

Automation and Jobs: When technology boosts employment

By James Bessen, Boston University School of Law

November, 2017

Abstract: Do industries shed jobs when they adopt new labor-saving technologies? Sometimes productivity-enhancing technology increases industry employment instead. In manufacturing, jobs grew along with productivity for a century or more; only later did productivity gains bring declining employment. What changed? Markets became saturated. While the literature on structural change provides reasons for the decline in the manufacturing share of employment, few papers can explain both the rise and subsequent fall. Using two centuries of data, a simple model of demand accurately explains the rise and fall of employment in the US textile, steel, and automotive industries. The model also predicts that computer technology should generate relatively greater job growth in nonmanufacturing industries today. Estimates show computer use is associated with declining employment in manufacturing industries, but not in other sectors.

JEL codes: J2, O3, N10

Keywords: Automation, technical change, sectoral growth, labor demand, manufacturing, computer technology, deindustrialization

Thanks for helpful comments from David Autor, Martin Fleming, Ricardo Hausman, Mike Meurer, Joel Mokyr, Bob Rowthorn, Anna Salomons and participants at the CID seminar at Harvard.

When an industry automates its production does its employment decline? There is widespread concern today that many jobs will be lost to new computer technologies. One recent paper concluded that new information technologies will put “a substantial share of employment, across a wide range of occupations, at risk in the near future” (Frey and Osborne 2013).


The example of manufacturing decline provides good reason to be concerned about technology and job losses. In 1958, the US broadwoven textile industry employed over 300 thousand production workers and the primary steel industry employed over 500 thousand. By 2011, broadwoven textiles employed only 16 thousand and steel employed only 100 thousand production workers.¹ Some of these losses can be attributed to trade, especially since the mid-1990s. However, overall since the 1950s in developed economies, most of the decline in manufacturing employment appears to come from technology and changing demand (Rowthorn and Ramaswamy 1999). Of course, jobs lost in one industry may be replaced with new jobs in another, but many fear that deindustrialization has heightened economic inequality in the US or caused unemployment in Europe and caused political instability in both. “Premature deindustrialization” has also become a concern in developing nations (Rodrik 2016).

Yet technical change does not *always* lead to declining employment in the affected industry. Figure 1 shows how textiles, steel, and automotive manufacturing all enjoyed strong employment growth during many decades that also experienced very rapid productivity growth. Despite persistent and substantial productivity growth, these industries have spent more decades with growing employment than with job losses. This “inverted U”

¹ These figures are for the broadwoven fabrics industry using cotton and manmade fibers, SIC 2211 and 2221, and the steel works, blast furnaces, and rolling mills industry, SIC 3312.

pattern appears to be quite general for manufacturing industries across many industries in the developed world and also in developing nations (Buera and Kaboski 2009, Rodrik 2016).

Figure 1. Production Employment in Three Industries



This presents an important puzzle. Why is automation associated with growing employment in some industries at some times, while at other times and in other industries jobs are lost? What determines whether technical change tends to increase employment in an industry or decrease it? The answer would seem to be important for understanding what affects the pace of deindustrialization in both the developed and developing worlds and for anticipating the impact of new technologies.

This paper uses century-long time series data on the US cotton textile, steel, and automotive manufacturing industries to explore what determines whether technology will increase or decrease employment. While a substantial literature has looked at structural change at the level of the manufacturing sector as a whole, the data for these individual industries allows a tighter identification of the interaction between technology, demand, prices, and income.² A simple model accounts for the inverted-U shape and the accompanying changes in demand. Despite abstracting away from several important factors that might affect industry employment, the model predicts the rise and fall of employment in these industries with reasonable accuracy: the solid line in Figure 1 shows that prediction. I then explore implications of this model. One implication is that computer automation should have a more positive effect on employment in nonmanufacturing industries than in manufacturing. I find support for this hypothesis using data on computer use across industries from 1984 through 2007.

Structural change

The inverted U pattern in Figure 1 is also seen in the relative share of employment in the whole manufacturing sector, shown in Figure 2. Logically, the rise and fall of the sector

² Papers empirically analyzing the sector shifts include Dennis and Iscan (2009), Buera and Kaboski (2009), Kollmeyer (2009), Nickell, Redding, and Swaffield (2008), and Rowthorn and Ramaswamy (1999).

Figure 2. Manufacturing Share of the Labor Force



Sources: US Bureau of the Census 1975; BLS Current Employment Situation. Labor force includes agricultural laborers

as a whole in this chart results from the aggregate rise and fall of separate manufacturing industries such as those in Figure 1. Yet explanations of this phenomenon based on broad sector-level factors face a challenge because individual industries show rather disparate patterns. For example, employment in the automotive industry appears to have peaked nearly a century after textile employment peaked. Data on individual industries are needed to analyze such disparate responses.

The literature on structural change provides two sorts of accounts for the relative size of the manufacturing sector, one based on differential rates of productivity growth, the

other based on different income elasticities of demand.³ Baumol (1967) showed that the greater rate of technical change in manufacturing industries relative to services leads to a declining share of manufacturing employment under some conditions (see also Lawrence and Edwards 2013, Ngai and Pissarides 2007, Matsuyama 2009).

But differences in productivity growth rates do not seem to explain the initial rise in employment. For example, during the 19th century, the share of employment in agriculture fell while employment in manufacturing industries such as textiles and steel soared both in absolute and relative terms. But labor productivity in these manufacturing industries grew faster than labor productivity in agricultural. Parker and Klein (1966) find that labor productivity in corn, oats, and wheat grew 2.4%, 2.3%, and 2.6% per annum from 1840-60 to 1900-10. In contrast, labor productivity in cotton textiles grew 3.0% per year from 1820 to 1900 and labor productivity in steel grew 3.0% from 1860 to 1900.⁴ Nevertheless, employment in cotton textiles and in primary iron and steel manufacturing grew rapidly then.

The growth of manufacturing relative to agriculture surely involves some general equilibrium considerations, perhaps involving surplus labor in the agricultural sector (Lewis 1954). But at the industry level, rapid labor productivity growth along with job growth must mean a rapid growth in the equilibrium level of demand—the amount consumed must increase sufficiently to offset the labor-saving effect of technology. For example, although labor productivity in cotton textiles increased nearly 30-fold during the nineteenth century, consumption of cotton cloth increased 100-fold. The inverted U thus seems to involve an interaction between productivity growth and demand.

³ Acemoglu and Guerrieri (2008) also propose an explanation based on differences in capital deepening.

⁴ My estimates, data described below.

A long-standing literature sees sectoral shifts arising from differences in the income elasticity of demand. Clark (1940), building on earlier statistical findings by Engel (1857) and others, argued that necessities such as food, clothing, and housing have income elasticities that are less than one (see also Boppart 2014, Comin, Lashkari, and Mestieri 2015, Kongsamut, Rebelo, and Xie 2001 and Matsuyama 1992 for more general treatments of nonhomothetic preferences). The notion behind “Engel’s Law” is that demand for necessities becomes satiated as consumers can afford more, so that wealthier consumers spend a smaller share of their budgets on necessities. Similarly, this tendency is seen playing out dynamically. As nations develop and their incomes grow, the relative demand for agricultural and manufactured goods falls and, with labor productivity growth, relative employment in these sectors falls even faster.

This explanation is also incomplete, however. While a low income elasticity of demand might explain late 20th century deindustrialization, it does not easily explain the rising demand for some of the same goods during the nineteenth century. By this account, cotton textiles are a necessity with an income elasticity of demand less than one. Yet during the 19th century, the demand for cotton cloth grew dramatically as incomes rose. That is, cotton cloth must have been a “luxury” good then. Nothing in the theory explains why the supposedly innate characteristics of preferences for cloth changed.

It would seem that the nature of demand changed over time. Matsuyama (2002) introduced a model where the income elasticity of demand changes as incomes grow (see also Foellmi and Zweimueller 2008). In this model, consumers have hierarchical preferences for different products. As their incomes grow, consumer demand for existing products saturates and they progressively buy new products further down the hierarchy. Given heterogeneous incomes that grow over time, this model can explain the inverted U pattern.

