ELSEVIER

Journal of Operations Management 18 (1999) 41–59

OPERATIONS MANAGEMENT

www.elsevier.com/locate/dsw

The economics of yield-driven processes

Roger E. Bohn [a,1], Christian Terwiesch [b,*]

[a] University of California, San Diego, CA, USA
[b] The Wharton School, University of Pennsylvania, Steinberg Hall–Dietrich Hall, Philadelphia, PA 19104-6371, USA

Received 5 February 1998; accepted 6 April 1999

Abstract

The economic performance of many modern production processes is substantially influenced by process yields. Their first effect is on product cost — in some cases, low-yields can cause costs to double or worse. Yet measuring only costs can substantially underestimate the importance of yield improvement. We show that yields are especially important in periods of constrained capacity, such as new product ramp-up. Our analysis is illustrated with numerical examples taken from hard disk drive manufacturing. A three percentage point increase in yields can be worth about 6% of gross revenue and 17% of contribution. In fact, an eight percentage point improvement in process yields can outweigh a US$20/h increase in direct labor wages. Therefore, yields, in addition to or instead of labor costs, should be a focus of attention when making decisions such as new factory siting and type of automation. The paper also provides rules for when to rework, and shows that cost minimization logic can again give wrong answers. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Production yields; Cost of quality; Product cost; Rework; Ramp-up; Location decisions; International operations

---

1. Introduction

Many modern production processes and services are driven by process yields. Not every unit of material that starts into the production process makes it to the end as a sellable, high quality product. Some ‘‘fall-out’’ along the way due to problems of various kinds. Often, some of the fall-out can be reworked, but always a fraction of it must be scrapped. This means that materials and effort go into making something that ultimately cannot be not sold.

The effect of yield losses on the economics of the product, factory, and business can be dramatic. The comprehensive Berkeley project on semiconductor manufacturing has documented many examples of integrated circuit factories with yields below 50% for years (Leachman, 1996). The impact of this is, crudely, that costs per good unit are multiplied by two compared with what they would be at 100% yield. The impact on profit is much greater.

---

[*] Corresponding author. E-mail: terwiesch@wharton.upenn.edu
[1] E-mail: rbohn@ucsd.edu.

PII: S0272-6963(99)00014-5

The main purpose of this paper is to analyze the economics of yield-driven production processes. Despite the widespread and important role of yields, their impact on economic performance is treated casually in management accounting systems, and has received little attention by operations management researchers. The result, we observe, is that some decisions are driven by analysis and intuition developed from inadequate models.

A secondary purpose of this paper is to compare the importance of yields with that of labor costs. Specifically, we show that under common conditions in ‘‘high-tech’’ industries, the impact of direct labor wage rates can be overshadowed by the effect of yields. Even eliminating direct labor entirely can have less effect on profit than modest changes in yield levels. Thus, yields matter when asking questions such as ‘‘Where to site the next factory?’’ and ‘‘Should we automate a process?’’

Our analysis is illustrated with examples from a high-tech industry, hard disk drives (HDDs). Disk drive production starts with the fabrication of key components (heads, media disks, and semiconductors). All of these fabrication processes are strongly yield-driven, i.e., much less than 100% of what goes ‘‘in’’ to the process comes ‘‘out’’ as good components. The components are then assembled in multi-step, labor- and testing-intensive processes. These assembly steps are also yield-driven. The industry is sensitive to yield issues, as illustrated by the following quotation. Nonetheless, it has not had good tools for quantifying their effects.

> It is how you can improve your yield that will get your productivity up. We are not in a business where you have a 99% yield. In many cases, there are initial yields on high-end products that are in the 50% range. So a 5% or 10% improvement in these yields is significant (Richard Downing, a senior VP of manufacturing at Seagate, quoted in Electronic Business Asia, Feb. 1997, p. 35).

Section 2 of this paper reviews the existing literature on yield-driven processes. Section 3 analyzes yields in multi-stage production process. Section 4 motivates our analysis by describing the yield-driven nature of the disk drive industry, and the yield-related decisions its managers must make. Section 5 examines the economics of rework and scrap in detail for a simple process. It concentrates on variable cost and output as the main effects of yield. Section 6 gives our conclusions and points at needed future research.

2. Prior research on yields

The subject of process yields has received considerable attention in various disciplines. We can group this research into four streams. First, engineering reports describe yield problems in specific industrial processes and provide technical solutions. Second, operations management and operations research models support production management of yield-driven processes. Typical concerns are inventories, inspection plans, order releases, scheduling and sequencing, and other issues related to production planning. Third, there is an organizational learning literature on how to improve yields and reduce ‘‘waste’’. Much of it is empirical- or case-based. Fourth, quality management research outlines a number of principles to reduce the ‘‘cost of quality’’. Yield losses correspond to internal quality problems, i.e., problems caught before goods leave the factory.

