KAUFFMAN
The Foundation of Entrepreneurship

KAUFFMAN DISSERTATION EXECUTIVE SUMMARY

Part of the Ewing Marion Kauffman Foundation's Emerging Scholars initiative, the Kauffman Dissertation Fellowship Program recognizes exceptional doctoral students and their universities. The annual program awards up to fifteen Dissertation Fellowship grants of $20,000 each to Ph.D., D.B.A., or other doctoral students at accredited U.S. universities to support dissertations in the area of entrepreneurship.

Since its establishment in 2002, this program has helped to launch world-class scholars into the exciting and emerging field of entrepreneurship research, thus laying a foundation for future scientific advancement. The findings generated by this effort will be translated into knowledge with immediate application for policymakers, educators, service providers, and entrepreneurs as well as high-quality academic research.

Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP
Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP
Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP

www.kauffman.org/kdfp

Dissertation Summary

Innovation and Industry Development: Lessons from the British Cotton Textile Industry During the U.S. Civil War

W. Walker Hanlon

May 16, 2012

This dissertation uses the large shock to the British cotton textile industry in the 19th century, caused by the U.S. Civil War (1861-1865), in order to address three long-running questions about technological progress and industry development. Chapter 1 considers the idea that a change in the availability of inputs to the production process can influence the direction of technological progress. This idea, called directed technical change, is a long-standing and influential line of research. Chapter 1 provides the first well-identified empirical test of the leading theory of directed technical change, due to Acemoglu (2002). Chapter 2 considers the possibility that the stock of existing knowledge related to a particular technology type may influence the level of innovation in that technology, which I call path dependence in innovation. In order to shed light on this relationship, I (1) highlight three important shortcomings of the approach used in previous empirical studies of path dependence in innovation, (2) suggest an approach which addresses these concerns, and (3) implement this approach in one empirical setting in order to generate empirical evidence that is not subject to these sources of bias. Chapter 3 considers the idea, suggested by Marshall (1890), that firms in different industries may benefit from locating near to one another, through localized inter-industry spillovers. Using a large, exogenous, and temporary shock to one industry, this chapter provides the first empirical evidence that temporary shocks, acting through inter-industry connections, can have a long-run impact on the geographic location of economic activity. All three of these chapters share a common empirical setting, and so I begin this summary by describing the setting, before dealing with each of the chapters in more detail.

Empirical Setting

The cotton textile industry was a large and important sector in the British economy during the 19th century. The industry was entirely dependent on imported raw cotton, most of which came from the U.S. South prior to the Civil War. The onset of the war sharply reduced the supply of Southern cotton to the British market, causing a severe downturn in the industry. In response to the shock, cotton textile producers turned to other sources of supply, chiefly India, but also Egypt, Brazil and others, to help meet their raw cotton needs. But cotton from these alternative suppliers, and India in particular, was very different from the U.S. cotton that British producers were used to spinning. As a result, British cotton textile producers were faced with a number of new challenges. The first two chapters of this dissertation describe how the cotton textile industry developed new technology in order to deal with these challenges, and what this response can tell us about the process of innovation. Chapter three then investigates the impact of the recession on other industries in the British economy.

This historical setting has a number of features which makes it a particularly good setting for investigating technological progress. First, the U.S. Civil War caused a shock which was both large and exogenous. The size of the shock ensures that the response will be large enough to clearly observe, while the exogenous nature of the shock means that it can be used as a natural experiment in order to uncover causal relationships. Second, the impact of the U.S. Civil War was largely industry-specific; while the impact of the war on the cotton textile industry was severe, most other sectors of the British economy were not directly impacted. This includes other textile industries based on wool, linen, and silk, which do not show any ill effects during the war. One advantage of this feature is that it allows me to control for other time-varying factors by comparing the cotton textile industry to these other similar industries. I will also be able to uncover evidence of inter-industry connections, since other industries will be affected primarily through their relationship with the cotton textile industry. Another feature of this shock is that, despite the magnitude, there was virtually no government intervention in the affected markets. This unique feature was due to the particularly strong free-market ideology that was dominant in Britain during this period.

