

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: The Economics of Artificial Intelligence: An Agenda

Volume Authors/Editors: Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-61333-8 (cloth); 978-0-226-61347-5 (electronic)

Volume URL: <http://www.nber.org/books/agra-1>

Conference Date: September 13–14, 2017

Publication Date: May 2019

Chapter Title: Artificial Intelligence and Jobs: The Role of Demand

Chapter Author(s): James Bessen

Chapter URL: <http://www.nber.org/chapters/c14029>

Chapter pages in book: (p. 291 – 307)

---

## Artificial Intelligence and Jobs

### The Role of Demand

James Bessen

---

There is widespread concern today that artificial intelligence technologies will create mass unemployment during the next ten or twenty years. 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 2017).




The example of manufacturing decline provides good reason to be concerned about technology and job losses. In 1958, the broadwoven textile industry in the United States employed over 300,000 production workers, and the primary steel industry employed over 500,000. By 2011, broadwoven textiles employed only 16,000, and steel employed only 100,000 production workers.<sup>1</sup> Some of these losses can be attributed to trade, especially since the mid-1990s. However, overall since the 1950s, most of the decline appears to come from technology and changing demand (Rowthorn and Ramaswamy 1999).

But the example of manufacturing also demonstrates that the effect of technology on employment is more complicated than a simple story of “automation causes job losses” in the affected industries. Indeed, figure 10.1 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

James Bessen is Executive Director of the Technology & Policy Research Initiative at Boston University School of Law.

For acknowledgments, sources of research support, and disclosure of the author’s material financial relationships, if any, please see <http://www.nber.org/chapters/c14029.ack>.

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.



**Fig. 10.1 Production employment in three industries**

with job losses. This “inverted-U” pattern appears to be quite general for manufacturing industries (Buera and Kaboski 2009; Rodrik 2016).<sup>2</sup>

The reason automation in textiles, steel, and automotive manufacturing led to strong job growth has to do with the effect of technology on demand, as I explore below. New technologies do not just replace labor with machines, but, in a competitive market, automation will reduce prices. In addition, technology may improve product quality, customization, or speed of delivery. All of these things can increase demand. If demand increases sufficiently, employment will grow even though the labor required per unit of output declines.

Of course, job losses in one industry might be offset by employment growth in other industries. Such macroeconomic effects are covered by other articles in this volume (chapter 13, chapter 9). This chapter explores the effect of technology on employment in the affected industry itself. The rise and fall of employment poses an important puzzle. While a substantial literature has looked at structural change associated with technology, I argue that the most widely accepted explanations for deindustrialization are inconsistent with the observed historical pattern. To explain the inverted-U pattern, I present a very simple model that shows why demand for these products was highly elastic during the early years and why demand became inelastic over time. This model forecasts the rise and fall of employment in these industries with reasonable accuracy: the solid line in figure 10.1 shows those predictions. I then explore the implications of this model for the future impact of artificial intelligence over the next two decades.

## 10.1 Structural Change

The inverted-U pattern in figure 10.1 is also seen in the relative share of employment in the whole manufacturing sector, shown in figure 10.2. Logically, the rise and fall of the sector as a whole in this chart results from the aggregate rise and fall of separate manufacturing industries such as those in figure 10.1. 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.<sup>3</sup> Baumol (1967) showed that the greater rate of technical change

2. Other 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).

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



**Fig. 10.2 Manufacturing share of the labor force**

Sources: US Bureau of the Census 1975; BLS Current Employment Situation.

Note: Labor force includes agricultural laborers.

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 nineteenth 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 agriculture. Parker and Klein (1966) find that labor productivity in corn, oats, and wheat grew 2.4 percent, 2.3 percent, and 2.6 percent per annum from 1840–1860 to 1900–10. In contrast, labor productivity in cotton textiles grew 3 percent per year from 1820 to 1900 and labor productivity in steel grew 3 percent from 1860 to 1900.<sup>4</sup> 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

4. My estimates, data described below.

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 thirtyfold during the nineteenth century, consumption of cotton cloth increased one hundred-fold. The inverted-U thus seems to involve an interaction between productivity growth and demand.

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 twentieth 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 nineteenth 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 Zweimüller 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 10.1.

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 the general problem of looking at the income elasticity of demand as the main driver of structural change: the data suggests that prices were often far more consequential for consumers than income. From 1810 to 2011, real gross domestic product (GDP) per capita rose thirtyfold, but output per hour in cotton textiles rose over eight hundredfold; inflation-adjusted prices correspondingly fell by three orders of magnitude. Similarly, from 1860 to 2011, real GDP per capita rose seventeenfold, 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 chapter 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 or 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.

