

**ECONOMIC
POLICY**



69th Economic Policy Panel Meeting

**4-5 April 2019
Tallinn**

Hosted by Eesti Pank

Automation and Jobs: When Technology Boosts Employment

James Bessen (Boston University School of Law)

The organisers would like to thank Eesti Pank for their support.
The views expressed in this paper are those of the author(s) and not those of the supporting organization.

Automation and Jobs: When Technology Boosts Employment

By James Bessen,
Technology & Policy Research Initiative, Boston University School of Law

Abstract: Will new technologies cause industries to shed jobs, requiring novel policies to address mass unemployment? 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? The elasticity of demand. Using data over two centuries for US textile, steel, and auto industries, this paper shows that automation initially spurred job growth because demand was highly elastic. But demand later became satiated, leading to job losses. A simple model explains why this pattern might be common, suggesting that today's technologies may cause some industries to decline and others to grow. Automation might not cause mass unemployment, but it may well require workers to make disruptive transitions to new industries, requiring new skills and occupations.

JEL codes: J2, O3, N10

Keywords: Automation, technical change, sectoral growth, labour demand, manufacturing, deindustrialization, industry life cycle

Thanks for helpful comments from Philippe Aghion, David Autor, Claude Diebolt, Martin Fleming, Richard Freeman, Ricardo Hausman, Mike Meurer, Joel Mokyr, Bob Rowthorn, Anna Salomons and participants at the ASSA and the SBBI and CID seminars at Harvard.

1. Introduction

There is widespread concern today that many jobs will be lost to new computer technologies as more human tasks can be performed by machines. A host of recent papers estimate that these technologies put anywhere from 9% of jobs to 47% at risk of automation in the near future (see, for example, Arntz et al. 2016, Frey and Osborne 2017). Some people argue (Ford 2015) that this expansion in the range of automatable tasks calls for new policies to address imminent mass unemployment.

But that inference is not warranted. Automation does not necessarily lead to a loss of jobs even in the affected industry. When major industries automate, their employment often rises rather than falls (see Figure 1). This might seem surprising because the recent experience in many manufacturing industries such as textiles and steel seems to suggest that automation leads to job losses in those industries. Some economic theory lends credence to that intuition. Baumol (1967) argued that faster productivity growth in manufacturing relative to other sectors has led to a declining share of employment in manufacturing; the same might pertain to fast productivity growth from information technologies today. However, some of these same manufacturing industries actually grew robustly at the same time as labour productivity for a century or more before declining. Figure 1 shows this “inverted U” pattern for the US cotton textile, primary steel, and automotive industries.

This pattern is important for two reasons. First, it shows that one cannot assume that productivity-improving technology necessarily leads to job losses. Second, while productive technology might or might not decrease aggregate employment, this pattern suggests it will likely have *disparate effects* on different industries at different times. Some industries will grow while others decline and this raises a distinct policy challenge: how to

support workers making transitions to new industries, new occupations with new skills, sometimes in new regions.

This paper explores why productivity-improving technology brought employment growth at some times but not at others, under what conditions such behaviour is likely to occur with today's technologies, and what this means for policy. The contributions of this paper are as follows. First, using long data series, I show that demand was initially highly elastic in the US cotton textile, primary steel, and automotive industries, but that it became inelastic as consumption grew. Second, I provide a novel model of demand satiation that can explain *why* demand became sharply less elastic. This model is able to closely account for the growth and subsequent decline in employment in the industries studied. Third, I show some general conditions that give rise to declining demand elasticities. The generality of these conditions and the historical experience suggest that industries today are likely to respond heterogeneously to new computer-based productivity-improving technologies; employment will rise in some and fall in others. Some evidence suggests that this is what is happening.

Finally, I draw implications for policy. Whether productivity-improving technology is increasing employment in some industries today can, and must, be determined empirically, of course. But this is a key policy implication: policymakers cannot assume that productivity-improving technology leads to unemployment in all the affected industries. If, instead, technology is increasing employment in some industries while decreasing it in others, then the policy challenge posed by new computer technologies may be less about ameliorating the impacts of mass unemployment and more about helping workers transition from some industries, some occupations, and some regions to others. The impacts of this sort of change may be no less disruptive, but the nature of policy responses would be

critically different. Some evidence suggests that automation, robotics, and information technologies are not associated with declining employment in many industries.

The Puzzle of the Inverted U


This paper argues that the key to understanding the changing employment response to labour productivity is the nature of demand. Empirical evidence shows that the non-constant elasticity of demand is a critical factor behind the patterns seen in Figure 1 and the paper develops a novel life-cycle model to show why the nature of demand changed in a common way across these industries. Common popular intuitions about automation ignore the role of demand and much economic theory about automation abstracts away from the dynamic role of demand.

A simple model illustrates the links between demand, labour productivity, and employment. Let demand for an industry's product be a function of price, p , and the wage, w , such that per capita demand $D = D(\frac{p}{w}, w)$. At the market-clearing equilibrium, domestic output, Y , equals consumption minus imports, so that $Y = (1 - I) \cdot N \cdot D(\frac{p}{w}, w)$, where I is the net import share and N is population (I assume for now that all consumers have the same demand function). Define labour productivity as industry output divided by industry employment, $A \equiv \frac{Y}{L}$, so that $L = \frac{Y}{A}$. For expositional simplicity, temporarily assume that the import share is independent of A . Then the response of log industry employment to log productivity is

(1)

$$\frac{\partial \ln L}{\partial \ln A} = \frac{\partial \ln D}{\partial \ln A} - 1.$$

Figure 1. Production Employment in Three Industries



Note: The solid line represents predicted employment based on the model developed below.

If demand increases sufficiently in response to productivity-improving technology, then employment will grow; otherwise, it will fall. The elasticity of demand with respect to labour productivity determines whether productivity-improving technology increases or decreases employment. A prime reason for demand to grow in response to new technology is that productivity growth will reduce prices in competitive markets. Below, I show how this elasticity is related to the price elasticity.

But what was the actual relationship between demand and labour productivity? Did the elasticity of demand change over time in a way that could explain the rise and fall in employment? Figure 2 shows labour productivity growth over time; Figure 3 shows per capita demand (final consumption) plotted against labour productivity, both on logarithmic scales. Note that for textiles and steel some recent observations, shown as empty circles, fail to control for import competition in downstream markets and, thus, no longer capture demand properly.¹ The data behind these figures are described below.

While each industry began growing rapidly at different times, each shows a strong, sustained, and relatively rapid growth of physical output per worker hour. Employment was indeed growing as new technology generated greater output per unit of labour. Figure 3 shows that consumption increased along with labour productivity, and all three industries exhibit a common concave pattern: demand increased rapidly along with productivity during the early years, but this relationship flattened out in later years. Demand satiated. Since the

¹ The consumption measure includes net imports; however, I am not able to correct for trade impacts in downstream industries to textiles and steel. 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 domestic apparel producers, and 3) apparel 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.

axes are plotted on log scales, the slope of the curve roughly represents the elasticity of demand with respect to productivity. This is not quite accurate because demand also depends on income levels, which also grew over this period. Below I estimate models that include both, but this figure conveys the essential intuition: demand was initially very elastic, allowing employment to grow with productivity; later, demand became inelastic and employment declined.