Chapter 1

The idea of directed technical change, first suggested by Hicks (1932), has been used to explain a diverse set of economic phenomena. It has been suggested that the increase in skilled workers in the U.S. in the 1970s caused skill-biased *directed technical change*, and that this directed technical change allowed the skill premium to increase in spite of the increase in the relative abundance of skilled workers (Acemoglu (1998), Kiley (1999)). Looking at an earlier period, several authors have suggested that a shortage of labor drove the development of labor-saving innovations which played an important role in industrialization in Britain and the U.S. (Habakkuk (1962), Allen (2009)). In the environmental literature, it has been pointed out that the impact of regulations that change the price of inputs, such as a carbon tax, will depend crucially on whether these changes generate directed technical change, and on the direction that this innovation takes (Acemoglu *et al.* (2012)).[1] Yet, despite the wide

[1] Related papers in the environmental literature include Porter (1991), Lanjouw & Mody (1996) and Jaffe & Palmer (1997). The idea of directed technical change has also been applied to consider the impact of high energy prices (Newell *et al.* (1999), Popp (2002)), the causes of cross-country productivity differences (Acemoglu & Zilibotti (2001), Caselli & Coleman (2006)), and agricultural productivity trends (Hayami &

application of this idea, several important predictions of leading directed technical change theories remain largely unexplored empirically.

Recently, Acemoglu (2002) presents an endogenous growth model that incorporates multiple inputs and uses it to consider the impact of a change in the supply of inputs, an exogenous parameter, on technological progress. The model provides two new results. First, Acemoglu shows that the direction of technical change depends crucially on the elasticity of substitution between inputs, represented by σ. When this elasticity is low ($\sigma < 1$), technical change will be directed towards technologies that augment the input which has become relatively scarce. In contrast, when the elasticity of substitution between inputs is high ($\sigma > 1$), technical change will be directed towards technologies that augment the input which has become relatively more abundant. This prediction stands in contrast to previous work, such as Hicks (1932), which hypothesized that technical change would always augment the more scarce input. The second major prediction of Acemoglu (2002) is that, when the elasticity of substitution between inputs is sufficiently high ($\sigma > 2$), technical change will be so strongly directed towards technologies that augment the more abundant input that the relative price of that input can increase. This *strong induced-bias hypothesis* may explain, for example, how an increase in the supply of skilled workers may increase the skill wage premium.

The shock to the British cotton textile industry generated a large exogenous shift in the relative cost of providing inputs to production that is used to identify the causal impact on the direction of technical change and input prices. This exogenous change is the first necessary ingredient for testing the theory. I construct an extensive data set covering the price and import quantity of raw cotton supplied by the U.S., India, Brazil, and Egypt before, during, and after the war. Observing data on prices and quantities for multiple inputs is the second crucial ingredient for testing the predictions of the theory, since they allow me to estimate the elasticity of substitution between inputs and track the impact on input prices. In order to track the impact on innovation rates, I gathered previously unexploited data on British patents containing a high level of detail on the types of new technologies being created. Using these patent data it is possible to track patterns of technical change in technologies related to particular inputs, which is a third necessary element for testing the theory. Previous studies have not incorporated all of these necessary elements, and as a result, the main predictions of Acemoglu (2002) have not yet been tested empirically. The novel contribution of this paper is to introduce a setting with these features and use it to test the predictions of the theory.

The analysis begins with a theory that extends the model of Acemoglu (2002). In his model, final goods are produced using inelastically supplied factors, and results are

Ruttan (1970), Olmstead & Rhode (1993)).

derived by varying relative factor supplies and observing the impact on technical change. I extend the model by incorporating elastically supplied inputs in place of the fixed factors. Generalizing the model in this way is important for making it empirically testable, since in practice we almost never observe truly fixed factors. In my model, inputs are supplied through trade and each input faces an input-specific transport cost. This connects the model to the empirical setting, where the U.S. Civil War can be thought of as a large increase in the transport cost of U.S. cotton. The main predictions of the model correspond to those found by Acemoglu (2002) which are described above. It also generates one new prediction: that input supply responses can act to magnify directed technical change and its impact on input prices. So, for example, when the elasticity of substitution between inputs is high, so that we expect technical change to be directed towards the input which becomes more abundant, this effect will be stronger the higher is the elasticity of input supply.

