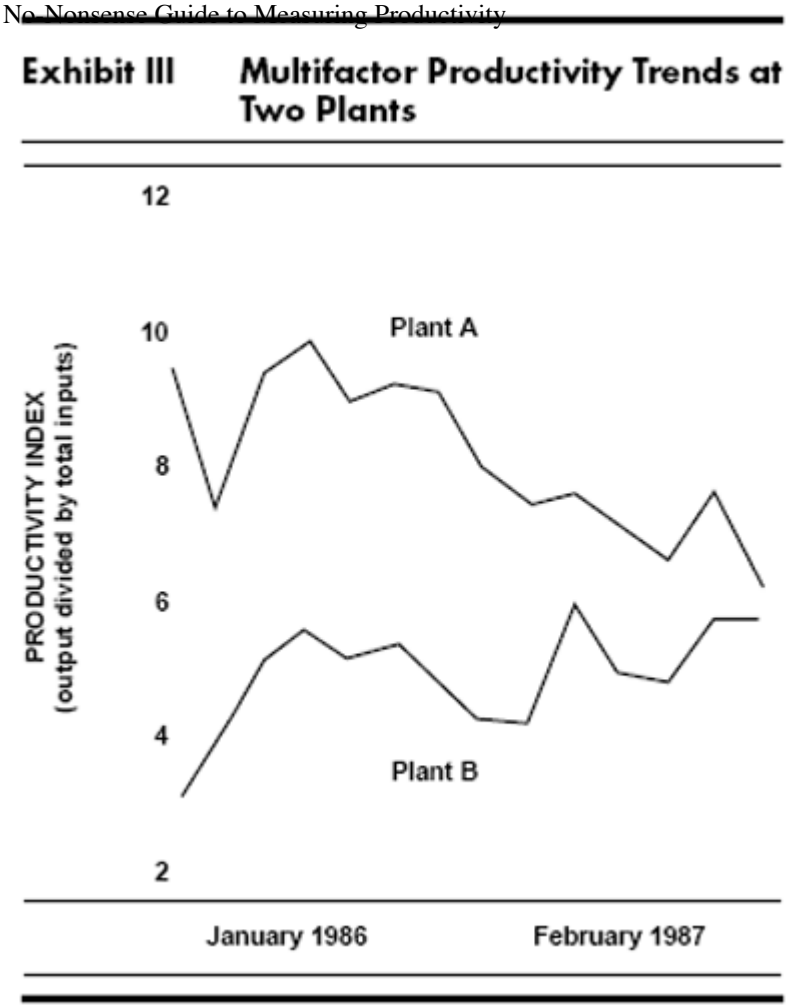

Exhibit III Multifactor Productivity Trends at Two Plants

Things are often not as they first appear with productivity data. One big manufacturer, after introducing a multifactor productivity index, discovered that its plants had suffered a significant productivity decline in the early 1980s. Bad management, right? Wrong. Demand for its product had fallen sharply during the period and, given the fixed inputs, overall productivity had declined. One plant had an especially large drop. Bad management, right? Wrong again. The plant, located in a rural area lacking skilled labor, treated skilled and semiskilled employees as fixed costs. To lay off these workers would be to lose them permanently to other employers. The large decline in productivity in response to a six-month downturn was thus evidence of good management; the employees should have been kept on. Productivity measurement raises issues and highlights changes, but it does not tell the whole story.

Proceed with Care—but Proceed

The complexities and ambiguities of productivity measurement should not discourage managers from using a system. Profit measures, after all, are also far from perfect, but we are accustomed to their shortcomings and have learned how to glean a wealth of insight from them. Managers should proceed with productivity measures, but with care. The seriousness

of a system's shortcomings depends on how it is used; if bonuses or promotions are based on certain measures, they had better be accurate. But this degree of accuracy is unnecessary for most applications.

Perhaps the most important use of productivity measurement is as an objective source of information about long-term operating trends. An index can draw attention to plants or departments experiencing unusual problems or uncommonly strong performance. Productivity comparisons can also inspire useful exchanges of ideas. Differences in the amount of vertical integration or subcontracting, accounting policies, and many other factors often obscure the relative productivity of companies. Nonetheless, if a business finds itself a lot less productive than a competitor, it probably has a real problem. Managers may insist that the productivity gap is overstated, and they may be right. They will be hard-pressed, however, to argue that it does not exist.

A version of this article appeared in the January 1988 issue of *Harvard Business Review*.

W. Bruce Chew is a senior consultant at Monitor Company, a global consulting firm in Cambridge, Massachusetts. He is a former associate professor of technology and operations management at the Harvard Business School in Boston, Massachusetts.