Demand satiation thus provides an intuitive explanation for the inverted U in employment. Below, I develop a model that can explain why this pattern of declining demand elasticity occurs and draws out policy implications.

Literature

The significance of this analysis relative to the literature is that it identifies the central role of demand in mediating the impact of productivity-improving technology on jobs. Both historical evidence and my theoretical model imply that, for significant industries, demand and productivity growth interact in a systematic way: at low, initial levels of productivity, demand elasticities are high; as productivity growth gradually satiates the market, demand elasticities fall sharply. Because industries today may have very different demand elasticities from one another, the impact of major new technologies on employment is likely to be highly heterogeneous. The impact will not be uniformly job-destroying, rather some industries will shrink, others will grow. Hence workers will need to transition to new industries, often requiring new skills, new organizations, and new locations. Facilitating these transitions and reducing their social cost thus becomes an important policy challenge.

Figure 2. Labour Productivity over Time


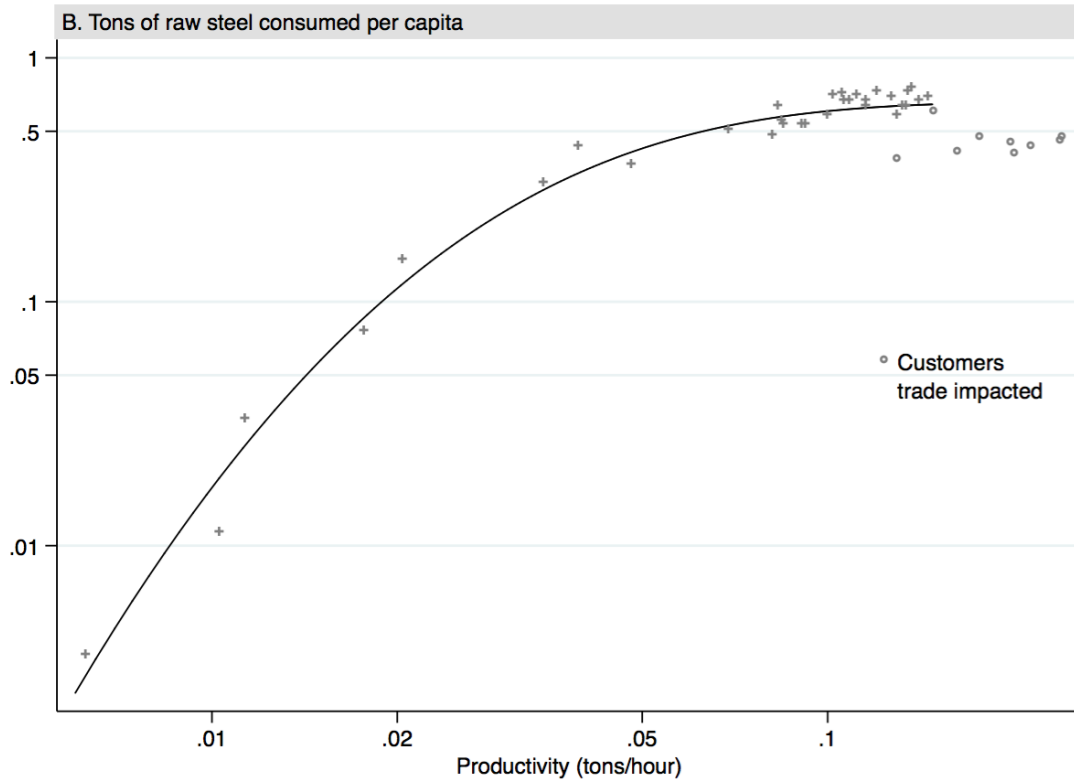
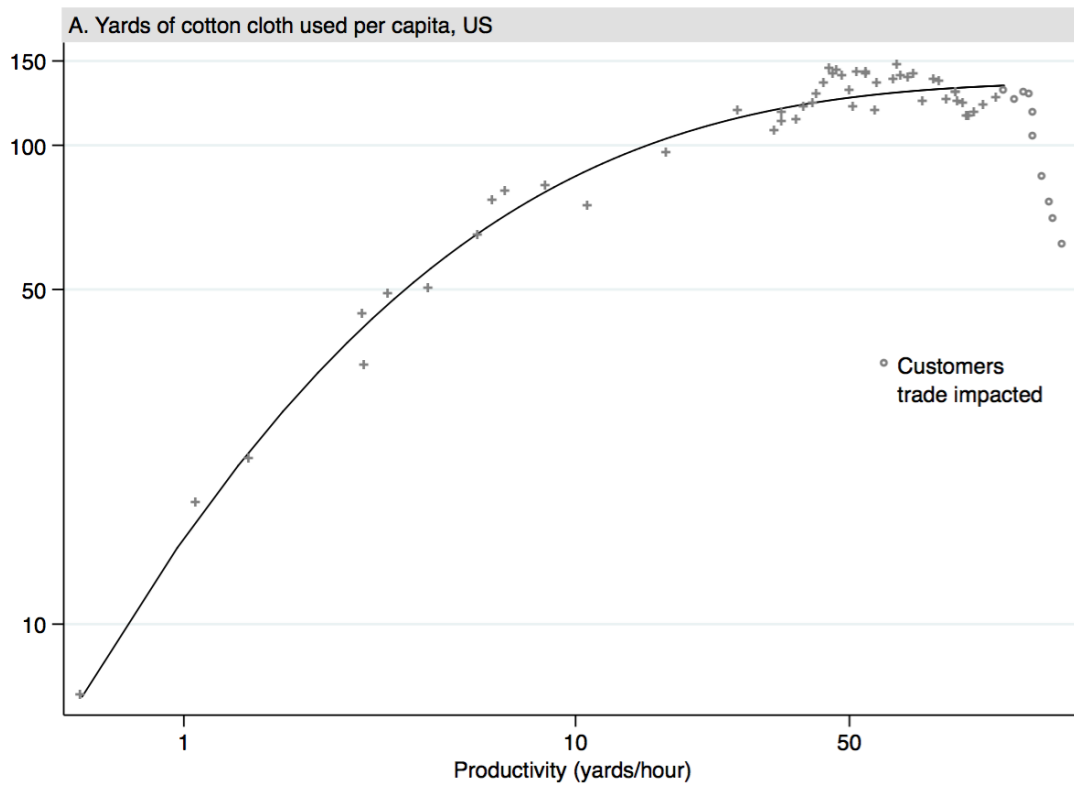
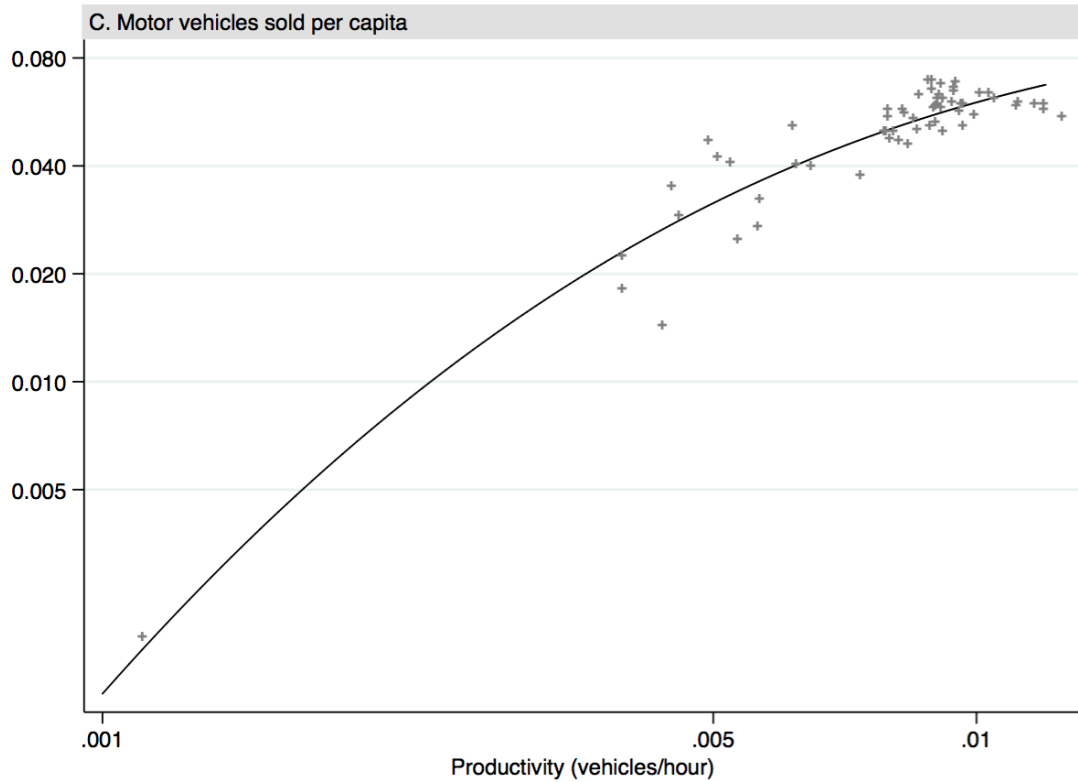