The predictions of the theory depend on the elasticity of substitution between inputs, and the elasticity of input supply for each input, so the first step in the empirical analysis is to estimate these elasticities. To estimate these parameters, I used data on the prices and quantities of the various cotton varieties, much of which I gathered from original sources. I estimate the elasticity of input supply for Indian, Brazilian, and Egyptian cotton by comparing the response of import quantities to the change in their prices, where I use the shock as an instrument for the increase in the cotton price. These estimates suggest that Indian cotton was the most responsive to the increase in prices caused by the Civil War, while Brazilian cotton was the least responsive. The elasticity of substitution between inputs is estimated by comparing the response of the price of each variety relative to the U.S. cotton price to the change in relative transport costs caused by the war, while taking into account the price elasticity of each input. The shock provides the necessary instrument for the change in transport costs. I find that the elasticity of substitution between Indian and U.S. cotton was well above 1 and likely to have been above 2. The elasticities of substitution for Egyptian and Brazilian cotton, relative to U.S. cotton, were also above 1, and likely above 2 as well. Given these, the predictions of the model are that (1) we should observe technical change that augments these alternative varieties, (2), the long-run price of these alternative inputs should increase, and (3), these changes should be more pronounced for Indian cotton than for the other varieties.

To track innovation patterns, I collected detailed British patent data which come from over 1500 pages of original printed documents. These data show that spinning technology patents, the largest category of textile-related technologies, increased by 18% during the Civil War period. In contrast, no increase is observed in non-textile technology patents. Within the spinning technology category, I can focus specifically on patents related to cotton, wool, linen, and silk. I observe a sharp increase in cotton-related spinning technology patents during the 1861-1865 period, on the order of 52-67%, but patents related to wool, linen, or silk show no similar increase. Data on patents of spinning technology subcategories allow me to look at these changes in even more detail. I find that the increase in cotton textile technology patents was driven by an increase in machines used to prepare the raw cotton for spinning, specifically gins, openers and scutchers – exactly the machines which were needed in order to take advantage of the Indian cotton. All of these results continue to hold when I focus only on high-quality patents using three different measures of patent quality. Thus, the patent data provide evidence of technical change which was directed towards those technologies which augment Indian cotton, an input which became relatively more abundant. This is consistent with the predictions of the theory, for Indian cotton, given my elasticity estimates. However, it runs counter to the hypothesis of Hicks (1932) that technical change would be directed towards the input that becomes more scarce.

Next, I consider the impact on input prices. While the price increased substantially for every cotton variety during the Civil War, my focus will be on what happened to the price of Indian, Brazilian, or Egyptian cotton *relative* to the U.S. cotton price during and after the war. In the absence of directed technical change, the model predicts that the relative price of each of these varieties should have fallen as it became relatively more abundant. On the other hand, the technical change directed towards augmenting Indian cotton may offset this, by increasing the demand for that variety. Graphing the relative price of Indian to U.S. cotton shows a decrease in the first two years of the Civil War, followed by a rapid rebound starting in 1863, around the time when the new technologies were becoming available. In contrast, for Brazilian or Egyptian cotton varieties, where I do not observe evidence of directed technical change, the price relative to U.S. cotton fell at the beginning of the war and remained low throughout the period in which these varieties remained abundant relative to the pre-war period. Econometric results strengthen this finding. I find that there was a significant decrease in the relative price of Brazilian and Egyptian cotton over the ten years following the onset of the war. In contrast, the relative price of Indian cotton was higher, on average, during this period. These results are consistent with the existence of strong induced bias acting on the relative price of Indian cotton, consistent with the prediction of the model for this variety given my elasticity estimates.

Chapter 2

How does the stock of knowledge in a particular type of technology affect the level of innovation in that technology type? Does current innovation "stand on the shoulders of giants", benefiting from the lessons learned from previous research? Or have previous innovators already harvested all of the "low hanging fruit", leading to decreasing returns along particular lines of research? The answer to these questions is central to our understanding of the process of innovation and economic growth, with important policy implications.

