

The Origins
of Efficiency
Brian Potter

Chapter 1
What Is a
Production
Process?

About *The Origins of Efficiency*

Efficiency is the engine that powers human civilization. It's the reason rates of famine have fallen precipitously, literacy has risen, and humans are living longer, healthier lives compared to preindustrial times. But where do improvements in production efficiency come from? In *The Origins of Efficiency*, Brian Potter argues that improving production efficiency—finding ways to produce goods and services in less time, with less labor, using fewer resources—is the force behind some of the biggest and most consequential changes in human history. The book is punctuated with examples of production efficiency in practice, including how high-yield manufacturing methods made penicillin the miracle drug that reduced battlefield infection deaths by 80 percent during World War II; the 100-year history of process improvements in incandescent light bulb production; and how automakers like Ford, Toyota, and Tesla developed innovative production methods that transformed not just the automotive industry but manufacturing as a whole. *The Origins of Efficiency* is a comprehensive companion for anyone seeking to understand how we arrived at this age of relative abundance—and how we can push efficiency improvements further into domains like housing, medicine, and education, where much work is left to be done.



Stripe Press
Ideas for progress
South San Francisco, California
press.stripe.com

1

What Is a Production Process?

If we are to understand how a production process becomes more efficient, we first need to understand how a production process works. To do this, let's take a look at a relatively simple process: the manufacture of glass bulbs used for incandescent lights in the late 1800s.

In 1880, Thomas Edison was awarded a patent for his electric incandescent light bulb, marking the beginning of the age of electricity. Although it was the result of thousands of hours of research that took place over decades, the ultimate design of Edison's light bulb was simple, consisting of just a few components: a filament, a thin glass tube in which the filament was mounted, a pair of lead-in wires, a base, and the glass bulb itself.¹

Until the 20th century, light bulbs were largely manufactured by hand. Workers would run the lead-in wires through the inner glass tube, attach the filament to the lead-in wires, and attach the glass tube to the bulb. A vacuum pump would then suck the air out of the bulb. Initially, this was done by connecting the pump to the top of the bulb, leaving a small tip of glass that had to be cut off. Later, tipless bulbs were developed that had the air removed from the bottom.²

Most of this manufacturing process was done in house by Edison's Electric Light Company, but the production of the glass bulb itself, known as a bulb blank, was outsourced. Edison placed his first order for bulb blanks with the Corning Glass Works company in 1880.³ The process of making the bulb blanks was fairly straightforward: Glassworkers would mix together sand, lead, and potassium carbonate, along with small quantities of niter, arsenic, and manganese oxide, place the mixture in a crucible, and melt it in a furnace into liquid glass. A worker called a gaffer would then gather a blob of glass on the end of a hollow iron tube and place the blob into a mold the shape of a light bulb. While the blob was still attached to the iron tube, the gaffer would blow into it to form the body of the bulb, then open the mold and cut the bulb from the end of the tube.⁴

We can draw this series of steps using a process flow diagram, a visual representation of how a process unfolds. See Figure 2 for an example of what the bulb blank process might look like.

Making bulb blanks is an example of what we'll call a *production process*—a series of steps through which input materials are transformed incrementally into a finished product. Each step in the process induces some change in the input material. The changed material is then passed on to the next step, which makes another change, and so

on, until the finished product comes out the other side. In the bulb blank process, sand, lead, and other chemicals are the inputs. These are gradually transformed by heat, chemical reactions, and physical manipulation until a finished bulb blank emerges at the other end.

In turn, this output might be the input to a subsequent process. Bulb blanks, for instance, would then be sent to Edison's factory to be assembled into complete light bulbs. Likewise, the input materials for the manufacture of bulb blanks were themselves the output of some other production process. Potassium carbonate, for example, was mined from potassium ore and then refined using the Leblanc process.⁵