It also corresponds, in a highly stylized way, to the sequence of growth across industries seen in Figure 1.

Yet there are two reasons that this model might not fit the evidence very well for individual industries. First, the timing of the growth of these industries seems to have much more to do with particular innovations that began eras of accelerated productivity growth than with the progressive saturation of other markets. Cotton textile consumption soared following the introduction of the power loom to US textile manufacture in 1814; steel consumption grew following the US adoption of the Bessemer steelmaking process in 1856, and Henry Ford's assembly line in 1913 initiated rapid growth in motor vehicles.

Second, there is a general problem of looking at the income elasticity of demand as the main driver of structural change: the data suggest that prices were often far more consequential for consumers than income. From 1810 to 2011, real GDP per capita rose 30-fold, but output per hour in cotton textiles rose over 800-fold; inflation-adjusted prices correspondingly fell by three orders of magnitude. Similarly, from 1860 to 2011, real GDP per capita rose 17 fold, but output per hour in steel production rose over 100 times and prices fell by a similar proportion. The literature on structural change has focused on the income elasticity of demand, often ignoring price changes. Yet these magnitudes suggest that low prices might substantially contribute to any satiation of demand. I develop a model that includes both income and price effects on demand, allowing both to have changing elasticities over time.

The inverted U pattern in industry employment can be explained by a declining price elasticity of demand. If we assume that rapid productivity growth generated rapid price declines in competitive product markets, then these price declines would be a major source of demand growth. During the rising phase of employment, equilibrium demand had to

increase proportionally faster than the fall in prices in response to productivity gains. During the deindustrialization phase, demand must have increased proportionally less than prices. Below I obtain estimates that show the price elasticity of demand falling in just this manner.

To understand why this may have happened, it is helpful to return to the origins of the notion of a demand curve. Dupuit (1844) recognized that consumers placed different values on goods used for different purposes. A decrease in the price of stone would benefit the existing users of stone, but consumers would also buy stone at the lower price for new uses such as replacing brick or wood in construction or for paving roads. In this way, Dupuit showed how the distribution of uses at different values gives rise to what we now call a demand curve, allowing for a calculation of consumer surplus.

This paper proposes a parsimonious explanation for the rise and fall of industry employment based on a simple model where consumer preferences follow such a distribution function. The basic intuition is that when most consumers are priced out of the market (the upper tail of the distribution), demand elasticity will tend to be high for many common distribution functions. When, thanks to technical change, price falls (income rises) to the point where most consumer needs are met (the lower tail), then the price and income elasticities of demand will be small. The elasticity of demand thus changes as technology brings lower prices to the affected industries and higher income to consumers generally.

I fit the model to actual demand data for the three industries with a lognormal specification that allows for changes in both the price elasticity of demand and the income elasticity of demand. The model estimates per capita demand accurately using only a single independent variable: labor productivity. I use the demand estimates to make the predictions of the actual rise and fall of employment in the textile, steel, and automotive industries shown in Figure 1.

This simple model of labor productivity and demand thus provides an explanation for the inverted U seen in manufacturing employment. The model implies that the characteristics of demand should determine whether the new technologies tend to increase or decrease employment. To test this, I compare the impact of computer technology on industry employment growth between manufacturing and nonmanufacturing industries. Inferring from the historical analysis that demand for manufactured products is less elastic on average than demand in other industries, the model suggests that computer technology should have a more positive impact on employment in nonmanufacturing industries. I find that from 1984 through 2007 computer technology use is associated with economically and statistically significant job losses in manufacturing industries but with a weak increase in employment in nonmanufacturing industries. This analysis supports Autor's (2015) explanation of why automation has not eliminated employment.

Model

Simple model of the Inverted U

Consider production and consumption of two goods—cloth and a general composite good—in autarky. The model will focus on the impact of technology on employment in the textile industry under the assumption that the output and employment in the textile industry are only a small part of the total economy.

Production

Let the output of cloth, $q = A \cdot L$, where L is textile labor and A is a measure of technical efficiency. Changes in A represent labor-augmenting technical change. Note that this is distinct from those cases where automation completely replaces human labor. Bessen

(2016) shows that such cases are rare and that the main impact of automation consists of technology augmenting human labor.

I initially assume that product and labor markets are competitive so that the price of cloth is

$$(1) \quad p = w/A,$$

where w is the wage. Below, I will test whether this assumption holds in the cotton and steel industries.

Then, given a demand function, $D(p)$, equating demand with output implies

$$D(p) = q = A \cdot L \quad \text{or}$$

$$(2) \quad L = D(p)/A.$$

We seek to understand whether an increase in A , representing technical improvement, results in a decrease or increase in employment L . That depends on the price elasticity of demand, ϵ , assuming income is constant. Taking the partial derivative of the log of (2) with respect to the log of A ,

$$\frac{\partial \ln L}{\partial \ln A} = \frac{\partial \ln D(p)}{\partial \ln p} \frac{\partial \ln p}{\partial \ln A} - 1 = \epsilon - 1, \quad \epsilon \equiv -\frac{\partial \ln D(p)}{\partial \ln p}$$

If the demand is elastic ($\epsilon > 1$), technical change will increase employment; if demand is inelastic ($\epsilon < 1$), jobs will be lost. In addition to this price effect, changing income might also affect demand as I develop below.

Consumption

Now, consider a consumer's demand for cloth. Suppose that the consumer places different values on different uses of cloth. The consumer's first set of clothing might be very valuable and the consumer might be willing to purchase even if the price were quite high. But cloth draperies might be a luxury that the consumer would not be willing to purchase

unless the price were modest. Following Dupuit (1844) and the derivation of consumer surplus used in industrial organization theory, these different values can be represented by a distribution function. Suppose that the consumer has a number of uses for cloth that each give her value v , no more, no less. The total yards of cloth that these uses require can be represented as $f(v)$. That is, when the uses are ordered by increasing value, $f(v)$ is a scaled density function giving the yards of cloth for value v . If we suppose that our consumer will purchase cloth for all uses where the value received exceeds the price of cloth, $v > p$, then for price p , her demand is

$$D(p) = \int_p^{\infty} f(z) dz = 1 - F(p), \quad F(p) \equiv \int_0^p f(z) dz$$

where I have normalized demand so that maximum demand is 1. With this normalization, f is the density function and F is the cumulative distribution function. I assume that these functions are continuous with continuous derivatives for $p > 0$.

The total value she receives from these purchases is then the sum of the values of all uses purchased,

$$U(p) = \int_p^{\infty} z \cdot f(z) dz.$$

This quantity measures the gross consumer surplus and can be related to the standard measure of net consumer surplus used in industrial organization theory (Tirole 1988, p. 8) after integrating by parts:

$$U(p) = \int_p^{\infty} z \cdot f(z) dz = \int_p^{\infty} z \cdot D'(z) dz = p \cdot D(p) + \int_p^{\infty} D(z) dz.$$

In words, gross consumer surplus equals the consumer's expenditure plus net consumer surplus. I interpret U as the utility that the consumer derives from cloth.⁵

The consumer also derives utility from consumption of the general good, x , and from leisure time. Let the portion of time the consumer works be l so that leisure time is $1 - l$. Assume that the utility from these goods is additively separable from the utility of cloth so that total utility is

$$U(v) + G(x, 1 - l)$$

where G is a concave differentiable function. The consumer will select v , x , and l to maximize total utility subject to the budget constraint

$$wl \geq x + pD(v)$$

where the price of the composite good is taken as numeraire. The consumer's Lagrangean can be written

$$\mathcal{L}(v, x, l) = U(v) + G(x, 1 - l) + \lambda(wl - x - p \cdot D(v)).$$

Taking the first order conditions, and recalling that under competitive markets, $p = w/A$, we get

$$\hat{v} = G_l \frac{p}{w} = \frac{G_l}{A}, \quad G_l \equiv \frac{\partial G}{\partial l}.$$

