



The March of the Techies: Job Polarization Within and Between Firms

James Harrigan, Ariell Reshef, Farid Toubal

► To cite this version:

James Harrigan, Ariell Reshef, Farid Toubal. The March of the Techies: Job Polarization Within and Between Firms. *Research Policy*, 2021, 50 (7), pp.104008. 10.1016/j.respol.2020.104008 . halshs-02973332

HAL Id: halshs-02973332

<https://shs.hal.science/halshs-02973332>

Submitted on 13 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

The March of the Techies: Job Polarization Within and Between Firms

James Harrigan

University of Virginia, Paris School of Economics, and NBER

Ariell Reshef

CNRS, Université Paris 1 Panthéon-Sorbonne,
Paris School of Economics and CEPII

Farid Toubal

ENS-Paris Saclay, CREST and CEPII

April 2020*

Abstract

Using administrative employee-firm-level data on the entire private sector from 1994 to 2007, we show that the labor market in France has polarized: employment shares of high and low wage occupations grew, while middle wage occupations shrank. At the same time, the share of technology-related occupations (“techies”) grew substantially. Aggregate polarization was driven mostly by changes in the composition of firms within industries. Within-firm adjustments and changes in industry composition were much less important. Polarization occurred mostly within urban areas, with roughly equal contributions of men and women. We study the role of technology adoption in shaping firm-level outcomes using a new measure of the propensity of a firm to adopt new technology: its employment share of techies. We find that techies were an important force driving aggregate polarization in France, as firms with more techies grew faster.

JEL classifications: J2, O3, D3, F1, F16, F66.

Keywords: Job polarization, technological change, STEM skills, techies, offshoring, firm level data.

*james.harrigan@virginia.edu, ariell.reshef@psemail.eu, ftoubal@ens-cachan.fr. This research is supported by a public grant overseen by the French National Research Agency (ANR) as part of the “Investissements d’avenir” program (ANR-10-EQPX-17 CASD). We also acknowledge financial support from the iCODE institute and from the ANR agency under the program ANR-16-CE26-0001-01. We thank Francis Kramarz for guidance with the data, and many seminar audiences in the United States and Europe for helpful comments.

1 Introduction

Job polarization—growth in the employment shares of high-wage and low-wage jobs at the expense of middle wage jobs—is one of the most striking phenomena in many advanced economies’ labor markets in the last several decades.¹ While job polarization has been mostly studied at aggregate levels, much less is known about the underlying firm-level mechanics that drive it. The firm-level dimension is important because firm-level decisions—for example, about investments in technological improvements and on global engagement—affect both overall firm employment growth and within-firm internal occupational changes, which together add up to aggregate changes.

We use administrative data for the entire private sector in France from 1994 to 2007 to show that France has indeed experienced job polarization: employment shares of high-wage managers and professionals increased; shares of middle-wage office workers and industrial workers fell; and shares of low-wage retail, personal service and unskilled manual workers increased. However, the picture that emerges is more complex than this simple relationship between wage ranks and changes in employment shares. For example, middle managers and technicians earn similar middle-income wages, yet employment in middle management occupations declined, while technicians increased their employment shares.

We then study the dimensions along which polarization happened: industries, firms, geography and gender. Using regressions motivated by theory, we also study whether technological change is associated with polarization through changing the composition of firms or by altering occupational shares within firms.

The use of matched firm-level data allows us to shed light on the dimensions along which job polarization occurs. The aggregate employment trends are visible across French urban and rural areas, at the industrial and firm level, and across gender. We compare how aggregate job polarization manifests in compositional changes of firms, industries, and departments—versus changes within these units.² As Cerina et al. (2017) do for the U.S, we also study whether polarization is more apparent for men than for women. We show that the main channel through which job polarization occurs is changes within industry in the composition of firms’ size, and not changes in occupational employment shares within firms. In other words, polarization didn’t happen because firms changed their employment mix, it happened because firms intensive in middle-wage jobs grew more slowly than firms intensive in high and low wage occupations. Indeed, even if firms had not changed their internal occupational composition at all, we would have seen almost the same polarization at the aggregate level, due to changes in firms’ sizes. This suggests that explanations that rely on substitution (either within local labor markets or

¹Job polarization has been documented in the United States (Autor et al. (2006), Autor et al. (2008), Firpo et al. (2011)), the United Kingdom (Goos and Manning (2007)), Germany (Spitz-Oener (2006), Dustmann et al. (2007)), and more generally in Europe (Goos et al. (2009) and Oesch (2013)). Polarization contrasts with earlier labor market developments, where changes in employment shares of middle-wage jobs were more modest, and the growth of high-wage jobs was at the expense of low-wage jobs. For example, in 1980s in the U.S., changes in employment shares are positively related to wages in the 1980s (Autor et al. (2008)).