Outside of the small number of things we can obtain directly from nature, all products of civilization are the result of some sort of production process—some series of transformations that take in raw materials, energy, labor, and information and produce goods and services. At first glance, services might seem far removed from the production of physical goods like cars or shoes, but the same basic model applies. A house cleaner, for example, goes through a specific series of steps—cleaning the bedrooms, then the bathrooms, then the kitchen—using various inputs—labor, electricity, cleaning products—to transform an input—a dirty house—into an output—a clean one. These processes might be comparatively simple, such as the production of light bulb blanks, or exceptionally complex, with hundreds or even thousands of steps. One 19th-century watch factory boasted that its watches “required 3,700 distinct operations to produce,” while a 1940s Cadillac—a relatively simple automobile by modern standards—required nearly 60,000 separate operations.⁶

Even everyday objects can mask a great deal of production complexity. In his book *The Toaster Project*, Thomas Thwaites disassembles a \$7 toaster to find that it contains 404 parts made up of more than a hundred different materials.⁷ And if we follow the chain of production further back, to the processes required to make the various input materials (and the processes to make the inputs for those processes, and so on), we find a sprawling mass of complexity for even the simplest products of civilization. In his famous 1958 essay “I, Pencil,” Leonard Read notes that a full accounting of the inputs required to make an ordinary pencil—the steel used to make the tools to harvest the cedar, the ships used to transport the graphite from Sri Lanka to

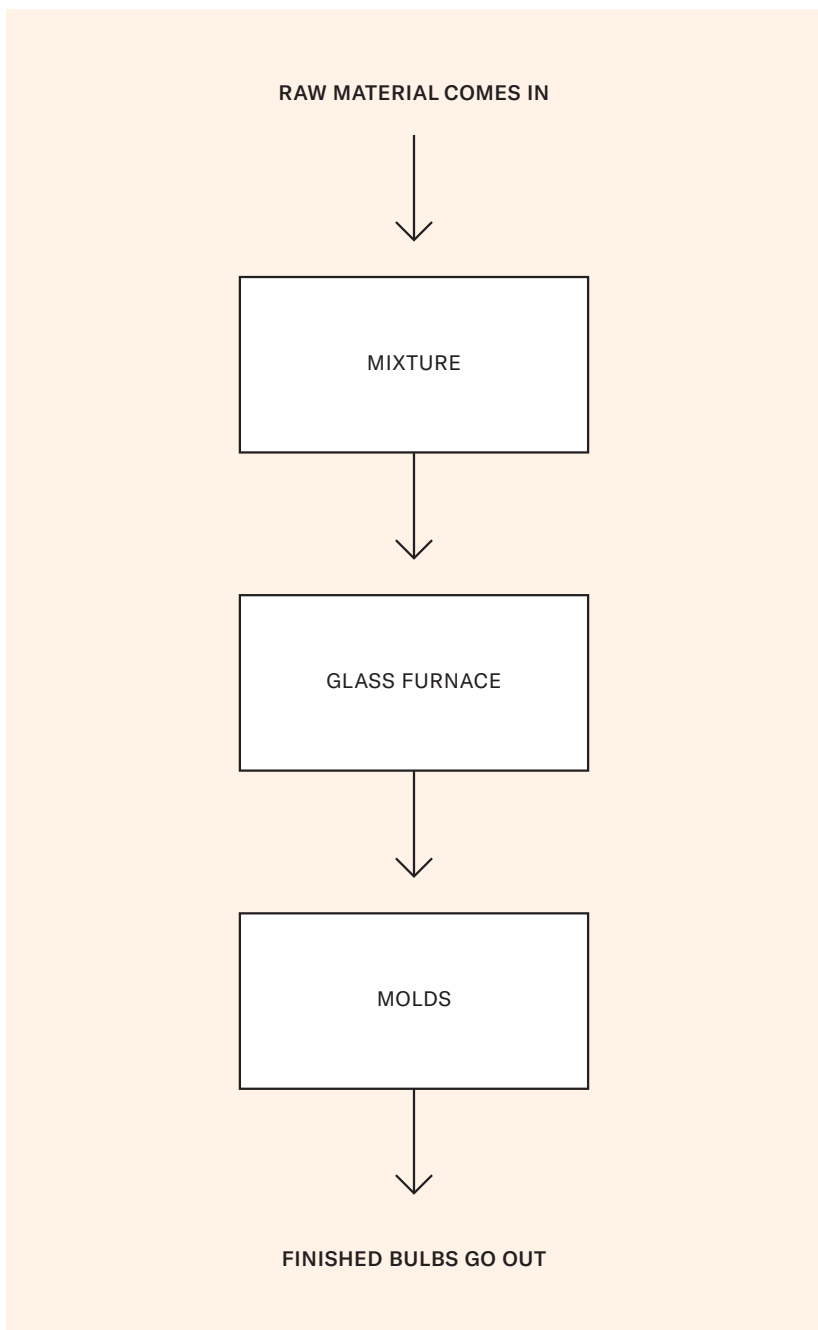


Figure 2. Process flow diagram of a bulb blank production process.