G_l represents the marginal value of leisure time and the second equality results from applying assumption (1). In effect, the consumer will purchase cloth for uses that are at least as valuable as the real cost of cloth valued relative to leisure time. Note that if G_l is constant, the effect of prices and the effect of income are inversely related. This means that the price

⁵ Note that in order to use this model of preferences to analyze demand over time, one of two assumptions must hold. Either there are no significant close substitutes for cloth or the prices of these close substitutes change relatively little. Otherwise, consumers would have to take the changing price of the potential substitute into account before deciding which to purchase. If there is a close substitute with a relatively static price, the value p can be reinterpreted as the value relative to the alternative. Below I look specifically at the role of close substitutes for cotton cloth, steel, and motor vehicles.

elasticity of demand will equal the income elasticity of demand. However, the marginal value of leisure time might very well increase or decrease with income; for example, if the labor supply is backward bending, greater income might decrease equilibrium G_l so that leisure time increases. To capture that notion, I parameterize $G_l = w^\alpha$ so that

$$(3) \quad \hat{v} = w^\alpha / A = w^{\alpha-1} p, \quad D(\hat{v}) = 1 - F(\hat{v}).$$

Elasticities

Using (3), the price elasticity of demand holding wages constant solves to

$$\epsilon = -\frac{\partial \ln D}{\partial \ln p} = \frac{\partial \ln D(\hat{v})}{\partial \ln \hat{v}} \frac{\partial \ln \hat{v}}{\partial \ln p} = \frac{pf(\hat{v})}{1 - F(\hat{v})} w^{\alpha-1}$$

and the income (wage) elasticity of demand holding price constant is

$$\rho = \frac{\partial \ln D}{\partial \ln w} = \frac{\partial \ln D(\hat{v})}{\partial \ln \hat{v}} \frac{\partial \ln \hat{v}}{\partial \ln w} = (1 - \alpha)\epsilon.$$

These elasticities change with prices and wages or alternatively with changes in labor productivity, A . The changes can create an inverted-U in employment. Specifically, if the price elasticity of demand, ϵ , is greater than 1 at high prices and lower than 1 at low prices, then employment will trace an inverted U as prices decline with productivity growth. At high prices relative to income, productivity improvements will create sufficient demand to offset job losses; at low prices relative to income, they will not.

A preference distribution function with this property can generate a kind of industry life cycle as technology continually improves labor productivity over a long period of time. An early stage industry will have high prices and large unmet demand, so that price decreases result in sharp increases in demand; a mature industry will have satiated demand so further price drops only produce an anemic increase in demand.

A necessary condition for this pattern is that the price elasticity of demand must increase with price over some significant domain, so that it is smaller than 1 at low prices but larger than 1 at high prices. It turns out that many distribution functions have this property.

This can be seen from the following propositions (proofs in the Appendix):

Proposition 1. Single-peaked density functions. If the distribution density function, f , has a single peak at $p = \bar{p}$, then $\frac{\partial \epsilon}{\partial p} \geq 0 \quad \forall p < \bar{p}$.

Proposition 2. Common distributions. If the distribution is normal, lognormal, exponential, or uniform, there exists a p^* such that for $0 < p < p^*$, $\epsilon < 1$, and for $p^* < p$, $\epsilon > 1$.

These propositions suggest that the model of demand derived from distributions of preferences might be broadly applicable. The second proposition is sufficient to create the inverted U curve in employment as long as price starts above p^* and declines below it.

Econometric specification

Below, I estimate the demand function using a lognormal distribution. To do this using aggregate data, the model, which describes the demand of an individual consumer, needs to be recast to describe aggregate consumer behavior. The distribution function can be recast as an aggregate distribution across both different consumers and different uses for each consumer. Also, an average wage now determines the equilibrium value of \hat{v} , so the distribution also reflects dispersion of wages across consumers.

Specifically, I estimate per capita demand, D ,

(4)

$$D = \gamma \left(1 - \Phi \left(\frac{-\ln A + \alpha \ln w - \mu}{\sigma} \right) \right) + \varepsilon$$

or

(5)

$$D = \gamma \left(1 - \Phi \left(\frac{\ln p/w + \alpha \ln w - \mu}{\sigma} \right) \right) + \varepsilon$$

where Φ is the standard normal cumulative distribution function and ε is an error term that captures, among other things, demand shocks and changing tastes. I estimate these equations using non-linear least squares (NLLS). I also estimate demand with a general polynomial form.

Data

Time series over a century in length often require combining data from different sources involving various adjustments. I describe the data sources and adjustments in detail in the Appendix. This section describes the main data series used in estimating employment in cotton textiles, steel, and automotive industries and in the computer technology analysis.

Production and demand

I use physical quantities to measure production and demand. For the textile industry, I measure output as yards of cotton cloth produced plus yards of cloth made of synthetic fibers from 1930 on. From 1958, I use the deflated output of the cotton and synthetic fiber broadwoven cloth industries (SIC 2211 and 2221). For the early years, I also included estimates of cotton cloth produced in households. For steel, I used the raw short tons of steel produced. For the motor vehicle industry I used the number of passenger vehicles and trucks produced each year.

To estimate per capita demand or consumption, I add net imports to the estimates of domestic production and divide by the population.

Note that these measures do not adjust for product quality.⁶ This approach avoids distortions that might arise from constructing quality adjusted price indices over long periods of time. It does mean that “true” demand and productivity are understated. However, this does not pose a significant problem for my analysis because I measure both without quality adjustments. The distribution function I estimate would, of course, be different if it were estimated with quality-adjusted data, but using unadjusted data allows for consistent predictions of employment.

Employment, prices, and wages

I count the number of industry wage earners or, from 1958 on, the number of production workers. For prices, I use the prices of standard commodities. For cotton textiles, I use the wholesale price for cotton sheeting. For steel, I use wholesale prices for steel rails. I do not have a similar commodity price for motor vehicles. The BLS does have a price index for the automotive industry, but this measure implicitly changes as the quality of vehicles improved. I need to use a commodity type price because my measures of output and consumption (cars and trucks) does not capture these quality improvements. For wages, I use the compensation of manufacturing production workers. This measure includes the value of employee benefits from 1906 on.

Labor productivity


I calculate labor productivity by dividing output by the number of production employees times the number of hours worked per year. I use industry specific estimates of hours if available and estimates of hours for manufacturing workers if not.

Over the sample periods, each industry exhibited rapid labor productivity growth. From 1820 to 1995, labor productivity in cotton textiles grew 2.9% per year; in steel, it grew

⁶ The deflators used from 1958 on in cotton are quality adjusted but the series closely matches the unadjusted output measure during the years when they overlap.

2.4% per year from 1860 to 1982; in motor vehicles, it grew 1.4% per year from 1910 through 2007. Figure 3 shows labor productivity for each industry on a log scale over time. Each industry exhibits steady productivity growth over long periods of time. Textiles and especially automotive show initially higher rates of growth; steel exhibits faster growth since the 1970s, likely the effect of steel minimills that use recycled steel rather than blast furnace production of iron.

Figure 3. Labor Productivity over Time



Computer technology

The data on computer use come from the Current Population Survey (CPS) of the Bureau of Labor Statistics. In select years, a supplemental survey asked whether the

respondent “directly used a computer at work?”⁷ The universe for these questions was employed persons of 15 years of age or older. I use these survey items as the measure of computer use. Because these measures are self-reported, they may undercount cases where information technology is embedded in hardware. I use the Census 1990 industry classification scheme where IPUMS crosswalks have converted industry codes to a consistent set of 227 industries (Ruggles 2015). I calculate the share of work hours performed by workers using computer technology for each industry for each sample year. Note that I assume that the use of computers means that industries are subject to an ongoing stream of productivity-enhancing technical changes as new hardware and software applications are installed.

Empirical Findings

Applying the model

The model above provides a parsimonious explanation for the inverted U shape of industry employment observed over time. But how well does this highly simplified model actually predict the patterns observed? A close fit would provide some support for its relevance. However, the model abstracts away from several considerations that might undermine efforts to fit the model, considerations that I discuss in this section. In general, I find that the model fits the data for employment and consumption rather well despite these concerns, except during the most recent decades of the textile and steel industries.

