

DISCUSSION PAPER SERIES

DP14970
(v. 3)

TECHNOLOGY ADOPTION AND PRODUCTIVITY GROWTH: EVIDENCE FROM INDUSTRIALIZATION IN FRANCE

Réka Juhász, Mara Squicciarini and Nico Voigtländer

**ECONOMIC HISTORY
MACROECONOMICS AND GROWTH**



TECHNOLOGY ADOPTION AND PRODUCTIVITY GROWTH: EVIDENCE FROM INDUSTRIALIZATION IN FRANCE

Réka Juhász, Mara Squicciarini and Nico Voigtländer

Discussion Paper DP14970
First Published 30 June 2020
This Revision 15 March 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Economic History
- Macroeconomics and Growth

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Réka Juhász, Mara Squicciarini and Nico Voigtländer

TECHNOLOGY ADOPTION AND PRODUCTIVITY GROWTH: EVIDENCE FROM INDUSTRIALIZATION IN FRANCE

Abstract

New technologies tend to be adopted slowly and - even after being adopted - take time to be reflected in higher aggregate productivity. One prominent explanation for these patterns is the need to reorganize production, which often goes hand-in-hand with major technological breakthroughs. We study a unique setting that allows us to examine the empirical relevance of this explanation: the adoption of mechanized cotton spinning during the First Industrial Revolution in France. The new technology required reorganizing production by moving workers from their homes to the newly-formed factories. Using a novel hand-collected plant-level dataset from French archival sources, we show that productivity growth in mechanized cotton spinning was driven by the disappearance of plants in the lower tail - in contrast to other sectors that did not need to reorganize when new technologies were introduced. We provide evidence that this was driven by the need to learn about optimal ways of organizing production. This process of 'trial and error' led to initially low and widely dispersed productivity, and - in the subsequent decades - to high productivity growth as knowledge diffused through the economy and new entrants adopted improved methods of organizing production.

JEL Classification: F63, O14

Keywords: Industrialization, Technology adoption, Firm productivity

Réka Juhász - rj2446@columbia.edu
Columbia University, NBER and CEPR

Mara Squicciarini - mara.squicciarini@unibocconi.it
Bocconi University and CEPR

Nico Voigtländer - nico.v@anderson.ucla.edu
UCLA Anderson, NBER and CEPR

Acknowledgements

We thank Sandy Black, Serguey Braguinsky, Bruno Caprettini, Diego Comin, Tarek Hassan, John Haltiwanger, Walker Hanlon, Michael Huberman, Felix Iglhaut, Amit Khandelwal, Markus Lampe, Thomas Le Barbanchon, Nils Lehr, Bob Margo, Joel Mokyr, Dávid Krisztián Nagy, Gianmarco Ottaviano, Tommaso Porzio, Diego Puga, Todd Schoellman, Eric Verhoogen and numerous seminar audiences at Bocconi, Boston University, Brown, CEU, Cologne, Columbia, CSEF, Leuven, Lund, the Minneapolis FED, Montreal University, Namur University, Penn State, Siena University, Yale, UCL, UCLA, Universidad de Los Andes, Uppsala and participants at the 2021 AEA meeting, the 2020 NBER Summer Institute Meetings, the 2019 SED meeting, the 1st Workshop of Economic History in Bologna, "The Transnational History of French Industrialisation Before 1914" conference, and the IGC-

Stanford conference on Firms, Development and Trade 2019 for helpful comments and suggestions. We thank Gaia Dossi, Johanna Gautier, Angelo Krüger, Lorine Labrue, Simon Margolin, Marie Michaille, Alberto Mola, and Sebastian Ottinger for excellent assistance throughout the construction of the dataset, Guillaume Daudin and Walker Hanlon for kindly sharing data, and Eugene White for sharing his expertise of hyperinflation during the French Revolution. Financial support from the National Science Foundation (Grant # 1628866) is gratefully acknowledged.

Technology Adoption and Productivity Growth: Evidence from Industrialization in France*

Réka Juhász
(Columbia, NBER, and CEPR)

Mara P. Squicciarini
(Bocconi and CEPR)

Nico Voigtländer
(UCLA, NBER, and CEPR)

First version: December 2019

This version: March 2021

Abstract

New technologies tend to be adopted slowly and – even after being adopted – take time to be reflected in higher aggregate productivity. One prominent explanation for these patterns is the need to reorganize production, which often goes hand-in-hand with major technological breakthroughs. We study a unique setting that allows us to examine the empirical relevance of this explanation: the adoption of mechanized cotton spinning during the First Industrial Revolution in France. The new technology required reorganizing production by moving workers from their homes to the newly-formed factories. Using a novel hand-collected *plant-level* dataset from French archival sources, we show that productivity growth in mechanized cotton spinning was driven by the disappearance of plants in the lower tail – in contrast to other sectors that did not need to reorganize when new technologies were introduced. We provide evidence that this was driven by the need to learn about optimal ways of organizing production. This process of ‘trial and error’ led to initially low and widely dispersed productivity, and – in the subsequent decades – to high productivity growth as knowledge diffused through the economy and new entrants adopted improved methods of organizing production.

JEL: F63, O14

Keywords: Industrialization, Technology Adoption, Firm Productivity

*We thank Sandy Black, Serguey Braguinsky, Bruno Caprettini, Diego Comin, Tarek Hassan, John Haltiwanger, Walker Hanlon, Michael Huberman, Felix Iglhaut, Amit Khandelwal, Markus Lampe, Thomas Le Barbanchon, Nils Lehr, Bob Margo, Joel Mokyr, Dávid Krisztián Nagy, Gianmarco Ottaviano, Tomasso Porzio, Diego Puga, Todd Schoellman, Eric Verhoogen and numerous seminar audiences at Bocconi, Boston University, Brown, CEU, Cologne, Columbia, CSEF, Leuven, Lund, the Minneapolis FED, Montreal University, Namur University, Penn State, Siena University, Yale, UCL, UCLA, Universidad de Los Andes, Uppsala and participants at the 2021 AEA meeting, the 2020 NBER Summer Institute Meetings, the 2019 SED meeting, the 1st Workshop of Economic History in Bologna, “The Transnational History of French Industrialisation Before 1914” conference, and the IGC-Stanford conference on Firms, Development and Trade 2019 for helpful comments and suggestions. We thank Gaia Dossi, Johanna Gautier, Angelo Krüger, Lorine Labrue, Simon Margolin, Marie Michaille, Alberto Mola, and Sebastian Ottinger for excellent assistance throughout the construction of the dataset, Guillaume Daudin and Walker Hanlon for kindly sharing data, and Eugene White for sharing his expertise of hyperinflation during the French Revolution. Financial support from the National Science Foundation (Grant # 1628866) is gratefully acknowledged.

[T]here were strong pairwise complementary relations between factory organization and machinery [...] employers needed to simultaneously determine the choice of technique, the level of worker effort, and the way incentives were set up and communications and decisions flowed through the firm hierarchy. [...] Factories were the repositories of useful knowledge ... but they were also the places in which experimentation took place. – Mokyr (2010, pp. 345-46)

1 Introduction

The diffusion of innovation is at the core of aggregate productivity growth in the long run. Yet, many technologies that ended up being widely adopted were slow to diffuse across firms (Griliches, 1957; Mansfield, 1961; Rosenberg, 1976).¹ This slow adoption is particularly puzzling given that new technology can provide a substantial boost to firm productivity (Syverson, 2011; Bloom, Eifert, Mahajan, McKenzie, and Roberts, 2013; Giorelli, 2019). There is also a second, well-documented puzzle: When major innovations such as information technology (IT) or electricity spread across firms, the widely expected boost in aggregate productivity has proved hard to document in the data. This prompted Robert Solow to remark in 1987 that “You can see the computer age everywhere but in the productivity statistics.”²

One prominent explanation for both puzzles is the need to modify and reorganize the production process when adopting major breakthrough technologies (David, 1990; Brynjolfsson, 1993; Brynjolfsson and Hitt, 2000; Hall and Khan, 2003; Brynjolfsson, Rock, and Syverson, 2018). Initially, many firms operate the new technology inefficiently – often because complementary organizational innovations are missing. If these challenges are indeed important during the early phase of technology adoption, we expect them to be reflected in a highly dispersed productivity distribution – due to a lack of standardized organizational knowledge that adopting firms can draw from. However, empirical evidence on this mechanism is scarce, as measuring the productivity distribution specific to new adopters is challenging for numerous reasons. First, standard data sources rarely make it possible to observe the use of specific technologies. Second, it is difficult to observe whether an adopting plant has also reorganized production. Third, productivity under the old and new technology are typically correlated.³

This paper shows how the need to reorganize production affects the productivity distribution of adopting plants in the short- and long-run. We bypass the typical challenges by studying a unique historical setting – the adoption of mechanized cotton spinning in France during the 19th century. Importantly, the macro-inventions that mechanized cotton spinning (the famous spinning jenny,

¹The more general observation that technology is often slow to diffuse is attributed to Rosenberg (1976). See Hall and Khan (2003) and Hall (2004) for surveys of the literature on technology diffusion. Comin and Hobijn (2010) document substantial lags in the adoption of new technologies, estimating that the variation in adoption lags across countries can account for at least one-quarter of per capita income disparities.

²New York Times, July 12, 1987, p. 36. Such productivity puzzles are not restricted to the introduction of computers: David (1990) documents similar trends following the diffusion of electricity earlier in the 20th century.

³See Brynjolfsson and Hitt (2000) for a discussion of some of these issues in the case of the IT revolution.

the water-frame, and the mule) went hand-in-hand with the need to reorganize production on a revolutionary scale. Prior to mechanization, workers produced in their homes in a cottage-industry setting. Adopting the new technologies required setting up factories from scratch, and moving workers there from their homes. This led to one of the most dramatic shifts in the organization of production in economic history (Mokyr, 2011).⁴ While the key elements of the new spinning technology itself were well-known across France (Horn, 2006), and multiple domestic producers supplied firms with the machinery (Chassagne, 1991), its adoption occurred in the absence of complementary knowledge on how the new cotton spinning plants should be organized (Pollard, 1965).

A number of features of this setting allow us to isolate the productivity distribution of adopters and study its evolution over a long time horizon (three decades), making headway on the typical empirical challenges faced by the literature. In particular, the sharp break in the *location* of production makes it possible to identify users of the new technology and verify that they had reorganized production in plants.⁵ This addresses the first two empirical challenges mentioned above. Moreover, we show that the first generation of mechanized cotton spinners did not typically have a background in the old technology, suggesting that productivity under the two technologies were not systematically related, which addresses the third challenge.

Our empirical analysis is based on a novel, hand-collected plant-level dataset from historical surveys covering three sectors (mechanized cotton spinning, metallurgy, and paper milling) at two points in time, around 1800 and in the 1840s. To help distinguish the effect of reorganization from broader trends such as general productivity growth, political and institutional change, or enhanced regional integration, we compare the evolution of the plant productivity distribution in mechanized cotton spinning to two comparison sectors (metallurgy and paper milling). This is similar in spirit to a difference-in-difference strategy. Crucially, in both comparison sectors, production was already organized in plants for centuries before the Industrial Revolution because of their reliance on water power and high-fixed-cost machinery. As a consequence, these sectors possessed fairly standardized knowledge about how to organize plant-based production. Moreover, during our sample period, all three sectors witnessed the arrival of new technologies that could be introduced fairly seamlessly into the *existing* organization of production.⁶ Thus, while all three sectors were adopting new technologies, only mechanized cotton spinning had to adapt to a radically reorganized production process.

⁴The new technology was organized in factories because of its reliance on inanimate power sources and the need to monitor workers more closely (Williamson, 1980; Szostak, 1989).

⁵Our data allow us to clearly distinguish plants operating mechanized spinning from earlier organizational forms, such as the pre-industrial putting-out system.

⁶In mechanized cotton spinning, there were improvements to the existing vintages of machinery and in preparatory processes (Allen, 2009). In paper milling, one part of the production process was mechanized (André, 1996) and in metallurgy, charcoal was replaced with coal as the source of fuel (Pounds and Parker, 1957).

We document four main findings for mechanized cotton spinning plants: 1) we observe a highly dispersed productivity distribution in the initial period (1806) relative to 1840; 2) we estimate that the industry underwent a substantial (82%) increase in plant productivity between 1806 and 1840 *after* mechanization had already been adopted; 3) this aggregate productivity growth was largely driven by the disappearance of plants in the lower tail of the distribution (which we refer to as ‘lower-tail bias’ of productivity growth); 4) the disappearance of the lower tail took place almost exclusively through plant exit and entry. Inefficient producers were replaced by more efficient entrants. In the comparison sectors, we also find a sizeable increase in plant productivity during the sample period (57% in metallurgy and 34% in paper milling). However, the lower-tail bias of productivity growth is unique to mechanized cotton spinning. In contrast, in the comparison sectors, the entire productivity distribution shifted right. Taken together, we interpret these results as suggestive of a link between the lower-tail bias of productivity growth and the feature unique to the mechanized cotton spinning industry – the need to reorganize production.

The second part of the paper examines why the need to reorganize production would lead to a lower-tail bias in productivity growth. Central to our argument is the fact that, at early stages of technology adoption, plants need to learn about optimal organizational forms. To fix ideas, we develop a simple framework in which plants learn about the optimal use of *multiple* inputs or tasks that in turn exhibit complementarities in the production function. We show that these features initially (when plants have little knowledge about the optimal ways to perform these tasks) lead to a fat lower tail in the plant productivity distribution. Over time, as plants learn about the efficient use of multiple inputs (tasks), the lower tail disappears. We present both historical and empirical evidence consistent with this.

According to the historical literature, there were two broad classes of challenges that early cotton spinning mills faced. First, they needed to contend with a range of issues related to mill layout and design. As Allen (2009, p.184) writes: “The cotton mill, in other words, had to be invented as well as the spinning machinery *per se*.” Second, a set of labor management innovations were required for setting up and operating spinning mills at a scale not seen elsewhere in the economy (Pollard, 1965).⁷

We examine the data for evidence consistent with the spread of organizational practices. We provide evidence for the spatial diffusion of knowledge during the early phase of industrialization, by showing that cotton plants located closer to high-productivity peers were themselves more productive. Strikingly, this spatial productivity pattern is not present in the comparison sectors, where plant-based production methods were more mature, nor in cotton spinning in the long-run,

⁷As Pollard (1965, p. 158) discusses, the managerial knowledge this required was to a large extent technical (i.e., how to coordinate and implement a division of labor given the specific production process) and for this reason, there was little scope to learn from other sectors. We test this formally with our data and show that there is no evidence that cotton spinning plants were learning from high-productivity plants in the comparison sectors.

once organizational knowledge had diffused. We show, using a rich set of controls and placebo tests, that these results are unlikely to be driven either by selection into productive locations or by omitted variables.

While these results are suggestive of a role for the spatial diffusion of knowledge, they do not distinguish between learning about the new technology itself (i.e., how to operate and maintain the new machines) and learning about efficient organizational forms. If plants were learning mostly how to operate the newly adopted technologies, we would expect incumbents to have an advantage relative to newer entrants. On the other hand, some aspects of organizational knowledge (such as mill layout and design) were costly to change once set up. Thus, if organizational knowledge diffused over time, we would expect later adopters to have an advantage in setting up their factories. Two features of the data point to the latter: First, cotton spinning plants that entered the market *later* had higher productivity during the initial phase of adoption. This holds even after controlling for newer capital vintages, and it does not hold for the comparison sectors or for cotton spinning in later periods. Second, the exit rate of plants in mechanized cotton spinning was substantially higher than in other sectors between 1800 and 1840, and buildings were also abandoned for use by the industry at higher rates. These findings suggest that knowledge of how to set up and organize cotton plants spread over time (and space), giving an edge to new entrants.

Finally, we examine alternative explanations that could account for our results. While our DID-style evidence based on the comparison sectors addresses many potential concerns, it is possible that some alternative channels affected mechanized spinning *differentially*. We control for a large set of these directly, showing that the lower-tail bias of productivity growth remains robust. For example, our results hold with region fixed effects, suggesting that the sorting of cotton plants into areas with better location fundamentals, better market access, or better input markets are unlikely to drive our results. Region fixed effects also make it unlikely that the Napoleonic Blockade, which had a regionally differential effect across France in mechanized cotton spinning (Juhász, 2018), drives our findings. We also account for plant-specific features such as size, quality upgrading, capital deepening, and life-cycle characteristics. Finally, our results hold when we exclude all plants that could be small ‘spinning workshops’ (which shared some, but not all characteristics of mature factory-based production).⁸

Related Literature and Contribution. Our paper is closely related to a literature on innovation and technology adoption in manufacturing – particularly the strand that has studied the productivity effects of the adoption of IT.⁹ Interestingly, some of the patterns we document for mechanized cotton spinning have been found in other settings. For example, Syverson (2011) discusses that the

⁸For a systematic overview of alternative mechanisms and the corresponding robustness, we refer the reader to the summary table in Appendix F.

⁹See Hall and Khan (2003) and Hall (2004) for an overview of the literature on technology diffusion. Brynjolfsson and Hitt (2000) and Syverson (2011) discuss the literature on the productivity effects of IT.

adoption of IT capital is associated with increased within-industry productivity dispersion. Foster, Grim, Haltiwanger, and Wolf (2018) provide empirical support for the argument by Gort and Klepper (1982) that periods of rapid innovation are associated with a surge in firm entry, followed by a period where experience with the new technology is accumulated, eventually leading to a shakeout where unsuccessful firms (or plants) exit. Brynjolfsson and Hitt (2000) conjecture that the surge in aggregate productivity in the 1990s was explained in part as a return on the large, intangible complementary organizational innovations that firms had undertaken in prior decades to make efficient use of IT. Our paper contributes to this literature in two ways. First, we provide evidence that these patterns are a more general feature of episodes of rapid technological change. Second, our unique setting allows us to more closely tie these patterns to the need to reorganize production that often accompanies major technological change. In particular, initial information disparities about the optimal organization of production can help to explain why breakthrough technologies tend to be adopted slowly and – even after being adopted – take time to be reflected in higher aggregate productivity. Along this dimension, our work relates to Atkin, Chaudhry, Chaudry, Khandelwal, and Verhoogen (2017), who also find that organizational barriers can impede technology adoption – albeit in the form of labor contracts.

In addition, our paper brings the insights of the firm productivity literature to the most important structural break in economic history – the First Industrial Revolution, which saw unprecedented growth in manufacturing productivity (Crafts, 1985; Crafts and Harley, 1992; Galor, 2011). So far, productivity growth during this period has been studied mostly at the country level, or – in some cases – at the aggregate sectoral level.¹⁰ Our paper is the first to study the contribution of plant dynamics to manufacturing productivity improvements during the Industrial Revolution. Our focus on the overall plant productivity distribution allows us to shed new light on how productivity growth evolved during this important period.¹¹ In particular, we isolate and track the productivity distribution of newly created adopting plants in cotton spinning. This goes beyond previous work (including with modern data), where major new technologies are typically introduced by existing producers, so that the changes in the productivity distribution reflect *both* the productivity differential of the new technology and subsequent gains due to organizational improvements. We are also the first to show that the extensive margin of plant entry and exit contributed decisively to productivity growth during the Industrial Revolution. Finally, we complement a rich historical literature by providing the first systematic empirical evidence for the importance of organizational innovations in driving productivity growth during the Industrial Revolution (Pollard, 1965; Sokoloff,

¹⁰Closely related, Clark (1987) studies cross-country differences in the productivity of mechanized cotton spinning, but only at the sectoral level. Braguinsky, Ohyama, Okazaki, and Syverson (2015) study the Japanese cotton spinning industry in the late 19th century and early 20th century. Rather than technology adoption, the paper focuses on the effects of acquisitions on acquired plants.

¹¹In related work, Braguinsky, Ohyama, Okazaki, and Syverson (2020) study how cotton firms grow by innovating vertically and horizontally in Japan’s Meiji period in the late 19th and early 20th century.

1984, 1986).

The paper is structured as follows. The next section discusses the historical context. Section 3 describes the data, while Section 4 presents and discusses our empirical results. Section 5 concludes.

2 Historical Background

Early nineteenth century France presents an ideal setting for our study of technology adoption. While England was the first country to industrialize, France was a close follower and structurally similar to England (Crafts, 1977; Voigtländer and Voth, 2006). The flagship inventions of the Industrial Revolution – most notably the spinning jenny – were developed in Britain. France adopted these widely during the first half of the 19th century (Nuvolari, Tortorici, and Vasta, 2021) and witnessed a similar acceleration in industrial output as Britain (Rostow, 1975).¹² By focusing on France, we thus study the effects of technology adoption in the context of an industrializing economy that was (at least initially) mostly adopting technology developed elsewhere.

In this section, we introduce the three sectors that we analyze in the paper, discussing the central features for the empirical analysis. We first examine the process of mechanization in cotton spinning and subsequent innovations during our sample period. We discuss the reasons why production had to be reorganized in the absence of organizational knowledge, and we show how the industry settled on, and disseminated best-practice organizational knowledge over a long time horizon. In the second part of the section, we introduce the two comparison sectors.

2.1 Mechanization and the Reorganization of Production in Cotton Spinning

Development of Mechanized Spinning and Subsequent Innovation. Cotton textiles was the flagship industry of the First Industrial Revolution, contributing one-quarter of TFP growth in Britain during the period 1780-1860 (Crafts, 1985). Cotton spinning is the process by which raw cotton fiber is twisted into yarn. Traditionally, this task was performed mostly by women in their homes, using a simple spinning wheel (see Figure A.1 in the appendix). With this old technology, each spinner was able to spin only one thread of yarn. The industry was rurally organized and generally centered around a local merchant-manufacturer who would supply spinners with the raw cotton, collect their output, take care of the marketing, and often also owned the spinning wheels (Huberman, 1996).

The breakthrough ‘macro-inventions’ in spinning were forged in Britain in the 1760s and 1770s, when three new vintages of machinery (the spinning jenny, the water-frame and the mule) were developed in quick succession (the left panel in Figure A.2 depicts the mule).¹³ These new

¹²Nuvolari et al. (2021) show that the French innovation system was capable of attracting and absorbing key British technologies. Below, we discuss the adoption of the most relevant technologies in our industries of interest. Appendix A.1 provides a more detailed discussion of the Industrial Revolution in France.

¹³Allen (2009) provides an in-depth discussion of each vintage.

machines made it possible to spin multiple threads simultaneously, as twist was imparted to the fibre not by using the workers' hands, but rather by using spindles. These innovations required production to move from workers' homes to the factory floor for two reasons. First, the machines typically used inanimate power sources leading to the concentration of production in one location.¹⁴ Second, mechanized production increased the need for monitoring workers as the machinery with which they worked was both more complex and more expensive (Williamson, 1980; Szostak, 1989; Mokyr, 2010).¹⁵

The productivity effect of these innovations was enormous. Allen (2009) estimates that the first vintage of the spinning jenny alone led to a threefold improvement in labor productivity. Correspondingly, the price of yarn declined rapidly in the late 18th century (Appendix Figure A.3), especially for the highest-quality yarn, where prices declined from 1,091 pence per pound to 76 pence per pound in real terms between 1785 and 1800 (Harley, 1998).

During the early decades of the 19th century, the industry was characterized by a steady stream of micro-inventions that improved these machines (Allen, 2009, p. 206).¹⁶ Evidence of this can be seen in Figure A.8 in the appendix, which shows a high level of British patenting in spinning – this sector was the third-most patent intensive out of 146 categories during this period. Importantly, the next major innovation, the self-acting mule (a completely new vintage of spinning machinery), did not spread widely until the 1840s, i.e., until after our sample period (Huberman, 1996). Thus, there was no major technology switching during our period of study, but rather a steady stream of inventions that improved existing vintages.

Adoption of Mechanized Spinning in France. Mechanized spinning was adopted with some lag in France. Efforts to adopt the technology had begun with state support during the *Ancien Régime*. By the beginning of our sample period in 1806, the large-scale expansion of the industry documented in Juhász (2018) had just begun. The technology was known throughout the country (Horn, 2006), and a number of domestic spinning machine-makers had been established (Chassagne, 1991).¹⁷ All three vintages (the spinning-jenny, the water-frame, and the mule) were used in France.