the factory, the agricultural equipment used to grow the castor beans to produce the lacquer—involves the work of millions of people all over the world.⁸

Five factors of the production process

Now that we have a basic model for how things get produced, we can add a bit of detail to the description, identifying five distinct factors of the production process. This slightly more regimented structure will be useful for pinpointing discrete sites of intervention that can improve the efficiency of a production process.

First is the transformation method itself. In bulb blank production, one transformation method is the process of blowing the glass bulbs. Of course, each transformation is itself made up of many steps (gathering the glass on a blowpipe, placing the mold around it, blowing while a worker holds it), which in turn might be made up of substeps (such as individual worker motions). Different situations will call for varying degrees of fidelity in describing a process—the scientific management movement of the early 20th century, for example, spent a great deal of time studying specific worker motions—but it will always be a simplified model that omits many details of what is actually occurring.

The idea of a well-defined transformation or series of transformations is something of a simplification, as there will inevitably be some degree of variation in the specific actions taken during a step. For a machine, this variation will be very small and occur in narrowly defined ways, but the farther we get from modern industrial production processes, the less true this becomes. A person might perform the same step slightly differently each time and modify their technique over time as they get more skilled. And craft production methods often require some degree of deciding what the next step should be. A glassblower blowing bulbs without a mold, for instance, will decide how hard to blow based on how they see the bulb taking shape.

Second, to understand how efficient a production process is, we need some idea of how much time the process takes. It obviously makes a big difference whether the bulb blank factory can produce 10 or 10,000 bulbs a day. Using bulb molds, three workers could produce about 150 bulbs per hour, or roughly 1,500 per day.⁹ This is called the *production rate*. Each step in the process will have its own rate, and these rates may differ from those of other steps. For example, filling the glass

crucibles might be done just once a week, even though glassworkers were producing bulb blanks daily.¹⁰

Third, to determine how much a given production process costs, we need to account for all the direct material inputs and outputs to the process. At the furnace step, raw materials go in and molten glass comes out. At the blowing step, molten glass goes in and a bulb comes out. Depending on how detailed we decide to be, we might also include inputs like the coal that fuels the furnace and outputs like the ash and smoke produced by the furnace. There are labor inputs as well. The blowing step, for instance, requires the labor of two or three workers to gather the glass, work the mold, and cut the finished bulb free.¹¹

We also need to account for the indirect inputs—things that aren't used directly by the process but are nevertheless necessary. A factory's rent can't be directly attributed to any particular operation within the factory, but the building is still an important input to the process. We can account for this cost by attributing some fraction of it to each step. Similarly, we can assign some fraction of the cost of the equipment, administration, insurance, and any other overhead costs to each step in the process. (The question of how best to assign these indirect costs is an involved area of accounting, but broadly speaking, these costs will be spread over the amount of output we produce.)

Taken together, the costs of the various inputs and outputs allow us to determine the cost to produce a bulb blank. The first blanks from Corning, for example, cost about 3¢ each.¹² For our purposes, the goal of any efficiency improvement is to reduce this cost. Or, to be more precise: An efficiency improvement is anything that is able to lower the cost per unit.

Fourth, to understand whether the process is efficiently arranged, we need to keep track of how much material is in the process at any given time. At any point, some material is actively being worked on and some is waiting to be worked on. In bulb blank production, once the raw materials had been added to the crucible, it might take a while before the glass was gathered by workers and blown into bulbs. If crucibles were filled once a week, there would be about half a week's worth of molten glass waiting to be turned into bulbs at any given time. Any material that isn't currently being worked on is considered to be in a buffer of available material. The total amount of material in the system—that is, the combination of what's in the

buffer and what's being worked on—is collectively known as *work in process*.