One concern is that the model assumes no substantial interference from close substitute products. That means that either there are no substitutes or that the productivity growth in substitutes is sufficiently slow that the effect of substitution can be taken as

⁷ I use the surveys of computer use in October of 1984, 1989, 1993, 1997, and 2003 and in September of 2001.

constant. Each industry did have substitutes, especially during the early years. However, it seems that these substitutes were fairly static technologically and were quickly overtaken. Cotton cloth competed with wool and linen. However, wool and linen were mainly produced within the household (Zevin 1971) and did not directly compete in most markets. In urban markets where they did compete, wool tended to be substantially more expensive per pound and its price declined only slowly compared to cotton.⁸ During the early years of the Bessemer steel process, steel rails were much more expensive than iron rails, but steel rails lasted much longer, making the higher price worth it for many uses. By 1883, the price of steel rails fell below the price of iron rails, eliminating the production of this substitute (Temin 1964 p. 222). And cars and trucks competed with horse drawn vehicles during the early years. However, here, too, production of horse drawn vehicles collapsed very quickly.⁹

It is also possible that new technologies introduce new substitutes or find new uses for commodities, changing the shape of the preference distribution function. Since the 1970s, steel may have faced greater competition from aluminum and other materials for use in cars and cans (Tarr 1988 p. 177-8), perhaps contributing to the poorer fit of the model then (see below). In any case, the close fit of the model overall suggests that substitution is not a significant problem.

Finally, in the stylized model, productivity equals wages over price, $A = w/p$. This is convenient for empirical analysis because it means the key independent variable can either be A or w/p . In practice, the productivity measure is preferable because it is less sensitive to

⁸ For example, in Philadelphia in 1820, wool was \$0.75 per pound while cotton sheeting was \$0.15 (US Bureau of the Census 1975).

⁹ The production of carriages, buggies, and sulkies fell from 538 thousand in 1914 to 34 thousand in 1921; the production of farm wagons, horse-drawn trucks, and business vehicles fell from 534 thousand in 1914 to 67 thousand in 1921 (US Bureau of the Census 1975).

demand shocks and because I have data over a longer time span for productivity and for motor vehicles.

For various reasons, A might diverge from w/p . For purposes of my empirical analysis, labor productivity can be used as the independent variable as long as it is proportional to wage over price, not necessarily equal.¹⁰ Figure 4 shows that this is largely the case for textiles and steel. The correlation coefficients are 0.98 for textiles and 0.88 for steel (0.98 through 1974 only). The only major deviation appears for steel from 1974 to 1982 when rising costs of energy and materials increased substantially as a share of output. In the curve-fitting exercise below, I use specifications both with A and with w/p .

Demand Curves

Before seeking to estimate the model, it is helpful to first examine the demand curves graphically. Figure 5 shows per capita demand (consumption) for each good against labor productivity, both on logarithmic scales. The solid line is simply a log version of equation (4) fit to the data.

In Figure 5A and 5B, the circles represent observations where the measure of demand fails to capture the effect of imports of downstream products. Demand needs to take trade into consideration and so I have calculated demand by adding net imports to the amount of product produced domestically. However, for textiles and steel further adjustment is needed because these are intermediate goods industries. The ultimate consumption good is produced by another industry and that good can be imported as well. For example, the consumption of textiles in the form of apparel includes: 1) apparel produced in the US with US cloth, 2) textiles that were imported to the US and used by

¹⁰ This will be true if the labor share of output is constant or if technical change is Hicks neutral.




Figure 5. Per Capita Consumption

domestic apparel producers, and 3) apparel that were produced outside the US using cloth also produced outside the US. Even after adjusting for imports of textiles, my measure of consumption misses the cloth imported in apparel made abroad.

For this reason, I can only estimate demand for those years where downstream imports are not too large. For textiles, I estimate demand through 1995; in 1996, imports comprised a third of apparel imports for the first time and have grown rapidly since. For steel, I estimate demand through 1982. After that, the largest steel-using industries, fabricated metal products and machinery excluding computers (SIC 34 and 35 excluding 357), show a large increase in import penetration. Between 1982 and 1987, the import penetration (net imports over domestic production) grew 10.5%. As the Figure shows, my measure of per capita consumption falls dramatically around these cutoff years. I also conducted my estimates using different cutoff years as a robustness check. Prediction of employment and the general inverted U pattern were not sensitive to small changes in the cutoff year.

Under the assumption that labor productivity is inversely proportional to price, the slope of the curve in the figure represents the price elasticity of demand. In each case, the price elasticity clearly decreases as labor productivity increases. As a first pass, Table 1 formally tests whether the elasticity of demand is constant by fitting the curve with a simple quadratic expression of the form

$$(6) \quad \ln D = \alpha + \beta \cdot \ln A + \gamma \cdot (\ln A)^2 + u.$$

Column 1 uses labor productivity as the base independent variable. Column 2 uses price over wage; the labor productivity variable has a somewhat better fit. In each case, the coefficient on the quadratic term is negative and highly significant, rejecting the null hypothesis that the price elasticity of demand is constant.

One concern is the possible endogeneity of price and labor productivity. A demand shock will affect both the equilibrium price and the error term, biasing the coefficient of price. A similar concern might relate to labor productivity if firms bring on less productive resources in response to positive demand shocks. In that case, the coefficient of labor productivity might be biased and, possibly, the predictions could be off as well. That problem might be mitigated for the data here because demand shocks are likely small compared to the large changes in demand over historical time frames. To test for endogeneity, Column 3 instruments the labor productivity variables in Column 1 using year and year squared as instruments in a GMM regression. Year should be independent of any demand shocks and the largely linear growth of log productivity seen in Figure 3 means that year should be a good instrument. The coefficients are quite similar for textiles and steel, but a bit different for auto. The table reports the probability value of a statistical test of the null hypothesis that all independent variables are exogenous.¹¹ The null hypothesis that the independent variables are exogenous cannot be rejected at the 5% level of significance, but it can be rejected at the 10% level for auto and cotton.

These tests thus raise some concern about endogeneity bias. However, my aim here is not unbiased coefficient estimates but accurate prediction. I compared the predictions of regressions of equation (6) with regressions that used year and year squared instead of the labor productivity variables. The predictions match fairly closely and are highly correlated. The correlation coefficients are 0.971 for textiles, 0.987 for steel, and 0.751 for automotive. The main disparity occurred in the automotive industry for 1910 when labor productivity was far below the 1920 value (thanks to Henry Ford). These tests provide some assurance

¹¹ The test is based on the difference in Sargan statistics; the specification used also shows a strong first stage regression; it is exactly identified.

that predictions from my model do not suffer substantially from the possible endogeneity of labor productivity.

Model estimates

Table 2 shows NLLS estimates of equations (4) in columns 1 and 2 and estimates of equation (5) for textile and steel in column 3. Columns 1 and 3 set $\alpha = 0$, excluding secondary income effects. All of the regressions have a good fit, although the regressions using labor productivity (columns 1 and 2) fit better than those using the ratio of prices to wages (column 3), probably because of the greater volatility of wholesale price data. Note that the model fits the data significantly better than estimates using a simple quadratic form in Table 1. None of the estimates in column 2 find a significant coefficient for α , suggesting that changes in the marginal value of leisure time are not important.¹² The lines in Figure 5 represent the predictions based on the column 1 regressions.

Using these predictions, I estimate the price elasticity of demand at each end of the estimation sample:

Cotton		Steel		Automotive	
Year	Elasticity	Year	Elasticity	Year	Elasticity
1810	2.13	1860	3.49	1910	6.77
1995	0.02	1982	0.16	2007	0.15

Demand was initially highly elastic, becoming highly inelastic.