## 10.2 Model

### 10.2.1 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 be  $q = A \cdot L$ , 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,  $\epsilon$ , assuming income is constant. Taking the partial derivative of the log of equation (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 is quite high. But cloth draperies might be a luxury that the consumer would not be willing to purchase unless the price is 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, 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.<sup>5</sup>

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

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  $v$  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.

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

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 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

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

### 10.2.2 Elasticities

Using equation (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,  $\epsilon$ , 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  $(\partial \epsilon / \partial p) \geq 0 \forall p < \bar{p}$ .*

**PROPOSITION 2.** *Common distributions.* *If the distribution is normal, log-normal, 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 the price starts above  $p^*$  and declines below it.

### 10.2.3 Empirical Estimates

This very simple model does not consider numerous factors that might influence demand. It does not consider the role of close substitutes or the effect of the business cycle on demand. New technology might create new products that generate new demand, altering the distribution, or new substitutes that decrease demand. Global trade might alter downstream industries, affecting the demand for intermediate goods such as cloth or steel. Nevertheless, the model appears to predict actual demand over a historical timeframe reasonably well.

Assuming that the preference distribution is lognormal, I estimate the per capita demand functions for these three commodities (see Bessen 2017 for details). The model fits the data quite closely, realizing  $R$ -squareds of .982 or higher. Using these predictions, I obtain very rough estimates of the price elasticity of demand at each end of the estimation sample (see table 10.1).

The demand was initially highly elastic but became highly inelastic.

Using estimated per capita demand, labor demand can be calculated incorporating population size, import penetration, labor productivity, and hours worked. These estimates 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 falls below the

**Table 10.1** Rough estimates of elasticity of demand

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

estimates when globalization takes a bite out of employment in textiles after 1995, and steel after 1982.

Thus, even though this overly simple model does not account for all of the factors that affect demand, it nevertheless provides a succinct explanation of the inverted-U in employment in these manufacturing industries.

### 10.3 Implication for AI

#### 10.3.1 The Importance of Demand

Although the model presented here appears to provide a good explanation for how demand mediated the impact of technology in the past, what is the relevance of this analysis for new technologies? There is, of course, no guarantee that AI or other new technologies will be applied in markets with preference distributions similar to those of the textile, steel, and automotive industries.

The relevance of this history is more general. Specifically, the responsiveness of demand is key to understanding whether major new technologies will decrease or increase employment in affected industries. Productivity-enhancing technology will increase industry employment if product demand is sufficiently elastic. 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. And demand will likely be sufficiently elastic if the technology is addressing large unmet needs affecting people with diverse preferences and uses for the technology. This situation corresponds to the upper tail of the distribution function. If, on the other hand, AI is targeted at more satiated markets, then jobs will be lost in the affected industries, although not necessarily in the economy as a whole.

The pace of change of a new technology is not sufficient by itself to determine the impact of that technology on employment. For example, a common view holds 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 for AI and this will cause unemployment (Ford 2015). However, my analysis highlights the importance of demand in mediating the impact of automation. If demand is sufficiently elastic and AI does not completely

replace humans, then technical change will create jobs rather than destroy them. In this case, a faster rate of technical change will actually create *faster* employment growth rather than job losses.

The demand response to AI is, of course, an empirical question and, therefore, an important part of the AI research agenda.

### 10.3.2 Research Agenda

To understand the interaction between AI and demand over the next ten or twenty years, empirical researchers will need answers to several specific questions.

First, to what extent will AI replace humans and to what extent will it, instead, merely augment human capabilities? That is, to what extent will AI completely automate occupations and to what extent will it, instead, merely automate some, but not all, tasks performed by an occupation. If humans are completely replaced, demand no longer affects employment because there isn't any demand for humans. In the past, despite extensive productivity growth, technology has almost always only partially automated work. Consider what happened to the 271 detailed occupations used in the 1950 census by 2010. Most occupations listed then still exist in some form (sometimes grouped differently) today. Some occupations were eliminated for a variety of reasons. In many cases, demand for the occupational services declined (e.g., boardinghouse keepers); in some cases, demand declined because of technological obsolescence (e.g., telegraph operators). This, however, is not the same as automation. In only one case—elevator operators—can the decline and disappearance of an occupation be largely attributed to automation. Nevertheless, this sixty-year period witnessed extensive automation; it was just mostly partial automation.