Figure 3. Per Capita Consumption





Note: The solid lines represent prediction from the model developed below.

In contrast to this analysis, recent work on automation abstracts away from demand heterogeneity. A recent theoretical literature, cognizant of new technologies, has focused on the impact of automation, defined as technical change that allows machines to replace humans performing specific production tasks (Acemoglu and Restrepo 2017, 2018a, 2018b, Hemous and Olsen 2016, Aghion et al. 2017). But these papers assume that the price elasticity of demand—or, equivalently, the elasticity of substitution between industries or between tasks—is constant and uniform across industries or tasks.²

² These papers use Dixit-Stiglitz aggregators to construct aggregate output from the output of individual industries or tasks. In this setting, the elasticity of substitution corresponds to a price elasticity of demand.

In some of these models, automation can increase the demand for labour, but that outcome depends on the relative productivity of the new technology. For example, Acemoglu and Restrepo (2017, 2018a) argue that “in contrast to some popular discussions, the new AI and robotics technologies that are more likely to reduce the demand for labour are not those that are brilliant and highly productive, but those that are ‘so-so.’” But the historical evidence shows that labour demand rises or falls with new technology even when productivity growth is roughly the *same*. The task-based models of automation miss the critically important role that demand elasticity can play in determining whether automation increases or decreases labour demand. “Brilliant” technologies will lead to declining employment if demand is inelastic and “so-so” technologies can lead to job growth if demand is highly elastic.

Moreover, by assuming constant and uniform demand elasticities, these models miss the dynamic interactions that create substantial heterogeneity between industries at different points in their life-cycles. This is also true of much of the larger literature on structural change and deindustrialization. For example, Baumol (1967) explains the decline of manufacturing’s share of employment as the result of higher productivity growth in manufacturing compared to services.³ But his model critically assumes inelastic demand for manufactured goods. Similarly, another literature following “Engel’s Law” attributes the decline in manufacturing employment to a low (and constant) income elasticity for manufactured goods; these goods are seen as “necessities” with income elasticities less than one.⁴ Yet these same goods apparently were luxuries with higher income elasticities in the

³ See also Lawrence and Edwards 2013, Ngai and Pissarides 2007, Matsuyama 2009.

⁴ 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,

past. This paper demonstrates substantial declines in the elasticity of demand, and it introduces a novel theory that explains the decline as a function of falling prices relative to wages. Because productivity growth reduces prices relative to wages, this model provides a theory of industry life cycles that applies under common conditions. While previous models provide a rationale for falling income elasticities, none has explored the link between productivity growth and dynamic demand elasticity.⁵

This paper also looks somewhat more broadly at technology than the recent papers on automation. The particular concern about automation is that, of all types of technological change, human-replacing automation most directly threatens to reduce employment. But any productivity-improving technology can reduce employment as in equation (1). This paper considers all technologies that increase output per worker, regardless of whether this increase is specifically achieved by task replacement or not. This category includes Hicks-neutral technical change, biased or labour-augmenting technical change, as well as automation. To be sure, much of the productivity-improving technology affecting the three industries studied here is automation per se. Bessen (2012) identifies the sources of output growth per worker in textile weaving and finds that most of the increase comes from task automation. In the framework used here, technology affects employment by increasing output per worker independently of how the capital-labour ratio changes.

Another concern about automation is that it might decrease labour's share of output, all else equal. In this paper, labour's share is taken as an empirical quantity rather than

Lashkari, and Mestieri 2015, Kongsamut, Rebelo, and Xie 2001 and Matsuyama 1992 for more general treatments of nonhomothetic preferences.

⁵ Matsuyama (2002) introduced a model where the income elasticity of demand for goods falls as incomes grow. See also Foellmi and Zweimueller (2008). Also, Banks et al. (1997) find that some commodities show a fall in income elasticity when examined in cross-sectional data.

assumed to either be fixed or to change in any prescribed way. In fact, the industry data on labour's share do not exhibit a secular trend prior to 1950 despite high levels of automation. And evidence suggests that technology does not play a major role in the recent decline in labour's share.⁶ In any case, the change in labour's share, while important, is separate from the question of the impact of technology on employment.

Another difference with the literature concerns the level of aggregation. The recent literature on automation mainly considers single-sector economies or economies where broad sectors are distinguished by the extent to which tasks are prone to automation. But industry-level analysis is more appropriate for this study because a high level of aggregation misses inter-industry transitions. In Figure 1, one can see that textile industry employment declined for many decades while steel and auto employment grew. The timing of growth in these industries varied because of historical industry-specific differences in technology.⁷ On net, workers made transitions between industries, but those transitions did not necessarily appear in the employment pattern for the manufacturing sector as a whole. To the extent that productivity-improving technologies have different employment impacts across industries at different times, broad sectoral analysis misses an important dimension of the disruption they may cause.