Fifth and finally, in evaluating a production process we need to make note of how the output of the process varies. While it's tempting to think of a step as producing the exact same output every time, there will inevitably be some variation. At times, the process may simply fail. For example, in some cases, the furnace would produce a batch of glass that was unsuitable for bulbs. In other cases, the crucibles that held the molten glass would crack, spilling the glass before it could be turned into bulbs.

But there will also be more subtle sources of variation. For instance, the composition of the glass and the thickness of the bulbs would differ slightly, perhaps imperceptibly, from bulb to bulb. No two bulbs were exactly alike. This discrepancy can be a natural outcome of the process, the result of a disparity in the inputs, or due to variation in the environment in which the process takes place. The quality of the bulb glass, for example, was greatly dependent on the quality of the chemicals used, how well they were mixed, and the temperature of the furnace.

One simple way of characterizing variation is in terms of yield—the fraction of inputs that are successfully transformed into outputs. A yield of 50 percent would characterize a process that is only successful half the time. An unsuccessful transformation might be a complete failure (a bulb falls on the floor and breaks) or one that is simply outside the range of acceptable tolerance (the glass on the bulb was slightly too thin).

In many cases, however, it will be useful to have a more detailed characterization of the variation in a process. In the production of light bulb filaments, very slight differences in temperature during the carburizing process resulted in the filament producing different amounts of light. Understanding how the resulting filaments varied was, therefore, necessary to determine how many bulbs of a given illumination could be produced. It might turn out that the variation in illumination could be described by a normal distribution with a particular mean and standard deviation, making it possible to track disruptions to the process by looking at whether values fell outside of the expected range. For now, we'll just note that variation is an important factor to consider without worrying about developing a certain measurement for it.

Looking at a single step in the process, we now have five factors that characterize it:

- 1 The transformation method itself. For example, the act of blowing molten glass into a mold.
- 2 The production rate. For example, how many molds the gaffers can fill in an hour.
- 3 The inputs and outputs, along with their associated costs. For example, the molten glass, the gaffer's wages, and wear and tear on the molds and blowpipes.
- 4 The size of the buffer. For example, how much molten glass is stored in the furnace waiting for the gaffer.
- 5 The variability of the output. For example, fluctuations in how fast the gaffer works and the thickness of the bulbs produced, or how often the gaffer drops and breaks a bulb.

This is, of course, a highly simplified model. For one thing, it omits the complexity of what specifically occurs during each step. For another, it suggests that these factors are steady over time, but in reality they will frequently be in flux. The variation in output may rise when a new worker starts, or at the end of the day when workers are tired, or over a long period of time as workers or managers grow complacent. Alternatively, variation may go down over time as workers gain experience and precision improves.

This model also doesn't include the many possible ways one step may influence another step, beyond how fast the step runs. The temperature of the glass furnace might influence how easy it is to blow the bulb into the mold, for example. Likewise, variation in one process may be a function of variation in some previous process. Bulbs breaking when the mold is removed, for instance, might be a function of inconsistent mixing of the ingredients or uneven temperature of the molten glass.

Finally, this model doesn't include any specifics about what is actually being produced. As we'll see later on, the form of the product and the method of production are intimately connected, and a change in one generally results in a change in the other.

Despite its various simplifications, however, this model gives us a useful way to structure our thinking about production processes and how they can be made more efficient.