In order to address these questions, existing empirical studies, such as Popp (2002) and

Aghion *et al.* (2011) construct measures of the rate of innovation and the stock of knowledge for different types of technologies over time using patent data.[2] They find evidence of a positive relationship between the lagged stock of knowledge for a particular technology type and rate of innovation in that technology, which they interpret as evidence in favor of what I will call *path dependence in innovation.*

While these previous studies have highlighted the importance of understanding path dependence in innovation, their empirical approach suffers from several important sources of bias. This study will (1) highlight three important shortcomings of the approach used in previous empirical studies of path dependence in innovation, (2) suggest an approach which addresses these concerns, and (3) implement this approach in one empirical setting in order to generate empirical evidence that is not subject to these sources of bias.

The first concern highlighted by this study has to do with the technology level at which to look for path dependence. For example, suppose that there is an increase in the stock of knowledge for a few technologies used in the cotton textile industry and we are interested in whether this generates path dependence in innovation. We may look for a subsequent increase in innovation in all cotton textile technologies, or an increase in only the types of technologies which experienced the initial increase, or an increase in innovation by only those firms which produced the initial technological improvements. In other words, there are potentially multiple levels at which we could look for path dependence in technology. This study begins by offering a theory showing that path dependence may exist at more than one of these technology levels. If we look for path dependence at only one level, then we may miss it at another level. Moreover, this problem cannot be solved by focusing only on the most disaggregated level (e.g. firms). This motivates an empirical exercise that focuses on multiple technology levels.

The second concern with previous studies is that, in the presence of persistent shocks that increases innovation in a particular firm or technology type, they will estimate a positive relationship between the lagged stock of knowledge and the innovation rate, even when no causal relationship exists. A third concern is that any trends which are specific to the cross-sectional unit of analysis will also generate a positive relationship between the lagged stock of knowledge and the innovation rate when using the empirical approach employed by previous studies.

In order to address these concerns, I suggest an approach that takes advantage of an exogenous temporary shock that affects the stock of knowledge related to a particular set of technologies. If there is path dependence in innovation, then this increase should lead to a higher rate of innovation in these technology types in the period just after the shock.

[2] Both of these studies are interested in energy-related technologies. Popp (2002) focuses on innovation in several categories of energy-efficient technologies, while Aghion *et al.* (2011) consider "clean" (not carbon emitting) and "dirty" (carbon emitting) automotive technologies.

Thus, if the shock is truly temporary, then observing a higher level of innovation in the post-shock period will provide causal evidence of path dependence in innovation. Also, in order to control for trends specific to particular technology types, this approach requires panel data where the time dimension is long enough to allow me to control for time-trends which are specific to the cross-sectional units of the data. Finally, this approach should be implemented at multiple technology levels.

I implement this methodology using the large exogenous shock to the cotton textile industry in 19th century Britain. As shown in Chapter 1, the cotton shortage led to a surge in innovation during the 1861-1865 period, as British inventors struggled to take advantage of Indian cotton. The result was a large increase in cotton textile technology innovation concentrated in two particular types of cotton textile machinery – gins and openers/scutchers – which were particularly important for dealing with Indian cotton. Because these effects were specific to one industry, cotton textiles, and to a particular subset of technologies within that industry, I can control for other time-varying factors by comparing the cotton textile industry to these other textile industries, and I am able to control for time-varying factors at the intermediate level by comparing technologies categories which were directly impacted by the shock to those that were not. The most important feature of this shock is that it is largely, though not entirely, temporary. British cotton imports quickly rebounded following the end of the Civil War, as did industry output. However, some changes did persist. Indian cotton remained a much larger fraction of imports in the post-shock period than it was prior to the war. The key point is that I am able to identify the expected direction of the bias that could be generated by these persistent effects, since they are continuations of the changes which drove the increase in innovation during the 1861-1865 period. While this is not a perfect empirical setting in which to implement the methodology, it approximates the ideal setting while allowing me to credibly assess the direction of the remaining bias. This is an improvement over previous studies which are not able to directly address these important sources of bias.