The challenging transition to factory-based production in cotton spinning. The transition of workers from their homes to the factory floor has been characterized as “one of the most dramatic sea changes in economic history” (Mokyr, 2010, p. 339). It fundamentally altered people's lives and, importantly for our setting, posed a host of challenges for the first generation of large-scale

¹⁴While the early spinning jenny was hand-powered, the water-frame and the mule relied on inanimate sources of power.

¹⁵This may explain why even spinning jennies were often housed in small workshops (Huberman, 1996).

¹⁶Our discussion focuses on mechanization of the spinning machinery. It is important to note, however, that preparatory processes were also mechanized in the late 18th century and improved upon throughout the 19th century. See Allen (2009) for a discussion.

¹⁷The spinning machinery itself was produced locally (using British blueprints) because of a ban on exporting machinery from Britain until 1843 (Saxonhouse and Wright, 2004).

factories.

In the case of cotton spinning, adopting the first generation of mechanized spinning machinery went hand-in-hand with the need to reorganize production as a plant. While cotton spinning was not the first sector to organize production in plants, the industry faced challenges for which a standard set of solutions did not exist at the time. Partly, this was because the knowledge required was largely technical and hence industry-specific (Pollard, 1965, p. 158). In addition, mechanized cotton spinning mills pioneered *flow production* – that is, the production of standardized goods in huge quantities at low unit costs by “arranging machines and equipment in line sequence to process goods continuously through a sequence of specialized operations” (Chapman, 1974, p. 470). This led to a finer division of labor and larger-scale plants than what had been seen before in other sectors, raising novel challenges (Chapman, 1974).

In what follows, we discuss the specific challenges of reorganizing production in mechanized cotton spinning. First, firms needed to contend with a range of issues related to mill layout and design. Allen (2009, p. 202) discusses some of the key challenges in developing the first mills in Britain: “[...] design issues emerged regarding the spatial location of the various machines, the flow of materials from one to the next, and the provision of power throughout a multi-story building.” The historical evidence suggests successful mill designs were observed and copied. Chapman (1970, p. 239) shows that early mills in England had a remarkably similar structure because plants quite literally copied the original design of the ‘Arkwright mills’ (the inventor of the water-frame). It took time for design defects to be improved; for example, contemporaries were aware of ventilation problems in the Arkwright-style mills, but continued to use the same layout regardless (Fitton and Wadsworth, 1958, p. 98). Moreover, fire hazards were a particularly pressing issue in the case of cotton spinning because of the highly flammable cotton dust (Langenbach, 2013). A process of trial and error eventually led to best-practice mill design that reduced fire hazards.¹⁸ Similarly, building structures needed to withstand the stress they faced from the vibrations of machines (Chassagne, 1991, p. 435). Iron rods with plates held beams to the masonry walls to prevent the vibrations of machines from shaking the walls apart (Langenbach, 2013).

Second, a large set of management innovations were required to run spinning mills efficiently at a scale not seen elsewhere in the economy.¹⁹ In the ‘Genesis of Modern Management,’ Pollard (1965, p. 160) described the development of efficient labor management practices as the primary management challenge facing early factories. There were three salient aspects of this for cotton spinning mills; i) how to get workers used to the independence of the domestic system to adapt to the rhythm and hierarchy of factory work, ii) how to coordinate and implement a fine division of

¹⁸We discuss the process of trial and error in minimizing fire hazards in cotton mills in Appendix A.3.1.

¹⁹Pollard (1965) and Mokyr (2010) contain a discussion of a wider set of challenges. Here, we focus on labor management challenges, which seem to have been the most important.

labor in a factory setting and, iii) how to solve monitoring problems.²⁰ Huberman (1996) estimates that it took two generations for efficient labor management practices to be developed in cotton spinning. Many firms in Britain initially experimented with crushing disciplinary measures for even the smallest infractions, which led to disastrously high turnover rates.²¹ Next, firms experimented with replacing male with female spinners in the hope that the latter could be more easily disciplined. Finally, around 1830, the industry settled on efficiency wages (Huberman, 1996).

The first generation of mechanized cotton spinners faced these design and management challenges all at once. Not only did it take time for best-practice solutions to emerge, but it also took time for this new body of knowledge to coalesce. It was not until 1835 that the first book on cotton factory management was published (Pollard, 1965). By the 1830s, the industry had reached maturity in Britain: “[in] a cotton mill [...] there was so much less scope for individual design, skill or new solutions to new problems, by 1830, at least, ... that little originality in internal layout was required from any but a handful of leaders” (Pollard, 1965, p. 90).

Isolating the Productivity Distribution of Adopters in Cotton Spinning. Mechanized spinning was operated in centralized locations (plants), while the old technology relied on home production. We can thus identify the users of the new technology (i.e., all *plants* observed in our data).²² Consequently, we are able to isolate the productivity distribution for plants that used the new, mechanized technology under the new organizational form and track the pattern of productivity dynamics along the entire distribution over a long time horizon.

The sharp break in organizational form under the new technology rendered experience with the old technology effectively useless. The type of skills necessary for successfully running mechanized cotton spinning factories were very different from those required under the old cottage-industry technology. Consistent with this, most owners who operated the new technology in France around 1800 did not have a background in handspinning. Table 1 presents information on the socioeconomic background of the owners of mechanized cotton spinning establishments for the early period of technology adoption (1785-1815). Based on these figures, the vast majority of owners were “traders, bankers and commercial employees” (62.5%). Within this category, most were in fact cloth merchants. In contrast, only a small fraction (10.2%) came from the production side, i.e., “workers and mechanics,” and three-quarters of this latter group were in fact highly skilled mechanics (Chassagne, 1991, p. 274). This highlights the importance of marketing skills (in the case of cloth merchants) or technical knowledge about the new technology (in the case of mechanics)

²⁰Appendix A.3.2 discusses each of these challenges in more detail.

²¹See Clark (1994) for a discussion of early factory discipline.

²²We follow Mokyr (2010, p. 339) in defining factory-based production as “the precise circumscription of work in time and space, and its physical separation from homes.” This definition is solely based on the *organization* of production; it does not rely on the use of inanimate sources of power. This is important because it allows us to refer to ‘factory’ or ‘plant’ production even when our data do not include specific information on power sources. Section 3.1 describes how we clean the data of a small number of observations that mix features of home and factory production.

as opposed to previous experience with handspinning in setting up cotton spinning factories. This suggests that the productivity under the old and new technology were not systematically related.

Importantly, there seems to have been little scope to bring industrial or managerial knowledge from other sectors that were already organized as plants, consistent with the notion that most knowledge was highly industry-specific. Of the entrepreneurs with a background in industry, most came from cotton printing (a downstream stage of production), and only a minuscule fraction came from other industries (Chassagne, 1991). This once again underscores the point that mechanized cotton spinners needed to develop the knowledge about running cotton mills from scratch.

Finally, the French needed to figure out many aspects of operating the technology efficiently themselves. This was partly because a lot of the learning that needed to be done was tacit (Mokyr, 2001, 2010), and partly because it wasn't until the 1830s that the British started codifying best practice in manuals.²³

The preceding discussion highlights the momentous challenges in reorganizing production in mechanized cotton spinning. To distinguish these from other, broader, trends at the time, we examine two sectors that did not need to reorganize production during this period – our ‘comparison sectors.’

2.2 Comparison Sectors: Metallurgy and Paper Milling

Metallurgy, the sector that supplied iron and steel to the rest of the economy, was a flagship industry of the Industrial Revolution. Paper milling – while not particularly important for other sectors – also underwent mechanization, which renders it a useful comparison sector. Despite the obvious differences in the production processes, metallurgy and paper milling share two important characteristics that differentiate them from mechanized cotton spinning.²⁴

First, both sectors had already organized production in plants since well before the Industrial Revolution. In metallurgy, plant production was mostly due to a reliance on high fixed-cost machinery such as the furnaces used both in smelting and refining. In paper milling, production was organized in plants because of the beating engine's reliance on water power.²⁵ The early start to plant-based production meant that these sectors had already accumulated industry-specific expertise in building design and labor management. For example, in paper milling, our detailed occupation data from 1794 shows that plants coordinated production across standardized occupational

²³Some British involvement in technology transfer took place *despite* the legal bans and the wars that occurred during this time period (Horn, 2006; Chassagne, 1991). However, while foreign technology transfer likely played *some* role in French learning, machines had to be produced domestically, workers were local, and some fundamental factors were different (such as the relative scarcity of coal).

²⁴We discuss the production process and technological change in the two comparison sectors in more detail in Appendix A.4.

²⁵The beating engine breaks down the raw input vegetable matter into cellulose fiber. The production process is described in more detail in Appendix A.4.

categories.²⁶

Second, while both were highly innovative sectors, in each, new technologies could be introduced within the existing organization of production.²⁷ The most prominent innovation in metallurgy was the switch from charcoal to coal, which could be introduced by modifying a plant's existing machines and ovens (see the illustration in Figures A.4 and A.5, and the corresponding discussion in Appendix A.4). In paper milling, the main technological innovation was the mechanization of forming paper with the Fourdrinier machine (one step in the production process). This invention still forms the basis of paper making today. Similar to metallurgy, this did not substantially alter the layout of the factory or other parts of the production process (see Figures A.6 and A.7, and the corresponding discussion in Appendix A.4).

In sum, metallurgy and paper milling underwent innovations similar to those observed in cotton spinning *after* the initial wave of mechanization, and integrating these did not require a substantial reorganization of production.

3 Data Construction

Our analysis is based on a novel plant-level dataset constructed from handwritten historical industrial surveys. The data have a panel-like structure covering three industries: mechanized cotton spinning, metallurgy, and paper milling. We observe plants in these sectors at two points in time: around 1800 and around 1840. Below, we discuss the main features of the data and the variables used in our analysis. We also examine key descriptive statistics important for the interpretation of our results. Appendix B provides a more detailed description of the data, including sources.

3.1 Industrial Surveys Around 1800

Our data from the turn of the 19th century are based on three industry-specific surveys that were conducted by the French government. The survey for paper milling was implemented in 1794 during the French Revolution; it contains data on 593 plants. The most important survey for our analysis – mechanized cotton spinning – was conducted by the Napoleonic regime in 1806,

²⁶See Appendix Table A.4 for occupational categories and their use across paper milling plants in France in 1794. Appendix C contains a more detailed discussion.

²⁷As Figure A.8 in the appendix shows, both industries witnessed the consistent arrival of patents during our sample period, although in both cases the level was somewhat lower than in spinning. In terms of patent intensity, metallurgy ranked ninth-highest, and paper milling twenty-first out of 146 categories during this dynamic period of innovation in Britain (as compared to spinning, which was third). It should be noted that this is a conservative comparison as spinning patents include those for all textile fibers, not just cotton.

covering 389 plants.²⁸ Finally, the survey in metallurgy in 1811 covers 477 plants.²⁹ Each of the three surveys provides hundreds of pages of handwritten returns that are available in the French National Archives in Pierrefitte-sur-Seine. Figures A.9, A.10, and A.11 in the appendix show sample pages from the three surveys. Although these data have not been systematically used for quantitative analyses,³⁰ the quality of French record-keeping in this period is well-known. The period is referred to as the “Golden Age of French Regional Statistics” (Perrot, 1975). Grantham (1997, p. 356) observes that “the quality equals that of any estimate of economic activity for a century to come.” Though the surveys were conducted at different points in time, we refer to their date henceforth as 1800.

Distinguishing Mechanized Cotton Spinning Taking Place in Plants. The rich data collected in the cotton spinning survey of 1806 allow us to identify production units that used the new, mechanized technology, organized in a central location, i.e., in a *plant* – as distinct from producers using handspinning technology. This distinction can be made with relatively high confidence because the 1806 survey specifically deals with *mechanized* cotton spinners. Thus, with the exception of a handful of cases, handspinning was not typically enumerated. We can identify these cases because establishments were asked about the vintage of mechanized capital that they used. We drop observations where handspinning wheels were reported. In addition, the 1806 survey also asked for the location of the plant. This helps us to filter out a few merchant-manufacturers who placed early vintages of the spinning jenny (i.e., a mechanized vintage) in workers’ homes. In these cases, the survey does not report one location for production, but many.³¹

²⁸In France, cotton spinning and weaving were generally not vertically integrated during this time period. Weaving, particularly in the early 19th century, was rurally organized. This implies less of an incentive to locate the workers in a common location, i.e. in a plant. Nevertheless, our dataset contains a few examples of vertically integrated spinning and weaving plants. We deal with these integrated plants in the following way. In the 1806 survey, enumerators were instructed to separately collect data for spinning and weaving activities (which is indicative of the lack of integration across these sectors in general). In the few cases where both took place under the same roof, we observe labor and output reported by activity and can thus estimate productivity separately for the spinning activities. In the 1840 survey (for which we only observe total labor and revenues), only 7% of plants that spun cotton yarn reported activities in both spinning and weaving. We follow the classification in Chanut, Heffer, Mairesse, and Postel-Vinay (2000) and use only plants that reported *exclusively* cotton spinning.

²⁹Bougin and Bourgin (1920) compiled an enormously rich overview of the metallurgy sector in 1788 using data from a wide variety of archival sources, including some recall data that was asked of plants in the 1811 survey. Unfortunately, since about 80% of plants do not report employment in Bougin and Bourgin (1920) for 1788, we cannot use these data in our baseline analysis. However, we do use the data as a validation check on plant survival.

³⁰The only exception that we are aware of is Juhász (2018), who uses the data from 1806 on the mechanized cotton spinning industry.

³¹Juhász (2018, Appendix pp. 21-22) contains a detailed description of the cleaning process for these data. Of the 626 entries, only 43 plants were dropped for these reasons. This is not because these types of production units were so insignificant, but rather because the survey was not designed to capture them, so they were typically *not* enumerated. Note that the number of observations in Juhász (2018) differs substantially from that used in this paper, because the former contains all plants active in the French Empire.

3.2 Industrial Census around 1840

The second period in our study is based on data from the first industrial census in France, conducted in 1839-47 and digitized by Chanut et al. (2000). For simplicity, we refer to these data as the ‘1840 census.’ While this census covers all manufacturing establishments, we only use data for cotton spinning (528 plants), metallurgy (839 plants), and paper milling (348 plants).³² Figure 1 shows the spatial distribution of plants around 1800 and in 1840 for the three industries.

The different surveys from around 1800 and 1840 contain remarkably rich information, although the exact set of variables varies from survey to survey. In what follows, we discuss only the variables used across all sectors, as well as additional, sector-specific variables when they are used in the empirical analysis.

3.3 Constructing Plant Productivity

Our main variable of interest is labor productivity measured at the plant level and defined as the log of revenues per worker. We use this measure in our baseline estimates because it can be constructed for all sectors and in both time periods. We will also examine TFP for mechanized cotton spinning, where this measure of productivity can also be constructed.³³

We face two challenges in constructing consistent productivity measures across plants and time. First, while the surveys for the three sectors around 1800 report output quantity (and some information on product-specific prices and quality), the census in 1840 directly reports plant-specific revenues (but not output quantities). To render productivity measures comparable over time, we have to construct revenues for 1800. Second, worker categories are not consistently reported across all plants in 1800 in metallurgy and paper milling. We discuss how we deal with each of these issues below.

Estimating Plant Revenues in 1800

In cotton spinning, the 1806 survey reports the quantity of yarn spun as well as the minimum and maximum count of yarn spun, where the count of yarn is the standardized measure of quality in the sector.³⁴ We construct plant-level revenue by multiplying the quantity of plant-level output by the price of the average quality of yarn produced by the plant. We use a schedule of prices

³²One potential concern with the 1840 census is that plants with less than 10 workers may have been systematically under-reported (Chanut et al., 2000). This is mostly relevant for paper milling, where plant size is the lowest. In robustness checks, we show that our baseline results hold even when only using plants with at least 10 workers in both periods. We also drop the handful of observations where multiple establishments were reported together in the survey.

³³Our revenue-based productivity measure reflects both product prices and quantities. It is thus potentially affected by changes in markups (Garcia-Marin and Voigtländer, 2019). However, this is unlikely to be quantitatively important because all three sectors in our analysis produced standardized, often intermediate products.

³⁴We use the (unweighted) average of the minimum and maximum count of the yarn produced by the plant as a proxy for its average output quality. The maximum and minimum count is the only information that plants provided on the quality of yarn that they produced.

for different counts of yarn reported by the French government.³⁵ In practice, the adjustment in price for the different qualities produced is not crucial for our results (which we confirm with an extensive battery of robustness checks). The reason for this is that the quality produced by the majority of firms is fairly similar. As Table A.2 in the appendix shows, the interquartile range for the average quality of yarn produced by the plants in our sample is 20 – 48. That is, most plants produced relatively low quality cotton yarn in this period (high count yarns typically start around 100), consistent with the British experience (Harley, 1998).

In metallurgy, the 1811 survey asked for the quantity of output produced (by product) as well as the price charged by the plant, by product type.³⁶ While product-specific output quantity is reported by all plants, the product-specific price is only reported by a subset of plants. We compute the average price for each product using the subset of plants where this information is available. We obtain plant revenues by multiplying product-specific plant output by the average price for each product and summing across products.

In paper milling, the 1794 survey reports the total quantity of paper products produced, but it does not provide plant-specific output prices. To construct revenues, we multiply plants' output quantity with the average price of paper products produced in the corresponding department, as reported in the *Tableaux du Maximum* – an extraordinary data source compiled in 1794 during the French Revolution that provides detailed data on goods prices and trade links across French regions. We use the department-specific price in order to accurately capture the product mix produced by plants across small geographies (see Appendix B.4 for detail). In robustness checks, we use the country-wide sectoral price.

Price deflators Finally, to compare revenues in the earlier periods and in 1840, for all three sectors, we deflate revenue data using the producer price index (PPI) for the respective survey years from Mitchell (2003). Appendix B.5 provides a discussion of these deflators. We note in passing that potential errors in the deflators would affect our estimates for *average* growth rates in the three sectors between 1800 and 1840, but they would not change the growth pattern across the plant distribution (e.g., the lower-tail bias in cotton spinning).

Constructing Consistent Labor Variables

In cotton spinning, the data provide consistent information on the number of workers employed by the plant.

In metallurgy, about 40% of the plants reported both 'internal' and 'external' labor in the 1811 survey, while the remainder of plants reported only total labor. Woronoff (1984, p. 138) describes

³⁵Source: Document number AN F12/533 from the French National Archives. In robustness checks, we use a single sector-level price, which we define based on the average quality of yarn produced across all cotton spinning plants.

³⁶The survey includes the following products: iron of first quality, iron of second quality, iron of third quality, steel using the cementation process, natural steel, and pig iron.

external labor as only having very loose ties to the plant. These workers did not typically work at the location of the plant, their work was not supervised by the manager, and their identity was often not even formally known to the manager. They performed tasks such as driving, collecting charcoal for the plant, or performing other jobs without belonging to the hierarchy or reporting to superiors in the chain of command. These types of workers were highly unlikely to be considered formal salaried employees of the plant in the 1840 census. The challenge is thus to construct a consistent measure of labor in 1811, given that approximately 60% of the observations report only total labor, with no indication of whether this includes external labor. For these plants, we need an estimate of the size of their internal labor force. We use a nearest neighbor matching algorithm to determine whether plants that only report total labor are more likely to be reporting internal labor only or the sum of internal and external labor.³⁷ When our algorithm suggests that the plant is reporting internal and external labor together, we estimate the number of internal workers by using the mean proportion of internal labor from all plants that report both types (the internal labor share is 20%).

In paper milling, the vast majority of plants only reported male labor in 1794. We impute the total number of employees in each plant by scaling male labor (reported by each plant) in 1794 by the average proportion of total employees to male employees in 1840 (where we observe both). The validity of this method hinges on the assumption that the proportion of male employees remained constant over time. We are able to check this using the subset of plants that report all types of workers in 1794. We find that the proportions are remarkably consistent.³⁸ Moreover, we show that our results are robust to using only male employees in both periods.

3.4 Linking Plants over Time

It is possible to link plants over time given that all surveys report the name of the owner and the location up to the commune, which is the lowest administrative unit in France.³⁹ We use two pieces of information to link plants over time. First, we match plants by their owner names in a given location in the respective industry.⁴⁰ Since the name of the owner may change even if the physical structure of the plant is the same, we also match by location in a second step. We match locations where there is *only* one plant in the respective sector in 1800 and where there is at least one plant active in the same sector in 1840. This turns out to be fairly common in the data. An obvious

³⁷We match each plant that reports only total labor to its nearest neighbor that reports internal and external labor, where matching is based on capital, output, and the stage of production. We then classify a plant as “reporting only internal labor” (“reporting total labor”) if its reported total labor is closer to the matched plant’s internal labor force (internal plus external labor force).

³⁸The proportion of total employees to male employees is 2.26 in 1794 for the subset of 20 plants that report all types of labor, while in 1840 it is 2.28 (among all plants).

³⁹In bigger cities such as Paris, the *arrondissement* is also reported.

⁴⁰We use a fuzzy string match to allow for differences in spelling as well as for different first names of owners, in cases where the plant was passed on within a family. Ambiguous matches were verified by hand.

concern is whether this ‘local matching’ indeed identifies the same plant. This is likely, given a fortuitous feature across all three of our industries: their reliance on water power. Only a small number of locations in a particular commune were suitable for setting up a water-powered mill, as rapid stream flow was needed to yield sufficient power. Moreover, the backwater created by one mill meant that another mill cannot be located in close proximity. Consistent with this, [Crafts and Wolf \(2014\)](#) argue that agglomeration in the cotton textile industry was not observed until steam became the common source of power in Britain. Consequently, our ‘local matching’ arguably identifies plants that have the same location within communes. Whether these were owned by the same entrepreneur (or their descendants), or whether they had passed on to a different owner is not crucial for our analysis.

One way to validate the assumptions underlying our ‘local matching’ is to examine how frequently communes with a single plant active in the sector in 1800 show up in 1840 with multiple plants active in the same sector. If this occurs frequently in the data, it would suggest that in fact there are multiple suitable locations for production in that sector for a particular commune. This is not the case in our data. For the vast majority of single-plant communes that we identify in the initial period, there either continues to be one plant or no plants in the subsequent survey. It is exceedingly rare (6% of cases in both paper milling and cotton spinning, and 8% in metallurgy) across all three surveys for single plant communes to ‘add’ additional plants (despite the large increase in the overall number of plants in metallurgy and cotton spinning). As an additional check on our methodology, it is also possible to compare plant survival in metallurgy to that reported in [Woronoff \(1984\)](#) for this sector over the period 1788-1811. If our strategy of ‘local matching’ led to too many plant matches over time, we would expect an exaggerated survival rate. The contrary is true. Our estimates of plant survival rates for the period 1811-1840 are well below (one half or less) those that we calculate for the period 1788-1811. This suggests that it is unlikely that we systematically overestimate plant survival.

3.5 Plant Survival Rates

Our main measure of plant survival is based on the combination of matching by owner name and ‘local matching’ that we described above. We define the survival rate as the percentage of plants from the initial period that survive into the later period. The numerator counts all plants that fulfill at least one of the following two conditions: i) the plant has the same owner in both periods; ii) there is only one plant in the respective sector in the location in the initial period and *at least* one plant in the same sector in 1840. The denominator is the sum of *all* plants in the given sector in the initial period. Note that this rate does not adjust for the fact that the number of plants located in communes that have only one plant varies across the three sectors in our sample.⁴¹ Thus, we

⁴¹ Among the 593 plants in paper milling in 1794, 218 (36.8%) were the only plants active in their commune in this sector. For cotton spinning in 1806, the proportion is 25.4% (99 out of the 389 plants), and in metallurgy in 1811, 69%

may mechanically find higher survival rates in a sector where single-plant communes are relatively more frequent. To address this issue, we also construct the ‘restricted sample’ survival rate as a robustness check. This measure is based solely on single-plant locations. The numerator counts the number of communes that had only one plant in the respective sector, in both the initial period and in 1840 (indicating plant survival). The denominator counts the number of communes that had a single plant in the respective sector in the initial period and either one or no plant in 1840 (indicating plant survival and plant exit, respectively).⁴²

3.6 Descriptive Statistics

Before moving to the empirical analysis, we note several important features of the data apparent from the descriptive statistics.⁴³ First, the scale of plants (measured by the number of employees) is striking for cotton spinning plants (see Table A.3 in the appendix). The average spinning plant in 1806 had 63 employees. This is in contrast to metallurgy and paper milling that had on average 20 and 13 employees respectively. This lends additional support to the argument that mechanized cotton spinners in France were facing labor management challenges that had not been encountered on a similar scale previously. Second, between 1806 and 1840, the mechanized cotton spinning industry expanded substantially, and this expansion was accompanied by an *increase* in the number of plants active in the market (from 372 in 1806, to 528 in 1840). That is, for every plant that exited the market, more than one new plant entered. As such, the results we present constitute more than a ‘shake-out’ of unsuccessful plants.