The paper is organized as follows. The next section briefly describes the data. Section 3 uses non-parametric methods to test whether the productivity elasticity of demand

⁶ Autor et al. (2017) find that most of the recent decline arises from changes between firms not changes within firms; firms with lower labour share have been growing faster.

⁷ 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. While each of these industries benefited from general purpose technologies such as steam and electric engines and machine tools, the periods of rapid growth began with industry-specific innovations.

declined over time from a level above one. Section 4 presents a simple model of demand satiation to explain why these elasticities declined and shows that this model applies under some rather general conditions. Section 5 shows that the model predicts the actual rise and fall of employment in the subject industries with reasonable accuracy. Section 6 discusses policy implications and Section 7 concludes.

2. 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. Prior to 1958, most of the data series are not available annually but typically occur at decadal intervals.

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 fibres from 1930 on. From 1958, I use the deflated output of the cotton and synthetic fibre 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.⁸ 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 US Bureau of Labour Statistics (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) do 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.

Because price data are limited, I also obtain data on labour’s share of output, the wage bill divided by the value of product shipped. Prices relative to wages can then be estimated from the labour share and labour productivity series, $s \cdot A = \frac{w}{p}$.⁹

⁸ 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.

⁹ $s \equiv \frac{wL}{pY} = \frac{w}{pA}$ where A is labour productivity, Y is output, L is labour, s is labour share of output, w is the wage, and p is the product price. It then follows that $s \cdot A = \frac{w}{p}$.

Labour productivity

I calculate labour 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 labour productivity growth. From 1820 to 1995, labour productivity in cotton textiles grew 2.9% per year; in steel, it grew 2.4% per year from 1860 to 1982; in motor vehicles, it grew 1.4% per year from 1910 through 2007. Figure 2 shows labour 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.

3. Non-parametric estimates of demand

Price elasticity and productivity elasticity

This section seeks to perform statistical tests on demand, $D(\frac{p}{w}, w)$, using a non-parametric expansion in $\frac{p}{w}$ and w . The statistical tests relate to the productivity elasticity, $\frac{\partial \ln D}{\partial \ln A}$, specifically testing the null hypotheses that 1) this elasticity is not constant or increasing, 2) that it was not greater than 1 during the early years of each industry, and 3) that it was not less than 1 during the later years. Rejections of these hypotheses confirm the key notions that demand elasticity is concave and that it declined from initial high levels to inelastic levels later.

In order to perform the tests, it is necessary to relate $\frac{\partial \ln D}{\partial \ln A}$ to $\frac{p}{w}$ and w . A useful relationship derives from the definition of the labour share of output:

(2)

$$s \equiv \frac{wL}{pY} = \frac{w}{pA} \text{ so that } \frac{w}{p} = As.$$

Assuming that $\frac{\partial w}{\partial D} = 0$,¹⁰ the elasticity of demand with respect to labour productivity can be written

(3)

$$\frac{\partial \ln D}{\partial \ln A} = \frac{\partial \ln D}{\partial \ln As} \frac{\partial \ln As}{\partial \ln A} = - \frac{\partial \ln D}{\partial \ln p/w} \frac{\partial \ln As}{\partial \ln A} = \epsilon \cdot \left(1 + \frac{\partial \ln s}{\partial \ln A} \right)$$

where ϵ is the price elasticity of demand while the partial derivative represents the influence of productivity on labour's share of output. In Acemoglu and Restrepo's model (2017, 2018a,b), automation reduces labour's share of output while increasing productivity. If so, then $\frac{\partial \ln s}{\partial \ln A}$ would be negative. I provide estimates of this term below, and it is not always negative nor is it large. In any case, (3) means that a large decline in the price elasticity of demand corresponds more or less to a large decline the elasticity of demand with respect to productivity. The productivity elasticity is also closely related to the income elasticity of demand.¹¹

¹⁰ This is because industry employment is a small part of total employment so that industry demand has a negligible effect on overall labour demand. See the model below.

¹¹ The income elasticity of demand is $\epsilon + \frac{\partial \ln D}{\partial \ln w}$. Hence a large decline in the price elasticity of demand might well be associated with a large decline in the income elasticity.

Testing demand

To estimate demand elasticities, I assume that the log of per capita demand can be captured as the sum of two quadratic forms, one in the price relative to the wage and the other in the wage:

(4)

$$\ln D\left(\frac{p}{w}, w\right) = \alpha + \beta_1 \ln \frac{w}{p} + \beta_2 \left(\ln \frac{w}{p}\right)^2 + \gamma_1 \ln w + \gamma_2 (\ln w)^2 + \varepsilon$$

This equation can be estimated separately for each industry over the observed years, using $A \cdot s$ in lieu of $\frac{w}{p}$ and measuring w as real GDP per capita. The top panel of Table 1 reports the regression coefficients. This non-parametric specification fits the data quite well with R-squareds of .979, .974, and .934 for textiles, steel, and autos respectively.

To check the robustness of these regressions, I repeated them using a measure of the production wage deflated by the Consumer Price Index (CPI). The results were broadly similar. Another concern was possible serial correlation. Annual data are not available until 1958, so serial correlation is not a realistic concern over most of the data. Nevertheless, I reran the estimates with AR(1) and AR(2) disturbances; the results were broadly similar. Another concern is possible co-integration of the dependent and independent variables giving rise to spurious correlation. Augmented Dickey-Fuller tests could not reject the null hypothesis of no co-integration (probability values of .989, .922, and .322 respectively). Demand might also be influenced by long term changes in tastes that could be correlated with prices and wages. Including trend terms in the regressions, I obtained very similar results, suggesting that the basic relationships in (4) are largely independent of secular trends.

Note that the estimate of $\hat{\beta}_2$ (the coefficient of $(\ln sA)^2$) is negative and highly significant in all three industries. This is sufficient to reject the null hypothesis that the demand curve is non-concave with respect to productivity.¹²

Note also that the terms in $A \cdot s$ are generally more significant statistically than the terms in w . They are also much more significant economically because the growth in labour productivity was much greater than the growth in wages. For example, from 1810 to 2011, real GDP per capita rose 30-fold, but output per hour in cotton textiles rose over 800-fold. Similarly, from 1860 to 2011, real GDP per capita rose 17-fold, but output per hour in steel production rose over 100 times. This comparison suggests that the pure income effects are relatively small compared to the productivity effects. This is important because productivity growth can vary substantially between industries, creating different demand patterns.

Assuming that demand can be written as a function of prices and wages in (4), the estimates can be used to calculate demand elasticities. To calculate the right-hand side of (3), we need estimates of ϵ and $\frac{\partial \ln s}{\partial \ln A}$. The former can be calculated from the regression coefficients:

$$\hat{\epsilon} = \frac{\partial \ln D}{\partial \ln w/p} = \hat{\beta}_1 + 2\hat{\beta}_2 \ln As.$$

I estimate the latter by regressing

$$\ln s = \delta \ln A + c + \epsilon.$$

¹² $\frac{\partial^2 \ln D}{\partial \ln A^2} = 2\hat{\beta}_2 \left(1 + \frac{\partial \ln s}{\partial \ln A}\right)$. As long as the partial derivative is not unreasonably large, the sign of this term depends on $\hat{\beta}_2$.