This same pattern is likely to be true for AI over the next ten or twenty years for the simple reason that although AI can outperform humans on some tasks, today's AI fails miserably at other tasks that humans perform. A casual review of current developments suggests that over the near term AI may be able to completely automate some jobs of drivers and warehouse workers, but most AI applications are targeted toward automating just some subset of tasks performed by specific occupations. Nevertheless, a more rigorous empirical investigation is needed to measure the extent to which AI is bringing or will bring complete versus partial automation.

To the extent that automation continues to be partial rather than complete in the near term, demand will be key. This raises a second question: To what extent will the effect of AI on demand and employment during the next ten or twenty years be similar to the effect that AI and computer automation generally had over the last several decades? Computers have been used to automate work in activities such as accounting and loan making since the 1950s. The first fully automatic loan application system was installed in 1972. In 1987, an artificial intelligence system was first put into commercial

operation in a system used to detect credit fraud. Since then, AI applications have been used to automate a variety of tasks in other industries and occupations, such as the electronic discovery of legal documents for litigation.

This means that we already have some evidence of the effects of AI and computer automation generally. It does not seem that computer automation or AI has so far led to significant job losses; the booming market for electronic discovery applications, for instance, has been associated with an increase in the employment of paralegals. A few studies have made estimates of the employment impact of computer technology (Gaggl and Wright 2017; Akerman, Gaarder, and Mogstad 2015), finding, if anything, a modest increase in employment following technology adoption.<sup>6</sup> Further studies could deepen our understanding of the impact of computer automation on employment, and how this impact differs across occupations and industries.

Also, we need to understand how AI applications in the near future will differ from those of the recent past. The model above provides a framework to analyze this question. In particular, to the extent that the new applications target the same services and industries as did the computer automation of the recent past, then we should expect the elasticity of demand to remain similar over the next ten or twenty years, perhaps with a modest decline. That is, the elasticity of demand is not likely to change very quickly. On the other hand, AI might introduce entirely new products and services that tap into otherwise unmet needs and wants. In this case, there may be new and unanticipated sources of employment growth. Research can help determine the extent of change in the sorts of applications, occupations, and industries affected by new AI applications that are also addressed by existing technologies. To the extent that AI creates wholly new applications, prediction will be more difficult. Indeed, in the past, predictions about technological unemployment have reliably failed to anticipate major new applications of technology and major new sources of demand.

A critical aspect of this research concerns the unevenness of the potential impact of AI. While AI might not create overall unemployment in the near future, it will likely eliminate jobs in some occupations while creating new jobs in others. The need to retrain and transition workers to new occupations, sometimes in new locations, might be highly disruptive even though the total employment rate remains high.

Finally, it is important to note that this analytical framework and research agenda are very much limited to the next ten or twenty years for two reasons. First, beyond a couple of decades, markets might well become saturated. Suppose, for example, that demand is highly elastic for many financial, health, and other services today so that information technology increases employment in these markets. If AI rapidly reduces costs or improves the

6. And, importantly, impacts that differed across skill groups.

quality of these services, the elasticity of demand will decline. That is, these markets might see the kind of reversals in employment growth seen in figure 10.1.

Second, in the future AI might very well be able to completely replace many more occupations. Then the effect of AI on demand will no longer matter for these occupations. For now, however, understanding how and where AI affects demand is critical to understanding employment effects.

## 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  $(\partial \epsilon / \partial p) \geq 0$  is

$$(10A.1) \quad \frac{f'}{f} + \frac{f}{1-F} \geq 0.$$

**PROPOSITION 1.** *For a single peaked distribution with mode  $\bar{p}$ , for  $p < \bar{p}$ ,  $f' \geq 0$  so that  $(\partial \epsilon / \partial p) \geq 0$ .*

**PROPOSITION 2.** *For each distribution, I will show that*

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

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

### Normal Distribution

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

where  $\varphi$  and  $\Phi$  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  $x > 0$ ,<sup>7</sup>

$$(10A.2) \quad x \leq \frac{\varphi(x)}{1-\Phi(x)}.$$

Applying this inequality, it is straightforward to show that (10A.1) holds for the normal distribution. This also implies that  $\lim_{p \rightarrow \infty} \epsilon = \infty$ . By inspection,  $\epsilon(0) = 0$ .

**Exponential Distribution**

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

Then,

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

so (10A.1) holds. By inspection,  $\epsilon(0) = 0$  and  $\lim_{p \rightarrow \infty} \epsilon = \infty$ .

**Uniform Distribution**

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

so that

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

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

**Lognormal Distribution**

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

so that

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

Canceling 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,  $\lim_{p \rightarrow 0} \epsilon = 0$ .