4 Empirical Analysis

In this section, we study the evolution of the plant productivity distribution in mechanized cotton spinning after the new technology had been adopted. Similar in spirit to a difference-in-difference estimation strategy, we contrast the observed patterns with those in the two comparison sectors – metallurgy and paper milling. This allows us to distinguish the unique feature in mechanized cotton spinning – the need to reorganize production – from other common factors that affected productivity growth in all three sectors.

4.1 The Pattern of Productivity Growth

We begin by examining average annual labor productivity growth. Column 1 in Table 2 shows that all three sectors experienced a significant increase in labor productivity between 1800 and 1840. This is consistent with the historical evidence that points to important innovations across all three

(329 out of 477 plants).

⁴²Based on this sample definition, we exclude plants that were the only ones in their commune in 1800, and where there was more than one plant in 1840. As discussed above, the number of these “uncertain” observations is very small across all sectors, which we consider a validation of our methodology.

⁴³See Appendix C for a more in-depth discussion.

sectors during this time period. The largest productivity gains were achieved in cotton spinning (2.4% per year), followed by metallurgy (1.9%) and paper milling (0.7%).⁴⁴ It is noteworthy that the large productivity increase in spinning reflects improvements *within* the new technology (as opposed to new vs. old technology) because all plants in this sector already used mechanized spinning in 1800.

In which part of the productivity distribution were these gains concentrated? Figure 2 plots the distribution of labor productivity in the three sectors at the beginning and at the end of our sample period, illustrating our main results. In cotton spinning, two features stand out. First, the initial dispersion in labor productivity was large in 1800 relative to that in 1840. Second, the productivity gains are almost exclusively concentrated in the lower tail – the lower tail disappeared over our sample period, while increases in productivity at the upper tail were modest. The contrast between cotton spinning and our two comparison sectors is striking. In metallurgy and paper milling, the entire productivity distribution shifted to the right between 1800 and 1840. Quantile regressions confirm this pattern. Columns 2-6 in Table 2 report these results for the three sectors, estimating regressions for productivity growth at different quantiles of the productivity distribution. Figure 3 displays the corresponding coefficients. In cotton spinning, the bias towards productivity growth in the lower tail is marked. Productivity growth at the 25th percentile was twice as large as that at the 75th percentile (3.3% per year relative to 1.65%), and the difference is fourfold between the 10th and the 90th percentile (3.9% and 1.0%, respectively). In the comparison sectors, the differences are more modest across the distribution. If anything, growth was concentrated in the upper tail. In both metallurgy and paper milling, the productivity growth at the 25th percentile was marginally *lower* than at the 75th percentile.

In summary, these set of results show that *after* the adoption of the new technology in mechanized cotton spinning, the industry witnessed major increases in productivity that were driven by a disappearance of the lower tail of the productivity distribution. Over time, the plant productivity distribution became less dispersed. This is in striking contrast to the patterns observed in the comparison sectors, where productivity growth occurred relatively evenly across the distribution. This constitutes suggestive evidence that the aspect unique to mechanized spinning – the need to reorganize production – was associated with the lower tail bias in productivity growth.

Output quality. Could the different pattern in cotton spinning be driven by output quality? Recall

⁴⁴Given that we discount revenues using price indices, all our productivity calculations reflect price-adjusted revenue-based productivity. To obtain the average annual growth rates between the two time periods, we first regress log output per worker $\ln(Y/L)$ on a dummy for 1840 in each sector. This coefficient measures the percentage growth in output per worker over the entire time period between the respective survey years. We then annualize these values (and corresponding standard errors) by dividing by the number of years between the surveys in each sector. Note that this method delivers average annual growth figures, not accounting for compound growth. In cotton spinning, the overall growth over the period 1806-40 amounts to 82% (2.42% per year x 34 years). In metallurgy, it is 57% (1.95% per year x 29 years), and in paper milling it is 34% (0.73% per year x 46 years).

that our data enable us to use quality-adjusted prices in 1800 for cotton spinning, while quality adjustments are not possible in the two comparison sectors. Panel B in Table A.5 in the appendix presents quantile regressions without quality-adjustments, i.e., using the same sector-level price across all plants in cotton spinning. The magnitude of the lower-tail bias is slightly smaller, but it remains striking. The difference in annual productivity growth between the 10th and the 90th percentile is 3.4% relative to 1.6%. Quality adjustment does not substantially alter the pattern of productivity growth because most plants produced yarn of a fairly similar (low) quality (see Table A.2).

Mark-up heterogeneity. Our productivity measure in 1800 is computed from plant-specific physical output (adjusted by sector-level prices), while the 1840 values are based on plant-specific revenues. Thus, the latter may also reflect markup heterogeneity across plants. To the extent that heterogeneity in markups leads to a more dispersed productivity distribution, this would only be a concern in 1840. That is, this data limitation – if quantitatively important – would work against our finding of a tightening productivity distribution.

Robustness to measuring productivity as TFP. Our baseline productivity measure is log output per worker. For cotton spinning, we can also compute TFP, using detailed data on physical machinery (number of spindles) – see Appendix E for detail. Panel C in Table A.5 confirms the lower-tail bias of productivity growth in cotton spinning using TFP.

Robustness to data construction choices in paper milling. Table A.6 in the appendix presents additional robustness checks on data construction choices we made in the paper milling industry, where we faced the most challenges in constructing comparable productivity measures. We show that our quantile regression results in paper milling are robust to i) using only male labor and ii) to pricing output in 1794 using the country-wide average price as opposed to the departmental price.

4.2 Mechanism: Learning About Best Practices in Factory-Based Production

Both the historical evidence presented in Section 2 and the different pattern of productivity growth found in mechanized cotton spinning suggest that the need to reorganize production may be driving the lower tail bias of productivity growth in mechanized cotton spinning. In this section, we first discuss a stylized framework that captures the key elements of the historical evidence. Then, we turn to examining the data for evidence compatible with the learning effects that reorganizing production would entail.

A Stylized Framework. We summarize the key features here; Appendix D presents the model and shows simulation results. The essential ingredients are: i) a production function that involves multiple *complementary* inputs (tasks); ii) independent, random draws in the efficiency associated with each input; and iii) the lower bound of each input’s efficiency draw increases over time (i.e., very bad draws disappear).

In cotton spinning, initially, very low input efficiency draws are possible. The complementarity across inputs implies that a low efficiency draw for only one input diminishes output substantially.⁴⁵ This gives rise to a fat lower tail of the productivity distribution. Over time, we let the lower bound for the input efficiency draws in cotton spinning increase. This reflects improved knowledge about organizing and handling each input – either due to learning within plants or due to knowledge diffusion across plants.⁴⁶ At the same time, we assume that the upper bound for input efficiency draws does not change, i.e., the technological frontier in mechanized cotton spinning remains the same. Intuitively, learning (within plants or from other producers) eliminates the worst mistakes in organizing production; but the best possible draws remain unchanged because the underlying technology does not change. This setup leads to productivity growth concentrated in the lower tail. Figure A.12 in the appendix illustrates the simulated productivity distributions.⁴⁷ The left panel of Figure A.12 shows the resulting evolution of the plant productivity distribution over time. The initially fat lower tail disappears in the second period – intuitively, very bad draws that pull down output towards zero (even if all other draws are high) have been eliminated. This productivity pattern mirrors the one for cotton spinning in the data (see Figure 2).

We also adapt our simulation to the comparison sectors. Since plant-based production in these sectors was already well-established around 1800, we choose a higher lower bound to reflect that learning had already occurred. To represent the technological progress that improved these technologies over time, we shift both the upper and lower bound outward in the second period. The corresponding simulation results are shown in the right panel of Figure A.12. They resemble the pattern observed for our control sectors, metallurgy and paper milling, where the whole productivity distribution shifted to the right.

While this stylized theoretical framework is not the only one that gives rise to lower-tail bias in productivity growth, it is a simple setup that represents the key features of the historical context. In what follows, we use this stylized framework to guide our discussion.

Spatial Diffusion of Knowledge. First, we shed light on *how* learning about organizational practices took place. The historical background discussed in Section 2 showed that plants copied successful designs and setups of the production process from each other. To examine this channel, we estimate whether a plant’s own productivity was higher in the proximity of other high-productivity

⁴⁵In the extreme complementarity case – a Leontief production function – output drops in proportion to the minimum efficiency draw, no matter how large the draws for the other inputs are.

⁴⁶While our stylized theory is agnostic as to which of these mechanisms dominated, the empirical results below suggest that the diffusion of organizational practices *across* plants was an important dimension.

⁴⁷We use a CES production function with three inputs and choose the elasticity of substitution across inputs to be 0.5, indicating a strong degree of complementarity. Efficiency in each of the inputs is drawn from a uniform distribution with support $[0, 1]$. Over time, the lower bound for each distribution increases to 0.3.

plants. We use the following specification:

$$\ln(Y/L)_{ij} = \beta_0 + \beta_1 Dist_{ij}^{p90} + FE_j + \epsilon_{ij} ,$$

where $\ln(Y/L)_{ij}$ is labor productivity (log output per worker) for plant i located in department j ; $Dist_{ij}^{p90}$ is log distance to the nearest plant (in the same sector) with productivity in the 90th percentile (in the distribution of *all* plants in the sector across France). Plants that are themselves in the top productivity decile are excluded from the sample to avoid introducing a mechanical relationship. Our preferred specification includes department fixed effects (FE_j) to absorb unobserved location characteristics that may make all plants in a given area more productive, irrespective of local spillovers. Thus, the coefficient of interest, β_1 , reflects the extent to which plants in the same department benefit from being located closer to a high-productivity plant (which may be located in the same or in another department). We do not interpret these correlations as causal effects, but as evidence that is compatible with spatial spillovers of knowledge. We estimate the specification separately for the three sectors, and in both time periods. Standard errors are clustered at the department level to account for spatial correlation.

Before presenting the results, we first examine the spatial distribution of high-productivity plants across our sectors and time periods. Figure A.13 in the appendix plots the spatial distribution of cotton spinning, metallurgy, and paper milling plants, distinguishing those in the 90th percentile of the productivity distribution. Unsurprisingly, some regions have a larger concentration of high-productivity plants than others. Due to our use of department fixed effects, these regional differences do not affect our results.

Figure 4 visualizes our baseline results, and Table A.7 in the appendix reports the corresponding regressions. To allow for direct comparability, we report standardized beta coefficients of $Dist^{p90}$ for all three sectors in the two periods. The estimated coefficient for cotton spinning in 1806 is negative, statistically significant and large in magnitude. A one standard deviation (std) increase in distance to a high productivity plant is associated with a 0.84 std decline in labor productivity. The pattern is much weaker in the two control sectors – the coefficients are less than one-third in magnitude as compared to cotton spinning. In addition, in 1840, there is at best a muted relationship between labor productivity and distance to top-plants in any of the three sectors.⁴⁸ Thus, proximity to high-productivity plants mattered the most in cotton spinning in 1806, i.e., in the period before knowledge about the optimal organization of production had spread widely. While this pattern is consistent with a learning mechanism, it could also be driven by a number of other forces. We turn to examining possible alternative explanations below.

⁴⁸The coefficient for mechanized cotton spinning is reduced to a quarter of its initial size, and it is significant only at 10%. In the other sectors, the coefficients of interest are reduced slightly, and they are no longer statistically distinguishable from zero.

While department fixed effects capture unobserved differences that vary at the department level, they cannot account for unobserved differences at a finer spatial level. To address this issue, we use four approaches. First, we control directly for prominent location fundamentals at the commune level such as the availability of fast-flowing streams (as a source of water power), proximity to coal (which matters for steam power), and the share of forest cover (which matters for access to charcoal – a major input in metallurgy).⁴⁹ Table A.8 shows the results for specifications that control for all of these location fundamentals. The pattern of the coefficients of interest does not change, and the estimated magnitudes remain very similar. Moreover, the location fundamentals themselves are mostly small and statistically insignificant. This is probably driven by the fact that the department fixed effects already account for the most important spatial differences in location characteristics. Second, our results could be affected by more general agglomeration externalities, as opposed to learning. In particular, our findings may be driven by high-productivity plants emerging (within departments) where the density of production is large due to agglomeration forces. We investigate this explanation by adding a control for the density of production at the commune level (measured as the log of total output in the sector, excluding a plant’s own output). Table A.9 in the appendix shows that controlling for the local density of production barely affects our results. The estimated coefficient on distance to high-productivity plants in cotton spinning in 1806 decreases slightly to -0.77 (se 0.18), while the distance coefficients in the other sectors and in 1840 remain relatively small. The coefficient on local production density itself is generally small, positive, and never statistically different from zero.

Appendix Table A.10 presents our third approach to probe for unobserved location fundamentals (within departments). We conduct a placebo exercise that examines whether plant productivity in 1800 was also related to the distance to high-productivity plants that only emerged later, i.e., plants in the top-90th percentile of productivity in 1840. Reassuringly, the estimated coefficient in cotton spinning is close to zero and statistically insignificant. In other words, productivity in cotton spinning in 1806 was not related to high-productivity locations three decades later. This suggests that the large estimated coefficient in our baseline specification is not driven by persistent location fundamentals within departments.

Fourth, we examine the extent to which the estimated distance coefficient in cotton spinning in 1806 may be driven by plant selection. It is possible that we estimate a large (negative) coefficient in cotton spinning in 1806 not because plants were learning from their high-productivity neighbors but rather because *ex-ante* high-productivity plants selected into ‘productive locations’ (i.e., chose

⁴⁹The sources for these data are as follows. Data on the stream flow of rivers are from EURO-FRIEND (<http://ne-friend.bafg.de/servlet/is/7397/>). These data report water flow rates from thousands of geocoded collection points across France. Data on charcoal sources are from the highly detailed ‘Cassini maps’ produced in the late 18th century that contain information on forest cover. These maps were geo-referenced by Vallauri, Grel, Granier, and Dupouey (2012). Finally, data on the location of coal deposits – both in France and near its borders – were geo-referenced from maps in Tarr and McMurry (1993) and Guiollard (1993).

to locate near existing high-productivity plants). Given that we observe plant age in cotton spinning in 1806, we can examine selection patterns. In Table 3, we show that our coefficient of interest is robust to controlling for log plant age (column 1) and to adding the interaction of plant age and distance to high productivity plants (column 2). This suggests that (conditional on department fixed effects) plant entry did not vary systematically with the location of high-productivity plants. A different, arguably more conservative approach is to estimate the coefficient of interest using only the subsample of plants that entered *before* the nearest high-productivity plant. The timing of entry of these plants rules out the type of selection described above. In column 3 of Table 3, we show that the coefficient on distance to high-productivity plants remains statistically highly significant, although it is somewhat smaller than in the baseline sample (-0.43, se 0.14). Columns 4 and 5 show that the coefficient of interest remains stable and statistically significant as controls for plant age (and its interaction with distance to high-productivity plants) are added in this subsample.

As a final placebo check, we examine whether there is evidence for similar types of learning *across* sectors. In particular, we estimate if proximity to high-productivity plants in the comparison sectors also mattered for mechanized cotton plants in 1806. Table A.11 in the appendix shows that there is no consistent pattern in the data. For metallurgy, the coefficient is negative and not statistically significant.⁵⁰ For paper milling on the other hand, the distance coefficient is actually *positive* and significant – that is, high-productivity paper mills crowded out productive cotton mills. A possible reason is that newer mechanized cotton spinning mills were competing for the same type of locations as paper milling plants (with access to fast-flowing streams). The cross-sector results are consistent with the historical record, showing no indication that early cotton spinning mills were able to learn from high-productivity plants in more mature sectors.

In summary, the consistently larger distance coefficient estimated in cotton spinning in 1806, in combination with a series of robustness checks, points to the spatial diffusion of knowledge as one mechanism through which learning across plants took place. However, these results cannot distinguish between the spatial diffusion of knowledge about the technology itself (for example, how to operate and maintain the machines efficiently) and organizational knowledge (i.e., how to design mills and organize workers within the plant). The next two exercises examine our data for patterns that would be more consistent with one or the other form of learning.

Plant Survival Across the Sectors. If an important component of learning in mechanized cotton spinning occurred along the organizational dimension of factory design, we would expect initially low plant survival rates relative to the other two sectors. Building design and layout are either sunk at the time of building or costly to change, implying that plants who got this wrong were

⁵⁰However, the coefficient is estimated with substantial noise. A likely reason for the noisy results is that metallurgy and cotton spinning mills were located relatively far from each other (see the maps in Figure A.13). The median cotton spinning mill was located 113 km from its nearest high-productivity peer in metallurgy, but only 58 km (68km) from the nearest high-productivity plant in paper milling (cotton spinning).

likely to exit the market. On the other hand, the inefficient operation of the new technology itself could be adjusted within an existing factory so that if anything, we would expect incumbents to have an edge over new entrants. The data speak in favor of the former. We find substantially larger exit rates in cotton spinning relative to the other two sectors. Table 4 reports plant survival rates over our sample period, using the two measures defined in Section 3.5 in each of the three sectors. Based on our baseline measure, survival rates in spinning (7%) were slightly lower than in paper milling (9%) and much lower than in metallurgy (34%). Note that the paper milling survey was conducted in 1794, more than 10 years earlier than the cotton spinning survey (1806). If we adjusted for the longer horizon in paper milling, the implied survival rates for 1806-1840 would be significantly higher. This implies that the differences in survival rates between cotton spinning and paper milling are probably higher than reflected in Table 4.

The ‘restricted sample’ survival rates test even more directly the extent to which building layout and design were key organizational challenges. By using only single-plant locations, we are in effect testing whether a building used for production in a particular industry in 1800 continued to be used in the latter period in the same industry (irrespective of who the owner was). The differences across the three sectors are even starker. By this measure, the survival rate in spinning is still about 7%, but it is much higher in the comparison sectors: 20% in paper milling and 49% in metallurgy. Thus, the low survival rate observed in cotton spinning means that many locations lost their (only) cotton mill. Owners that invested in a cotton spinning mill with poor layout had to exit the market, and the structure of the mill was not subsequently used by other firms in cotton spinning.⁵¹

Table 5 examines the extent to which plants that eventually exited the market by 1840 had lower initial productivity around 1800, as compared to surviving plants. Consistent with large exit rates, this pattern is strong in cotton spinning. Exiting plants were 53% less productive than survivors, and this difference is highly statistically significant. This pattern is much less pronounced in the comparison sectors: exiting plants were about 15% less productive in both sectors, and this number is not statistically different from zero.

In summary, plant survival rates show a pattern that is more consistent with learning about best practice organizational challenges. In particular, the evidence presented here is consistent with first-generation plants facing important design challenges. Those that got it wrong were particularly unproductive, leading to a highly dispersed productivity distribution. These plants eventually exited the market, and for a subset of our sample, we can confirm that the same building

⁵¹Lower survival rates in spinning could also be consistent with the fact that the industry moved towards steam power to a larger extent than other sectors and hence moved away from water power. However, note in Table A.14 that even in spinning, steam power seemed to be used *in addition to* water power: 66% of cotton spinning plants still used water power as a source in 1840. Moreover, cotton mills using steam power were somewhat less productive (see Table A.15), suggesting that in France plants did not face a strong profit incentive to move away from water power.

was not used by another plant in the industry.

Age Profile of Plant Productivity. We now examine whether the age profile of plant productivity is more consistent with learning about technology or best-practice organizational methods. If learning about the technology is the dominant dimension, we expect older plants to have accumulated more experience and hence have a productivity advantage. On the other hand, if organizational practices are more important, younger plants have a larger pool of knowledge to draw from when setting up the plant and deciding on mill design and the corresponding division of labor.

We exploit the richness of our data to test this in both 1800, when best practice mill design was still evolving, and in 1840, when according to Pollard (1965), the industry had reached maturity – at least in Britain. The 1806 survey in cotton spinning contains the year of foundation of plants. This allows us to compute a dummy for ‘young’ plants, defined as below-median age (with the median age in 1806 being three years). We first examine whether plant age is systematically correlated with productivity in cotton spinning. Column 1 in Table 6 shows that the unconditional association is strongly positive: ‘Young’ plants were 58% more productive in 1806. This could be driven by mechanisms other than the one we have in mind. For example, new entrants may have used the most recent vintage of capital, leading to higher physical productivity (Foster, Haltiwanger, and Syverson, 2008). To address this issue, we control for several important plant characteristics in columns 2-7 of Table 6. These include the average quality (count) of the yarn spun, the capital intensity of the plant (measured as log spindles per worker), the number of workers in the plant, and the vintage of machinery (a binary variable for using different vintages of machinery – these are not mutually exclusive categories)⁵² The coefficient of interest remains large and highly significant when we add these controls one-by-one.⁵³

Next, we turn to the productivity-age pattern in 1840. While the data for this second period are generally more comprehensive, we do not observe plant age. However, we can perform a similar – albeit weaker – test based on the comparison of surviving and entrant plants: In Table 7 we regress log output per worker on an indicator for whether the plant was an ‘entrant’ in 1840 (as opposed to a surviving plant by our definition from Section 3.5). The coefficient on the ‘entrant’ dummy thus reflects the average productivity differential for plants that entered between the initial survey year (1806) and 1840. Best-practice mill design evolved over this period, and it had largely converged by 1840 (Pollard, 1965). Correspondingly, we find that ‘young’ plants were not more productive; the coefficient is in fact negative, but not statistically different from zero. This holds also when we control for the use of water power, steam power, any other power source (wind or animal power used by a small subset of plants), and for the number of workers.

For comparison, we perform an additional exercise in the appendix, using data on metallurgy

⁵²The three different vintages of machinery are the spinning jenny, the water-frame (throstle), and the mule-jenny.

⁵³We do not include all controls together because of multicollinearity concerns.

plants in both periods (Tables A.12 and A.13). For this sector, in both periods, the best measure of plant age that we observe is a binary indicator of plant survival from 1788 to 1811, and for survival from 1811 to 1840 (the latter being the same procedure as for cotton spinning in Table 7). Overall, the results do not point to younger plants in metallurgy having a strong productivity advantage: In 1811, ‘young’ metallurgy plants were only marginally more productive, and in 1840 ‘young’ metallurgy plants were somewhat *less* productive.⁵⁴

In summary, the evidence presented in this section is in line with plants experimenting and learning slowly about best-practice organizational methods. Knowledge seems to have diffused spatially. The exit patterns and age profile of productivity are more consistent with learning about organizational practices as opposed to learning how to operate the new technology itself.

4.3 Robustness to Alternative Explanations

In the final part of this section, we consider additional alternative mechanisms that could also explain our results. We examine these along a number of different dimensions.

4.3.1 Common Patterns across all Three Sectors

First, recall that we observe the lower-tail bias in productivity growth only in cotton spinning, and not in the other two sectors. For this reason, it is unlikely that factors that also affected the comparison sectors in similar ways can explain our findings. For example, the fact that mechanized spinning experienced innovations during our sample period seems unlikely to explain the lower-tail bias, as both comparison sectors also witnessed the introduction of new technologies. In a similar vein, improvements to power sources (notably water power, which remained the dominant source of power in cotton spinning) affected the other two sectors similarly.⁵⁵ Finally, all three sectors used a mix of different vintages of capital, suggesting that the choice of capital in itself is not likely to account for the lower-tail bias of productivity growth in cotton spinning. In fact, below, we show that the lower-tail bias of productivity growth remains intact if we remove all plants using the earliest vintage of the spinning jenny.⁵⁶

⁵⁴In metallurgy in 1811, ‘young’ plants were 22% more productive when no controls are added (as compared to almost 60% for the same specification in cotton spinning). Once we control for plant size (log workers) in column 2 of Table A.12, the productivity advantage disappears. Note that the metallurgy survey has sparser information on this dimension; plant size is the only control that can be added in 1811. For paper milling, information on plant age or recent entry is not available for the early period. We thus cannot perform the comparison for this sector.

⁵⁵Appendix Table A.14 shows that all three sectors relied heavily on water power in the 1840s.