There are a number of engineering articles and technical reports describing methods of dealing with yield-driven production processes, especially in the semiconductor industry. For example, IEEE Transactions on Semiconductor Manufacturing has several articles per issue related to yields or ‘‘defects’’. The emphasis is on methods, concepts, and tools that will improve yields by detecting diagnosing and solving specific problems. Examples include methods of defect classification (Breaux and Kolar, 1996), yield-loss modeling (Stamenkovic et al., 1996), in-line product inspection (Wang et al., 1996), statistical software to analyze process control data (Burggraaf, 1996), and expert systems to provide estimates on quality of certain batches (Khera et al., 1994). This literature is vital to continued technological progress in these industries. As new products and processes push the state of the art, yields fall, and new cycles of yield improvement are needed.

The random nature of yield-driven processes and the resulting challenges for managing production have attracted a number of operations management researchers. Most of this literature takes the production technology, and thus, the yield problems, as given and provides models supporting standard production decisions such as how to manage work-in-process and congestion (e.g., Chen et al., 1988), inspection plans and quality improvement (e.g., Barad and Bennett, 1996), scheduling and sequencing (e.g., Ou and Wein, 1995), and other issues related to production planning (e.g., Denardo and Tang, 1997).

A smaller group within the operations management literature argue that the overall yields of a production process can be improved by effective management of the process. Proposed methods for yield improvement include inspection policies for quick feedback on the quality of the process (e.g., Tang, 1991), keeping the work in progress level low (e.g., Wein, 1992), and effectively combining items from different batches (e.g., Seshadri and Shanthikumar, 1997). In contrast to the engineering literature, these papers focus on improving various performance measures, including yields, without really changing the underlying production technology. This makes them more general across processes, but limits their potency.

The third stream of yield research is at the intersection between production management and organizational research, especially organizational learning, and has contributed some in-depth empirical studies on yield improvement. Mukherjee et al. (1995) categorized various quality projects undertaken at a major manufacturer of wire cord depending on the type of learning approach taken in the projects. A follow-up study (Lapré et al., 1996) links these quality projects to waste reduction (yield improvement). Bohn (1995a; b) looks at factors which influence the speed of yield improvement in semiconductor manufacturing. Kantor and Zangwill (1991) give a theoretical model of waste reduction learning. Like the engineering literature, the organizational literature has little to say on the economic value of yield improvement. For the most part, yield improvement is implicitly treated as a way to reduce costs without looking at other effects.

Finally, under the ‘‘cost of quality’’ paradigm as outlined in (Juran and Gryna, 1993), yield losses are viewed as part of internal failure costs and thus, as one of the main drivers of the costs of quality. Juran and Gryna emphasize the need to assign economic values to these quality costs, to make them easier to understand for top management decision-makers. The cost of quality approach is valuable in its recognition of hidden effects from quality problems, and its emphasis on quantifying them. For example, this approach would show that when first-pass yields get high enough, in-process inspection can be eliminated, which has various desirable effects. However, one of the main benefits of yield improvement is ignored, namely the improvement in effective capacity and output.

In the quality literature, yield loss is the extreme form of a defect — the product is unsalable. Therefore, much of the quality improvement literature is applicable to yield improvement. Probably most important are the tools and concepts of statistical process control to yield monitoring and improvement. Again, this is most active for semiconductors; see the survey/tutorial by Spanos (1992). Typical issues include how to detect a defective machine quickly, what inspection policies to set, and how to modify SPC tools such as control charts to cope with the huge amount of data produced by automated semiconductor manufacturing lines.

Although the literature reviewed above has substantially improved our understanding of yield issues in production processes, none of it has provided the basic economic analysis of how yields matter. We attempt to extend the literature in three directions:

- we assign concrete economic values to yield issues (Juran and Gryna, 1993)
- we do not take yields as given, rather, we concentrate on the value of improvement
- we look beyond the cost impacts of yield improvement.

This article can be viewed as an effort to evaluate the value of yield improvement.

3. Multi-stage yield-driven production processes

In this section, we discuss production processes consisting of a sequence of sub-processes, of which at least one has yield below 100%. Although defects can occur anywhere, they are detected mainly at test points. An important question in designing processes with yield losses is the positioning of tests or inspections. Tests are costly, and can sometimes reduce yields themselves. There are various formulations of where to put them. Common rules are to position them before expensive or irreversible operations, at the end of modules in modular subassembly, after low-yield operations (to avoid adding more value to bad units), or immediately after operations targeted for process improvement (to provide fast feedback for learning).

At each test point, items are classified into ‘‘good items’’ and various categories of ‘‘defective items’’. Whereas good items can continue processing at the next operation, defective items are removed from the line. They can then be either reworked or scrapped.

3.1. *Yields and rework*

Rework means that some operations prior to the defect detection point must be redone, or defects must be otherwise repaired. Thus, rework changes the capacity utilization profile of the process. In analyzing the influence of yields (and rework) on process capacity, we need to distinguish between bottleneck and non-bottleneck machines. If rework involves only non-bottleneck machines with a large amount of idle time, it has a negligible effect on the overall process capacity.