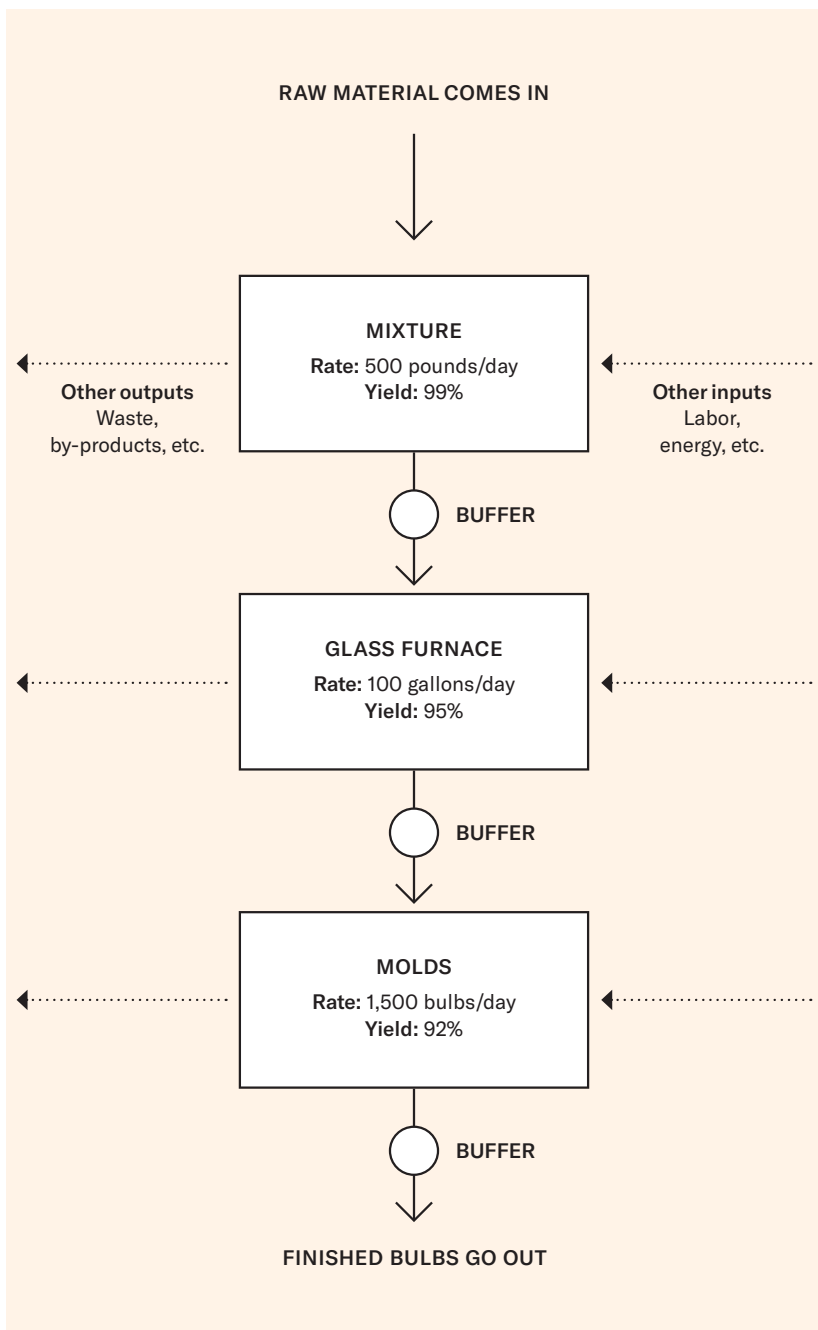


Figure 3. Process sketch of a bulb blank process showing inputs, outputs, buffers, production rates, and yields.

Improvements to the process

The goal of any efficiency improvement is to minimize the costs of producing something. If we're running a bulb blank factory, we want to figure out how to produce those bulb blanks as cheaply as possible, which means using the fewest, lowest-cost inputs we can. The way to do this is to change one or more of these five factors.

First, we can change the transformation method itself, to one that requires fewer resources. The very first bulb blanks produced by Corning didn't use molds but were produced using a much slower free-hand method, which entailed manually rolling out tubes of glass. Changing the bulb-blowing process to the mold method greatly increased output and decreased the labor required for each bulb, such that workers went from producing 165 bulbs on the first day to 150 an hour.¹³

Second, we can try to improve the rate of production and take advantage of economies of scale—the fact that per-unit costs tend to fall as production volume rises. Glass furnaces in the bulb blank factory ran continuously, because starting a furnace cold took a great deal of time (24 hours or more) and was very likely to damage the crucibles. The furnaces were, therefore, burning coal regardless of whether glass was being blown and bulbs produced.¹⁴ Similarly, the rent needed to be paid whether the factory was producing bulb blanks or not. For these reasons, a factory that manufactured bulbs continuously over 24 hours would have lower unit costs than a factory that only operated for eight hours a day (and, in fact, some glass manufacturers did run continuously for this reason).¹⁵

Third, we can try to reduce variation in the process. The quality of the glass was dependent on the temperature of the furnace: Variations in temperature would result in glass that would break after a short period of use. Reducing temperature variation would, therefore, result in more glass within acceptable bounds, producing a higher yield.