⁵⁶In this regard, the metallurgy sector is a helpful comparison as we observe the use of different vintages of capital in 1811. Some plants used the so-called ‘direct technology’ widely known in France as the Catalan forge, while others used the ‘indirect technology’ which separates the production process into smelting and refining. The direct technology was gradually replaced by the indirect technology during our sample period (Pounds and Parker, 1957). See Appendix A for a discussion.

4.3.2 Regional Differences

It is important to highlight the robustness of our results to using only within-region variation. Table A.16 shows that the lower-tail bias of productivity growth in mechanized cotton spinning remains intact if we add fixed effects for 22 French regions.⁵⁷ The lower-tail bias of productivity growth is more muted, but still striking. Productivity growth is almost twice as high at the lowest relative to the highest decile.

This makes it unlikely that the sorting of plants into areas with better location fundamentals or different growth potential over time is driving our results. Moreover, as plants within regions arguably had access to similar input markets (for raw cotton, machinery, or labor), it is unlikely that these types of forces account for our findings. In addition, the early 19th century was a turbulent period of French history, and some parts of France were more affected by wars and revolts than others (we deal explicitly with the various effects of the Napoleonic Blockade below). The inclusion of region fixed effects once again absorbs the possible differential exposure of French regions to these upheavals and institutional changes. In what follows, we examine other prominent mechanisms that may have played a role at a finer geographic level, i.e., potentially even within regions.

Market Integration. Could increased market integration explain our results in cotton spinning? As the French economy became more integrated over time, it is possible that lower-productivity plants faced tougher competition and had to exit the market.⁵⁸ However, we do not observe a lower-tail bias in productivity growth in the two comparison sectors. Thus, for market integration to explain our results, it must have affected the cotton textiles sector differentially.

The historical background and available data speak against such a mechanism. Cotton yarn (and textiles more generally) are high value-to-weight products, which makes them more easily tradable over long distances than iron and steel or paper. This suggests that cotton spinners may already have faced tougher competition through more integrated markets in the early 1800s. Thus, we would not necessarily expect cotton spinning to be the most affected sector by increased market integration after 1800. Consistent with this reasoning, we present evidence for relatively high market integration in cotton yarn in the late 18th century. We use data in 1794 from Daudin (2010) on the number of districts across France that reported consuming cotton textiles, iron, or paper products from any district in a given department.⁵⁹ Intuitively, higher market integration

⁵⁷Regions are larger than departments: There are 22 regions in France and 86 departments. Our data in the three sectors do not have sufficiently many observations to include *department* fixed effects, i.e., to estimate meaningful productivity distributions within departments.

⁵⁸Market integration arguably increased during our sample period both for policy reasons such as the abolition of internal barriers to trade during the French Revolution (Daudin, 2010), and because of infrastructure improvements that reduced transport costs such as the introduction of railways in the late 1820s.

⁵⁹Districts are administrative units that stayed in place for approximately five years, from 1790 to 1795 – with each department including from a minimum of 3 to a maximum of 10 districts.

means lower price differentials across departments, which in turn implies that highly productive areas could dominate the market throughout France. Consequently, we can infer high market integration from the data if we observe that a few (presumably highly productive) departments sold to many other departments, while the majority of departments produced no output, or did so only for local consumption. Figure A.14 in the appendix shows that this pattern is particularly strong in cotton textiles. Many departments produced mostly for themselves if at all (these are the zeros and ones), while a few departments supplied cotton to a large number of districts. The top tercile of departments exported cotton textiles to 30 or more districts. In the two comparison sectors, there is less specialization and less evidence for market integration: Fewer departments report not supplying to anyone (particularly in paper), and the top decile of departments supplied only to 6 (paper) and 7 (iron) districts in total. This suggests that cotton textile plants were already competing in a bigger market than the comparison sectors around 1800.

The appendix presents a further robustness check that probes the extent to which market integration may explain our results. Table A.17 controls for market access in our quantile regressions.⁶⁰ The coefficients of interest change only marginally, and the lower-tail bias of productivity growth in spinning remains strong.

Machine Quality. Machine production and even maintenance was typically in the hands of external regional suppliers.⁶¹ Given that we find the lower-tail bias of productivity growth also within regions, it is unlikely that heterogeneity in machine quality was an important driver.

Effects of the Napoleonic Blockade on the Cotton Spinning Sector. Juhász (2018) shows that temporarily higher trade protection from British competition shifted the location of the mechanized cotton spinning industry within France. Since our results hold within regions, where the pattern of protection was very similar, it is unlikely that they are affected by the Napoleonic Blockade. Figure A.15 presents further evidence that varying trade protection does not drive our results, by splitting the sample into plants in northern and southern regions in France (corresponding to the main dimension along which protection varied). The productivity distributions in the north and south are remarkably similar, and in both regions, productivity growth until 1840 was due to a disappearing lower tail.

Could the blockade explain the lower plant survival rates observed in cotton spinning relative to the other sectors? In Table A.18 we split mechanized cotton spinning plants into the same northern and southern regions and report survival rates separately. Indeed, consistent with Juhász (2018), survival rates are lower in the south than in the north (which experienced a relative increase in trade protection during the blockade). However, survival rates remain low relative to the other two sectors even when we narrow our sample to only northern plants.⁶² This suggests that an important

⁶⁰Market access is computed as the sum of inverse distance-weighted city population in 1800.

⁶¹We discuss the market structure of domestic machine suppliers in more detail in Appendix A.2.

⁶²Note that the raw baseline survival rate for northern cotton spinning plants is now slightly higher than in paper

part of low survival rates in mechanized cotton spinning are not driven by the uneven effects of the Napoleonic blockade across France.

Finally, we note that France also experienced high raw cotton prices (the input to producing yarn) during the blockade (which ended in 1813) as documented in Juhász (2018). However, high input prices are unlikely to explain the high exit rate of plants in the sector for the period 1806 – 1840, as mechanized cotton spinning activity increased dramatically at the aggregate level during the period of the Napoleonic Blockade (Juhász, 2018).

4.3.3 *Early Spinning Workshops*

Another potential concern is that our results may be driven by the disappearance of small cotton spinning plants (or ‘workshops’). While our baseline sample filters out small establishments that used *handspinning* wheels (see Section 3.1), it may include relatively small plants that operated early vintages of mechanized spinning jennies and did not necessarily need inanimate sources of power. This small-scale setup may have been inherently different from the larger-scale factories powered by inanimate power sources. While the move to factory-based production was swift, systematic differences of smaller mechanized cotton workshops could account for the lower-tail bias of productivity growth in this sector. Our data does not differentiate between these two types of plants (we observe the capital vintage, but not the power source in the 1800 data). However, we can examine the extent to which our results could be driven by these forces in a number of ways.

First, we adopt a stricter definition of ‘factory-production’ and use only those plants that had more than 10 employees. This should exclude the majority of the smaller workshops that may have been organized as factory-based production along some but not all dimensions. Table A.19 in the appendix show that the lower-tail bias in productivity growth is robust to using only larger plants, and that it remains unique to mechanized cotton spinning. In other words, small plants are not responsible for the fat lower tail in 1806. This finding also addresses the concern that smaller plants may have been under-sampled in 1840.

Next, we implement an even more conservative definition of factory production. We drop all plants from the 1806 survey that used the earliest vintage of machinery – the spinning jenny. These were the types of machines that could, in principle, have been operated also in small workshops without inanimate power sources. We note first that there are not many plants in our data that reported using spinning jennies (see Table A.2 in the appendix). Table A.20 contains the quantile regression results. The lower-tail bias of productivity growth remains striking. Thus, it is unlikely that early spinning workshops that share some, but not all features of factory-based production

milling (11% and 9%, respectively). Recall, however, that the cotton spinning survey was conducted much later than paper milling (1806 as opposed to 1794). Moreover, the restricted sample survival rate in the north is half that measured in paper milling (10% and 20%, respectively). Taken together, this suggests that even comparing only cotton spinning plants in the north of France to the comparison sectors (calculated for all of France), survival rates for northern cotton spinning plants were low.

drive the lower-tail bias of productivity growth.

4.3.4 *Changes at the Plant Level*

Finally, we consider explanations that rely on changes in the characteristics of plants.

Plant Scale. First, our results could also be driven by increasing plant scale. This is unlikely, as all sectors witnessed an increase in plant scale. Indeed, as the results in Table A.21 show, controlling for the number of workers (at the plant level) does not alter our findings.

Capital Deepening. Over time, spinning machines were equipped with more spindles, and hence less labor was needed to produce a unit of output. Table A.22 shows that the lower-tail bias of productivity growth remains robust and similar in magnitude when we control for the capital-labor ratio at the plant level (measured as spindles per employee). This suggests that neither differential capital deepening nor changes in the vintages of capital drive our results.

Quality Upgrading. We assess the extent to which the lower-tail bias in mechanized cotton spinning could be driven by quality upgrading across plants over time. The concern is that if plants shifted their product mix to higher-quality products over time (or more precisely, plants producing higher-quality yarn entered the market), this could explain our result. We have no way of directly testing for this because we do not observe any information about the quality of output produced in the 1840 census. However, we can use our data in 1806 to examine the extent to which plants producing low-quality yarn account for the fat lower tail in that period. Unsurprisingly, quality is positively correlated with our measure of labor productivity (see column 2 in Table 6). However, quality does not account for a significant part of our core result (see Panel B of Table A.5): The lower-tail bias of productivity growth in cotton spinning continues to hold even when using prices in 1806 that do not account for quality differences across plants.⁶³ The negative gradient is not as steep when we do not adjust for quality, but productivity growth at the 10th percentile remains twice that of productivity growth at the 90th.

Finally, quality could still drive our results indirectly if it led to higher sales and thus larger plant size. To examine this possibility, we estimate the quantile regressions *not* adjusting for quality differences in prices across plants in 1806, and controlling for the number of employees (Table A.23). The lower-tail bias of productivity growth continues to hold. Taking all pieces of evidence together, it seems unlikely that quality upgrading alone could explain our results.

Age profile of plants. The median age of plants in mechanized spinning in 1806 was strikingly small (3 years). This is consistent with historical accounts that the period witnessed the birth of mechanized spinning in France. Could it be that young plants are always more dispersed in terms of productivity, and that this drives the fat lower tail? Our earlier results in Table 6 render this interpretation unlikely. We have shown that younger plants were substantially *more* productive

⁶³As highlighted in Section 3, the majority of plants produced similar, low-quality yarn.

than older ones. Thus, if anything, a predominance of younger plants would tend to lead to heavier *upper* tail.

Altogether, the findings of this section show multiple pieces of evidence that point to an important role for reorganizing production in explaining the unique lower-tail bias of productivity growth in mechanized cotton spinning. While this is not the only possible explanation, we have shown that numerous prominent alternative mechanisms are incompatible with the data. This leaves organizational challenges in technology adoption as the most promising candidate to explain the observed patterns in the data.

5 Conclusion

The unique setting examined in this paper allows us to shed new light on important open questions in the technology adoption literature. First, our findings speak directly to why the *aggregate* productivity effect of major technological breakthroughs such as IT and electricity may be hard to pin down in the data. Based on our results, the full effects of a new technology may take significant time to materialize, as firms still need to learn how to organize production efficiently. In our context, adopting mechanized cotton spinning required producers to reorganize production from households to factories. Our results suggest that initially, many plants operated the new technology in combination with inefficient complementary organizational practices. This led to a widely dispersed productivity distribution and relatively low aggregate productivity. Observers estimating the productivity effect of switching from handspinning to mechanized spinning would significantly underestimate the long-run aggregate productivity gain if they only looked at the initial data around 1800.

Second, our results also shed light on the slow adoption of major new technologies. When there is uncertainty about how to operate a new technology efficiently, and the organizational knowledge – once acquired – is observable to competitors, firms face a strategic incentive to delay adoption. The high exit rates observed in cotton spinning relative to other sectors, alongside the higher productivity observed for younger plants in 1806, suggest that plants that entered later were at an advantage. If firms understood the significant uncertainty they face when setting up a spinning mill at early stages of adoption, they had an incentive to delay the switch to the new technology in order to take advantage of the learning externalities generated by other early adopters.

In summary, our unique setting allows us to speak to a dimension of productivity growth that is usually hidden. Productivity differences across firms (or plants) reflect both the underlying technology and the complementary organizational practices with which the respective technology is used. Both features play important roles in firms' decisions to adopt new technologies: What are the potential productivity gains of a new technology (i.e., its frontier), and is the organizational knowledge needed to achieve these gains (i.e., operate at the frontier) readily available? Separating

these features empirically is difficult because of data limitations. Our results suggest that the need to reorganize production is an important dimension of technology adoption. Approaching the frontier of a new technology via organizational improvements can take a long time, and it can explain some of the salient features in the adoption of major innovations.

Finally, our paper provides a first look at how the unprecedented growth in manufacturing productivity during the First Industrial Revolution played out at the plant level. We show that in mechanized cotton spinning – the flagship industry of the period – a substantial proportion of productivity growth materialized along the extensive margin of plant exit and entry. Our results suggest that throughout this process, organizational innovations (alongside the traditionally emphasized technological ones) were an important driver of productivity growth. Future research – building on the increasing availability of historical data – should examine whether these findings constitute a common feature of the structural transformation from agriculture to modern, factory-based manufacturing. Our paper lays the groundwork by using comparative historical analysis to deepen our understanding of why the diffusion of innovation is often a complex and slow process.

References

- Allen, R. C. (2009). *The British Industrial Revolution in Global Perspective*. Cambridge University Press.
- André, L. (1996). *Machines à papier: Innovation et transformations de l'industrie papetière en France : 1798-1860*. Editions de l'Ecole des hautes études en sciences sociales.
- Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017). Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan. *Quarterly Journal of Economics* 132(3), 1101–1164.
- Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2013). Does Management Matter? Evidence from India. *Quarterly Journal of Economics* 128(1), 1–51.
- Bougin, H. and G. Bourgin (1920). *L'Industrie Sid'érurgique en France au D'ebut de la R'evolution*. Paris: Imprimerie Nationale.
- Braguinsky, S., A. Ohyama, T. Okazaki, and C. Syverson (2015). Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry. *American Economic Review* 105(7), 2086–2119.
- Braguinsky, S., A. Ohyama, T. Okazaki, and C. Syverson (2020). Product Innovation, Product Diversification, and Firm Growth: Evidence from Japan's Early Industrialization. NBER Working Paper #26665.
- Brynjolfsson, E. (1993). The Productivity Paradox of Information Technology. *Communications of the ACM* 36(12), 66–77.
- Brynjolfsson, E. and L. M. Hitt (2000). Beyond Computation: Information Technology, Organizational Transformation and Business Performance. *Journal of Economic Perspectives* 14(4), 23–48.
- Brynjolfsson, E., D. Rock, and C. Syverson (2018). The Productivity J-curve: How Intangibles Complement

- General Purpose Technologies. NBER Working Paper, no. 25148.
- Chanut, J., J. Heffer, J. Mairesse, and G. Postel-Vinay (2000). Editions de l'école des hautes études en sciences sociales, Paris.
- Chapman, S. (1970). Fixed Capital Formation in the British Cotton Industry, 1770-1815. *Economic History Review* 23(2), 235–253.
- Chapman, S. (1974). The Textile Factory Before Arkwright: A Typology of Factory Development. *Business History Review* 48(4), 451–478.
- Chassagne, S. (1991). *Le coton et ses Patrons: France 1760-1840*. Editions de l'Ecole des Hautes Etudes en Sciences Sociales.
- Clark, G. (1987). Why Isn't the Whole World Developed? Lessons from the Cotton mills. *The Journal of Economic History* 47(1), 141 – 173.
- Clark, G. (1994). Factory Discipline. *The Journal of Economic History* 54(1), 128–163.
- Comin, D. and B. Hobijn (2010). An Exploration of Technology Diffusion. *American Economic Review* 100(5), 2031–59.
- Crafts, N. (1977). Industrial Revolution in England and France: Some Thoughts on the Question, 'Why Was England First'? *Economic History Review* 30(3), 429–441.
- Crafts, N. (1985). *British Economic Growth During the Industrial Revolution*. Oxford: Oxford University Press.
- Crafts, N. and K. Harley (1992). Output Growth and the British Industrial Revolution: A Restatement of the Crafts-Harley View. *Economic History Review* 45(4), 703–730.
- Crafts, N. and N. Wolf (2014). The Location of the UK Cotton Textiles Industry in 1838: A Quantitative Analysis. *Journal of Economic History* 74(4), 1103–1139.
- Daudin, G. (2010). Domestic Trade and Market Size in Late-Eighteenth-Century France. *Journal of Economic History* 70(3), 716–743.
- David, P. (1990). The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *American Economic Review* 80(2), 355–361.
- Fitton, R. and A. Wadsworth (1958). *The Strutts and the Arkwrights, 1758-1830: A Study of the Early Factory System*. Manchester University Press.
- Foster, L., C. Grim, J. Haltiwanger, and Z. Wolf (2018). Innovation, Productivity Dispersion, and Productivity Growth. NBER Working Paper 24420.
- Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review* 98(1), 394–425.
- Galor, O. (2011). *Unified Growth Theory*. Princeton University Press.
- Garcia-Marin, A. and N. Voigtländer (2019). Exporting and Plant-Level Efficiency Gains: It's in the Measure. *Journal of Political Economy* 127(4), 1777–1825.
- Giorcelli, M. (2019). The Long-Term Effects of Management and Technology Transfers. *American Economic Review* 109(1), 1–33.
- Gort, M. and S. Klepper (1982). Time Paths in the Diffusion of Product Innovations. *Economic Journal* (367), 630–653.
- Grantham, G. (1997). The French Cliometric Revolution: A Survey of Cliometric Contributions to French Economic History. *European Review of Economic History* 1(3), 353–405.

- Griliches, Z. (1957). Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica* 25(4), pp. 501–522.
- Guiollard, P. (1993). *Les chevalements des houillères françaises*. Pierre-Christian Guiollard.
- Hall, B. (2004). Innovation and Diffusion. NBER Working Paper 10212.
- Hall, B. and B. Khan (2003). Adoption of New Technology. NBER Working Paper 9730.
- Harley, C. (1998). Cotton Textile Prices and the Industrial Revolution. *Economic History Review* 51(1), 49–83.
- Horn, J. (2006). *The Path Not Taken: French Industrialization in the Age of Revolution, 1750-1830*. Cambridge MA: MIT Press.
- Huberman, M. (1996). *Escape from the Market*. Cambridge University Press.
- Juhász, R. (2018). Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade. *American Economic Review* 108(11), 3339–76.
- Langenbach, R. (2013). Better than Steel? Part 2: Tall Wooden Factories and the Invention of “Slow-burning” Heavy Timber Construction. [Article URL accessed June 20, 2020](#).
- Mansfield, E. (1961). Technical Change and the Rate of Imitation. *Econometrica* 29(4), 741–766.
- Mitchell, B. (2003). *International Historical Statistics. Europe, 1750-2000*. Palgrave Macmillan.
- Mokyr, J. (2001). The Rise and Fall of the Factory System: Technology, Firms, and Households since the Industrial Revolution. *Carnegie-Rochester Conference Series on Public Policy* 55, 1–45.
- Mokyr, J. (2010). *The Enlightened Economy: An Economic History of Britain 1700-1850*. Yale University Press.
- Mokyr, J. (2011). The Economics of Being Jewish. *Critical Review* 23(1-2), 195–206.
- Nuvolari, A., G. Tortorici, and M. Vasta (2021). British-French Technology Transfer from the Revolution to Louis Philippe (1791-1844): Evidence from Patent Data . CEPR working paper 15620.
- Perrot, J. (1975). *L'âge d'Or de la Statistique Régionale Française (An IV-1804)*. Paris: Société des Etudes Robespierriéristes.
- Pollard, S. (1965). *The Genesis of Modern Management: A Study of the Industrial Revolution in Great Britain*. Harvard University Press.
- Pounds, N. and W. Parker (1957). *Coal and Steel in Western Europe: The Influence of Resources and Techniques on Production*. Bloomington: Indiana University Press.
- Rosenberg, N. (1976). Factors Affecting the Diffusion of Technology. *Explorations in Economic History* 10(1), 3–33.
- Rostow, W. (1975). *How it All Began: Origins of the Modern Economy*. Methuen and Co., London.
- Saxonhouse, G. and G. Wright (2004). Technological Evolution in Cotton Spinning. In D. Farnie and D. Jeremy (Eds.), *The Fibre that Changed the World: The Cotton Industry in International Perspective, 1600-1990s*, pp. 129–152. Oxford University Press.
- Sokoloff, K. (1984). Was the Transition from the Artisanal Shop to the Nonmechanized Factory Associated with Gains in Efficiency? Evidence from the US Manufacturing Censuses of 1820 and 1850. *Explorations in Economic History* 21(4), 351–382.
- Sokoloff, K. (1986). Productivity Growth in Manufacturing During Early Industrialization: Evidence from the American Northeast, 1820-1860. In *Long-Term Factors in American Economic Growth*, pp. 679–736.

University of Chicago Press.

Syverson, C. (2011). What Determines Productivity? *Journal of Economic Literature* 49(2), 326–365.

Szostak, R. (1989). The Organization of Work: The Emergence of the Factory Revisited. *Journal of Economic Behavior and Organization* 11(3), 343–358.

Tarr, R. and F. McMurry (1993). *New Geographies: Second Book*. Macmillan Company.

Vallauri, D., A. Grel, E. Granier, and J. Dupouey (2012). *Les Forêts de Cassini. Analyse Quantitative et Comparaison Avec les Forêts Actuelles*. Marseille: WWF/INRA.

Voigtländer, N. and H. Voth (2006). Why England? Demographic Factors, Structural Change and Physical Capital Accumulation During the Industrial Revolution. *Journal of Economic Growth* 11(4), 319–361.

Williamson, O. (1980). The Organization of Work: A Comparative Institutional Assessment. *Journal of Economic Behavior and Organization* 1(1), 5–38.

Woronoff, D. (1984). *L'industrie sidérurgique en France pendant la Révolution et l'Empire*. Paris: École des hautes études en sciences sociales.