These estimates are found in panel B of Table 1. The estimates of δ for textiles and autos are small and positive; the estimate for autos is modest and negative. Interpreting the change in labour's share over time as a linear function of labour productivity requires the heroic assumption that nothing else affects labour's share. That is almost certainly not true. For instance, the decline in the auto industry, which occurred entirely after 1980, might well be related to the dramatic decline in union membership since then. In any case, the elasticities estimated using (3) are not too sensitive to these estimates of δ . Using the alternative assumption that labour's share is fixed does not change the pattern of results.

Panel C in Table 1 shows estimates of the elasticities at different years and the probability values of a series of F tests on these various quantities. The first row shows estimates of the elasticities for early years in each industry, 1820, 1870, and 1910 respectively for textiles, steel, and auto. All of the estimates are significantly greater than one. The next two rows repeat the exercise for 1950 (1951 for autos). All of the elasticities declined from their earlier values and all are less than one, both economically and statistically.

Estimates of elasticities based on (4) might suffer because the quadratic specification might not fit the true relationship well in all ranges. Note, in particular, that the estimate of the price elasticity of steel is negative in 1950. The bottom panel of Table 1 provides estimates of price elasticities using a specification based on the model developed below for comparison. This specification, from Table 2, Column 1, fits the data better than does equation (4), and it produces more reasonable, although generally similar, estimates of the productivity elasticities of demand.

In summary, non-parametric estimates of demand functions reject the view that the elasticities of demand with respect to price and income for three major manufactured goods have been less than one and constant. These industries began with elastic demand that

became inelastic over time. By (1) this implies that productivity growth increased employment at first and then decreased it.

4. A Model of Demand and Technical Change

Simple model of the Inverted U

These findings reject some common assumptions about how technology leads to employment loss and about the nature of demand elasticity. But it also raises a question: why was it that demand in three major industries followed a similar pattern? This section develops a model to explain why demand elasticities began high and then declined, how that change generated the inverted U in employment, and how general this phenomenon might be with an eye to understanding its relevance to current conditions.

Consider production and consumption of two goods, cloth, quantity y , and a general composite good, quantity x , 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. The model aims to sketch out how industry-specific productivity growth and general income growth can affect demand, including conditions where these trends give rise to an inverted U in employment.

Consumption

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 . Now suppose our consumer will purchase cloth for all uses where the value received exceeds a threshold, $v > \bar{v}$. This threshold will be determined by utility maximization subject to a budget constraint. Demand is then a function of this threshold:

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

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 $\bar{v} > 0$ and that $f(\bar{v}) > 0$ in this domain.¹³

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

$$U(\bar{v}) = \int_{\bar{v}}^{\infty} z \cdot f(z) dz.$$

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). In that setting, $\bar{v} = p$, the price. After integrating by parts and assuming $D(\infty) = 0$, we obtain

¹³ The last condition is necessary to ensure that the indifference curve between cloth and the general good, x , is convex.

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

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.¹⁴

The consumer also derives utility from consumption of the general good, x . Assume that the utility from this good is additively separable from the utility of cloth so that total utility is

$$U(\bar{v}) + G(x)$$

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

(5)

$$w \geq x + pD(\bar{v})$$

where the price of the general good is taken as numeraire and w is the consumer's wage (all consumers are workers). The consumer's Lagrangean can be written

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

Taking the first order conditions,

$$G_x = \lambda, \quad \hat{v}(p, w) = p \lambda = p \cdot G_x(\hat{x}(p, w))$$

where the subscript designates a derivative and the "hats" indicate optimal solutions.

¹⁴ 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 w 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.

Production

Let there be three sectors, one producing cloth, one producing good x , and one producing an investment good, quantity I . Each sector is composed of many firms in competitive markets. The aggregate output of cloth, $Y = Y(L, K; t)$, where L is textile labour, K is capital, t captures the state of technology, and $Y(\cdot)$ is a constant returns production function that is continuous and differentiable. Capital and the general good, x , are produced using simpler production functions

$$X = a \cdot L_x, \quad I = a \cdot L_I, \quad N = L + L_x + L_I$$

where X is the aggregate output of x , N is population (or workforce), L_x and L_I are the workforce size in the x and I production sectors, and a is a measure of general productivity that increases over time. Taking the price of good x and the investment good as numeraire, aggregate profits of each sector are

(6)

$$\pi_y = p \cdot Y - w \cdot L - r \cdot K, \quad \pi_x = X - w \cdot L_x, \quad \pi_I = I - w \cdot L_I.$$

Firms in each sector employ a fraction of the total labour and capital and earn the same fraction of profits. The first order profit maximizing conditions imply

(7)

$$w = a = p \cdot \frac{\partial Y}{\partial L}.$$

Assuming that competitive markets generate zero profits in each sector and equating aggregate income ($wN + rK$) with aggregate output ($pY + X + I$), it is straightforward to show that per capita consumption expenditures equal w , as in the individual budget constraint (5).

Finally, since I am concerned here just with the determinants of the demand for cloth, I do not specify the savings function and the dynamic growth path of capital. Equations (6) and (7) will hold at each point in time. Also, I assume that textile consumption (or steel or autos) is very small compared to total consumption, $pY \ll X$. The consumption of cotton textiles, steel, and motor vehicles never exceeded a few percent of national income during the entire period studied. Note that this implies that $\frac{\partial \hat{x}}{\partial p} \approx 0$ and $X \approx wN$ so that each individual's consumption of x is $\hat{x} \approx w$. Taking (2), rearranging, substituting from (6) and equating total demand and output of cloth ($Y = N \cdot D$),

(8)

$$p = \frac{a}{sA}, \quad \hat{L} = \frac{N \cdot D(\hat{v}(p, w))}{A}.$$

Price elasticity of demand and the Inverted U

Equation (1) shows that $\frac{\partial \ln D}{\partial \ln A}$ determines the relationship between employment growth and productivity growth. Specifically, if this elasticity is greater than 1 at high prices (relative to wages) and lower than 1 at low prices, then employment will trace an inverted U as prices decline with productivity growth. Equation (3) says that this elasticity is determined by the price elasticity of demand, ϵ , and the elasticity of labour's share of output with respect to productivity.

I will now show that this pattern can occur under some fairly general conditions.

Begin with the price elasticity of demand. Assuming, as above, that $\frac{\partial \hat{x}}{\partial p} = 0$, the price elasticity of demand is

(9)

$$\epsilon(p, w) = -\frac{\partial \ln D}{\partial \ln p} = -\frac{\partial \ln D(\hat{v})}{\partial \hat{v}} \frac{p}{\partial p} \frac{\partial \hat{v}}{\partial p} = \frac{pf(\hat{v}(p, w))}{1 - F(\hat{v}(p, w))} G_x(\hat{x}(p, w)).$$

Holding the wage constant, the price elasticity will change with the price depending on the nature of the preference function, $F(\cdot)$. The following proposition holds that for common distribution functions, the price elasticity of demand will be greater than 1 at sufficiently high prices and less than 1 at sufficiently low prices (see Appendix for proofs):

Proposition. Holding the wage constant and assuming $\frac{\partial \hat{x}}{\partial p} = 0$,

1. Single-peaked density functions. If the distribution density function, f , has a single peak at $p = \bar{p}$, then $\frac{\partial \epsilon}{\partial p} \geq 0 \quad \forall p < \bar{p}$.
2. Common distributions. If the preference distribution is normal, lognormal, exponential, or uniform, there exists a p^* such that for $0 < p < p^*$, $\epsilon < 1$, and for $p^* < p$, $\epsilon > 1$.