These predicted levels of per capita demand can also be used to estimate industry production employment by dividing domestic demand (total demand divided by $1 + \text{import}$

¹² Of course, leisure time increases dramatically over this historical period studied.

penetration) by the annual output per production worker. Measuring labor productivity as output per production worker-hour this is¹³

$$\frac{\text{Demand per capita} \cdot \text{Population}}{1 + \text{Import penetration}} \cdot \frac{1}{\text{Labor productivity} \cdot \text{Hours worked/year}}$$

These estimate are shown as the solid lines in Figure 1. The estimates appear to be accurate over long periods of time. There are notable drops in employment during the Great Depression and excess employment in motor vehicles during World War II. Finally, employment drops sharply for the years when my measure of consumption fails in textiles (after 1995) and steel (after 1982). It appears that this simple model using a lognormal distribution of preferences provides a succinct explanation of the inverted U in employment in these industries.

Implications

Trade vs. technology in manufacturing job losses

The model provides an estimate of the impact of technology on industry employment as mediated by demand. We can use this to understand how much technology contributed to the loss of manufacturing jobs compared to other factors including trade. Using predicted industry employment without the correction for imports, actual and predicted employment changes can be compared:

¹³ For 1820 and before, I also subtract the estimate of labor performed in households.

Year	Cotton & synthetic textiles Production workers (1000s)			Steel Production workers (1000s)		
	Actual	Predicted	Tech share	Actual	Predicted	Tech share
1950	350	336		550	514	
1995	120	134				
2011	16	98		100	153	
Job losses 1950 - 2011	334	238	71%	450	361	80%
Job losses 1995 - 2011	104	36	35%			

In both textiles and steel, most of the jobs have disappeared since 1950 and most of the losses can be attributed to growth in labor productivity without compensating growth in demand. Other factors, including trade, the recession, and, perhaps changing tastes, can account for only 29% of the job losses in textiles and 20% in steel.

That said, technology does not account for much of the more recent losses especially in textiles. Since 1995, only about one third of the job losses can be explained by growing labor productivity. Most of the loss in textile manufacturing jobs appears to be the result of the collapse of the domestic apparel industry in the face of heavy global competition. Note that the effect of trade on textile and steel manufacturing jobs does not appear to be mainly about imports of textiles and steel—the import penetration in textiles was 8.9% in 2005 and in steel it was 5.2%. Instead, the main impact of trade appears to have come through its effect on downstream industries.

The direction and rate of job changes

A naïve view holds that more rapid productivity growth will be more likely to create job losses. The analysis in this paper suggests, to the contrary, that demand determines whether productivity growth eliminates or increases jobs. Moreover, if demand is elastic,

then more rapid productivity growth will actually lead to *faster* employment growth, all else equal. While the sign of the employment effect depends on the elasticity of demand and not on the rate of productivity growth, the rate of change depends on both. For instance, although productivity grew faster in cotton textiles than in auto, employment grew faster in auto; the distribution of preferences for motor vehicles was much more concentrated (small σ) than the distribution for textiles.

One might expect that computer technology would have a different effect on job growth across industries depending on industry demand elasticity. Assuming that the historical process of deindustrialization means that manufacturing industries have less elastic (more satiated) demand than most other industries on average, then computer technology should have a relatively more negative impact on employment in manufacturing industries, all else equal.¹⁴

Table 3 performs a simple test of this hypothesis. From equation (2) it follows that, all else equal,

$$\Delta \ln L = (\epsilon - 1) \cdot \Delta \ln A.$$

I assume that labor productivity growth in industry i at time t is proportional to the share of workers using computers in that industry, C_{it} . Then, estimating

$$\Delta \ln L_{it} = \beta_{nonmanu} \cdot C_{it} \cdot I(nonmanu.) + \beta_{manu} \cdot C_{it} \cdot I(manu.) + \delta_t + \varepsilon_{it}$$

the hypothesis can be tested as a test of $\hat{\beta}_{nonmanu} > \hat{\beta}_{manu}$.

Column 1 shows that the share of the industry workforce using computers is positively associated with employment growth in nonmanufacturing industries, but is

¹⁴ To the extent that information technology allows manufacturers to create new products, then the technology might tap into new sources of demand and thus be more elastic. There is some evidence that information technology is used to create new products (see for example Bartel, Ichniowski, and Shaw 2007). New product varieties might provide a reason that the association between IT and employment growth in manufacturing is only weakly negative.

significantly associated with job loss in manufacturing. The difference in the coefficients is highly significant. The bottom panel of the table calculates the effect at the sample mean; computer use is associated with about 1% per annum faster employment growth on average in nonmanufacturing industries, but is associated with a loss of jobs of about 3% per year on average in manufacturing. These estimates do not take into account the possible impact of trade as was done above for the historical analysis. To consider the impact of trade, I repeated the regression after adjusting the dependent variable for the rate of change in the log import penetration for manufacturing industries.¹⁵ The coefficients were very close to the original estimates.

Column 2 adds two additional controls for occupational characteristics that might be related to industry growth. Industries that send work offshore would tend to grow more slowly. Jensen and Kletzer's (2010) developed an index for the likelihood that the work of occupations can be sent offshore. I use the weighted average over the occupations in each industry. Also, industries might grow at different rates depending on the relative supply of different type of labor they require. The Bureau of Labor Statistics O*NET database (version 17.0) rates occupations on the extent to which they require a college diploma or advanced degree. I include a weighted average of these ratings over the occupations in each industry. The additional controls are significant, but coefficients of computer use are similar.

Perhaps other industry characteristics might be associated with both computer use and industry growth. Column 3 repeats the regression but with a full set of industry fixed effects. The difference between the coefficients is still highly significant, although both coefficients are now negative. The standard errors are also substantially larger, not surprising, given the short panel length.

¹⁵ That is, I added $\Delta \ln output / (output + imports - exports)$.

Computer use might also be endogenous if, say, faster growing industries tend to also use computers more. For example, perhaps rapidly growing industries have greater difficulty managing resources and they rely more on computer systems to help manage. To check for endogeneity, I instrument computer use with a measure of how sedentary each industry's workforce is. The notion is that desktop computers can be more easily applied to sedentary jobs. I obtain the variable from the 1977 edition of the Dictionary of Occupational Titles.¹⁶ This variable is correlated with industry computer use in the sample (correlation coefficient of 0.43). However, because the 1977 estimates were made before computers were widely adopted, the instrument is likely independent of computer use; that is, computers could not have had much influence on whether jobs were sedentary then. The measure also appears to be independent of other unobserved variables that might be associated with employment growth. I performed t-tests comparing employment growth rates between sedentary and non-sedentary occupations for each decade. During the 1960s and 1970s, the difference in employment growth between these groups of occupations was not significant (probability values of 0.409 and 0.286 respectively). Thus there does not appear to have been a significant association between the instrumental variable and employment growth before computers became widely adopted. After 1980, however, the picture changes and the sedentary occupations grow significantly faster during the 1980s, 1990s, and from 2000 through 2013 (probability values of 0.000, 0.009, and 0.009 respectively). These tests suggest that while the instrument has been correlated with employment growth since computers were widely adopted, it was not correlated before then. This pattern suggests that sedentary

¹⁶ Experts rated each occupation for the level of physical exertion involved ranging from 1 (sedentary) to 5 (very heavy). England and Kilbourne (2013) have mapped the 14,000 occupations in the DOT to Census detailed occupation codes, averaging them to this higher level of aggregation. I classified each occupation with a mean score less than 2 as sedentary and then calculated weighted averages across all occupations for each industry in the 1980 Census 5% sample.

work does not influence employment growth except through computer use, satisfying the exclusion restriction.

Columns 4 and 5 perform separate instrumental variables estimations for non-manufacturing and manufacturing industries using this measure of sedentariness as well as year dummies as instruments.¹⁷ For nonmanufacturing industries, computer use is associated with industry employment growth that is both economically and statistically significant; in manufacturing, the association is not significantly different from zero. Given the larger standard errors, however, the difference between these sectors is not statistically significant ($P = .124$). Also, a test of the null hypothesis that computer use is exogenous cannot be rejected for the nonmanufacturing sample ($P = .267$), but can be rejected for manufacturing ($P = .005$).