FIGURES

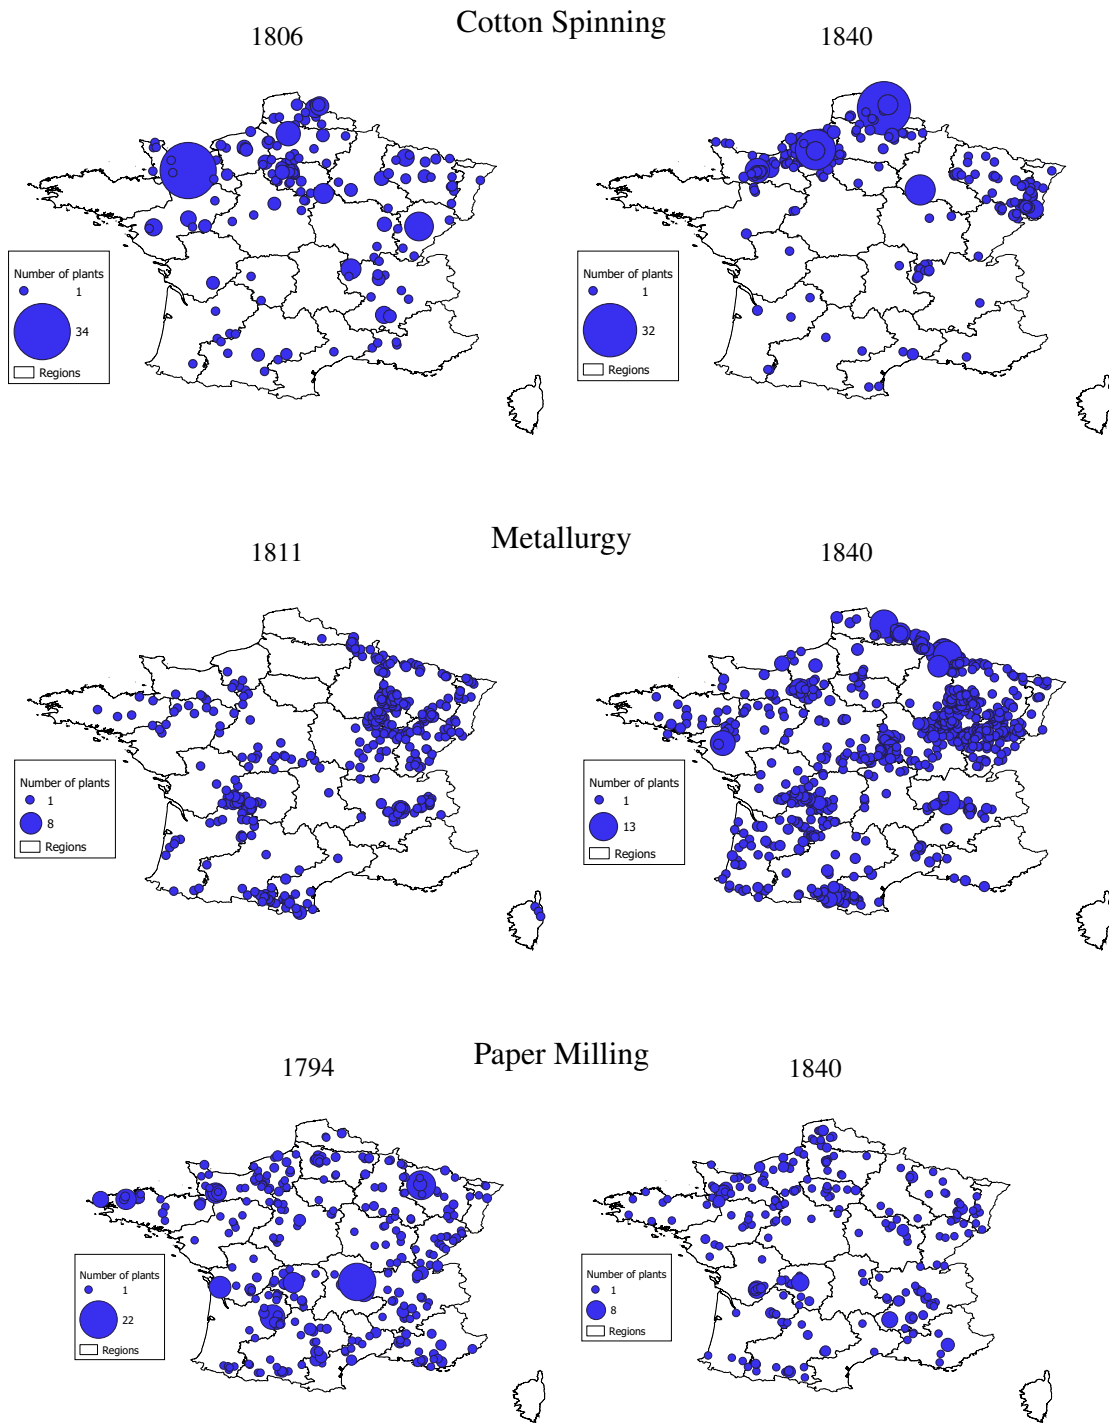


Figure 1: Spatial Distribution of Plants Across France in the Three Sectors

Note: The figure shows the spatial distribution of plants in cotton spinning (top), metallurgy (middle), and paper milling (bottom). Dot sizes reflect the number of plants per commune.




Figure 2: Changes in the Productivity Distributions in the Three Sectors

Notes: The figure shows the distribution of $\log(\text{output per worker})$ for the three sectors at the beginning of our sample period (around 1800) and in 1840. Productivity growth in spinning was mainly due to a disappearing lower tail. In contrast, in metallurgy and paper milling, the whole distribution shifted to the right.




Figure 3: Productivity Growth at Different Quantiles of the Distribution

Notes: The figure visualizes the results of quantile regressions for growth in $\log(\text{output per worker})$ for the three sectors, estimated at each decile. Productivity growth in spinning was concentrated in the lower tail of plant productivity. In contrast, in metallurgy and paper milling, productivity growth occurred relatively evenly across the distribution.




Figure 4: Proximity to High-Productivity Plants

Notes: The figure shows that proximity to high-productivity plants mattered the most in mechanized cotton spinning at the beginning of our sample period (around 1800), when the technology had just been introduced in France. The figures plot the standardized beta coefficients of $Dist^{0.90}$, which measures the log distance to the closest plant (in the same sector) with productivity in the 90th percentile. The dependent variable is $\log(\text{output per worker})$. All regressions include department fixed effects (see Table A.7). Whiskers indicate 90% confidence intervals.

TABLES

Table 1: Background of Owners of Mechanized Cotton Spinning Establishments

<i>Owners active between 1785-1815</i>	
Traders, bankers, and commercial employees	62.2%
Nobility or administrator pre-1789	10.1%
Workers and mechanics	10.8%
Industrialists	9.5%
Liberal profession	4.1%
Other	3.4%

Notes: Data are from Chassagne (1991, p. 274). The author collected data on the owners from a variety of archival sources. The sample covers 148 owners of mechanized cotton spinning establishments.

Table 2: Annual Productivity Growth (in %) at Different Quantiles of the Distribution

	(1) Average	(2) 0.1	(3) 0.25	(4) 0.5	(5) 0.75	(6) 0.9	(7) N
Spinning (1806-1840)	2.420*** (0.154)	3.917*** (0.204)	3.293*** (0.229)	2.234*** (0.151)	1.651*** (0.167)	1.014*** (0.297)	868
Metallurgy (1811-1840)	1.949*** (0.185)	1.700*** (0.415)	1.776*** (0.271)	1.787*** (0.236)	2.025*** (0.187)	2.465*** (0.226)	1296
Paper milling (1794-1840)	0.734*** (0.111)	0.713*** (0.157)	0.681*** (0.137)	0.779*** (0.098)	0.759*** (0.137)	0.726*** (0.256)	868

Notes: The table reports the average annual productivity growth (in %) between the initial sample period (around 1800) and 1840 for the three sectors (column 1), and annual productivity growth estimated at different quantiles (columns 2-6). Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Testing for Spatial Selection of New Plants in Cotton Spinning in 1806

Dependent variable: log(Output per worker)					
	(1)	(2)	(3)	(4)	(5)
			Only plants entering before high-productivity plants		
$Dist^{p90}$ (~ 1800)	-0.818*** (0.129)	-0.869*** (0.126)	-0.425*** (0.144)	-0.386*** (0.141)	-0.482** (0.188)
Plant age (in 1806)	-0.049 (0.084)	-0.187 (0.123)		-0.146 (0.127)	-0.388* (0.203)
Plant Age* $Dist^{p90}$ (~ 1800)		0.211 (0.191)			0.370 (0.247)
Department FE	✓	✓	✓	✓	✓
R ²	0.57	0.57	0.66	0.66	0.66
N	284	284	175	175	175

Notes: The table shows that our results for proximity to high-productivity plants (see Figure 4) are not confounded by new, more productive plants entering near existing high-productivity plants. $Dist^{p90}$ (1806) is the log distance to the nearest plant in cotton spinning with productivity in the 90th percentile in 1806. Plant age is the log number of years since the plant had been established, reported in 1806. Standard errors (clustered at the departmental level) in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Survival Rates Across Sectors

	Spinning	Metallurgy	Paper milling
Period	1806-1840	1811-1840	1794-1840
Survival rate	7.5%	34%	9%
Number of plants	389	477	593
Restricted sample survival rate	6.5%	49%	20%
Number of plants	93	303	218

Notes: The “survival rate” is defined as the percentage of plants from the initial period that survived to the later period based on matching either by name or location (see Section 3.5 for details). The “restricted sample survival rate” adjusts for the fact that different sectors have single-plant communes to a varying degree. It is based on the subset of plants located in communes that had only one plant in the initial period and that had either exactly one plant in 1840 (‘survival’) or no plant at all in the 1840 data (‘exit’).

Table 5: Productivity of Exiting Relative to Surviving Plants

Dependent variable: log(Output per worker)			
	(1)	(2)	(3)
	Spinning	Metallurgy	Paper Milling
Exit dummy	-0.533*** (0.165)	-0.139 (0.087)	-0.179 (0.150)
R ²	0.03	0.005	0.004
N	340	457	520

Notes: Exit is a dummy variable equal to one for plants that existed in the initial period and that had exited the market by 1840 (based on the baseline survival rate – see Section 3.5). In cotton spinning, there were 340 plants in 1806 with information on output and labor, and 314 of these had exited by 1840. In metallurgy, there were 457 plants with data to compute productivity in 1811, and 292 had exited by 1840. In paper milling, there were 520 plants with information on output and labor in 1794, 473 of which had exited by 1840. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Cotton Spinning in 1806: Productivity and Plants' Age Profile

Dependent variable: log(Output per worker)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Young plant	0.575*** (0.088)	0.374*** (0.079)	0.543*** (0.083)	0.575*** (0.089)	0.493*** (0.085)	0.608*** (0.086)	0.534*** (0.085)
log(Yarn quality)		0.673*** (0.074)					
Spinning jenny			-0.626*** (0.087)				
Water Frame				-0.003 (0.092)			
Mule jenny					0.481*** (0.086)		
log(Workers)						0.107*** (0.025)	
log(Spindles per worker)							0.336*** (0.070)
R ²	0.11	0.32	0.20	0.11	0.18	0.14	0.17
N	340	323	340	340	340	340	340

Notes: The table shows that mechanized cotton spinning plants that had just entered the market by 1806 had significantly higher productivity. 'Young' is a dummy variable equal to one for cotton spinning plants with below-median age (with the median age in 1806 being three years). log(Yarn quality) is the log (unweighted) average of the minimum and maximum count of the yarn produced by the plant. Spinning-jenny, water-frame and mule-jenny are binary indicators equal to one for plants using the earliest (spinning-jenny), intermediate (water-frame), and latest (mule-jenny) vintage of spinning machinery, respectively. The number of spindles is a standard measure of a spinning machine's production capacity, irrespective of vintage. Robust standard errors in parentheses. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table 7: Cotton Spinning in 1840: Productivity and Plants' Age Profile

Dependent variable: log(Output per worker)					
	(1)	(2)	(3)	(4)	(5)
Entrant 1840	-0.053 (0.099)	-0.060 (0.099)	-0.072 (0.100)	-0.056 (0.099)	-0.116 (0.098)
Water power		0.062 (0.050)			
Steam power			-0.093** (0.047)		
Other power				0.172 (0.140)	
log(Workers)					-0.157*** (0.028)
R ²	0.00	0.00	0.01	0.00	0.07
N	528	528	528	528	528

Notes: The table shows that in 1840, when the mechanized cotton spinning technology had reached maturity, new entrant plants did not have a productivity advantage over existing plants anymore. Entrant 1840 is a binary indicator equal to one for plants that entered the market after 1806. Water power, steam power, and other (wind or animal) power are binary indicators equal to one for plants using the respective source of power. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Online Appendix

Technology Adoption and Productivity Growth: Evidence from Industrialization in France

Réka Juhász Mara P. Squicciarini Nico Voigtländer

A Historical Background: Additional Detail

This section provides additional information about industrialization in France in the first half of the 19th century, followed by further detail on the three sectors that we examine: mechanized cotton spinning as well as metallurgy and paper milling.

A.1 The Industrial Revolution in France

An earlier literature “derogated the economic development of France as a story of retardation or relative backwardness” (O’Brien and Keyder, 1978, p. 194). This view has been largely revised in recent decades, leading to a new consensus that – similar to Britain – economic growth also accelerated in France in the mid-18th century (Rostow, 1975).¹ Horn (2006, p.10) writes that “[i]n an astonishing number of sectors, French entrepreneurs of the 1780s competed successfully with their English counterparts” and Daudin (2010) shows that French producers had access to well-integrated domestic markets (larger than those in Britain), leading to substantial specialization in production across space.

In the initial phase of industrialization, France largely depended on the adoption of British technology: In all three sectors that we examine in this paper, the new technologies arrived from Britain, most famously the British spinning jenny.² The fact that France *adopted* the major new technologies from Britain in this period makes the setting an ideal one in which to examine technology adoption.³

¹The “retardation view” of French industrialization was defended by traditionalists who criticized the emerging cliometric approach and its quantitative methods and data. Numerous studies following the work of French historian Jan Marczewski gave credence to the cliometric approach (most prominently Maddison, 2001). These studies weakened – and eventually eliminated – the idea that French economic growth had stagnated in the 19th century. The consensus view that emerged holds that French growth had in fact been substantial (Crouzet, 2003).


²Industrial espionage became widespread and, despite the attempts of the British government to block it, detailed reports and descriptions of English technology were sent across the Channel (Harris, 1998; Bradley, 2010). Additionally, “industrially minded” people in Britain and France entertained an intense correspondence on scientific and technological advances (Mokyr, 2005). Correspondingly, upper-tail human capital played an important role in industrialization (Squicciarini and Voigtländer, 2015).

³The adoption of new technologies from Britain was the primary source of innovation in the early 19th century. However, as French industrialization proceeded, “technological progress became indigenous, built in to the economy, so that ... France became at mid-[19th]century a centre of invention and diffusion for modern technologies” (Crouzet, 2003, p.234).

A.2 Mechanized Cotton Spinning

In Section 2.1, we discussed the development of mechanized cotton spinning in Britain as well as its adoption in France. Figure A.1 provides an illustration of how cotton spinning was traditionally performed, mostly by women in their homes, using a simple spinning wheel. With this technology, each spinner was able to spin only one thread of yarn. The invention of the spinning jenny by James Hargreaves in 1765 made it possible to spin multiple threads simultaneously, as twist was imparted to the fibre by using spindles rather than by the workers' hands. Throughout the 1760s – 1770s, Richard Arkwright and Samuel Crompton developed two subsequent vintages: the water-frame and the water-powered spinning mule, respectively (see Figure A.2, left panel). The mechanization of preparatory processes was also well-underway prior to the 19th century. These new technologies entailed a move from home-based to factory-based production (right panel in Figure A.2). This was partly due to the machines' reliance on inanimate power sources, and partly to an increased need to monitor workers more closely (Williamson, 1980; Szostak, 1989).⁴

The Spinning Wheel



Home Spinning




Figure A.1: Old Handspinning Technology

These innovations had enormous productivity effects. The first vintage of the spinning jenny alone led to a threefold improvement in labor productivity (Allen, 2009). As a consequence, the price of yarn declined in the late 18th century, especially for the highest-quality yarn. This can be seen in Figure A.3 showing price data for three different qualities of yarn: 18, 40 and 100 count yarn.⁵ While we observe a striking price decline for all counts, this trend was most pronounced

⁴The spinning jenny was typically hand-powered.

⁵The count is an industry-wide standard that refers to the length per unit mass, implying that higher counts are

Water-Powered Spinning Mule



Spinning Mule Operated in a Mill




Figure A.2: New Mechanized Technology in Cotton Spinning

for the highest-quality variety, where prices dropped from 1,091 pence per pound to 76 pence per pound in real terms between 1785 and 1800.

Historical Evidence About Machinery Producers. As we discussed in Section 2, machines were mainly produced domestically in France. Importantly, master mechanics and builders were typically not employees of the firm – they were paid by the factories to install, maintain and repair equipment (Cookson, 1997). Technologically complex tasks were ‘outsourced’ to engineers (Mokyr, 2010). This suggests that plants had access to broadly similar markets for the capital equipment within the same regions.

No Large-Scale Switch to Steam-Power. In contrast to Britain, mechanized cotton spinners in France did not switch from water to steam power to a large extent, owing to the fact that France was not particularly well-endowed with coal (Cameron, 1985).⁶ As such, we can think of the power-source used as remaining mostly constant over the time period. Moreover, improvements to the technology used to operate water-wheels should have a similar effect on productivity growth in paper-milling, one of our comparison sectors, as this was also reliant on water-power.

finer. Harley (1998) collected price data for three different qualities of yarn from British sources: 18, 40 and 100 count yarn (the count is an industry-wide standard that refers to the length per unit mass, implying that higher counts are finer). While all counts saw striking price declines (see Figure A.3), this trend was most pronounced for the finest, highest-quality varieties. Machine spinning had the largest impact on the fine high-quality yarn, which British hand-spinners had not been able to effectively produce and to which the mule-jenny (a subsequent vintage of the machine introduced in the late 18th century) was well-suited (Riello, 2013). Our data include information on the type of yarn produced, allowing us to account for quality differences across plants.

⁶This is confirmed in our data for 1840, showing that the majority of cotton spinning plants were still using water power (see Table A.14 in the appendix).




Figure A.3: Price of Different Counts of Yarn in Britain

Notes: Data are from Harley (1998), who collected prices for three different qualities of yarn: 18, 40 and 100 count yarn. The count is an industry-wide standard that refers to the length per unit mass (implying that higher counts are finer). Finer count yarns are used to produce higher-quality cloth, while lower counts are used to produce heavier, cheaper cloth.

A.3 Additional Details Regarding the Challenging Transition to Factory-Based Production in Mechanized Cotton Spinning

In Section 2, we discussed some of the key challenges regarding the move to factory-based production in mechanized cotton spinning. Here, we provide additional evidence and examples.

A.3.1 Building design challenges

We illustrate the trial and error process of overcoming building design challenges using the example of constructing buildings better able to withstand fires. Cotton textile mills introduced the so-called “fire-proof building” in Britain in the late 18th century, which entailed leaving no timber surfaces exposed by using cast-iron columns instead of wood (Johnson and Skempton, 1955). However, it quickly became apparent that fireproof mills were not indeed fireproof, because “steel or wrought iron, when heated, will fail by buckling or bending very much sooner than the equivalent beam of post or wood” (Boston Mutual Fire, 1908, p. 3). US textile mills developed what became known as “slow-burning mills” in the 1820s, recognizing that fires could not be prevented but their effects could be minimized by better mill design. Partly, this entailed moving back to using wood: “Timber posts offer more resistance to fire than either wrought-iron, steel, or cast iron pillars, and in mill construction are preferable in many respects (Boston Mutual Fire, 1908, p. 3). Chassagne (1991, p. 340) posits that early 19th century French mills consisted of multiple buildings and covered vast spaces (as opposed to building vertically) partly to minimize the fire hazard.

A.3.2 Labor management challenges

Cotton spinning plants needed to develop organizational and management practices for running spinning mills at a scale not typically seen elsewhere in the economy. Here, we provide additional evidence regarding labor management challenges.⁷ There were three salient aspects of this for cotton spinning mills as described in the main text; i) how to get workers to adapt to the rhythm of factory work, ii) how to coordinate labor in a factory setting and, iii) how to solve monitoring problems.

First, from the workers' side, the move to factory-based production fundamentally altered both the location and the nature of work (Clark, 1994). Under the factory system, the employer "dictated when workers worked, their conduct on the job and that they steadily attend to their assigned tasks." (Clark, 1994, p. 128). Following instructions, showing up to work on time, or getting along with other employees was a challenge for the first generation of factory workers used the high degree of independence typical of the domestic system (Pollard, 1965). As Pollard (1965, p. 181) writes, "What was needed was regularity and steady intensity in place of irregular spurts of work; accuracy and standardization in place of individual design; and care of equipment and material in place of pride in one's tools". Simply put, an industrial labor force needed to be created where none had existed before (Mokyr, 2010).

Second, coordination of labor was crucial, as flow production meant that one worker could hold up the entire production line. This is strikingly illustrated by Karl Marx, who quoted a large cotton manufacturer, Henry Ashworth: "When a laborer lays down his spade, he renders useless for that period, a capital worth eighteen pence. When one of our people leaves the mill he renders useless a capital that cost £100,000" (Clark, 1994, p. 129).

Third, monitoring worker effort, much of which was hard to observe (Huberman, 1996), was another novel aspect of factory work. Huberman (1996, p. 11) describes the need for monitoring in mechanized cotton spinning: "If there were multiple breakages of yarn on the larger machines, the mule had to come to a complete stop to piece the broken threads. There was also doffing, when the reels were full of spun cotton, the mule had to be stopped and the reels removed. Finally, there was cleaning. At all times, the spinners could expend effort as they were motivated to, and without proper supervision or incentives they could disguise how hard they could in fact work." This created a strong need for monitoring, so that even early hand-powered machines (a particular vintage of spinning machinery) were housed in the "garrets of cottages and later in sheds" (Huberman, 1996, p. 11).

As we discuss in the main text, overcoming these challenges proceeded via a slow process of trial and error. The industry eventually settled on efficiency wages in the 1830s in Britain

⁷Pollard (1965) and Mokyr (2010) provide more general discussions of other management challenges facing firms at the time.

(Huberman, 1996).

A.4 Comparison Sectors: Metallurgy and Paper Milling

This section discusses our two comparison sectors during the First Industrial Revolution in France in more detail. We give an overview of the production process in each, and discuss new technologies developed during our sample period.

Metallurgy


Iron was a flagship product of the Industrial Revolution. It was used for railways, steamships, and for machines. The fundamental process of producing iron has remained the same over centuries. Despite large productivity improvements achieved during the Industrial Revolution (Allen, 2009), the actual changes to methods of production were modest.

Iron is extracted from iron ore in a process called ‘smelting’ – freeing the iron by combining carbon with the oxygen of the ore under heat. The difficulty comes from the fact that the iron also needs to be separated from other metallic substances in iron ore. This is achieved by controlling the heat of the furnace so that most of the foreign matter separates out with the lowest-possible expenditure of fuel. In the Medieval period, the production process of iron used ‘direct’ technology. Smelting with this technology produced malleable iron directly from iron ore in a bloomery (a type of furnace) where the temperature was low enough for the iron not to melt. The product of this technology is wrought iron. This process is referred to as ‘direct’ because iron was produced in a near-finished condition in a single process. One vintage of this technology, the Catalan forge, survived into our study period and beyond in certain parts of France (Pounds and Parker, 1957).

Starting in the late Middle Ages, the direct technology began to be gradually replaced by an ‘indirect’ technology, which consists of two steps: smelting and refining. Smelting is similar to the direct technology. A blast furnace – with temperatures high enough for the complete fusion of the metal – is used to produce an intermediate product – pig iron. However, as pig iron is too brittle to be used in many applications, it needs to be refined one more time on a hearth. This second stage is known as refining. The blast furnace first appeared in Europe in the 15th century, but it was not widely adopted until the 17th or 18th century (Pounds and Parker, 1957). Figure A.4 illustrates the blast furnace and the organization of an 18th century metallurgy plant (foundry).

Prior to the Industrial Revolution, both stages of the indirect process relied on charcoal as the source of fuel. The key innovation during the Industrial Revolution was the switch from charcoal to coal, through a series of gradual improvements in the period 1700-1850. The change in the type of fuel required modifications to the blast furnace, but the new technology “merely replaced earlier, recognizably similar, though less ‘efficient’ methods” (Pollard, 1965, p. 101). In particular, a coal-based blast furnace required slightly larger furnace sizes and a switch from water to steam power. Such modifications could be made to existing blast furnaces (Pounds and Parker, 1957).

18C Charcoal Iron Blast Furnace



Organization of 18C Metallurgy Plant





Figure A.4: ‘Old’ 18th Century Charcoal-Based Technology in Metallurgy

Allen (2009) estimates that the cost of producing pig iron using coal decreased by 75% during this period in Britain. Figure A.5 presents illustrations of the ‘new’ coal-based technology, showing that the organization of metallurgy plants remained practically unchanged.

In refining, the move from charcoal to coal was achieved by the puddling process. Pig iron was melted in a reverberatory furnace fueled by coal while stirring, or ‘puddling’ the molten mass until the free carbon in pig iron was oxidized and the mass reduced to malleable form as bar iron. Similar to smelting, the development of puddling proceeded gradually, and the technologies were not meaningfully different to others already in use. Pounds and Parker (1957, p. 35) characterize the development of puddling in the following way: “Cort’s [the inventor] invention was less the introduction of entirely new processes than a synthesis of practices that were already familiar. The reverberatory furnace, burning any fuel that might be available, had long been known. (...) Rolling mills were also known (...). The lawsuits with which Cort was faced served to emphasize that some of his contemporaries regarded his claim to be an inventor as remarkably thin.”

Importantly, in contrast to cotton spinning, the new, coal-based technologies were adopted within the existing plant-based organization of production. Therefore, the plant productivity distributions that we observe for metallurgy reflect patterns of productivity growth for an industry undergoing innovation without the need to reorganize production.

19th Century Coal Blast Furnace



Organization of 19C Metallurgy Plant





Figure A.5: 'New' 19th Century Coal-Based Technology in Metallurgy

Paper Milling

In contrast to metallurgy and cotton spinning, paper milling was not a flagship industry of the Industrial Revolution and its output was not particularly important for other sectors. However, paper making was one of the few manufacturing activities that was organized as plant-based production from well before the Industrial Revolution because of its reliance on water power. Moreover, it too underwent mechanization during our study period. For these reasons, paper making serves as a useful second comparison sector to cotton spinning.