Fourth, we can try to decrease the costs of our inputs. Replacing the hand-blowing process with the bulb mold process not only reduced the amount of labor required but also enabled the factory to use less expensive labor, since the molding process required less skill.¹⁶

Fifth, we can try to reduce our work in process by decreasing the size of our buffers. Work in process is material that has been paid for but hasn't yet been sold—it's an investment that has yet to yield a return. If glass furnaces were filled with new glass once a week,

on average there would be half a week's worth of glass simply sitting idle in the factory.¹⁷ If the crucibles were instead half the size and filled twice a week, on average there would only be a quarter of a week's supply of glass in the crucibles, reducing work in process by 50 percent.

We also have one more option available to us: We can try to delete an entire step in the process. This will, obviously, remove all of its associated costs. If, for instance, it becomes possible to buy premixed glass powder, we no longer need to perform the mixing step ourselves—our input materials can instead go directly to the glass crucibles.

These are the options available to improve the efficiency of a process. So, what does this suggest about what an extremely efficient process looks like?

-
- It's a process with no buffers. Material moves smoothly from one step to the next without any waiting or delay, and material tied up in the process is minimized.
-
- It's a process with no variability. The process works every time and always produces exactly what it's supposed to, at exactly the time when it's needed. More generally, the output of the process is as close to perfectly predictable as possible.
-
- It's a process with no unnecessary or wasteful steps. Every step is contributing value, and no steps can be eliminated.
-
- It's a process with inputs that are as cheap as possible and no wasted outputs. Either all inputs are successfully transformed, or the ancillary outputs are repurposed elsewhere.
-
- It's a process that acts at as large a scale as the technology and market will allow. Fixed costs are spread over as much output as possible, and the process takes maximum advantage of scale effects.
-
- It's a process that uses transformation methods that require as few inputs as possible, at the limits of what production technology will allow.

This sort of production process is sometimes called a *continuous flow* process—it continuously transforms inputs into outputs without any delays, downtime, waiting, unnecessary steps, or unneeded inputs. A steady stream of inputs goes in, and a steady stream of completed products swiftly and smoothly comes out.

One way of thinking about a continuous flow process is that it's like driving on the highway. In the city, there's the constant stop-and-start

of traffic lights and waiting behind other cars. But on the highway, the flow of cars is consistent and uninterrupted, as one car smoothly follows another. In practice, it's often not possible to achieve a true continuous-flow process, just as it's not always possible for traffic to flow perfectly smoothly on the highway. The technology may not allow it, or the size of the market may not justify the cost of the equipment required. There are any number of reasons why continuous flow may not be achievable. But when it *is* possible, it results in the production of enormous volumes of incredibly inexpensive goods.

To see what a continuous flow process looks like in practice, let's look at how the light bulb manufacturing process evolved in the century after Edison.

In 1891, just over a decade after Edison's invention, the US was producing 7.5 million incandescent bulbs per year.¹⁸ By the turn of the 20th century, that figure had climbed to 25 million. But production was still largely manual, and the cost of light bulbs, though falling, was still high. In 1907, a 60-watt light bulb cost \$1.75, or about \$54 in 2022 dollars.¹⁹

In 1912, Corning introduced the first semiautomatic machine for blowing light bulbs, called the Empire E machine. Though it still required workers to manually gather the molten glass, the machine could produce bulbs at a rate of 400 per hour, over twice as fast as the manual mold method.²⁰ This was followed by General Electric's fully automatic Westlake machine, as well as Corning's Empire F. In 1921, a Westlake machine could manufacture over 1,000 bulbs an hour. By the 1930s, improved Westlake machines could produce 5,000 bulb blanks an hour.²¹