I look for path dependence in innovation at the industry level by comparing patenting patterns in the cotton textile industry to the wool, linen, and silk industries. I find that despite the sharp increase in cotton textile related patents during the 1861-1865 period, there is no evidence of a continuing higher level of innovation after 1865. This result is strengthened because I expect that any changes caused by the Civil War which persisted after 1865 should cause an increase in innovation in the post-shock period. Thus, this paper provides clear evidence against the existence of path dependence in innovation at the industry level. In contrast, when I use the approach offered in previous studies I find strong evidence in favor of path dependence in innovation. Next, I look for path dependence at the intermediate level using patents of eight specific types of textile technologies. Two of these eight technology types – gins and openers/scutchers – were specifically related to using Indian cotton inputs. Thus, these two technology types experienced a large increase in patents during the 1861-1865 period. The analysis focuses on whether higher levels were sustained in the years after 1865, while using the remaining six technologies as controls. At this level I find some evidence that there was a sustained higher level of patents in these categories in the three years after 1865. The fact that these results differ from those obtained at the industry level highlights the importance of considering multiple technology levels.

Chapter 3

Marshall (1890) suggested that firms in different industries may benefit from locating near to one another, through localized inter-industry spillovers. He identified three channels through which these benefits could flow: input-output linkages, labor market pooling, or technology spillovers. Since then, inter-industry spillovers have been incorporated into theories of industrial linkages and development (Hirschman (1958), Ciccone (2008)), industrial clusters (Porter (1990)), the benefits of urban economies (Jacobs (1969)), the benefits of trade and FDI (Young (1991), Rodriguez-Clare (1996)), and the geography of economic activity (Krugman & Venables (1995)). Policy makers, too, have been influenced by these ideas. The existence of inter-industry spillovers is one of the prime motivations for local industrial policies, such as the tax incentives offered to firms by U.S. municipalities, or the widespread use of special industrial zones in developing countries. The cost of these policies can run into the hundreds of millions of dollars for single U.S. municipalities.[3]

Despite this interest, and widespread application, empirical evidence on the role of inter-industry connections in influencing the geographic location of economic activity is sparse. This is due largely to the difficulty of measuring the patterns of connections between industries, though these measures are improving. In an important recent study, Greenstone *et al.* (2010) show evidence of localized productivity spillovers between plants that share similar labor or technology pools. In a similar vein, Javorcik (2004) finds evidence of spillovers from FDI firms to upstream suppliers. More related to the current study is Ellison *et al.* (2010), which uses measures of input-output connections, labor force similarity, and technological spillovers, to provide evidence that the underlying pattern of connections between industries is influencing the geographic location of industries. While these are important contributions, they do not address the most relevant question for policy: can *temporary changes* in the local availability of inter-industry spillovers, such as those created by economic shocks or policy interventions, have long-term impacts on the geographic location of economic activity?

This study takes the next step, by utilizing a large temporary external shock that altered the spillovers available to certain industries in certain locations, in order to provide

[3] See Greenstone & Moretti (2004).

causal evidence that changes in the availability of inter-industry spillovers can influence the long-run distribution of economic activity. My focus will be on how this affected the location of those industries *related to* (sharing spillovers with) the cotton textile industry. An important feature of this setting is that some locations were severely impacted, while other, economically similar locations, were left nearly untouched. This study compares outcomes in towns from two industrial counties in the north of England, Lancashire and Yorkshire. Lancashire was the heart of Britain's cotton textile industry at the time. Yorkshire, lying just to the east, was similar to Lancashire in many ways. The key difference between these two counties was that, while towns in Yorkshire also had large textile industries, Yorkshire producers focused primarily on wool-based textiles (woollens & worsted) rather than cotton. Thus, while towns in Lancashire were severely affected by the cotton shortage, towns in Yorkshire were not negatively impacted. Comparing industry growth rates from towns in these two counties thus allows us to better identify the impact of the shock.

The basic hypothesis that I test is that the shock negatively affected employment and employment growth in industries more closely related to the cotton textile industry, in locations more severely impacted by the shock, by reducing the spillover benefits available to these industries. Importantly, I focus on impacts occurring in the years and decades after the Civil War had ended and the cotton textile industry had rebounded. The hypothesis is motivated by a two-country dynamic Ricardian trade model building on work by Young (1991) and Matsuyama (1992). In the model, technology growth is driven by localized learning-by-doing spillovers within and between industries. A negative shock to one industry reduces the spillover benefits enjoyed by related industries, in the location in which the affected industry is concentrated. The result is a reduction in technology growth in the related industry, in the more severely affected location relative to the less affected location. Moreover, the model describes channels through which a loss of spillovers in one period can affect employment and employment growth in related industries in future periods. The model is used to derive the empirical specification used in the analysis.