The labour share is the remaining consideration. Clearly, if the labour share of output is constant, then the proposition implies that $\frac{\partial \ln D}{\partial \ln A}$ will be greater than one for a sufficiently large initial price and will be less than one for a sufficiently small price; the inverted U property will hold. The property will also hold if labour's share of output does not change drastically, as has been true historically. The inverted U property will only fail to hold in the unusual case where $\frac{\partial^2 \ln s}{\partial \ln A^2} \gg 0$. Thus this proposition suggests that the model of demand derived from distributions of preferences can account for the inverted U curve in employment under fairly general circumstances as long as price starts above p^* and declines below it.

5. Parametric estimates of the Model

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. For a substitute to pose a problem for this empirical exercise, it would have to replace a substantial share of the uses of the target product over a substantial period of time. Otherwise, the substitute could not produce more than a temporary deviation from the level of consumption of the target product predicted by the model. In fact, each industry studied here did have substitutes, especially during the early years, but 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.¹⁵ 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

¹⁵ For example, in Philadelphia in 1820, wool was \$0.75 per pound while cotton sheeting was \$0.15 (US Bureau of the Census 1975).

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.¹⁶ In all three industries, the substitution that took place might not have caused a major deviation from the model.

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 aluminium 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).

Another concern is that the distribution of preferences changes over time, for instance, as income inequality changes. Also, product quality changes over time, distorting consumption measures that are not adjusted for quality. In addition, the model does not take into account time patterns of consumption for consumer durables (auto) and investment goods (some steel). In any case, despite all these potential problems, the model fits the data reasonably well.

Parameterizing the model

In order to investigate the model empirically, it is helpful to provide a flexible functional form for G_x :

(10)

¹⁶ The production of carriages, buggies, and sulkies fell from 538,000 in 1914 to 34,000 in 1921; the production of farm wagons, horse-drawn trucks, and business vehicles fell from 534,000 in 1914 to 67,000 in 1921 (US Bureau of the Census 1975). Bicycles and motorbikes do not appear to have substantially replaced large numbers of motor vehicles.

$$G_x(\hat{x}(p, w)) = \hat{x}^{\alpha-1}, \quad 0 \leq \alpha < 1. \text{ }^{17}$$

Recalling that the assumption that $\hat{x} \approx w$, (7) and (8) yield

(11)

$$\hat{v}(p, w) \approx \frac{p}{w} \cdot w^\alpha = \frac{w^\alpha}{s \cdot A}.$$

The first expression presents \hat{v} as the product of the ratio of price to the wage—as is commonly specified in indirect utility functions—and a pure income term. The second expression presents \hat{v} as the product of labour productivity in textiles and an income term (the labour share of output is approximately constant during most of the sample period).

Then, choosing a lognormal specification for $F(\cdot)$, per capita demand, D , can be written

(12)

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

or

(13)

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

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. γ , μ and σ are parameters to be estimated. Finally, the per capita demand function is defined above as the demand of a

¹⁷ Making $G(x) = \frac{1}{\alpha}x^\alpha$ if $\alpha > 0$ or $G(x) = \ln x$ if $\alpha = 0$.

single individual. It is straightforward to re-conceptualize D as an average over all individual consumers.

Model estimates

This section estimates the log version of (12) and (13). The dependent variable is the log of consumption per capita as seen in Figure 3. The solid line in that figure is simply a log version of equation (12) fit to the data (with $\alpha = 0$ as discussed below).

In order to measure consumption, it is necessary to make adjustments to production figures in order to account for trade. All figures are adjusted for net imports; however, I lack complete data on imports in downstream industries. 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; for steel, I estimate demand through 1982.¹⁸ As a robustness check, I used different cutoff years, but small changes in the cutoff year did not change coefficient estimates significantly.

Table 2 shows NLLS estimates of equations (12) in columns 1 and 2 and estimates of equation (13) 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 labour 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 better than estimates using the non-parametric specification in

¹⁸ In 1996, imports comprised a third of apparel consumption for the first time and have grown rapidly since. After 1982, 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, per capita consumption falls dramatically around these cutoff years. Because the consumption data become unrepresentative after these years, I estimate the model only for prior years.

Table 1. None of the estimates in column 2 find a significant coefficient for α and the NLLS regression for the auto industry failed to converge for this specification.¹⁹ The lack of a significant estimate of α may be because of lack of statistical power, but it suggests that per capita demand is close to a simple function of $A \cdot s$, as was seen in the non-parametric estimates. Assuming that $\alpha \approx 0$ makes for a simple interpretation of Figure 3. This means that the pure income effect is negligible, so the figure represents the functional relationship between demand and productivity (ignoring changes in labour's share) and the slope thus represents the elasticity.

These predicted levels of per capita demand can also be used to estimate industry production employment by dividing domestic demand (total demand divided by $1 +$ import penetration) by the annual output per production worker. Measuring labour productivity as output per production worker-hour, this is²⁰

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

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 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.

¹⁹ This may be due to insufficient curvature in the data to identify the role of α .

²⁰ For 1820 and before, I also subtract the estimate of labour performed in households.

6. Policy Implications

A strong implication of the above analysis is that the productivity of new productivity-improving technologies is not sufficient to understand their impact on employment. Over 18 recent studies predict job losses from new automation technologies.²¹ These predictions are based on estimates of the abilities of technology to replace humans, that is, the productivity of the technology. As noted above, recent economic theory on automation has also largely focused on the productivity-improving effect of new technology, abstracting away from non-constant or heterogeneous demand.

Yet there are important reasons why demand factors also need to be incorporated in the analysis. The historical evidence shows that demand elasticities can be high, especially when automation is only beginning to take over the work of humans. Moreover, the theoretical model provides reason to think that some industries today, especially industries that have had little automation to date, may also experience highly elastic demand. If so, automation may increase employment in these industries even if the technology is “brilliant” and brings large productivity advances. That is, the rate of productivity growth determines the pace of employment change, but the elasticity of demand determines the *sign*.

Moreover, both the historical evidence and the model estimates suggest that although demand elasticities change substantially over time, this change may occur slowly. This is important because it means that for many industries, the demand elasticity in ten or twenty years will be similar to the demand elasticity today. Hence, if we want to understand

²¹ Erin Winick, “Every study we could find on what automation will do to jobs, in one chart,” *MIT Technology Review*, January 25, 2018, <https://www.technologyreview.com/s/610005/every-study-we-could-find-on-what-automation-will-do-to-jobs-in-one-chart/>.

the impact of productive new technologies on an industry in ten or twenty years, we need to predict productivity growth and also, to estimate the elasticity of demand.