In summary and consistent with my hypothesis, computer use is associated with lower and typically negative employment growth in manufacturing industries. In other sectors, computer use is associated with employment growth in most estimates. These results correspond to several recent papers finding that information technology increases employment for some groups and does not appear to reduce net employment, including some causal studies. Gaggl and Wright (2014) find that ICT tended to raise employment in wholesale, retail, and finance industries, but had no statistically significant effect on other sectors, including manufacturing. Akerman, Gaarder, and Mogstad (2015) find that Internet technology increased employment of skilled workers and had no effect on unskilled. Bessen (2016) finds that computers tend to increase occupational employment modestly overall,

¹⁷ I estimate the two samples separately in order to have a sufficient number of exogenous variables relative to endogenous variables.

with job losses in low wage occupations. Autor, Dorn, and Hansen (2015) find that local markets susceptible to computerization are not more likely to experience employment loss.

Conclusion

A simple model explains the rise and subsequent fall of manufacturing employment in the face of ongoing productivity growth. Productivity-enhancing technology will increase industry employment if product demand is sufficiently elastic. Technical change reduces the labor required needed to produce a unit of output, but it also reduces prices in competitive markets. If the price elasticity of demand is greater than one, the increase in demand will more than offset the labor saving effect of the technology.

Understanding the responsiveness of demand is thus key to understanding whether major new technologies will decrease or increase employment in affected industries. This paper proposes that industry employment dynamics can be analyzed by deriving demand from a distribution of preferences. For many distribution functions, the elasticity of demand declines as price declines and productivity grows. In particular, a parsimonious model using a lognormal distribution fits the demand curves well for cotton textiles, steel, and motor vehicles over long periods of time.

This model generates an industry life cycle explanation for the inverted U pattern of industrialization/deindustrialization seen in manufacturing employment. At high initial prices, industries have large unmet demand that is highly elastic. Productivity improvements give rise to robust job growth. Over time and with ongoing productivity gains, prices progressively decline until most demand is met and the price elasticity of demand is quite low. Then further productivity gains bring reduced employment.

This view implies that major new technologies today should increase employment if they improve productivity in markets that have large unmet needs. That appears to be the case with computer technology. I find, consistent with the analysis, that computer technology use is associated with employment growth in nonmanufacturing industries, but in manufacturing, where a couple centuries of rapid productivity growth mean that demand is more satiated, computer technology is associated with employment declines.

This model challenges a popular view that faster technical change is more likely to eliminate jobs. Some people argue that because of Moore's Law, the rate of change will be fast in new information technologies and this will cause unemployment (Ford 2015). However, my model shows that if demand is elastic, faster technical change will, instead, create *faster* employment growth. Faster technical change will, however, also hasten the day when demand is no longer so elastic and deindustrialization sets in.

This analysis raises a number of other questions. For one, it would be helpful to understand what factors shape the preference distribution functions. For instance, in the model, the pace of industrialization/deindustrialization is affected by the variance of the distribution of preferences. Nations with greater income equality might have more homogenous preferences and hence a narrower distribution (smaller standard deviation). A narrower distribution of preferences, in turn, implies more rapid employment growth during industrialization. In this way, income inequality might slow the pace of economic development. Correspondingly, income inequality might also affect the pace of deindustrialization as markets mature. Another area for investigation concerns trade. The model in this paper abstracts away from the affect of imports on demand. Clearly, imports might play a role in decelerating industrial development in exposed economies, heightening patterns of "premature deindustrialization" (Rodrik 2017).

Appendix

Propositions

To simplify notation, let the wage remain constant at 1. Then

$$\epsilon(p) = \frac{p f(p)}{1 - F(p)}$$

so that

$$\frac{\partial \epsilon(p)}{\partial p} = \frac{f' p}{1 - F} + \frac{f^2 p}{(1 - F)^2} + \frac{f}{1 - F} = \epsilon \left(\frac{f'}{f} + \frac{f}{1 - F} + \frac{1}{p} \right)$$

Note that the second and third terms in parentheses are positive for $p > 0$; the first term could be positive or negative. A sufficient condition for $\frac{\partial \epsilon}{\partial p} \geq 0$ is

(A1)

$$\frac{f'}{f} + \frac{f}{1 - F} \geq 0.$$

Proposition 1. For a single peaked distribution with mode \bar{p} , for $p < \bar{p}$, $f' \geq 0$ so that

$$\frac{\partial \epsilon}{\partial p} \geq 0.$$

Proposition 2. For each distribution, I will show that

$$\frac{\partial \epsilon}{\partial p} \geq 0, \quad \lim_{p \rightarrow 0} \epsilon = 0, \quad \lim_{p \rightarrow \infty} \epsilon = \infty.$$

Taken together, these conditions imply that for sufficiently high price, $\epsilon > 1$, and for a sufficiently low price, $\epsilon < 1$.

a. Normal distribution

$$f(p) = \frac{1}{\sigma} \varphi(x), \quad F(p) = \Phi(x), \quad \epsilon(p) = \frac{p}{\sigma} \frac{\varphi(x)}{(1 - \Phi(x))}, \quad x \equiv \frac{p - \mu}{\sigma}$$

where φ and Φ are the standard normal density and cumulative distribution functions respectively. Taking the derivative of the density function,

$$\frac{f'}{f} + \frac{f}{1-F} = -\frac{x}{\sigma} + \frac{\varphi(x)}{\sigma(1-\Phi(x))}.$$

A well-known inequality for the normal Mills' ratio (Gordon 1941) holds that for $x > 0$,¹⁸

(A2)

$$x \leq \frac{\varphi(x)}{1-\Phi(x)}.$$

Applying this inequality, it is straightforward to show that (A1) holds for the normal distribution. This also implies that $\lim_{p \rightarrow \infty} \epsilon = \infty$. By inspection, $\epsilon(0) = 0$.

b. Exponential distribution

$$f(p) \equiv \lambda e^{-\lambda p}, \quad F(p) \equiv 1 - e^{-\lambda p}, \quad \epsilon(p) = \lambda p, \quad \lambda, p > 0.$$

Then

$$\frac{f'}{f} + \frac{f}{1-F} = -\lambda + \lambda = 0$$

so (A1) holds. By inspection, $\epsilon(0) = 0$ and $\lim_{p \rightarrow \infty} \epsilon = \infty$.

c. Uniform distribution

$$f(p) \equiv \frac{1}{b}, \quad F(p) \equiv \frac{p}{b}, \quad \epsilon(p) = \frac{p}{b-p}, \quad 0 < p < b$$

so that

$$\frac{f'}{f} + \frac{f}{1-F} = \frac{1}{b-p} > 0.$$

By inspection, $\epsilon(0) = 0$ and $\lim_{p \rightarrow b} \epsilon = \infty$.

¹⁸ I present the inverse of Gordon's inequality.

d. Lognormal distribution

$$f(p) \equiv \frac{1}{p\sigma} \varphi(x), \quad F(p) \equiv \Phi(x), \quad \epsilon(p) = \frac{1}{\sigma} \frac{\varphi(x)}{(1 - \Phi(x))}, \quad x \equiv \frac{\ln p - \mu}{\sigma}$$

so that

$$\frac{\partial \epsilon(p)}{\partial p} = \epsilon \left(\frac{f'}{f} + \frac{f}{1 - F} + \frac{1}{p} \right) = \epsilon \left(-\frac{1}{p} - \frac{x}{p\sigma} + \frac{\varphi}{p\sigma(1 - \Phi)} + \frac{1}{p} \right).$$

Cancelling terms and using Gordon's inequality, this is positive. And taking the limit of

Gordon's inequality, $\lim_{p \rightarrow \infty} \epsilon = \infty$. By inspection $\lim_{p \rightarrow 0} \epsilon = 0$.