In Europe, paper making had traditionally taken place in mills. Production consisted of several stages. First, the vegetable matter (the raw material) was broken down into cellulose fiber, which involved a water-powered stamping machine (see Figure A.6, left panel). Next, it was formed into thin, wet sheets by a skilled worker, called a vatman (Figure A.6, right panel). It was then dried and – depending on its intended use – finished in different ways. Each of these steps was performed in a different section or room of the mill with a marked division of labor by function and gender. The only step of the production process that required water power was the washing, breaking, and beating (stamping) down of rags into fiber. The machine that performed all of these tasks (known as the washing, breaking or stamping engine) consisted of an oval, wooden tub containing a water-powered revolving roll and was operated by a skilled workman or an engineer. This stage of the production process and the work of the engineer determined to a large extent the quality of the paper that could be produced. Moreover, it was because of this technology that production was

Sketch of Fourdrinier Machine



Fourdrinier Machine in a Plant




Figure A.7: New Technology in Paper Milling Plants: The Foudrinier Machine

A.5 Evidence on Innovation Across the Three Sectors Using British Patent Data

Here, we provide evidence that all three sectors we study were highly innovative during our sample period. In Figure A.8, we show that patenting activity was consistently high through the 1800-1840 period using data on the number of British patents by category. Spinning was the third-most patent intensive industry among 148 categories, while metallurgy and paper milling were ninth and twenty-first respectively.




Figure A.8: Number of British Patents, 1800-1840

Notes: ‘Spinning’ refers to the patents related to all textile fibers, not just cotton. Data were kindly shared by Walker Hanlon based on work on patenting in Hanlon (2020).

Moreover, Table A.1 shows that in all three sectors, innovations were broad-based in the sense that they covered different parts of the production process. In combination with the historical

evidence on the nature of innovations, we conclude that the three sectors were undergoing technological change of a similar type *after* the adoption of mechanized cotton spinning.

Table A.1: Patents, 1800-1847

Sector	'Core'	Other	Total
Spinning	176	313	489
Metallurgy	100	143	243
Paper	89	33	122

Notes: The table reports the number of patents divided between different stages of the production process. 'Core' patents refer to innovations in the main part of the production process, while 'Other' refers to innovations in preparatory or finishing stages. Data were kindly shared by Walker Hanlon based on work on patenting in Hanlon (2020).

MINISTÈRE DE L'INTÉRIEUR.

Maire de Forges
des Mines de fer, les Forges et les Usines

arrondissement de Soissons — Département de l'Aisne

OBJET DES QUESTIONS

EN 1789.	DEPUIS LE 1 ^{er} JANVIER 1811 JUSQU'AU 31 JUIN 1811.	OBSERVATIONS
1. Quel a été le prix de quintal métrique en deux cas indiqués par les deux premières questions, comparé avec votre usine?	610. 00	Le 1 ^{er} et 2 ^e cas se produisent dans la même usine. L'augmentation est de 100 p. 100.
2. Combien y a-t-il eu de hauts-fourneaux en activité dans votre usine?	0	Il n'y a eu aucun haut-fourneau en activité.
3. Quel a été le nombre de fers de forge en activité dans votre usine?	2 1/2	Il y a eu 2 1/2 fers de forge en activité.
4. Combien y a-t-il eu de fers de forge alternativement par la haute dans votre usine?	0	Il n'y a eu aucun fer de forge alternativement par la haute.
5. Quel a été le nombre de fers de forge alternativement par la basse dans votre usine?	0	Il n'y a eu aucun fer de forge alternativement par la basse.
6. Combien de quintaux métriques, en force ou en poids, ont été produits par votre usine?	8. 00	Il y a eu 8 quintaux métriques produits.
7. Quelle a été la quantité de quintaux métriques que votre usine a produits en fer forgé, deuxième qualité, en poids?	0	Il n'y a eu aucun quintal métrique produit en fer forgé, deuxième qualité.
8. Combien de quintaux métriques ont été produits par votre usine en fer forgé, troisième qualité, en poids à froid ou à chaud?	0	Il n'y a eu aucun quintal métrique produit en fer forgé, troisième qualité.
9. Combien y a-t-il eu de fers de forge à froid en activité dans votre usine?	0	Il n'y a eu aucun fer de forge à froid en activité.
10. Quelle quantité de quintaux métriques d'acier de cimentation a été produite par votre usine?	0	Il n'y a eu aucun quintal métrique d'acier de cimentation produit.
11. Combien de quintaux métriques d'acier travaillé ont été produits par votre usine?	0	Il n'y a eu aucun quintal métrique d'acier travaillé produit.
12. A quel prix est retenu le quintal métrique d'acier travaillé dans votre usine?	0	Il n'y a eu aucun quintal métrique d'acier travaillé produit.
13. Quel a été le prix de quintal métrique d'acier de cimentation dans votre usine?	0	Il n'y a eu aucun quintal métrique d'acier de cimentation produit.
14. Combien a coûté une corde de hauts-fourneaux dans votre usine, et quelle a été la longueur, la hauteur et la largeur de cette corde?	0	Il n'y a eu aucune corde de hauts-fourneaux produite.
15. Combien a coûté à votre usine l'entretien de la mine de fer?	0	Il n'y a eu aucun quintal métrique produit.
16. Combien a coûté à votre usine le transport de la mine de fer?	0	Il n'y a eu aucun quintal métrique produit.
17. Quel a été le prix de la mine d'acier dans votre usine?	0	Il n'y a eu aucun quintal métrique produit.
18. Quel nombre d'ouvriers a travaillé dans votre usine?	0	Il n'y a eu aucun ouvrier travaillé.

Figure A.10: Sample Page from the Metallurgy Survey, 1811

Département des Ardennes
District de Reims

Etat Des Papiers Millés En La Commune d'Aspach

papier

Liste de la papeterie de Darnay, dans le Département des Ardennes, sous le nom de la papeterie de Darnay, dirigée par M. le Citoyen Lefebvre, ancien député de la Convention.

Noms et surnoms des ouvriers	Lieu de la naissance			Age	Jours de travail	nombre		Délivrance	Observations
	Communes	Districts	Départements			des	des		
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	1 ^{er} jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	2 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	3 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	4 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	5 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	6 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	7 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	8 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	9 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	10 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	11 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	12 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	13 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	14 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	15 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	16 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	17 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	18 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	19 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	20 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	21 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	22 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	23 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	24 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	25 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	26 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	27 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	28 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	29 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	30 ^e jour de la semaine
André Lefebvre	Reims	Reims	Ardennes	35 ans	250	1	250	1	31 ^e jour de la semaine

Figure A.11: Sample Page from the Paper Milling Survey, 1794

used and the number of machines,⁸ physical output produced (reported in kilograms), the maximum and minimum quality of yarn that the plant spun (measured by the count of the yarn), the date of foundation, the name of the owner, and the location of the plant (up to the commune level). As discussed in Juhász (2018), the number of spindles used is imputed for plants that do not report this information, based on the vintage of capital used and other plant characteristics.

B.2 The Metallurgy Survey in 1811

The 1811 survey on the metallurgy industry, discussed in detail in Woronoff (1984), covers information on all branches of metallurgy. The survey covers labor employed, types of physical capital used, types of output produced in physical quantities and the corresponding price, as well as some information on raw materials. In addition, the survey provides information on the owner of the plant and on its location (up to the commune level).

B.3 The Paper Milling Survey in 1794

The 1794 survey covers plants active in paper milling. The survey provides particularly detailed information on the workforce, i.e., names and tasks of the employees, as well as their age and the time spent working in that plant. Moreover, it contains information on the number of tanks used (a measure of capital), the quantity of paper produced, and the location of the plant (up to the commune level).

B.4 Construction of Prices for 1794 in Paper Milling

The *Tableaux du Maximum* that we use to estimate the price of paper products reported in the 1794 survey was compiled by the revolutionary French government. It reports prices of a standardized list of goods in 1790. Compiling the *Tableaux* was part of the strategy of the revolutionary French government to fight hyperinflation by setting a price ceiling for products. To this end, the government asked each French district to send a list of all goods produced and imported, together with their respective prices increased by one third.⁹ In the case of paper products, for each district, we have information on the different types of paper sold and their price (for a total of 852 entries). However, information on the unit of measurement for which the price is reported is only observed for about 30% of our observations. To impute the missing information, we identify clusters of products that were similar, based on the product names. We then assign product-district observations with missing unit-of-measurement information the average price per kilogram in that cluster. For each department, we then construct the average price using all products listed for that department. The nine clusters are ‘gris,’ ‘puissant,’ ‘moyen,’ ‘bulle,’ ‘pate superfin,’ ‘carre fin double,’ ‘carre fin d’impression,’ ‘raisin fin,’ ‘ecu fin,’ and ‘pot.’

⁸The survey asked for throstles (water-frame) and mule-jennies, but many plants also reported spinning jennies. Juhász (2018) thus constructed a third, separate category to capture these.

⁹Districts were administrative units that were in place from 1791 to 1795, i.e., during the revolutionary period. In 1800, they were replaced by arrondissements.

B.5 Further Detail on Price Deflators

For all three sectors, we deflate revenues from the census using the PPI in 1840, which is equal to 0.88.¹⁰ The PPI in 1800 is 1.18. The PPIs in 1806 and in 1811 are 1.25 and 1.68, respectively – and we use these to deflate revenues in cotton spinning and in metallurgy in the corresponding exact years of the survey. The PPI data from Mitchell (2003) are reported only starting in 1800. In the case of paper milling – where prices are reported for 1790 – we need to make an assumption about how 1790 prices relate to prices in 1800. Based on our reading of the literature on the timing of hyperinflation during the French Revolution, we assume that 1790 prices were the same as prices in 1800.¹¹

C Descriptive Statistics

In this section, we describe key descriptive statistics for the three industries covered in our data.

C.1 Summary Statistics From the 1806 Cotton Spinning Survey

Table A.2 shows summary statistics for the mechanized cotton spinning plants covered in the 1806 survey. A few points stand out. First, most plants produced fairly similar, low quality yarn, consistent with the British experience (Harley, 1998).¹² The interquantile range for the quality of yarn produced is 20 – 48. Second, plant age is fairly low – the median plant was only three years old at this time. This suggests that indeed, we are capturing the industry at a fairly early stage of wide-scale technology adoption. Third, while we do capture some plants that use the early vintage of the spinning jenny, the vast majority of plants use more modern vintages, in particular the mule jenny.

C.2 Plant Scale

We also examine plant scale, and the number of plants in each industry. Table A.3 provides an overview, reporting plant size, measured by the number of workers. A few points stand out. First, as early as 1806, cotton spinning plants were strikingly large. The average spinning plant in this

¹⁰Deflation, defined here as a decline in the aggregate price level, was common in the 19th century. Borio and Filardo (2005) show that France, for example, experienced deflation in the CPI during more than one-third of the years between 1801 and 1879. Indeed, in this period, inflation tended to be a mean-reverting process.

¹¹The high inflation during the revolutionary period was accompanied by the increasing issuing of *assignats* – a type of paper money. Initially, *assignats* were linked to gold – one *louis d'or* (a gold coin) was equal to 100 *assignats* in 1791. Thereafter, the value of *assignats* depreciated: In 1794, the same *louis d'or* corresponded to about 500 *assignats*. After a period of hyperinflation (1793-1797), the *assignats* were removed from the market (White, 1991; Sargent and Velde, 1995). In 1800, Napoleon founded the *Banque de France* and then eliminated paper money (Lefebvre, 2011), while trying to increase the stock of metallic money. While there is no systematic price index linking the 1790 and 1800 prices, it seems that by 1800, prices reverted back to their 1790 level (we are grateful to Eugene White for these insights on hyperinflation during the French Revolution). Moreover, for one town (Château-Gontier), we also observe a (wheat) price index every year from 1790 to 1800: This was 202 in 1790, it increased to 336 in 1794, and went back to 209 in 1800 (Hauser, 1985). This evidence from Château-Gontier further supports our assumption that 1790 prices were approximately the same as prices in 1800.

¹²High count yarns typically start around count 100 (Harley, 1998).

Table A.2: Summary Statistics: Cotton Spinning Plants in 1806

Sector	mean	(1)	(2)	(3)	(4)	(5)
		sd	median	25perc	75perc	N
Avg. quality of yarn	35	(19.8)	28	20	47.7	346
Plant's Age	6.6	(7.9)	3	2	10	383
Total num. of workers	63	(101)	30	10	70	372
Total num. of machines	14.8	(23)	8	4	16	348
Num. of spinning jennies	2.6	(8.8)	0	0	0	385
Num. of water-frames	4.1	(11.3)	0	0	4	349
Num. of spinning-mules	7.8	(18.6)	1	0	8	350
Spindles	1602	(3487.4)	674	220	1640	389
Spindles per worker	28.2	(19.3)	22.3	14.6	35.4	372

Notes: The table reports summary statistics for mechanized cotton spinning firms in 1806. The number of observations differ across the columns as not all plants responded to all questions. Data sources: See Section 3.

period had 63 employees (the median was 30). Despite the recent introduction of mechanized cotton spinning in France, plants were already larger than in the two comparison sectors, which had a longer tradition of factory-based production. Plants in metallurgy (reported in 1811) had on average 20 workers; paper milling plants had on average 13 employees.¹³

Second, Table A.3 also shows that between 1806 and 1840, the number of mechanized cotton spinning plants active in the market expanded markedly (from 372 in 1806, to 528 in 1840). This is important, as it suggests that our results, which show a disappearance of the lower tail of the productivity distribution, were driven by more than simply the ‘shake-out’ of unsuccessful plants. In fact, each exiting plant was replaced on average by more than one new entrant.

C.3 Occupations and Division of Labor in Paper Milling in 1794

The data from the 1794 paper milling survey contain highly detailed information on occupations in each plant. In Table A.4, we examine the extent to which paper milling plants in France used standardized occupational categories to implement a division of labor. The table clearly shows that despite the relatively small scale of plants (measured by their number of employees), work within

¹³One caveat with making this comparison is that the paper milling survey dates from 1794. Thus, plant size may have grown by 1806 – the year of the cotton spinning survey. In addition, we had to extrapolate the overall number of workers in paper milling in 1794 (including women and children – see Section 3.3). However, it is unlikely that true plant scale would have been very different in 1806. This is because even in 1840, the average plant size in paper milling was only 43 (including women and children, which are reported in this year). We can thus be confident that paper milling plants in 1806 were substantially smaller than cotton plants. Finally, there is a concern that the 1840 census did not enumerate all plants with less than 10 employees.

Table A.3: Summary Statistics: Plant Scale Across the Three Sectors

Sector	year	(1)	(2)	(3)	(4)	(5)	(6)
		mean	sd	median	10perc	90perc	N
Cotton spinning	1806	63	(101)	30	4	150	372
	1840	112	(148)	72	28	210	528
Metallurgy	1811	20	(23)	11	4	46	457
	1840	57	(114)	22	7	135	839
Paper milling	1794	13	(19)	11	5	23	550
	1840	43	(58)	19	5	112	348

Notes: The table reports summary statistics on the number of workers per plant in the three sectors covered by our analysis. The year of the first survey varies across the sectors, while information in 1840 is available for all sectors. Data sources: See Section 3.

the vast majority of plants was divided amongst highly standardized occupational categories. The three occupations listed most frequently were found in the majority of plants. In fact, only 15% of plants report no specific occupational categories. This underscores the point that even before the turn of the 19th century, paper milling plants across France employed a similar division of labor. This is consistent with a common set of industry-wide ‘best-practices’ for organizing plant-based production.

Table A.4: Division of Labor in Paper Milling

Profession	(1)	(2)
	Share of firms	Number of firms
Coucher	75%	410
Remover	67%	366
Engineer	64%	353
Loftman	29%	158
Vatman	9%	49
Total number of firms: 550		

Notes: The table shows the share of firms (column 1) and the number of firms (column 2) reporting, as part of their workforce, one of the five main occupational categories: coucher, remover, engineer, loftman, and vatman. The total number of paper milling firms is 550. Only 81 firms (14.7 % of the total) do not report any specific occupational category.

D A Stylized Theoretical Framework

This section describes a stylized theoretical framework that generates the main pattern in our data: a strong lower-tail bias in productivity growth in cotton spinning, and a more balanced shift of the productivity distribution in the comparison sectors. For simplicity, we focus on a partial equilibrium setting where the economy-wide expenditure for a sector is given. This expenditure can reflect both domestic and international demand. Firms within a sector produce differentiated products, reflecting differences in output varieties as well as spatially segmented markets. Firms randomly draw their productivity (based on a combination of complementary inputs), and they have to pay a fixed cost to produce. Firms with negative profits exit the market.

D.1 Production

We chose a production function that features complementarity across multiple inputs (tasks). For tractability, we constrain inputs to capital (K), material (M), and labor (L). Firm i draws input-specific efficiencies $\gamma_{K,i}$, $\gamma_{M,i}$, and $\gamma_{L,i}$ from independent uniform distributions. Output Y_i is produced by a CES production function:

$$Y_i = A \left(a_K (\gamma_{K,i} K_i)^{\frac{\sigma-1}{\sigma}} + a_M (\gamma_{M,i} M_i)^{\frac{\sigma-1}{\sigma}} + a_L (\gamma_{L,i} L_i)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad (\text{D.1})$$

where A is a sector-wide productivity parameter and a_j reflect the average (across all firms) importance of input $j \in \{K, M, L\}$ in production.

For simplicity, we focus on a static, one-period problem. Firms incur a fixed cost F , which affects their decision to remain in the market at the end of the production period (as discussed below). Firms minimize the total variable cost of production $\tau_K K_i + \tau_M M_i + \tau_L L_i$, where τ_j are the given prices of inputs $j \in \{K, M, L\}$. Solving the cost minimization problem yields the firm-specific marginal cost of production:

$$c_i = \frac{1}{A} \left(a_K^\sigma \left(\frac{\tau_K}{\gamma_{K,i}} \right)^{1-\sigma} + a_M^\sigma \left(\frac{\tau_M}{\gamma_{M,i}} \right)^{1-\sigma} + a_L^\sigma \left(\frac{\tau_L}{\gamma_{L,i}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \quad (\text{D.2})$$

Firm i 's labor demand is given by

$$L_i = Y_i \left(\frac{a_L c_i}{\tau_L} \right)^\sigma (A \gamma_{L,i})^{\sigma-1}, \quad (\text{D.3})$$

where the total demand for firm i 's output, Y_i , is determined below. In our simulation, we use this equation to compute log output per worker, $\ln(Y_i/L_i)$.¹⁴

¹⁴We are interested in log output per worker for surviving firms only. In what follows, we outline the conditions for firms that exit the market.

D.2 Consumption

Demand for a sector's products Y_i is given by

$$Y = \left(\sum_{i=1}^I Y_i^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}, \quad (\text{D.4})$$

where ϵ is the elasticity of substitution, and I is the total number of firms in the respective sector. We take the economy-wide expenditure for the sector as given and denote it by E . The price index that follows from (D.4) is given by:

$$\mathbf{P} = \left(\sum_{i=1}^I P_i^{1-\epsilon} \right)^{\frac{1}{1-\epsilon}}, \quad (\text{D.5})$$

where P_i is the price of product i . Firms maximize profits and charge a markup over the marginal costs in (D.2). Thus, their output price is given by

$$P_i = \frac{\epsilon}{\epsilon-1} c_i \quad (\text{D.6})$$

Consumer optimization also implies that the demand for firm i 's output is given by

$$Y_i = \frac{1}{P_i} \left(\frac{\mathbf{P}}{P_i} \right)^{\epsilon-1} \cdot E \quad (\text{D.7})$$

where $E = \mathbf{P}Y$ is the total expenditure for the sector's output.

D.3 Firms' Market Exit Decision

Firm i 's profits are given by

$$\Pi_i = (P_i - c_i) Y_i - F = \frac{1}{\epsilon-1} c_i Y_i - F \quad (\text{D.8})$$

Following (D.2) and (D.7), profits depend on firm i 's draw of input efficiencies $\gamma_{ji}, j \in \{K, M, L\}$. If these are such that $\Pi_i < 0$, firm i exits the market at the end of the period.¹⁵

D.4 Simulation Results

To simulate the model, we use the following parameters: For the importance of the three inputs, $a_K = a_M = a_L = \frac{1}{3}$; $\sigma = 0.5$, reflecting complementarity across inputs; and $\epsilon = 2$, so that the individual products within a sector are substitutes. Regarding the input efficiency draws in cotton

¹⁵We assume that the fixed costs cannot be recovered once the firm received its draw of γ_{ij} (i.e., once the production facility has been put in place). Since $\frac{1}{\epsilon-1} c_i Y_i > 0$, even firms with negative profits produce until the end of the period, recovering part of their fixed cost, and then exit the market.

spinning, we initially draw all three parameters γ_{ji} , $j \in \{K, M, L\}$ from a uniform distribution with bounds $[0, 1]$. Thus, some input draws are initially close to zero, which, in combination with the complementarity across inputs, yields the fat lower tail of the firm productivity distribution. We choose the parameter A together with the overall expenditure E and fixed costs F so that the firm survival rate at the end of the period is about 0.925, which over a horizon of 34 periods leaves about 7% surviving firms, as in our data. Finally, to reflect learning, we increase the lower bound of the uniform distribution of γ_{ji} draws to 0.3, while keeping the upper bound unchanged at 1. Thus, learning within the cotton spinning sector eliminates very low input efficiency draws, but it does not change the technological frontier. This reflects our unique empirical setting, where all firms that we observe in the data already operated the modern, mechanized cotton spinning technology.

The results of the simulation for cotton spinning in the initial and post-learning periods are shown in the left panel of Figure A.12.¹⁶ The model delivers the lower-tail bias in productivity growth that we documented for the cotton spinning plants.

Finally, we adapt our simulation to the comparison sectors. Since these used technologies that were already well-established around 1800, we choose a lower bound that reflects that learning had already occurred. We thus use a uniform distribution with bounds $[0.3, 1]$ at the beginning of the period (i.e., the same as for cotton spinning after learning). To reflect incremental technological progress that improved the technology over time, we shift both the upper and lower bound outward in the second period, using the support interval $[0.4, 1.1]$. Notice the contrast to cotton spinning, where productivity growth occurred without an outward shift of the technological frontier. The results of the simulation are shown in the right panel of Figure A.12. They reflect the pattern observed for our comparison sectors, metallurgy and paper milling (see Figure 2 in the paper).

¹⁶The figure plots $\ln(Y_i/L_i)$ for the subset of firms that do not exit the market at the end of the one-period simulation.




Figure A.12: Simulated Productivity Distributions

Notes: The left panel shows simulation results for the productivity distributions in cotton spinning in the initial period (dashed line), and after learning eliminated very low input efficiency draws (solid line). Learning does not shift the upper bound of the input efficiency distribution, i.e., the technological frontier is unchanged. The right panel shows the simulated productivity distributions for a comparison sector, where technological progress shifts the input efficiency distribution to the right, thus raising both its lower and upper bound.

E Additional Results

In this section, we provide details on the robustness checks referenced in the main text.

E.1 Cotton Spinning: Output Quality and TFP

In Section 4.1 we showed that the productivity gains in mechanized cotton spinning are almost exclusively concentrated in the lower tail of the productivity distribution. The lower tail disappeared over our sample period, while increases in productivity at the upper tail of mechanized spinning were modest. This is evident in Table 2, displaying the results of quantile regressions that use quality-adjusted sector-level prices. Panel B of Table A.5 shows that these results are robust when we do not adjust for quality differences in the count of yarn spun by plants, but rather we use the *same* sector-level price across all plants in cotton spinning. In particular, we use the price of the yarn count that the average plant produces. The difference in productivity growth is more muted, as we would expect, but the lower-tail bias remains striking.

Panel C of Table A.5 presents results for total factor productivity (TFP) instead of output per worker in the productivity regressions of mechanized cotton spinning. To estimate TFP, we use data on the labor employed by the firm (measured as number of workers) and proxy for the capital stock with the number of spindles – a standard measure of production capacity in the industry. We regress the revenue of the firm on a constant, log labor, and the log number of spindles of the plant, separately for each year. This allows for the capital and labor shares to be time-varying. Log TFP

for each plant i in a given year t is thus the regression constant plus the residual of the regression. We find that the lower-tail bias of productivity growth in mechanized cotton spinning is robust to estimating TFP.