There is evidence to suggest that new technologies are not uniformly decreasing employment but, rather, have disparate effects across industries. Several recent papers find that information technology increases employment for some groups and does not appear to reduce net employment, except, perhaps, in manufacturing. 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. Mann and Püttmann (2017) find that automation increases jobs in services but decreases them in manufacturing. Bessen (2016) finds that computers tend to increase occupational employment modestly overall, 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.

Studies of automation also find little evidence of strong unemployment effects. A study of automation in the Netherlands finds that although a substantial share of employees work at firms that automate each year, relatively few of those affected leave their employer as a result (Bessen et al. 2019). Cirera and Sabetti (2019), studying 53 developing countries, find that automation alone does not significantly affect employment, although it may reduce employment gains from product innovations. Evidence on the impact of industrial robots is more mixed. Graetz and Michaels (2015) and Dauth et al. (2017) find no effect on employment and a positive effect on wages; Acemoglu and Restrepo (2017) find a negative

effect on both. Other studies look more generally at the effects of productivity growth and innovation on employment.²²

Thus, the limited evidence available to date suggests that the impact of these new technologies will likely be negative in some industries and neutral or positive in others. These findings are entirely consistent with heterogeneous responses across industries reflecting different demand elasticities. While new robotic and information technologies surely bring significant productivity gains, demand mediates their impact on employment.

This point raises a second major implication of the analysis: the main impact of automation in the near future may be to cause a major reallocation of jobs even if it does not permanently eliminate large numbers of jobs. That is, employment will fall in some industries and grow in others because some industries have satiated demand while others have large unmet needs and elastic demand. This kind of change can be highly disruptive nevertheless. Workers switching industries often need new skills, and they may need to move to new occupations and sometimes new geographical locations and these transitions may require periods of unemployment.

For example, during the Industrial Revolution automation was disruptive even though it did not create mass unemployment. Numerous observers, including Marx, predicted that new technologies would eliminate work, for instance, in textiles. Instead, employment grew, but the dislocations brought about were nevertheless substantial, contributing to 19th century political upheaval: workers had to acquire new skills, new types of vocational training organizations were needed to provide these skills, and there were parallel demands to change and broaden the education system, providing universal

²² See Autor and Salomons (2018) and Dosi and Mohnen (2018) for an overview of recent papers.

education; workers had to adjust to new types of work organization, moving from farm and workshop to highly-disciplined factories; new labour markets had to develop that recognized the new skills; many workers had to migrate to urban areas to obtain work.

The social challenges posed by new technologies today are, of course, quite different. Nevertheless, if it is true that the pace of technological change is accelerating and if the impact is varied across industries, then many workers will need new occupations, new skills, new ways of working in organizations, and, perhaps, new locations. It is not surprising that we see experiments with new types of apprenticeship and work-study programs, new forms of skill certification for skills learned on the job, new forms of training such as MOOCs, and a general rising discussion of the need for “lifelong learning” incorporated into education systems. Other changes could reduce barriers to employee mobility across occupations (occupational licensing) or locations (land use, zoning).

In addition, workers may need income support in order to make such transitions. While relatively few workers have been displaced by automation in the Netherlands so far, the evidence shows that those displaced workers suffer extended unemployment and income losses that are not offset by the current safety net, including unemployment, welfare, and disability payments (Bessen et al. 2019). Nor are most of these losses offset by early retirement or self-employment. This suggests a need for at least some form of temporary income support beyond the current safety net. While discussion of specific policy proposals is beyond the scope of this paper, the main idea here is that the range of policies needed to address large scale worker transitions is quite different from those aimed at dealing with the supposed imminent rise of permanent unemployment.

Conclusion

This paper, using empirical evidence and a theoretical model, argues that the elasticity of demand with respect to labour productivity substantially affects whether productivity-improving technological change will increase or decrease industry employment. While the pace of productivity growth determines the magnitude of employment change, the elasticity of demand determines the sign of the effect. Moreover, both historical evidence and the model provide reason to expect substantial heterogeneity across industries in their response to new productivity-improving technologies. If so, this implies that a critical policy challenge of new automation technologies is the disruptive transitions for workers as some industries grow while others shrink. Whether automation brings a decline in aggregate employment depends not just on these individual industry responses, but also on spillovers, effects on downstream producers, and general equilibrium effects on labour demand.²³ Regardless of the aggregate effect, this paper argues that new productivity-improving technologies will likely bring a disruptive reallocation.

Much of the literature on automation and deindustrialization generally abstracts away from considerations of demand, assuming that demand is inelastic and static. Using long time series, this paper looks at the changes in the elasticity of demand for the US cotton textile, primary steel, and motor vehicle industries. A non-parametric analysis rejects the hypotheses that these elasticities were always less than one and were static. Demand was highly elastic during the early years of automation in these industries, but it became highly inelastic by the mid-20th century.

²³ See Autor and Salomons (2018) for some estimations of these effects.

A theoretical model explains why demand elasticity changed and also why this change altered the direction of employment growth: during the early years, there were large unmet needs, so the price reduction brought by productivity-improving technology spurred strong demand growth. Demand grew faster than labour productivity, resulting in increased employment. Later, however, demand was satiated; ongoing automation still reduced prices, but these reductions did little to spur demand, resulting in job losses. Moreover, the paper shows that this pattern of declining elasticities holds for common distribution functions, suggesting it may be relevant to understanding the impact of new technologies today. This investigation thus highlights the importance of demand heterogeneity for understanding the impact of new technologies, but it also suggests how to analyse that heterogeneity.

At least since the Industrial Revolution, observers of automation have incorrectly predicted the imminent demise of demand for labour. For example, in 1930, Keynes (1930), anticipating continued productivity growth, predicted that in 100 years his grandchildren would enjoy a fifteen-hour workweek. Now that we are close to that 100-year mark, the average workweek for OECD nations is 34 hours. Yet Keynes was right about productivity growth. In the US, the 1930 level of mean GDP per capita could be realized in 15 hours on average by 1977. What Keynes did not grasp was the depth of human wants and desires, that is, the depth of consumer demand. The reason we don't work 15 hours is that we choose to demand more goods and services that technology has made cheaper and better. A similar underestimation of demand lies behind many other failed predictions of automation-induced mass unemployment.

Of course, this time is different. Now white collar and professional tasks are being taken over by machines and perhaps the pace and breadth of change are greater than past episodes of automation. However, not *everything* is different. Human needs and wants are not

very different and are not likely to be very different in 10 or 20 years. Because demand is not likely to change rapidly, empirical analysis of current demand elasticities and of the current impacts of productivity-improving technologies are critical to inform policy regarding the expected new wave of technology. And demand may not only be elastic enough to mitigate unemployment effects, but it is likely to have different effects across industries as well.

Appendix

Propositions

To simplify notation, let $G_x = 1$. Then, keeping wages constant,

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

so that

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

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 $x > 0$,²⁴

(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

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

²⁴ I present the inverse of Gordon's inequality.

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 $\lim_{p \rightarrow b} \epsilon = \infty$.