Historical data sources

I obtain data on production employees for cotton and steel from Lebergott (1966, see also US Bureau of the Census 1975) through 1950, and from 1958 on from the NBER-CES manufacturing database for SIC 2211 and 2221 (broadwoven fabric mills, cotton and manmade fibers and silk) and SIC 3312 (primary iron and steel). The former measures the number of wage earners while the more recent series measure production employees. I find that these series are reasonably close for overlapping years. For 1820 in cotton, I estimate 5,600 full time equivalent workers producing in households, using estimates of household production and Davis and Stettler's (1966) estimates of output per worker. For the auto industry, I use the BLS Current Employment Statistics series for motor vehicle production workers from 1929 on. For 1910 and 1920, I obtained the number of employees in the motor vehicle industry from the 1% Census samples (Ruggles et al. 2015) and prorated those figures by the ratio of BLS production workers to Census industry employees for 1930.

Weekly hours data for motor vehicles also come from the BLS from 1929 on. For earlier years and for cotton and steel before 1958, I use Whaples (2001) before 1939, linearly interpolating for missing year observations. From 1939 to 1958 I use the BLS Current Employment Statistics series for manufacturing production and nonsupervisory personnel. In cotton and steel, I use the NBER-CES data for production hours from 1958 on (this comes from the BLS industry data).

For cotton production, I begin with Davis and Stettler's (1966, Table 9) estimates of yards produced per man-year for 1820 and 1831 multiplied by the estimate of the number of cotton textile wage earners for those years (I assume productivity was the same in 1830 and 1831). For 1820, I estimate that an additional 9.6 million yards were produced in households based on data from Tryon (1917). From 1830 on, Tryon's estimates indicate little cotton cloth was produced at home. From 1840 through 1950, I use estimates of the pounds of cotton consumed in textile production times three yards per pound (US Bureau of the Census 1975 and Statistical Abstracts, various years). This ratio is the historically used rule of thumb, but I also found that it applies reasonably well to a variety of twentieth century test statistics. While some cotton is lost in the production process (5% or less typically), these losses changed little over time. From 1930 on, I also include the weight of manmade fibers

consumed in textile production. From 1958 on I found that the deflated output of SIC 2211 and 2221 in the NBER-CES tracked the pounds of fiber consumed closely for the ten years when I had measures of both. I used the average ratio for these years to estimate yards of cloth produced based on the NBER-CES real output from 1958 on. For steel, my output measure is the short tons of raw steel produced (Carter 2006). From 1913 through 1950, I measure motor vehicle production using the NBER Macrohistory Database series on passenger car and truck production. I obtained a figure for 1910 production from Wikipedia.¹⁹ From 1951 on, I use car and truck production figures from the Ward's Automotive Yearbook, prorated to match the NBER series.

For consumption of motor vehicles, I use the Ward's Automotive series on sales of passenger cars and trucks. For cotton and steel, I add net imports to domestic production. For cotton from 1820 through 1950, I use the net dollar imports of cotton manufactures divided by the price of cloth. From 1820 through 1860, I use Sandberg's (1971) estimate of the price of British imports; from 1860 through 1950, I use the price of cotton sheeting (see below). From 1958 on, I use import penetration ratios from Feenstra (1958 through 1994) and Schott (1995 on). For steel, I use Temin's (1964, p. 282) estimates for steel rail imports from 1860 through 1889. I use the Feenstra and Schott import penetration estimates from 1958 on; I ignore steel imports between 1890 and 1957.

For prices, I use the series on cotton sheeting from 1820 through 1974 (Carter 2006, Cc205); for steel I use series for the price of steel rails, splicing together separate series for Bessemer, open hearth, standard, and carbon steel (Carter 2006, CC244-7).

References

- Acemoglu, Daron, and Veronica Guerrieri. "Capital deepening and nonbalanced economic growth." *Journal of political Economy* 116.3 (2008): 467-498.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. "The skill complementarity of broadband internet." *The Quarterly Journal of Economics* (2015), 1781–1824.
- Autor, David H. "Why are there still so many jobs? The history and future of workplace automation." *The Journal of Economic Perspectives* 29.3 (2015): 3-30.
- Autor, David, David Dorn, and Gordon H. Hanson. "Untangling trade and technology: Evidence from local labour markets." *Economic Journal* 125.584 (2015): 621-46.
- Bartel, Ann, Casey Ichniowski, and Kathryn Shaw. "How does information technology affect productivity? Plant-level comparisons of product innovation, process improvement, and worker skills." *The Quarterly Journal of Economics* 122.4 (2007): 1721-1758.
- Baumol, William J. "Macroeconomics of unbalanced growth: the anatomy of urban crisis." *The American economic review* (1967): 415-426.

¹⁹ https://en.wikipedia.org/wiki/U.S._Automobile_Production_Figures.

- Becker, Randy, Wayne Gray, and Jordan Marvakov. NBER-CES Manufacturing Industry Database, 2016, <http://www.nber.org/data/nberces.html>
- Bessen, James E. "How computer automation affects occupations: Technology, jobs, and skills." (2016).
- Boppart, Timo. "Structural Change And The Kaldor Facts In A Growth Model With Relative Price Effects And Non-Gorman Preferences." *Econometrica* (2014): 2167-2196.
- Buera, Francisco J., and Joseph P. Kaboski. "Can traditional theories of structural change fit the data?." *Journal of the European Economic Association* 7.2-3 (2009): 469-477.
- Carter, Susan B. *Historical statistics of the United States: Earliest times to the present.* Cambridge University Press, 2006.
- Clark, Colin, *The Conditions of Economic Progress.* (1840) London: Macmillan.
- Comin, Diego A., Danial Lashkari, and Martí Mestieri. Structural change with long-run income and price effects. No. w21595. National Bureau of Economic Research, 2015.
- Davis, Lance E., and H. Louis Stettler III. "The New England Textile Industry, 1825–60: Trends and Fluctuations." In Dorothy Brady, ed., *Output, employment, and productivity in the United States after 1800.* NBER, 1966. 213-242.
- Dennis, Benjamin N., and Talan B. İşcan. "Engel versus Baumol: Accounting for structural change using two centuries of US data." *Explorations in Economic history* 46.2 (2009): 186-202.
- Dupuit, Jules. "De la mesure de l'utilité des travaux publics" 1844, *Annales des ponts et chaussées*, v.8 (2 sem), p.332-75
- Engel, Ernst (1857). "Die Productions- und Consumtionsverhältnisse des Königreichs Sachsen". *Zeitschrift des statistischen Bureaus des Königlich Sächsischen Ministerium des Inneren.* 8–9: 28–29.
- Feenstra, Robert, "U.S. Imports and Exports by 4-digit SIC Industry, 1958-94," <http://cid.econ.ucdavis.edu/usixd/usixd4sic.html>.
- Foellmi, Reto, and Josef Zweimüller. "Structural change, Engel's consumption cycles and Kaldor's facts of economic growth." *Journal of monetary Economics* 55.7 (2008): 1317-1328.
- Gaggl, Paul, and Greg C. Wright. "A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive." Working Paper (2014).

- Gordon, Robert D. "Values of Mills' ratio of area to bounding ordinate and of the normal probability integral for large values of the argument." *The Annals of Mathematical Statistics* 12.3 (1941): 364-366.
- Jensen, J. Bradford, and Lori G. Kletzer. "Measuring tradable services and the task content of offshorable services jobs." *Labor in the new economy*. University of Chicago Press, 2010. 309-335.
- Kollmeyer, Christopher. "Explaining Deindustrialization: How Affluence, Productivity Growth, and Globalization Diminish Manufacturing Employment 1." *American Journal of Sociology* 114.6 (2009): 1644-1674.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie. "Beyond balanced growth." *The Review of Economic Studies* 68.4 (2001): 869-882.
- Lawrence, Robert Z., and Lawrence Edwards. "US employment deindustrialization: insights from history and the international experience." *Policy Brief* 13-27 (2013).
- Lebergott, Stanley. "Labor Force and Employment 1800-1960." In Dorothy Brady, ed., *Output, Employment and Productivity in the United States After 1800*. National Bureau of Economic Research, *Studies in Income and Wealth* 30 (1966).
- Lewis, W. Arthur. "Economic development with unlimited supplies of labour." *The manchester school* 22.2 (1954): 139-191.
- Matsuyama, Kiminori. "Agricultural productivity, comparative advantage, and economic growth." *Journal of economic theory* 58.2 (1992): 317-334.
- Matsuyama, Kiminori. "Structural change in an interdependent world: A global view of manufacturing decline." *Journal of the European Economic Association* 7.2-3 (2009): 478-486.
- Matsuyama, Kiminori. "The rise of mass consumption societies." *Journal of political Economy* 110.5 (2002): 1035-1070.
- NBER Macrohistory Database, Chapter 1, Production of Commodities.
<http://www.nber.org/databases/macroyhistory/contents/chapter01.html>.
- Ngai, L. Rachel, and Christopher A. Pissarides. "Structural change in a multisector model of growth." *The American Economic Review* 97.1 (2007): 429-443.
- Nickell, Stephen, Stephen Redding, and Joanna Swaffield. "The uneven pace of deindustrialisation in the OECD." *The World Economy* 31.9 (2008): 1154-1184.
- Parker, William N., and Judith LV Klein. "Productivity growth in grain production in the United States, 1840-60 and 1900-10." *Output, employment, and productivity in the United States after 1800*. NBER, 1966. 523-582.