Table A.5: Alternative Productivity Measures in Cotton Spinning

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Average		At the following quantiles:				N
		0.1	0.25	0.5	0.75	0.9	
PANEL A: Baseline (Table 2)							
Spinning (1806-1840)	2.420*** (0.154)	3.917*** (0.204)	3.293*** (0.229)	2.234*** (0.151)	1.651*** (0.167)	1.014*** (0.297)	868
PANEL B: Using prices not quality-adjusted							
Spinning (1806-1840)	2.373*** (0.138)	3.381*** (0.285)	2.828*** (0.199)	2.105*** (0.193)	1.829*** (0.160)	1.628*** (0.188)	868
PANEL C: Using TFP							
Spinning (1806-1840)	2.845*** (0.050)	3.233*** (0.080)	3.107*** (0.072)	2.834*** (0.056)	2.647*** (0.083)	2.317*** (0.072)	868

Notes: Panel A reproduces the specification in Table 2. In Panel B, the dependent variable is log(output per worker) computed using prices that are *not* adjusted for quality. In Panel C, the dependent variable is total factor productivity. Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.2 Paper Milling: Labor Input and Price Index

Next, we check the robustness of our results to other data construction choices we made. Table A.6 shows that the results in paper milling are robust to using only male labor in both periods to construct productivity (Panel B). Moreover, the results are also similar when we use the average, country-wide price to estimate revenues in 1794, as opposed to departmental ones.

Table A.6: Alternative Productivity Measures in Paper Milling

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Average	At the following quantiles:					N
		0.1	0.25	0.5	0.75	0.9	
PANEL A: Baseline (Table 2)							
Paper milling (1794-1840)	0.734*** (0.111)	0.713*** (0.157)	0.681*** (0.137)	0.779*** (0.098)	0.759*** (0.137)	0.726*** (0.256)	868
PANEL B: Using only male labor							
Paper milling (1794-1840)	0.611*** (0.116)	0.408** (0.160)	0.408*** (0.154)	0.538*** (0.127)	1.020*** (0.136)	0.684*** (0.221)	868
PANEL C: Using country-wide average price							
Paper milling (1794-1840)	0.841*** (0.111)	0.729*** (0.159)	0.750*** (0.142)	0.912*** (0.099)	0.841*** (0.140)	0.984*** (0.254)	868

Notes: Panel A reproduces the specification of Table 2. In Panel B, the dependent variable is log(output per worker) computed using male labor only. In Panel C, labor productivity is computed using the average price for all plants in the country. Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.3 Robustness: Spatial Diffusion of Knowledge

In this section, we asked how information about best practice organizational methods diffused through the economy. We suggested that plants copied from successful experimenters and provided evidence supporting the spatial diffusion of knowledge. Figure 4 in the paper plots the estimated coefficients that show that proximity to high-productivity plants mattered in cotton spinning in 1806, while this pattern was much weaker for the later period and for the two comparison sectors. Table A.7 reports the corresponding regressions. Figure A.13 displays the spatial distribution of plants in mechanized cotton spinning, metallurgy, and paper milling, distinguishing those in the 90th percentile of the productivity distribution.

Next, we show that our evidence on the spatial diffusion of knowledge is robust to a series of potentially confounding explanations. Importantly, the inclusion of department fixed effects across all specifications already captures department-level unobserved characteristics. However, it does not account for unobserved differences at a more disaggregated level. Table A.8 addresses this possibility, by controlling directly for some key location fundamentals at the commune level such as the availability of fast-flowing streams, proximity to coal, and the share of forest cover. The pattern of the coefficients of interest are very similar to those observed in Table A.8. Moreover, the coefficients of the location fundamentals are small and insignificant – perhaps because the department fixed effects already account for the most important spatial differences in location fundamentals.

Another possible concern is that our results may be affected by more general agglomeration ex-

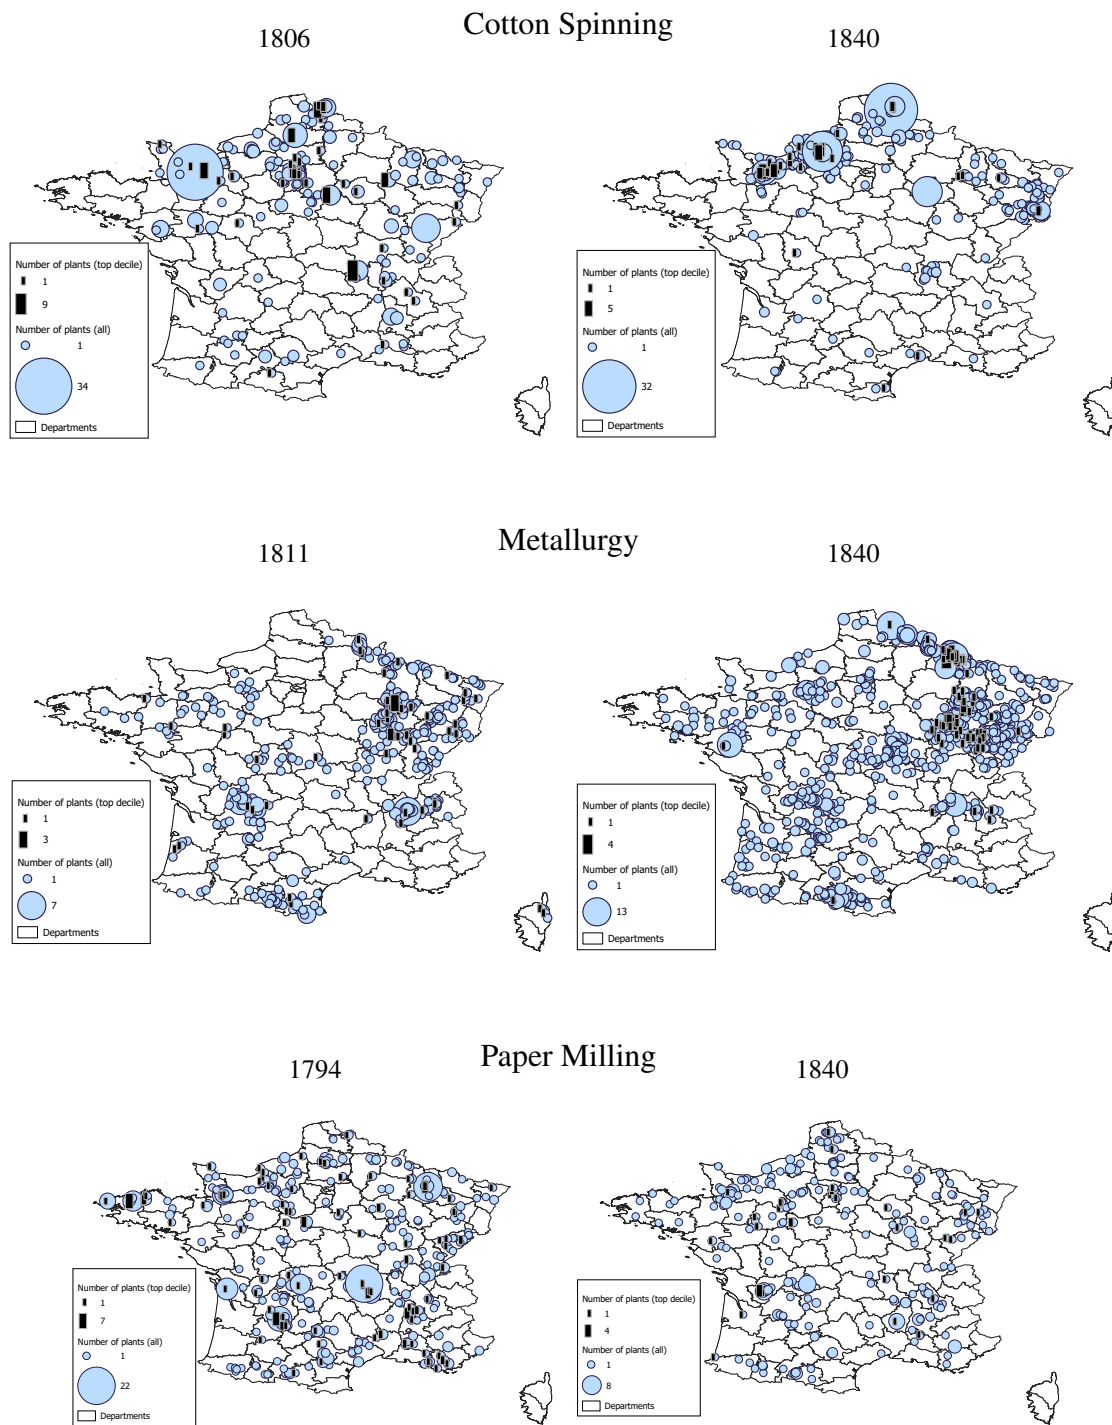


Figure A.13: Spatial Distribution of Plants Across France in the Three Sectors

Note: The figure shows the spatial distribution of plants in cotton spinning (top), metallurgy (middle), and paper milling (bottom). The figure distinguishes plants in the 90th percentile of the productivity distribution (black columns) from all other plants in a commune (light dots).

Table A.7: Proximity to High-Productivity Plants

Dependent variable: log(Output per worker)						
	(1)	(2)	(3)	(4)	(5)	(6)
	Spinning		Metallurgy		Paper milling	
	1806	1840	1811	1840	1794	1840
$Dist^{p90} (\sim 1800)$	-0.841*** (0.135)		-0.249*** (0.085)		-0.213* (0.125)	
$Dist^{p90} (1840)$		-0.178* (0.100)		-0.115 (0.107)		-0.182 (0.127)
Department FE	✓	✓	✓	✓	✓	✓
R ²	0.56	0.15	0.37	0.27	0.30	0.46
N	290	471	377	746	460	312

Notes: $Dist^{p90} (\sim 1800)$ and $Dist^{p90} (1840)$ measure the log distance to the nearest plant in the same sector with productivity in the 90th percentile in 1800 and in 1840, respectively. The number of observations in these specifications is smaller than the full sample size as all plants that belong to the 90th percentile are excluded. Standard errors (clustered at the departmental level) in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

ternalities, as opposed to learning. In particular, our findings could be driven by high-productivity plants emerging (within departments) where the density of production is large due to agglomeration forces. To address this possibility, Table A.9 controls for the density of production at the commune level. This is measured as the log of total output in the sector, excluding a plant's own output. The main pattern holds.

Next, Table A.10 performs a placebo exercise and examines whether plant productivity in 1800 was also related to the distance to high-productivity plants in the top-90th percentile of productivity in 1840. The estimated coefficient in cotton spinning is close to zero and insignificant, implying that productivity in cotton spinning in 1806 was not related to high-productivity locations three decades later. This suggests that our results are not driven by persistent location fundamentals within departments.

Finally, we examine whether there is evidence consistent with learning across sectors. We are mostly interested in whether mechanized cotton spinning plants located closer to high productivity metallurgy and paper milling plants in 1800 were more productive. Table A.11 shows that there is no consistent pattern in the data that would point to this type of learning. The coefficient on distance to high productivity metallurgy plants (column 1) is not statistically different from zero and very noisily estimated. This is arguably because cotton spinning plants were located relatively far away from metallurgy plants. In paper milling, the coefficient of interest is actually *positive* and significant – that is, there is a positive association between between the productivity of cotton

Table A.8: Proximity to High-Productivity Plants – Controlling for Location Fundamentals

Dependent variable: log(Output per worker)						
	(1) Spinning		(3) Metallurgy		(5) Paper milling	
	1806	1840	1811	1840	1794	1840
$Dist^{p90}$ (~ 1800)	-0.882*** (0.116)		-0.258*** (0.080)		-0.199 (0.128)	
$Dist^{p90}$ (1840)		-0.191* (0.103)		-0.123 (0.099)		-0.190 (0.118)
Access to high streamflow	-0.115 (0.283)	0.270** (0.113)	-0.039 (0.159)	0.245 (0.215)	-0.169 (0.247)	-0.182 (0.204)
Proximity to coal	-0.011 (0.201)	-0.084 (0.321)	-0.249 (0.185)	0.074 (0.155)	0.165 (0.357)	-0.160 (0.226)
Share forest area	-1.337*** (0.484)	0.380 (0.342)	-0.144 (0.314)	0.004 (0.390)	0.338 (0.504)	0.288 (0.861)
Department FE	✓	✓	✓	✓	✓	✓
R ²	0.58	0.16	0.38	0.28	0.30	0.47
N	290	471	369	746	460	312

Notes: $Dist^{p90}$ (~ 1800) and $Dist^{p90}$ (1840) measure the log distance to the closest plant in the same sector with productivity in the 90th percentile in 1800 and in 1840, respectively. Access to high streamflow is a binary variable that takes the value of one if a plant's nearest data collection point for river discharge has streamflow in the top quartile of the distribution. Proximity to coal is a binary indicator that takes the value of one if a location is within the bottom quartile of plant locations in terms of distance to the nearest coalfield. Share forest area measures the forest area over the total area of the commune where the plant is located (using data on forest cover from the late 18th century). The number of observations in these specifications is smaller than the full sample size as all plants that belong to the 90th percentile are excluded. Standard errors (clustered at the departmental level) in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.9: Proximity to High-Productivity Plants – Controlling for Local Production Density

Dependent variable: log(Output per worker)						
	(1)	(2)	(3)	(4)	(5)	(6)
	Spinning		Metallurgy		Paper milling	
	1806	1840	1811	1840	1794	1840
$Dist^{p90}$ (~ 1800)	-0.771*** (0.175)		-0.245*** (0.085)		-0.178 (0.127)	
$Dist^{p90}$ (1840)		-0.150 (0.112)		-0.129 (0.100)		-0.179 (0.135)
Production density	0.020 (0.021)	0.007 (0.014)	0.003 (0.012)	-0.004 (0.008)	0.018 (0.015)	0.001 (0.019)
Department FE	✓	✓	✓	✓	✓	✓
R ²	0.57	0.15	0.37	0.27	0.30	0.46
N	290	471	377	746	460	312

Notes: $Dist^{p90}$ (~ 1800) and $Dist^{p90}$ (1840) measure the log distance to the closest plant in the same sector with productivity in the 90th percentile in 1800 and in 1840, respectively. The number of observations in these specifications is smaller than the full sample size as all plants that belong to the 90th percentile are excluded. Standard errors (clustered at the departmental level) in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.10: Proximity to High-Productivity Plants – Distance Placebo in 1840

Dependent variable: log(Output per worker)			
	Spinning	Metallurgy	Paper milling
	1806	1811	1794
	(1)	(2)	(3)
$Dist^{p90}$ (1840)	-0.053 (0.235)	-0.235 (0.160)	0.192 (0.143)
Department FE	✓	✓	✓
R ²	0.55	0.30	0.23
N	321	415	507

Notes: $Dist^{p90}$ (1840) measures the log distance to the closest plant in the same sector with productivity in the 90th percentile in 1840. Standard errors (clustered at the departmental level) in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

spinning plants and distance to high productivity paper milling plants, which is the opposite of what we would expect if learning from the more mature comparison sectors was an important feature of the data. It instead suggests that high productivity paper milling plants crowded out productive cotton spinners.

Table A.11: Proximity of Cotton Spinning Plants to High-Productivity Plants in Metallurgy and Paper Milling in 1800

Dependent variable: log(Output per worker)		
	Spinning-Metallurgy	Spinning-Paper
	(1)	(2)
$Dist^{p90}$ (~ 1800)	-0.461 (0.351)	0.399*** (0.125)
Department FE	✓	✓
R ²	0.56	0.58
N	321	321

Notes: $Dist^{p90}$ (~ 1800) measures the log distance of cotton spinning plants to the closest plant in metallurgy (col. 1) and in paper milling (col. 2) with productivity in the 90th percentile in 1800. Standard errors (clustered at the department level) in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.4 Plant Age and Productivity in Metallurgy

Following the results from the spatial regressions, we turned to examining the data for patterns consistent with i) learning about the new technology or ii) learning about best practice organizational methods. One piece of evidence that points to the latter mechanism is the higher productivity of younger firms in mechanized cotton spinning in 1806. Here, we investigate whether a similar pattern holds in metallurgy, where similar data are available.

Tables A.12 and A.13 examine the relationship between age profile and productivity for metallurgy plants. For both periods, the best measure of plant age available is a binary indicator of plant survival from 1788 to 1811, and from 1811 to 1840. The results do not point to younger plants in metallurgy having a strong productivity advantage. Table A.12 (column 1) shows a positive association between ‘young’ metallurgy plants in 1811 and labor productivity, which, however, becomes smaller in magnitude and insignificant once we control for plant size (column 2). At the same time, in 1840 ‘young’ metallurgy plants were associated with somewhat *lower* productivity, although the negative coefficient is statistically significant only when we control for the number of workers (Table A.13).

Table A.12: Metallurgy in 1811: Productivity and Plants' Age Profile

Dep. variable: log(Output per worker)		
	(1)	(2)
Young 1811	0.226*	0.101
	(0.118)	(0.117)
log(Workers)		-0.313***
		(0.051)
R ²	0.01	0.10
N	448	448

Notes: Young 1811 is a binary indicator equal to one for plants that entered the market after 1788. Robust standard errors in parentheses. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

Table A.13: Metallurgy in 1840: Productivity and Plants' Age Profile

Dependent variable: log(Output per worker)					
	(1)	(2)	(3)	(4)	(5)
Entrant 1840	-0.084	-0.029	-0.080	-0.078	-0.144**
	(0.077)	(0.077)	(0.078)	(0.077)	(0.065)
Water power		0.327***			
		(0.062)			
Steam power			-0.045		
			(0.076)		
Other power				-0.193**	
				(0.090)	
log(Workers)					-0.373***
					(0.027)
R ²	0.00	0.03	0.00	0.01	0.24
N	839	839	839	839	839

Notes: Entrant 1840 is a binary indicator equal to one for plants that entered the market after 1811. Water power, steam power, and other (wind or animal) power are binary indicators equal to one for plants using the respective source of power. Robust standard errors in parentheses. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

E.5 Power Sources in the Three Sectors

In Section 4.2 in the main text, we showed that the survival rate in spinning was lower than in the comparison sectors. This is consistent with a mechanism in which entrepreneurs that invested in cotton spinning mills with poor layout had to exit the market, and the mill was not subsequently used by other cotton spinning entrepreneurs. However, it could also be consistent with the fact that the industry adopted steam power (and moved away from water power) more than the other sectors. Table A.14 and Table A.15 suggest that this was not the case. Table A.14 shows that even in spinning, water remained the prominent source of power until the end of our sample period in 1840.¹⁷ Moreover, Table A.15 shows a weak *negative* association between labor productivity and the use of steam power, confirming that in France, plants did not face a strong profit incentive to move away from water power.

Table A.14: Sources of Energy in the Three Sectors (1840)

Sector	Share of plants using:				Number of plants
	Water power	Steam power	Other power	Any power	
	(1)	(2)	(3)	(4)	(5)
Cotton spinning	0.66	0.39	0.02	0.93	528
Metallurgy	0.64	0.16	0.09	0.77	839
Paper milling	0.85	0.12	0.02	0.88	348

Notes: The table reports the share of plants in each of the three sectors using the source of power reported in the header. In column (3) 'other power' refers to wind and animal power.

¹⁷This is a well-known aspect of the French setting; see Cameron (1985) for a discussion.

Table A.15: Productivity and the Use of Steam Power in Cotton Spinning (1840)

Dependent variable: log(Output per worker)				
	(1)	(2)	(3)	(4)
Steam power	-0.090** (0.046)		-0.082* (0.048)	-0.087* (0.046)
Water power		0.060 (0.050)	0.017 (0.053)	
Other power				0.144 (0.142)
R ²	0.01	0.00	0.01	0.01
N	528	528	528	528

Notes: Water power, steam power, and other (wind or animal) power are binary indicators equal to one for plants using the respective source of power. Robust standard errors in parentheses. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

E.6 Robustness: Region Fixed Effects

In Section 4.3, we examined a set of alternative explanations that may account for the lower-tail bias observed in mechanized cotton spinning. One key robustness check studies the extent to which our main result holds *within* regions. Table A.16 shows that the lower-tail bias, while more muted, remains striking when we add a binary indicator for each of 22 regions in France. This suggests that local fundamentals, or differential access to input markets do not account for our results. Moreover, they also show that any uneven effect of political and institutional changes, or wartime disruptions cannot fully account for our results.

Table A.16: Productivity Growth by Quantiles – Controlling for Region FE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Average	At the following quantiles:					N
		0.1	0.25	0.5	0.75	0.9	
Spinning (1806-1840)	2.028*** (0.158)	2.941*** (0.455)	2.352*** (0.208)	1.982*** (0.168)	1.943*** (0.218)	1.659*** (0.191)	844
Region FE	✓	✓	✓	✓	✓	✓	
Metallurgy (1811-1840)	1.766*** (0.181)	2.012*** (0.211)	1.317*** (0.273)	1.781*** (0.158)	1.838*** (0.139)	1.786*** (0.214)	1243
Region FE	✓	✓	✓	✓	✓	✓	
Paper milling (1794-1840)	0.785*** (0.118)	0.928*** (0.099)	0.827*** (0.132)	0.730*** (0.098)	0.614*** (0.120)	0.664*** (0.106)	853
Region FE	✓	✓	✓	✓	✓	✓	

Notes: The table reports the average annual productivity growth (in %) between the initial sample period (around 1800) and 1840 for the three sectors, as well as annual productivity growth estimated at various quantiles. Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.7 Robustness: Market Integration

Among the spatially differential effects that could account for our results, we address a particularly prominent one more specifically – that of market integration. While the robustness of the main findings to the inclusion of region fixed effects make it unlikely that market integration can account for our results, we present several additional pieces of evidence underscoring why this is the case. Figure A.14 uses data in 1794 from Daudin (2010) on the number of districts across France that reported consuming cotton textiles, iron, or paper products from any district in a given department. The figure shows that many departments produced cotton textiles mostly for themselves, while a few departments supplied cotton to a large number of districts – consistent with a high degree of market integration already in 1800. On the other hand, in the two comparison sectors, there is less specialization and less evidence for market integration. This suggests that cotton textile plants were already competing in bigger markets than the comparison sectors around 1800.

Moreover, Table A.17 controls directly for market potential. Reassuringly, when adding these controls, the lower-tail bias of productivity growth remains a robust feature of the mechanized cotton spinning sector. In fact, the magnitude of the lower tail-bias remains very stable suggesting that market integration has little explanatory power.

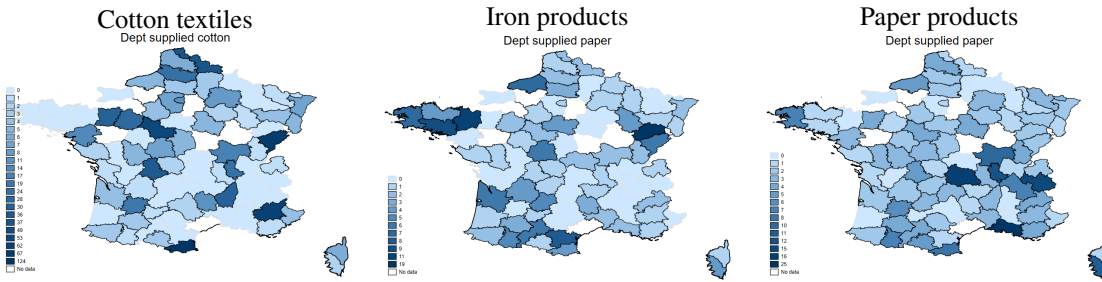


Figure A.14: Market Integration for the Three Sectors in 1794

Notes: Data source: Daudin (2010). The figure shows the extent of market integration in cotton textiles (left), iron (middle), and paper products (right). Market integration is measured as the number of districts across France that reported consuming cotton textiles, iron, or paper products from any district in a given department. A higher number for a department indicates that firms from that department sell their products to many distinct locations across France. Data are from Daudin (2010).