This intuition is perhaps best illustrated using an example from the empirical setting, provided by the Engine & Machinery industry (E&M). This was an important industry in Britain at this time, and one that was connected to the cotton textile industry in Lancashire, as well as the wool textile industry in Yorkshire, through all three of Marshall's channels. In fact, Marshall himself used the textile and engineering industries to illustrate the possibility of labor market pooling benefits. There is also evidence suggesting that textile machine makers learned from the nearby textile producers that they supplied. Engine and machine makers in Lancashire and Yorkshire competed, both to supply customers in these locations, as well as in the important export market. The data show that the E&M industry had a similar growth path in the two locations prior to the shock, but that E&M firms in Yorkshire towns gained an advantage relative to Lancashire producers in the following decades, allowing them to expand employment more rapidly. This suggests that the recession in the cotton textile industry had persistent impacts on distribution of the E&M industry across locations.

The empirical strategy used to test this hypothesis involves using panel data with two cross-sectional dimension, allowing us to compare impacts across time (pre vs. post shock periods), locations (towns with higher vs. lower shock intensity), and industries (more vs. less related to cotton textiles). The primary data are drawn from the British Census of Production, which were gathered from original sources. These data provide employment disaggregated into 171 industry groups, spanning nearly the entire private sector economy, available for every ten years from 1851-1891. Thus, I have multiple observations in both the pre- and post-shock period, and are able to observe effects up to 25 years after the end of the recession. These data are available for 11 principal towns, 6 in Lancashire, and 5 in Yorkshire, providing the geographic dimension of the analysis. Additional data from local Poor Law boards, which were the primary source of funds for unemployed workers during this period, allow us to measure the severity of the shock in each town. I find that the share of cotton textiles in a town's employment in 1851 is a good predictor of the severity of the shock in each location. Thus I am able to strengthen my identification strategy by using each town's cotton textile employment in 1851 as an instrument for shock intensity.

An important input into the analysis is a measure of the pattern of connections between the cotton textile industry and other industries. Two measures are used. The first is based on the degree to which each industry is geographically coagglomerated with the cotton textile industry, following work by Ellison & Glaeser (1997). Ellison *et al.* (2010) show that this measure is related to measures of input-output linkages, labor market pooling potential, and technology spillovers. The second measure is an intermediate goods input-output matrix based on Thomas (1987).

The results suggest that industries more closely related to the cotton textile industry, based on the geographic coagglomeration measure, suffered lower employment and employment growth, in more severely affected towns, in the post shock period. The impacts are estimated while controlling for aggregate industry-level and town-level shocks in each year, as well as time-invariant industry-location factors. The impact of the shock on employment and employment growth continues to appear through the 1881-1891 period. The implication is that inter-industry connections can play an important role in affecting the geographic location of economic activity across locations. Moreover, temporary changes in the availability of these connections appear to have the potential to generate long-lasting effects. There appear to be no persistent effects related to the intermediate goods input-output matrix, suggesting that the observed effects are being driven by other types of inter-industry connections.

References

Acemoglu, Daron. 1998. Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *Quarterly Journal of Economics*, **113**(4), 1055 – 1089.

Acemoglu, Daron. 2002. Directed technical change. *Review of Economic Studies*, **69**(4), 781–809.

Acemoglu, Daron, & Zilibotti, Fabrizio. 2001. Productivity Differences. *Quarterly Journal of Economics*, **116**(2), 563–606.

Acemoglu, Daron, Aghion, Philippe, Bursztyn, Lonardo, & Hemous, David. 2012. The Environment and Directed Technical Change. *American Economic Review*, **102**(1), 131–166.

Aghion, Philippe, Dechezlepretre, Antoine, Hemous, David, Martin, Ralph, & Van Reenen, John. 2011 (July). *Carbon Taxes, Path Dependency and Directed Technical Change : Evidence from the Auto Industry*. Working paper.