- Rodrik, Dani. "Premature deindustrialization." *Journal of Economic Growth* 21.1 (2016): 1-33.
- Rowthorn, Robert, and Ramana Ramaswamy. "Growth, trade, and deindustrialization." *IMF Staff papers* 46.1 (1999): 18-41.
- Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. *Integrated Public Use Microdata Series: Version 6.0* [Machine-readable database]. Minneapolis: University of Minnesota, 2015.
- Sandberg, Lars. "A Note on British Cotton Cloth Exports to the United States: 1815-1860," *Explorations in Economic History*, 9 (1971): 427-428.
- Schott, Peter K. "U.S. Manufacturing Exports and Imports by SIC or NAICS Category and Partner Country, 1972 to 2005," http://faculty.som.yale.edu/peterschott/sub_international.htm.
- Tarr, David G. "The steel crisis in the United States and the European Community: Causes and adjustments," in Robert E. Baldwin, Carl B. Hamilton and Andre Sapir, editors, *Issues in US-EC trade relations*. University of Chicago Press, 1988. 173-200.
- Temin, Peter. "Iron and Steel in 19th century America." *An Economic Inquiry* (1964). Massachusetts: MIT Press.
- Tirole, Jean. *The theory of industrial organization*. MIT press, 1988.
- Tryon, Rolla Milton. "Household manufactures in the United States, 1640-1860." (1917).
- United States. Bureau of the Census. *Historical Statistics Of The United States, Colonial Times To 1970*. No. 93. US Department of Commerce, Bureau of the Census, 1975.
- United States. Bureau of the Census. *Statistical Abstract of The United States*, various years.
- United States. Department of Labor. *Dictionary of Occupational Titles*, Fourth edition, 1977.
- Ward's Automotive Yearbook, various years, Detroit: WardsAuto.
- Whaples, Robert. "Hours of Work in U.S. History". *EH.Net Encyclopedia*, edited by Robert Whaples. August 14, 2001. URL <http://eh.net/encyclopedia/hours-of-work-in-u-s-history/>
- Zevin, Robert Brooke. "The growth of cotton textile production after 1815." Robert Fogel and Stanley Engerman, *Reinterpretation of American Economic History*, New York: Harper and Row (1971).

Tables

Table 1. Regressions on Log Per Capita Demand

	1	2	3
<u>Cotton</u>			
Ln labor productivity	1.05 (.04)***		1.04 (.04)***
(Ln labor productivity) ²	-0.13 (.01)***		-0.13 (.01)***
Ln <i>p/w</i>		-0.96 (.06)***	
(Ln <i>p/w</i>) ²		-0.19 (.02)***	
Ln real GDP/capita			
Number observations	52	31	52
Adjusted R-squared	0.981	0.954	0.981
Exogeneity P-value			0.096
<u>Steel</u>			
Ln labor productivity	-2.89 (.29)***		-3.19 (.37)***
(Ln labor productivity) ²	-0.65 (.05)***		-0.69 (.06)***
Ln <i>p/w</i>		1.44 (.60)**	
(Ln <i>p/w</i>) ²		-0.25 (.06)***	
Ln real GDP/capita			
Number observations	35	34	35
Adjusted R-squared	0.977	0.968	0.977
Exogeneity P-value			0.134
<u>Automotive</u>			
Ln labor productivity	-3.55 (.78)***		-5.26 (1.22)***
(Ln labor productivity) ²	-0.45 (.07)***		-0.62 (.12)***
Ln real GDP/capita			
Number observations	61		61
Adjusted R-squared	0.904		0.891
Exogeneity P-value			0.073

Note: Robust standard errors in parentheses. ***= significant at 1%; ** = significant at 5%; * = significant at 10%. Constant term not shown. Column 3 is an instrumental variables GMM estimation using year and year-squared to instrument the labor productivity terms. The reported probability value is for the null hypothesis that these variables are exogenous.

Table 2. Regressions of Per Capita Demand

Independent variable	1 Labor productivity	2 Labor productivity	3 Price / wage
A. Cotton cloth, 1820 - 1995			
μ	-1.74 (.10)***	-1.51 (2.35)	-1.49 (.79)*
σ	1.29 (.15)***	1.36 (.69)*	2.04 (.58)**
γ	132.94 (3.2)***	133.09 (3.48)**	184.42 (49.15)**
α		-0.06 (.58)	
Observations	52	52	37
R-squared	0.993	0.993	0.990
B. Raw steel, 1860 – 1982			
μ	3.24 (.14)***	2.31 (.80)***	3.60 (.86)***
σ	0.77 (.18)***	0.56 (.23)**	1.46 (.40)***
γ	0.68 (.05)***	0.69 (.05)***	1.32 (.64)**
α		0.26 (.22)	
Observations	35	35	116
R-squared	0.982	0.982	0.958
C. Motor vehicles, 1910 – 2007			
μ	5.30 (.06)***	7.32 (1.63)***	
σ	0.49 (.10)***	1.61 (1.17)	
γ	64.60 (4.55)***	81.13 (24.79)***	
α		-0.80 (.68)	
Observations	61	61	
R-squared	0.984	0.985	

Note: Non-linear least squares estimates of equation (4) in columns 1 and 2 and equation (5) in column 3. Robust standard errors in parentheses; ***= significant at 1%; ** = significant at 5%; * = significant at 10%.

Table 3. Annual Growth Rate of Industry Hours Worked

	1	2	3	4	5
	OLS	OLS	FE	IV	IV
Sample	All	All	All	Non-mfg.	Mfg.
Computer use x nonmanufacturing	3.00 (1.89)	2.30 (2.87)	-5.23 (5.10)	6.05 (1.71)***	
Computer use x manufacturing	-9.09 (4.37)**	-10.99 (4.70)**	-23.54 (5.89)***		0.49 (3.19)
Manufacturing	2.73 (1.82)	4.09 (1.97)**			
Offshoreability		-2.57 (1.10)**			
College required		7.00 (2.43)***			
No. observations	1299	1299	1299	799	464
R-squared	0.024	0.028	0.138	0.016	0.012
Prob $\hat{\beta}_{nonmanu} = \hat{\beta}_{manu}$	0.010	0.006	0.003	0.124	
Percentage contribution to annual employment growth					
Nonmanufacturing	1.27	0.97	-2.22	2.53	
Manufacturing	-3.48	-4.20	-9.00		0.19

Note: Errors, clustered by industry, in parentheses. ***= significant at 1%; ** = significant at 5%; * = significant at 10%. Data are for 1984, 1989, 1993, 1997, 2001, and 2003 (extending to 2007 for the dependent variable). The dependent variable is the annual growth rate in hours worked from the observation year to the next year in the sample trimmed of 1% tails. Year dummies not shown. The fixed effects regression is over 227 detailed industries. The regressions in columns 4 and 5 are instrumented using year dummies and the share of workers in sedentary jobs based on the Dictionary of Occupational Titles of 1977 (see text) and the 1980 Census. Probability values are shown for F tests that the manufacturing coefficient equals the nonmanufacturing coefficient.