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 fibres 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 fibres 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 fibre 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.²⁵ 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). I obtain data on the labour share of output from various Census of Manufactures and the Annual Survey of Manufactures after 1958.

²⁵ "U.S. Automobile Production Figures," https://en.wikipedia.org/wiki/U.S._Automobile_Production_Figures (accessed Jan. 30, 2019).

References

- Acemoglu, Daron, and Pascual Restrepo. Robots and Jobs: Evidence from US Labor Markets. No. w23285. National Bureau of Economic Research, 2017.
- Acemoglu, Daron, and Pascual Restrepo. Artificial Intelligence, Automation and Work. No. w24196. National Bureau of Economic Research, 2018a.
- Acemoglu, Daron and Pascual Restrepo. 2018b. "The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment." *American Economic Review*, forthcoming.
- Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. Artificial Intelligence and Economic Growth. No. w23928. National Bureau of Economic Research, 2017.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. "The skill complementarity of broadband internet." *The Quarterly Journal of Economics* (2015), 1781–1824.
- Autor, David, David Dorn, and Gordon H. Hanson. "Untangling trade and technology: Evidence from local labour markets." *Economic Journal* 125.584 (2015): 621-46.
- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. The fall of the labor share and the rise of superstar firms. National Bureau of Economic Research, 2017.
- Autor, David, and Anna Salomons, "Is automation labour-displacing? Productivity growth, employment, and the labour share," *Brookings Papers on Economic Activity*, 2018 forthcoming.
- Banks, James, Richard Blundell, and Arthur Lewbel. "Quadratic Engel curves and consumer demand." *Review of Economics and statistics* 79, no. 4 (1997): 527-539.
- Baumol, William J. "Macroeconomics of unbalanced growth: the anatomy of urban crisis." *The American economic review* (1967): 415-426.
- Bessen, James. "More machines, better machines... or better workers?." *The Journal of Economic History* 72, no. 1 (2012): 44-74.
- Bessen, James E. "How computer automation affects occupations: Technology, jobs, and skills." (2016).
- Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge, 2019. "Automatic Reaction: What happens to workers at firms that automate?" CBS working Paper.
- Boppart, Timo. "Structural Change And The Kaldor Facts In A Growth Model With Relative Price Effects And Non-Gorman Preferences." *Econometrica* (2014): 2167-2196.

- Carter, Susan B. *Historical statistics of the United States: Earliest times to the present*. Cambridge University Press, 2006.
- Cirera, Xavier, and Leonard Sabetti. "The effects of innovation on employment in developing countries: evidence from enterprise surveys." *Industrial and Corporate Change* (2019).
- 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.
- Dauth, Wolfgang, Sebastian Findeisen, Jens Südekum, and Nicole Woessner. "German robots-the impact of industrial robots on workers." (2017).
- 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.
- Diebolt, Claude. "La théorie de la sous-consommation du cycle des affaires de Emil Lederer", *Economie Appliquée. An International Journal of Economic Analysis*, Vol. 50, No. 1, 1997, pp. 27-50.
- Dosi, Giovanni, and Pierre Mohnen. "Innovation and employment: an introduction." *Industrial and Corporate Change* (2018).
- 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.
- Frey, Carl Benedikt, and Michael A. Osborne. "The future of employment: how susceptible are jobs to computerisation?." *Technological forecasting and social change* 114 (2017): 254-280.
- 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.
- Graetz, G., and G. Michaels. "Robots at work. forthcoming." *Review of Economics and Statistics* (2015).
- Hémous, David, and Morten Olsen. "The rise of the machines: Automation, horizontal innovation and income inequality." (2016).
- Keynes, John Maynard. "Economic possibilities for our grandchildren." In *Essays in persuasion*, pp. 321-332. Palgrave Macmillan, London, 2010.
- 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. "Labour 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).
- Mann, Katja and Lukas Püttmann, "Benign Effects of Automation: New Evidence From Patent Texts" working paper, 2017.
- 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.
- 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 Labour. *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. Non-parametric Regressions and Estimates of Demand Elasticity

A. Dependent Variable: ln D (log demand per capita)						
	1 Cotton Textiles		2 Primary Steel		3 Auto	
ln sA	0.71 **	(0.10)	-7.19 **	(1.50)	-12.43 **	(1.32)
(ln sA) ²	-0.13 **	(0.02)	-0.75 **	(0.15)	-0.92 **	(0.09)
ln w	-0.71	(0.87)	8.68 *	(4.03)	4.51 **	(1.61)
(ln w) ²	0.03	(0.05)	-0.39	(0.21)	-0.21 *	(0.08)
N	52		35		61	
R-squared	0.979		0.974		0.934	
B. Dependent Variable: ln s (log labour share of output)						
ln A	.016	(0.010)	.066 **	(0.023)	-.348 **	(0.051)
N	52		35		61	
R-squared	.053		.202		.442	
C. Non-parametric Estimates of Demand Elasticity, $\frac{\partial \ln D}{\partial \ln A}$						
	Year	Elasticity	Year	Elasticity	Year	Elasticity
H0: elasticity=1	1820	1.3	1870	2.2	1910	2.4
	P =	0.008	P =	0.015	P =	0.000
H0: elasticity=1	1950	0.2	1950	-1.2	1951	0.4
	P =	0.000	P =	0.000	P =	0.000
D. Model Estimates of Demand Elasticity, $\frac{\partial \ln D}{\partial \ln A}$						
	1820	1.7	1870	2.9	1910	4.4
	1950	0.2	1950	0.4	1951	1.0

Note: Robust standard errors in parentheses. Constant term not reported. The wage, w , is real GDP per capita. The hypothesis tests are two-tailed F tests of the null hypothesis. Note that $s \cdot A = w/p$. The model estimates come from the specification of Table 2, Column 1.

Table 2. Regressions of Per Capita Demand

Independent variable	1 Labour productivity	2 Labour productivity	3 Price / wage
A. Cotton cloth, 1820 - 1995			
μ	-0.24 (0.10)**	0.40 (1.36)	-1.49 (.79)*
σ	1.43 (0.15)***	1.64 (0.29)***	2.04 (.58)**
γ	134.60 (3.80)***	135.17 (3.66)***	184.42 (49.15)**
α		-0.16 (0.32)	
Observations	52	52	37
R-squared	0.993	0.993	0.990
B. Raw steel, 1860 – 1982			
μ	5.04 (0.18)***	4.85 (3.38)	3.60 (.86)***
σ	0.83 (0.18)***	0.79 (0.50)	1.46 (.40)***
γ	0.66 (0.05)***	0.66 (0.07)***	1.32 (.64)**
α		0.05 (0.90)	
Observations	35	35	116
R-squared	0.981	0.981	0.958
C. Motor vehicles, 1910 – 2007			
μ	7.31 (0.04)***		
σ	0.30 (0.06)***		
γ	59.40 (2.78)***		
Observations	61		
R-squared	0.977		

Note: Non-linear least squares estimates of equation (12) in columns 1 and 2 and equation (13) in column 3. Robust standard errors in parentheses; ***= significant at 1%; ** = significant at 5%; * = significant at 10%.