Allen, Robert. 2009. The Industrial Revolution in Miniature: The Spinning Jenny in Britain, France, and India. *The Journal of Economic History*, **69**(4), 901–927.

Caselli, Francesco, & Coleman, II, Wilbur John. 2006. The world technology frontier. *American Economic Review*, **96**(3), 499–522.

Ciccone, Antonio. 2008. *Linkages*. Basingstoke: Palgrave Macmillan.

Ellison, G., Glaeser, E., & Kerr, W. 2010. What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns. *American Economic Review*, **100**(3), pp. 1195–1213.

Ellison, Glenn, & Glaeser, Edward L. 1997. Geographic concentration in US manufacturing industries: A dartboard approach. *Journal of Political Economy*, **105**(5), 889–927.

Greenstone, Michael, & Moretti, Enrico. 2004. *Bidding for Industrial Plants: Does Winning a 'Million Dollar Plant' Increase Welfare?* MIT Department of Economics Working Paper No. 04-39.

Greenstone, Michael, Hornbeck, Richard, & Moretti, Enrico. 2010. Identifying Agglomeration Spillovers: Evidence from Million Dollar Plants. *Journal of Political Economy*, **118**(3), pp. 536–598.

Habakkuk, HJ. 1962. *American and British Technology in the Nineteenth Century: Search for Labor Saving Inventions*. Cambridge University Press.

Hayami, Yujiro, & Ruttan, V. W. 1970. Factor Prices and Technical Change in Agricultural Development: The United States and Japan, 1880-1960. *Journal of Political Economy*, **78**(5), 1115–1141.

Hicks, J. R. 1932. *The Theory of Wages*. New York: Macmillan.

Hirschman, A.O. 1958. *The Strategy of Economic Development*. New Haven, CN: Yale University Press.

Jacobs, Jane. 1969. *The Economy of Cities*. New York: Vintage Books.

Jaffe, Adam B., & Palmer, Karen. 1997. Environmental Regulation and Innovation: A Panel Data Study. *The Review of Economics and Statistics*, **79**(4), 610–619.

Javorcik, Beata S. 2004. Does Foreign Direct Investment Increase the Productivity of Domestic Firms? In Search of Spillovers through Backward Linkages. *The American Economic Review*, **94**(3), 605–627.

Kiley, Michael T. 1999. The Supply of Skilled Labour and Skill-Biased Technological Progress. *Economic Journal*, **109**(458), 708–724.

Krugman, Paul, & Venables, Anthony. 1995. Globalization and the Inequality of Nations. *Quarterly Journal of Economics*, **110**(4), pp. 857–880.

Lanjouw, Jean Olson, & Mody, Ashoka. 1996. Innovation and the international diffusion of environmentally responsive technology. *Research Policy*, **25**(4), 549 – 571.

Marshall, Alfred. 1890. *Principles of Economics*. New York: Macmillan and Co.

Matsuyama, Kiminori. 1992. Agricultural Productivity, Comparative Advantage, and Economic Growth . *Journal of Economic Theory*, **58**(2), 317–334.

Newell, Richard G., Jaffe, Adam B., & Stavins, Robert N. 1999. The Induced Innovation Hypothesis and Energy-Saving Technological Change. *The Quarterly Journal of Economics*, **114**(3), 941–975.

Olmstead, Alan L., & Rhode, Paul. 1993. Induced Innovation in American Agiculture: A Reconsideration. *Journal of Political Economy*, **101**(1), 100–118.

Popp, David. 2002. Induced Innovation and Energy Prices. *American Economic Review*, **92**(1), 160–180.

Porter, Michael E. 1990. *The Competitive Advantage of Nations*. New York, NY: The Free Press.

Porter, Michael E. 1991. America's Green Strategy. *Scientific American*, **264**, 168.

Rodriguez-Clare, Andrs. 1996. Multinationals, linkages, and economic development. *American Economic Review*, **86**(4), 852–873.

Thomas, Mark. 1987. *An Input-Output Approach to the British Economy, 1890-1914*. Ph.D. thesis, Oxford University.

Young, Alwyn. 1991. Learning By Doing and the Dynamic Effects of International Trade . *Quarterly Journal of Economics* , **106**(2), 369–405.