In many cases, however, rework is severe enough to make a machine a bottleneck (or, even worse, rework needs to be carried out on the bottleneck machine). As the capacity of the bottleneck equals the capacity of the overall process, all capacity invested in rework at the bottleneck is lost from the perspective of the overall process.

A second complication related to rework, which affects bottleneck and non-bottleneck machines, is related to the amount of variability in the process. A yield of 90% means not that every 10th item is bad, but that there is a 10% chance that a given item is bad. Thus, yield losses increase variability, which is the enemy of capacity. The best stochastic case is that yields are Bernoulli, i.e., that the process has no memory. Suppose that bad items at an operation are immediately reworked by repeating the operation. Even if the actual processing time of the operation is itself deterministic, the yield losses force items into multiple passes, and thus make the effective processing time for a *good* item a random variable. Hopp and Spearman (1996) (Section 12.3) show for this case that the variability (measured by the squared coefficient of variation) in the effective processing time increases linearly with $(1-y)$.

Capacity losses due to variability can be partially compensated by allowing WIP after each operation with yields below 100%. The larger these buffers, the more the capacity-reducing impact of variability can be reduced. However, additional WIP increases costs, lead times, and throughput times; it also can hurt problem detection and solution, thereby reducing yields.

3.2. *Yields and scrap*

Scrap occurs when bad items are discarded. Final output is correspondingly reduced. Rework is generally preferable, but sometimes, it is technically infeasible or uneconomic. An economic comparison of scrap and rework is given in Section 5.

Strictly speaking, scrap is a special form of rework, where the rework loop includes all operations between the defect generating machine and the beginning of the process. The impact of scrap losses on system capacity are even stronger than in the rework case, since additional capacity must be added at *all* stations upstream of yield test points, with the most capacity needed at the start of the process. It does not matter where the defective unit is actually created, only where it is detected. In order to get 100 good parts at the end of the process, more than $100/y$ must be started at the beginning, where y is the cumulative yield all the way through the process. Further, the stochastic variation in load is felt at *all* stages downstream of the yield loss, not just at the stages involved in the rework loop.

This points to the importance of capacity planning in yield-driven processes. If yields and resulting rework requirements are known at the time a line is laid out and remain roughly constant, then capacity planning and

Table 1
Summary of yield effects on cost

	Rework is done	**Scrap**
Material-related costs	Incremental material to replace bad components	All material up to failed test is lost
Labor-related costs	Rework labor	All labor up to failed test is lost
Capacity-related costs	More capacity needed in the rework loops of process	More capacity needed at all stages upstream of failed tests
Variability-related costs	WIP cost to buffer variability	WIP still needed but less effective; more capacity needed to counteract
	Lead time variability in make to order systems	Extra large lots needed in make-to-order systems
		Line never perfectly balanced; more capacity needed to counteract

line balancing is done by increasing the capacity at each station enough to handle its anticipated yield-caused extra load. With scrap, it takes the form of increasing the capacity at all upstream stations enough that they can keep up with demand at the end of the process. Usually, however, yields are neither known accurately in advance nor are they constant over time. Instead, the aggregate yield shows both a positive trend (learning) and a week-by-week variation which cannot be buffered out economically, even by finished goods inventory. Therefore, once a process starts up, the actual capacity at each stage usually will be ‘‘sub-optimal’’ by static criteria.

A related complication arises in make-to-order situations with scrap. To respond to a customer order of N units, we must start N/y at the beginning to compensate for the expected yield losses. This approach would work fine, if yields were deterministic. However, since they are not, the production scheduler has to trade off the costs of making too much against the cost of making too little. Mathematically this is a newsboy-type problem (Table 1).

3.3. Cost and value at different stages of the process

In addition to its effect on capacity, yields determine the value that a good unit of WIP has at various stages in the process. This information is, for example, important in deciding where to concentrate process improvement efforts. A two-point yield improvement has different value at different places in the process.

The value of a good unit of WIP also help to decide whether it is more economical to scrap a defective item or to rework it. For example, suppose that after a test a defective item can be reworked for a labor cost of US$10, with a 90% chance of success and a 10% chance that the item must be scrapped. Is it better to pay for rework, or to scrap the item? Clearly if x is the value of a good item at that point, the decision rule is to rework if $10 < 0.9x$. However, determining x is not trivial.

At the beginning of the process, the value of a good item equals the cost of raw materials. At the end of the process, the value is given by the marginal revenue from a good item that can be sold. The value of a good item increases as it moves through the process, even if no additional material is being added. Let y_n be the yield at the nth stage. If there are no binding capacity constraints, the value leaving stage n is approximately $1/y_n$ times the sum of the value entering stage n and the variable costs at stage n.

This gives two different ways to calculate value: cost-based working forward, and price-based working backwards. The two will be equivalent if there is no binding capacity constraint, and differ if there is one. The discontinuity in value comes at the bottleneck operation(s). After the bottleneck, value is based on selling price; before the bottleneck, it is based on cost. An analogous effect will be formally discussed in Section 5. It can have surprising consequences when cheap raw materials are transformed into expensive products (e.g., semiconductors).

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

Smarter legal research.