The Westlake machine, though speedy, was largely a faster, mechanized version of the existing method for hand-blowing bulbs. It consisted of a large rotating drum with a series of iron tube arms mounted to it; as the machine rotated, the arms would lower into a glass furnace, gather a glob of molten glass, and swing it into a mold, after which air would be blown into it to form the bulb. Then, in 1926, a new type of machine for manufacturing bulb blanks was introduced: the Corning ribbon machine. Unlike previous machines, which largely duplicated the manual bulb-blowing process, the ribbon machine used a different mechanism for forming the bulbs. Instead of gathering a blob of glass on an iron rod, molten glass was poured onto a conveyor belt, which produced a continuous ribbon of molten glass (giving the

machine its name). The glass would sag through holes in the belt, forming a bowl shape. As the conveyor moved, a mold attached to a second conveyor below would snap shut around the bowl-shaped glass and air would be blown in from above, forming the shape of the bulb. The formed bulbs were then released and carried away by conveyor belt.²²

What had previously been a process with many small stops and starts became an uninterrupted, continuous flow. Glass poured onto the conveyor, sagged through the holes, and was repeatedly transformed into finished bulbs, one after the other, without any delays or waiting. Every step was perfectly synchronized.

The ribbon machine was extraordinarily complex and required constant intervention to keep operational.²³ But it could produce bulb blanks in truly staggering volumes. The first ribbon machine produced 16,000 bulbs an hour—over three times faster than the Westlake machines. By 1930, an improved ribbon machine could produce 40,000 bulbs an hour.²⁴

The ribbon machine represented the final evolution of incandescent bulb blank production. It produced bulbs in such enormous quantities that by the early 1980s, fewer than 15 ribbon machines were needed for the entire world's supply of light bulbs. By then, machine improvements had increased the production volume to nearly 120,000 bulbs an hour, or 33 bulbs every second.²⁵

Similar improvements took place in the rest of the light bulb manufacturing process, though none were quite so dramatic as the ribbon machine. In the late 19th and early 20th centuries, machines were developed to attach the inner tube to the outside bulb, mount the filament to the tube, make and then insert the lead-in wires, and seal the bulb.²⁶ Enhanced vacuum pumps were developed to evacuate bulbs much more quickly—Edison's original pumps took five hours to produce a vacuum in a bulb—and they did so automatically.²⁷

By the 1920s, most steps in the bulb manufacturing process had been automated, but they were largely performed by separate machines. Large volumes of in-process bulbs would accumulate between workstations, creating severe storage problems.²⁸ Starting in 1921, these steps were rearranged into groups, or cells, so one machine would smoothly feed another at synchronized rates.²⁹ Work in process was greatly reduced, storage requirements fell, and output per worker nearly doubled.³⁰ By 1930, the major manufacturing innovations were

complete, and by 1942, finished bulbs could be produced by a work cell at a rate of 1,000 per hour.³¹

As a result of these improvements, the cost of a light bulb plummeted. By 1942, the cost of a 60-watt bulb had fallen to 10¢. Over this same period, bulb efficiency, the amount of light emitted per watt, also improved, nearly doubling from 1907 to 1942. Combined with cheaper electricity, the cost per lumen dropped 98.5 percent between 1882 and 1942.³²

Other parts of the light bulb-making process benefited from the same types of improvements: new production technology that required fewer inputs, increased economies of scale, reduced variability, minimization of buffers, and the elimination of unnecessary steps. As with bulb blanks, these processes gradually evolved toward a continuous, uninterrupted transformation of material.

Of course, such gains are not restricted to light bulbs. Any production process that can be described as a series of sequential steps can be made more efficient in the exact same ways. As we'll see throughout this book, these types of improvements have resulted in increased efficiency in everything from steelmaking to cargo shipping. Over the next several chapters, we'll take a closer look at each of the five factors of a production process and how they can contribute to increased production efficiency.



About the author

Brian Potter is the author of the Construction Physics newsletter and a senior infrastructure fellow at the Institute for Progress. He writes about the technology and economics of building construction with a focus on improving productivity and reducing costs. He previously managed an engineering team at Katerra, a SoftBank-backed construction startup, and has 15 years of experience as a structural engineer. He has a bachelor's in civil engineering from Georgia Tech and a master's in systems engineering from University of Central Florida.



Stripe Press
Ideas for progress
South San Francisco, California
press.stripe.com