7. I present the inverse of Gordon's inequality.



## References

- Acemoglu, Daron, and Veronica Guerrieri. 2008. "Capital Deepening and Nonbalanced Economic Growth." *Journal of Political Economy* 116 (3): 467–98.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. 2015. "The Skill Complementarity of Broadband Internet." *Quarterly Journal of Economics* 130 (4): 1781–824.
- Baumol, William J. 1967. "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review* 57 (3): 415–26.
- Bessen, James E. 2016. "How Computer Automation Affects Occupations: Technology, Jobs, and Skills." Law and Economics Research Paper no. 15-49, Boston University School of Law.
- . 2017. "Automation and Jobs: When Technology Boosts Employment." Law and Economics Research Paper no. 17-09, Boston University School of Law.
- Boppart, Timo. 2014. "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences." *Econometrica* 82 (6): 2167–96.
- Buera, Francisco J., and Joseph P. Kaboski. 2009. "Can Traditional Theories of Structural Change Fit the Data?" *Journal of the European Economic Association* 7 (2–3): 469–77.
- Clark, Colin. 1940. *The Conditions of Economic Progress*. London: Macmillan.
- Comin, Diego A., Danial Lashkari, and Martí Mestieri. 2015. "Structural Change with Long-Run Income and Price Effects." NBER Working Paper no. 21595, Cambridge, MA.
- Dennis, Benjamin N., and Talan B. İşcan. 2009. "Engel versus Baumol: Accounting for Structural Change Using Two Centuries of US Data." *Explorations in Economic History* 46 (2): 186–202.
- Dupuit, Jules. 1844. "De la Mesure de l'utilité des Travaux Publics." *Annales des Ponts et Chaussées* 8 (2 sem): 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.
- Foellmi, Reto, and Josef Zweimüller. 2008. "Structural Change, Engel's Consumption Cycles and Kaldor's Facts of Economic Growth." *Journal of Monetary Economics* 55 (7): 1317–28.
- Ford, Martin. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic Books.
- Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114 (2017): 254–80.
- Gaggl, Paul, and Greg C. Wright. 2017. "A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive." *American Economic Journal: Applied Economics* 9 (3): 262–94.
- Gordon, Robert D. 1941. "Values of Mills' Ratio of Area to Bounding Ordinate and of the Normal Probability Integral for Large Values of the Argument." *Annals of Mathematical Statistics* 12 (3): 364–66.
- Kollmeyer, Christopher. 2009. "Explaining Deindustrialization: How Affluence, Productivity Growth, and Globalization Diminish Manufacturing Employment 1." *American Journal of Sociology* 114 (6): 1644–74.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie. 2001. "Beyond Balanced Growth." *Review of Economic Studies* 68 (4): 869–82.

- Lawrence, Robert Z., and Lawrence Edwards. 2013. "US Employment Deindustrialization: Insights from History and the International Experience." Policy Brief no. 13–27, Peterson Institute for International Economics.
- Lewis, W. Arthur. 1954. "Economic Development with Unlimited Supplies of Labour." *Manchester School* 22 (2): 139–91.
- Matsuyama, Kiminori. 1992. "Agricultural Productivity, Comparative Advantage, and Economic Growth." *Journal of Economic Theory* 58 (2): 317–34.
- . 2009. "Structural Change in an Interdependent World: A Global View of Manufacturing Decline." *Journal of the European Economic Association* 7 (2–3): 478–86.
- . 2002. "The Rise of Mass Consumption Societies." *Journal of Political Economy* 110 (5): 1035–70.
- Ngai, L. Rachel, and Christopher A. Pissarides. 2007. "Structural Change in a Multi-sector Model of Growth." *American Economic Review* 97 (1): 429–43.
- Nickell, Stephen, Stephen Redding, and Joanna Swaffield. 2008. "The Uneven Pace of Deindustrialisation in the OECD." *World Economy* 31 (9): 1154–84.
- Parker, William N., and Judith L. V. Klein. 1966. "Productivity Growth in Grain Production in the United States, 1840–60 and 1900–10." In *Output, Employment, and Productivity in the United States after 1800*, edited by Dorothy Brady, 523–82. Cambridge, MA: National Bureau of Economic Research.
- Rodrik, Dani. 2016. "Premature Deindustrialization." *Journal of Economic Growth* 21 (1): 1–33.
- Rowthorn, Robert, and Ramana Ramaswamy. 1999. "Growth, Trade, and Deindustrialization." *IMF Staff Papers* 46 (1): 18–41.
- Tirole, Jean. 1988. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press.
- United States Bureau of the Census. 1975. Historical Statistics of the United States, Colonial Times to 1970, no. 93. US Department of Commerce, Bureau of the Census.