Table A.17: Productivity Growth by Quantiles – Controlling for Market Access

	(1) Average	(2) 0.1	(3) 0.25	(4) 0.5	(5) 0.75	(6) 0.9	(7) N
Spinning (1806-1840)	2.444*** (0.157)	3.978*** (0.221)	3.028*** (0.203)	2.170*** (0.158)	1.697*** (0.167)	1.251*** (0.313)	844
Market access	0.349*** (0.095)	0.401*** (0.138)	0.397*** (0.097)	0.229*** (0.075)	0.279*** (0.104)	0.483*** (0.186)	
Metallurgy (1811-1840)	1.951*** (0.189)	1.438*** (0.431)	1.687*** (0.287)	1.881*** (0.232)	2.161*** (0.213)	2.488*** (0.248)	1242
Market access	0.136 (0.198)	0.979** (0.431)	0.190 (0.409)	-0.378 (0.320)	-0.114 (0.142)	-0.421* (0.236)	
Paper milling (1794-1840)	0.710*** (0.110)	0.735*** (0.164)	0.685*** (0.135)	0.759*** (0.098)	0.743*** (0.136)	0.409 (0.259)	853
Market access	0.680*** (0.187)	0.209 (0.205)	0.314 (0.271)	0.433** (0.171)	0.537* (0.307)	1.775*** (0.664)	

Notes: The table reports the average annual productivity growth (in %) between the initial sample period (around 1800) and 1840 for the three sectors, as well as annual productivity growth estimated at various quantiles. Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** p<0.01, ** p<0.05, * p<0.1.

E.8 Robustness: Trade Protection and Napoleonic Blockade

Another alternative explanation that relies on spatial reallocation is the potential confounder of the Napoleonic Blockade. While robustness to the inclusion of region fixed effects rule out that this can explain the full lower-tail bias of productivity growth, we now examine this alternative explanation in more detail. First, Figure A.15 splits the sample into plants into the northern and southern regions of France (corresponding to the main dimension along which protection varied). It shows that the productivity distributions in the two areas are remarkably similar, suggesting that varying trade protection does not drive our results.




Figure A.15: Spinning: Productivity growth in the ‘North’ and ‘South’ of France

Notes: Northern communes are the set of plants with above-median latitude in 1806. Southern communes are those located in below-median latitude communes. Median latitude is time invariant and defined based on the sample of 1806 plants.

In addition, we also examine plant survival rates separately by northern and southern regions. Table A.18 shows that survival rates remain low when we only focus on plants in the northern parts of France. These are the locations where the blockade protected plants from British competition. Indeed, we find that survival rates are higher in the northern parts of the country consistent with Juhász (2018), yet they remain low relative to the other sectors.

Table A.18: Plant Survival Rates in Cotton Spinning in the North and South of France

	'North'	'South'
Survival rate	11%	4%
Number of plants	202	184
Restricted sample survival rate	10%	5.4%
Number of plants	31	56

Notes: The “survival rate” is defined as the percentage of plants from the initial period that survive to the later period based on matching either by name or location (see Section 3.5 for detail). The “restricted sample survival rate” adjusts for the fact that different sectors have single-plant communes to a varying degree. It is based on the subset of plants located in communes that have only one plant in the initial period and that either do not show up in the 1840 data or they show up with only one plant. Northern communes are the set of plants with above-median latitude in 1806. Southern communes are those located in below-median latitude communes. Median latitude is time invariant and defined on the sample of 1806 plants.

E.9 Robustness: Excluding Small Plants and Those That Use Spinning Jennies

Another concern is that our results may be driven by the disappearance of small cotton spinning workshops. To address this, we first adopt a stricter definition of plant production and use only those that have more than 10 employees. This should exclude the majority of smaller workshops that may only be organized as factory-based production along some, but not all dimensions. Table A.19 show that the lower-tail bias in productivity growth is robust to using only larger plants. The magnitudes remain similar to those in the baseline specification the lower-tail bias of productivity growth remains unique to cotton spinning plants.

As a final check on the effect of small workshops, we adopt an even more conservative approach and drop all plants that report using a spinning jenny (even if they use other, newer vintages of capital in addition). Table A.20 shows that the lower-tail bias of productivity growth remains striking.

Table A.19: Productivity Growth by Quantiles – Plants with at Least 10 Workers

	(1) Average	(2) 0.1	(3) 0.25	(4) 0.5	(5) 0.75	(6) 0.9	(7) N
		At the following quantiles:					
Spinning (1806-1840)	2.261*** (0.177)	3.917*** (0.227)	3.191*** (0.258)	2.179*** (0.170)	1.612*** (0.240)	0.309 (0.292)	777
Metallurgy (1811-1840)	1.990*** (0.235)	1.751** (0.759)	1.578*** (0.405)	1.523*** (0.275)	2.029*** (0.235)	1.845*** (0.243)	905
Paper milling (1794-1840)	1.245*** (0.136)	1.024*** (0.225)	1.086*** (0.162)	1.186*** (0.121)	1.434*** (0.159)	1.289*** (0.274)	507

Notes: The table reports the average annual productivity growth (in %) between the initial sample period (around 1800) and 1840 for the three sectors, as well as annual productivity growth estimated at various quantiles. Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.20: Productivity Growth by Quantiles – Excluding Plants that used Spinning Jennies

	(1) Average	(2) 0.1	(3) 0.25	(4) 0.5	(5) 0.75	(6) 0.9	(7) N
		At the following quantiles:					
Spinning (1806-1840)	1.983*** (0.169)	3.165*** (0.303)	2.678*** (0.256)	2.097*** (0.187)	1.213*** (0.228)	0.383 (0.248)	792

Notes: The table reports the average annual productivity growth (in %) between 1806 and 1840 in cotton spinning, as well as annual productivity growth estimated at various quantiles. Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.10 Robustness: Controlling for Plant Scale

We also address a set of alternative explanations concerning plant level characteristics. One concern is that our findings may be driven by increasing plant scale. For this reason, Table A.21 controls for the number of workers at the plant level. Our results continue to hold.

Table A.21: Spinning Productivity by Quantiles – Controlling for Number of Workers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Average	At the following quantiles:					N
		0.1	0.25	0.5	0.75	0.9	
Spinning (1806-1840)	2.427*** (0.162)	3.941*** (0.231)	3.427*** (0.243)	2.292*** (0.165)	1.836*** (0.185)	0.974*** (0.304)	868
log(Workers)	-0.006 (0.063)	-0.073 (0.092)	-0.072 (0.090)	-0.048 (0.058)	-0.169** (0.071)	-0.257** (0.115)	
Metallurgy (1811-1840)	2.852*** (0.177)	3.296*** (0.438)	2.539*** (0.275)	2.552*** (0.214)	2.916*** (0.184)	2.507*** (0.202)	1296
log(Workers)	-1.219*** (0.082)	-1.338*** (0.143)	-1.183*** (0.093)	-1.193*** (0.099)	-1.105*** (0.081)	-1.066*** (0.083)	
Paper milling (1794-1840)	0.808*** (0.125)	0.744*** (0.139)	0.627*** (0.157)	0.780*** (0.112)	0.955*** (0.161)	1.505*** (0.188)	868
log(Workers)	-0.100* (0.060)	0.206*** (0.066)	0.111 (0.073)	-0.002 (0.046)	-0.124* (0.073)	-0.432*** (0.051)	

Notes: The table reports the average annual productivity growth (in %) between the initial sample period (around 1800) and 1840 for the three sectors, as well as annual productivity growth estimated at various quantiles. We control for the number of workers at the plant level across all specifications. Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.11 Robustness: Capital Deepening

Table A.22 examines whether capital deepening can account for our results. It controls for the capital-labor ratio in cotton spinning plants (measured as the log of the number of spindles per employee at the plant level). The table shows that the lower-tail bias of productivity growth remains robust and similar in magnitude to the pattern in the baseline specifications.

Table A.22: Spinning Productivity by Quantiles – Controlling for Capital Deepening

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Average	At the following quantiles:					N
		0.1	0.25	0.5	0.75	0.9	
Spinning (1806-1840)	1.960*** (0.167)	3.555*** (0.267)	2.930*** (0.247)	1.966*** (0.178)	1.254*** (0.190)	0.755*** (0.281)	868
K/L	0.522*** (0.075)	0.374*** (0.082)	0.467*** (0.085)	0.389*** (0.068)	0.379*** (0.090)	0.542*** (0.141)	

Notes: The table reports the average annual productivity growth (in %) between 1806 and 1840 in cotton spinning (column 1), as well as annual productivity growth estimated at various quantiles (cols 2-6). Column 7 reports the number of observations. *K/L* is the capital-labor ratio in the plant, measured as the log of the number of spindles per employee.) Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.12 Robustness: Quality Upgrading

Finally, we also considered quality upgrading as a potential confounder in Section 4.3. Table A.23 addresses the concern that even if quality is not directly driving the lower-tail bias, it could still affect it indirectly through higher sales and larger plant size. To address this, we estimate quantile regressions using prices *not* adjusted for quality in 1806, *and* controlling for plant size using the number of employees. The lower-tail bias of productivity growth in cotton spinning remains strong.

Table A.23: Annual Productivity Growth (in %) at Different Parts of the Distribution—Using prices not quality-adjusted and controlling for number of workers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Average	At the following quantiles:					N
		0.1	0.25	0.5	0.75	0.9	
Spinning (1806-1840)	2.517*** (0.143)	3.488*** (0.295)	2.886*** (0.223)	2.420*** (0.198)	1.994*** (0.141)	1.657*** (0.218)	868
log(Workers)	-0.140** (0.058)	-0.055 (0.103)	-0.127 (0.087)	-0.210*** (0.068)	-0.200*** (0.057)	-0.211** (0.090)	

Notes: The table reports the average annual productivity growth (in %) between 1806 and 1840 in cotton spinning, as well as annual productivity growth estimated at various quantiles. Column 7 reports the number of observations. Robust standard errors in parentheses. Notation for statistical significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

F Summary of Robustness Checks

<u>Alternative Mechanisms</u>	
<i>Our argument is that the lower-tail bias in mechanized cotton spinning is due to the reorganization of production. Here, we test for alternative (potentially confounding) mechanisms. Importantly, factors that affected the comparison sectors in similar ways (e.g., economy-wide effects, the introduction of innovations during our sample period, and improvements in power sources) are unlikely to explain our findings. Thus, we consider confounders that are either specific to cotton spinning or that affected this sector differentially.</i>	
Type of evidence: Historical/Empirical	<u>Description of Evidence</u>
A. Regional Differences	
E: Table A.16	Different growth potential of French regions, access to better domestic suppliers of machines, or plants sorting into areas with better location fundamentals (such as fast-flowing streams) could potentially be driving our results. Similarly, wars and revolts (we deal explicitly with the various effects of the Napoleonic Blockade below) affected some regions of France more than others. Table A.16 performs the quantile regressions including region fixed effects and shows that our results are robust to using only within-region variation.
<i>A.1 Market Integration</i>	
E: Figure A.14, Tables A.17 and A.16	As the French economy became more integrated, lower-productivity firms might have faced stronger competition and exited the market. However, this should have affected mechanized cotton spinning differentially, as we do not observe the lower-tail bias in the comparison sectors. Figure A.14 shows the extent of market integration in the three sectors and suggests that cotton plants were already competing in a larger market around 1800. In addition, Table A.17 performs the quantile regressions controlling for market potential. The results on the lower-tail bias hold.
<i>A.3 Machine Quality</i>	
H: Section 2.1 and Appendix Section A.2	The quality of machines available to producers could potentially explain our findings. However, a large body of historical evidence (reported in Section 2.1 and Appendix Section A.2) argues that machine production and maintenance was external to most plants. This suggests that plants within the same region had access to the same suppliers. As our results hold when using within-region variation, it is unlikely that low quality machines can account for the lower-tail bias.
<i>A.2 The Napoleonic Blockade</i>	
E: Figure A.15 and Table A.16	Varying trade protection during the Napoleonic Blockade affected the location of mechanized cotton spinning plants raising the concern that the blockade may explain the lower-tail bias. Table A.16 shows that the results hold within region, where the pattern of protection was very similar. In addition, Figure A.15 splits the sample in northern and southern regions (the main dimension along which protection varied) and shows that in both regions, productivity growth until 1840 was due to a disappearing lower tail.

Alternative Mechanisms (continued)

Our argument is that the lower-tail bias in mechanized cotton spinning is due to the reorganization of production. Here, we test for alternative (potentially confounding) mechanisms. Importantly, factors that affected the comparison sectors in similar ways (e.g., economy-wide effects, the introduction of innovations during our sample period, and improvements in power sources) are unlikely to explain our findings. Thus, we consider confounders that are either specific to cotton spinning or that affected this sector differentially.

Type of evidence: Historical/Empirical	<u>Description of Evidence</u>
B. Early Spinning Workshops	
E: Section 3, Tables A.19 and A.20	Our results may be driven by the disappearance of small cotton spinning plants. This concern is partly mitigated by the fact that our baseline sample excludes small establishments using handspinning wheels (see Section 3). However, it may still be the case that there were systematic differences between small plants operating early vintages of mechanized spinning jennies (that did not necessarily need inanimate sources of power) and larger-scale factories. We test for this in two ways: first, we adopt a stricter definition of ‘factory-production’ and use only plants with more than 10 employees (see Table A.19); second, we drop plants using the earliest vintage of machinery, the spinning jenny (see Table A.20). In all cases, the lower-tail bias remains striking.
C. Changes at the Plant Level	
<i>Plant Scale</i>	
E: Tables A.3 and A.21	Our findings could be driven by increasing plant scale. First, this is unlikely as all sectors witnessed an increase in plant scale (see Table A.3). In addition, Table A.21 shows that the results are robust to controlling for the number of workers (at the plant level).
<i>Capital Deepening</i>	
E: Table A.22	Productivity growth in cotton spinning could be driven by technological improvements of the mechanized machinery after 1800. One important dimension of improving technology in cotton spinning was the increased capital per unit of labor (measured as spindles per worker). Table A.22 controls for the capital-labor ratio at the plant level (measured as the number of spindles per worker). The lower-tail bias for productivity growth remains robust.
<i>Quality Upgrading</i>	
E: Tables A.5 and A.23	The lower-tail bias could be driven by quality upgrading across plants over time. Table A.5 (Panel B) uses prices not quality-adjusted and shows that our results hold. In addition, if quality led to higher sales and thus larger plant size, it could still drive our results indirectly. Table A.23 estimates the quantile regressions not adjusting for quality differences in prices across plants and controlling for plant size. The lower-tail bias is robust to this specification.

Alternative Mechanisms (continued)

Our argument is that the lower-tail bias in mechanized cotton spinning is due to the reorganization of production. Here, we test for alternative (potentially confounding) mechanisms. Importantly, factors that affected the comparison sectors in similar ways (e.g., economy-wide effects, the introduction of innovations during our sample period, and improvements in power sources) are unlikely to explain our findings. Thus, we consider confounders that are either specific to cotton spinning or that affected this sector differentially.

Type of evidence: Historical/Empirical	<u>Description of Evidence</u>
<i>Age Profile of Plants</i>	
E: Tables 6	The median age of mechanized cotton spinning plants in 1806 is only 3 years. One concern is that our results are driven by the fact that cotton spinning plants are very young and inexperienced, thus operating the new technology at varying levels of efficiency and leading to the lower-tail bias. Table 6 shows exactly the opposite: younger plants are substantially more productive than older ones.

Evidence in Line with Learning About Factory-Based Production

We argue that learning about the efficient organization of factory-based production can explain the lower-tail bias of productivity growth in mechanized cotton spinning. Here we summarize the evidence in line with our proposed mechanism.

Type of evidence: Historical/Empirical	<u>Description of Evidence</u>
<i>Learning About Best-Practice Methods</i>	
H: Section 2	Section 2 documents in details how early adopters in mechanized cotton spinning needed to engage in trial and error along multiple dimensions. This process led to the development of best-practice methods for operating the new technology efficiently.
<i>Spatial Diffusion of Knowledge</i>	
E: Figure 4 and Table A.7	Figure 4 and Table A.7 show that plants located closer to other plants (in the same sector) with productivity in the 90th percentile were themselves more productive, and this relationship is strong only for cotton spinning plants and only during the initial period of technology adoption. This evidence is in line with spatial spillovers of knowledge in mechanized cotton spinning in 1800.
<i>Plant Survival Across Sectors</i>	
E: Table 4	Table 4 shows that plant survival rates in mechanized cotton spinning was lower than in our comparison sectors. This is consistent with a mechanism in which owners of a cotton spinning mill faced considerable challenges along the organizational dimension of factory design. Those that invested in mills with poor layout had to exit the market.
<i>Plant Exit and Productivity</i>	
E: Table 5	Table 5 shows that in mechanized cotton spinning exiting plants were substantially less productive than those that survived. This is consistent with large organizational challenges and low initial guidance in switching to factory-based production in cotton spinning. It also suggests that learning about the new technology is not the dominant dimension of learning, as we would expect this type of learning to give incumbents who have accumulated more experience an advantage.
<i>Age Profile of Plant Productivity</i>	
E: Tables 6, 7, A.12, and A.13	Table 6 shows that 1800, cotton spinning plants that entered the market later had higher productivity. This is in line with the argument that knowledge about the optimal organization of mechanized cotton spinning diffused slowly. On the other hand, if learning about technology was the dominant dimension, older plants that had accumulated more experience should have had a productivity advantage. In addition, we show that younger spinning plants were not more productive in 1840 (Table 7) and that in metallurgy, young plants were as productive as older ones in both periods (Tables A.12 and A.13). This is consistent with the idea that later entrants in mechanized cotton spinning could draw from a better pool of knowledge about organizing production. This process was muted in cotton spinning in 1840, when best practice about had diffused and in metallurgy, where plant-based production methods had been established much earlier.

Alternative Explanations for the Spatial Diffusion of Knowledge

Our argument, supported by the historical evidence (see Section 2.1), is that plants copied organizational solutions by observing successful experimenters nearby. Empirically, we find support for the spatial diffusion of knowledge. Plants located closer to other high-productivity plants were themselves more productive, and this relationship is strong only for cotton spinning plants and only during the initial period of technology adoption. Here we show that this result holds when performing a series of robustness checks.

Type of evidence: Historical/Empirical	<u>Description of Evidence</u>
<i>Location Fundamentals</i>	
E: Table A.8	Our results could be affected by prominent location fundamentals, not captured by department fixed effects (such as the availability of fast-flowing streams, proximity to coal, and the share of forest cover). Table A.8 controls for these factors and shows that the pattern of proximity to high-productivity plants holds.
<i>Agglomeration Externalities</i>	
E: Table A.9	Our findings could be driven by high-productivity plants emerging (within departments) where the density of production is large due to agglomeration forces. Table A.9 controls for the local density of production and shows that our results hold.
<i>Unobserved Location Fundamentals</i>	
E: Table A.10	If there are unobserved location fundamentals (within departments) not captured by our controls (in Tables A.8 and A.9), they could still confound our results. Table A.10 performs a placebo exercise and studies whether plant productivity in 1800 was related to distance to plants in the top-90th percentile of productivity in 1840 (i.e., plants that emerged later). The coefficient is close to zero and statistically insignificant, suggesting that our results are not driven by persistent unobserved location fundamentals within departments.
<i>Plant Selection</i>	
E: Table 3	Another concern is that ex-ante high-productivity plants selected into ‘productive locations’ (i.e., chose to locate near existing high-productivity plants). Table 3 shows that the results hold when controlling for log plant age, adding the interaction of plant age and distance to high productivity plants, and when focusing only on the subsample of plants that entered before the nearest high-productivity plant.
<i>Similar Types of Learning Across Sectors</i>	
E: Table 3	One concern could be that there were similar types of learning across sectors and that 1806 cotton spinning plants closer to high productivity metallurgy or paper milling plants were also more productive. Table A.11 shows that this was not the case, suggesting that early cotton spinning mills were unlikely to learn from high productivity plants in these more mature sectors.

References

- Allen, R. C. (2009). *The British Industrial Revolution in Global Perspective*. Cambridge University Press.
- André, L. (1996). *Machines à papier: Innovation et transformations de l'industrie papetière en France : 1798-1860*. Editions de l'Ecole des hautes études en sciences sociales.
- Borio, C. and A. Filardo (2005). Looking Back at the International Deflation Record. *North American Journal of Economics and Finance* 15(3), 287–311.
- Boston Mutual Fire (1908). *Report No. 5: Slow-Burning or Mill Construction*. Boston Manufacturers Mutual Fire Insurance Co. 3rd ed., rev. and enl. Boston: The Company. .
- Bradley, M. (2010). Examples of Industrial and Military Technology Transfer in the Eighteenth Century. *Documents Pour L'histoire Des Techniques* 19, 87–95.
- Cameron, R. (1985). A New View of European Industrialization. *Economic History Review* 38(1), 1–23.
- Chassagne, S. (1991). *Le coton et ses Patrons: France 1760-1840*. Editions de l'Ecole des Hautes Etudes en Sciences Sociales.
- Clark, G. (1994). Factory Discipline. *The Journal of Economic History* 54(1), 128–163.
- Cookson, G. (1997). Family Firms and Business Networks: Textile Engineering in Yorkshire, 1780–1830. *Business History Review* 39(1), 1 – 20.
- Crouzet, F. (2003). The Historiography of French Economic Growth in the Nineteenth Century. *Economic History Review* 62(2), 215–242.
- Daudin, G. (2010). Domestic Trade and Market Size in Late-Eighteenth-Century France. *Journal of Economic History* 70(3), 716–743.
- Hanlon, W. W. (2020). The Rise of the Engineer: Inventing the Professional Inventor During the Industrial Revolution. Manuscript.
- Harley, C. (1998). Cotton Textile Prices and the Industrial Revolution. *Economic History Review* 51(1), 49–83.
- Harris, J. (1998). *Industrial Espionage and Technology Transfer: Britain and France in the Eighteenth Century*. Aldershot: Ashgate.
- Hauser, H. (1985). *Recherches et documents sur l'histoire des prix en France de 1500 à 1800*. Slatkine.
- Horn, J. (2006). *The Path Not Taken: French Industrialization in the Age of Revolution, 1750-1830*. Cambridge MA: MIT Press.
- Huberman, M. (1996). *Escape from the Market*. Cambridge University Press.
- Johnson, H. and A. Skempton (1955). William Strutt's Cotton Mills, 1793-1812. *Transactions of the Newcomen Society* 30(1).
- Juhász, R. (2018). Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade. *American Economic Review* 108(11), 3339–76.
- Lefebvre, G. (2011). *Napoleon*. Taylor and Francis.
- Maddison, A. (2001). *The World Economy. A Millennial Perspective*. Paris: OECD.
- McGaw, J. (1987). *Most Wonderful Machine: Mechanization and Social Change in Berkshire Paper Making, 1801-1885*. Princeton University Press.
- Mitchell, B. (2003). *International Historical Statistics. Europe, 1750-2000*. Palgrave Macmillan.
- Mokyr, J. (2005). The Intellectual Origin of Modern Economic Growth. *Journal of Economic History* 65(2),

285–351.

- Mokyr, J. (2010). *The Enlightened Economy: An Economic History of Britain 1700-1850*. Yale University Press.
- O'Brien, P. and C. Keyder (1978). *Economic Growth in Britain and France, 1780-1914: Two Paths to the 20th Century*. London: Allen and Unwin.
- Pollard, S. (1965). *The Genesis of Modern Management: A Study of the Industrial Revolution in Great Britain*. Harvard University Press.
- Pounds, N. and W. Parker (1957). *Coal and Steel in Western Europe: The Influence of Resources and Techniques on Production*. Bloomington: Indiana University Press.
- Riello, G. (2013). *Cotton: The Fabric that Made the Modern World*. Cambridge University Press.
- Rostow, W. (1975). *How it All Began: Origins of the Modern Economy*. Methuen and Co., London.
- Sargent, T. and F. Velde (1995). Macroeconomic Features of the French Revolution. *Journal of Political Economy* 103(3), 474–518.
- Squicciarini, M. and N. Voigtländer (2015). Human Capital and Industrialization: Evidence from the Age of Enlightenment. *Quarterly Journal of Economics* 30(4), 1825–1883.
- Szostak, R. (1989). The Organization of Work: The Emergence of the Factory Revisited. *Journal of Economic Behavior and Organization* 11(3), 343–358.
- White, E. (1991). Measuring the French Revolution's Inflation: the Tableaux de dépréciation. *Histoire et mesure* 6(3–4), 245–274.
- Williamson, O. (1980). The Organization of Work: A Comparative Institutional Assessment. *Journal of Economic Behavior and Organization* 1(1), 5–38.
- Woronoff, D. (1984). *L'industrie sidérurgique en France pendant la Révolution et l'Empire*. Paris: École des hautes études en sciences sociales.