²France is divided into 101 departments. Our analysis is restricted to the 94 departments that comprise mainland France. The median mainland department had a population of just above 500 thousand in 2007.

national industries) are missing an important dimension of the mechanics of polarization. Moreover, as Autor (2019) notes for the U.S., we find that job polarization is primarily an urban phenomenon.

These results lead us to ask what factors explain changes of firms' sizes and changes in occupational shares. Most explanations for job polarization focus on the "routinization hypothesis": new and cheaper information and communication technologies (ICT) perform "routine", codifiable tasks that would otherwise be performed by middle-wage workers. At the same time, ICT has been found to be particularly complementary to "non-routine", non-codifiable tasks (e.g., analysis, decision making) that are typically performed by high-wage workers.

In France, the price of ICT capital dropped by roughly 30% in our sample. We develop a simple model which shows how a drop in the cost of ICT capital affects within-firm occupational shares and relative firm size. In the model, variation in initial levels of ICT intensity across firms leads to different responses to a common drop in the price of ICT, with more ICT-intensive firms seeing their costs drop more. As a result of this cost-reduction effect, firms that are more ICT-intensive grow relative to less ICT-intensive firms. The effect of falling ICT prices on within-firm occupational polarization depends on how ICT substitutes for or is complementary to different types of labor. We show that under a broad set of parameters, more ICT-intensive firms will tend to experience greater occupational polarization.

We evaluate the predictions of the model, in particular how changes depend on initial ICT intensity, by estimating regressions using weighted least squares (WLS) and two-stage least squares (W2SLS). We use the initial share of techies in the firm as a proxy for initial ICT intensity, as suggested by the model. In addition to the role of routinization and technological progress, many authors have linked the phenomenon of polarization with globalization.³ The data allow us to measure the importer and exporter status of the firm. We find that export access to new markets increases employment, while firm level imports have ambiguous effects. Access to imported intermediate inputs will increase a firm's competitiveness, which will tend to raise employment, but may also directly substitute for some tasks formerly done by workers in the firm, thus reducing employment of some occupations. Concerning the composition of employment within firm, we expect a greater effect of international trade on occupations that are more directly exposed to it. Imports may explain the decline of mid-level wage jobs because they substitute the tasks associated with jobs that can be carried out by a less costly workforce abroad. Since these tasks are to be performed at long distance, they are more likely to be offshored if they require less face-to-face interpersonal interaction (Blinder (2006)). International trade can also favor high-wage workers by increasing demand for non-routine cognitive tasks, those associated with within-firm internal occupational changes, or those related to management and communication between the firm's affiliates that are located in different countries.

A novel contribution of our paper is our focus on the firm-level employment share of techies: workers with STEM skills and experience. We focus on techies because of their central role in planning, installing, and maintaining information and computer technology (ICT) and other technologies, as well as in training and assisting other workers in

³Important contributions including Autor and Dorn (2013), Autor et al. (2013), Autor et al. (2015), Goos et al. (2014), and Malgouyres (2017) link polarization with trade.

the use of technology. These roles make techies the crucial link between economy wide technological progress and firm level technology adoption. As noted by Tambe and Hitt (2014), “the technical know-how required to implement new IT innovations is primarily embodied within the IT workforce” in a firm. Brynjolfsson and Hitt (2003) observe that “In recent years, companies have implemented thousands of large and small innovations in software applications, work processes, business organization, supply-chain management, and customer relationship management”. Implementation of these innovations was undoubtedly mediated by techies, who are a good measure of the “organizational capital” that Brynjolfsson and Hitt (2003) argue is crucial to ICT adoption. The richness of our data allows us to make a distinction between two different categories of techie workers: highly-educated technical managers and engineers on the one hand, and lower-ranked technicians on the other.

Of course, not all technological change comes through ICT investment, and our measure of techies includes engineers, technology managers, and technicians more broadly. For example, Helper and Kuan (2018) survey the auto parts industry and find that innovation occurs through the efforts of firms’ engineers and technicians, although without explicit R&D expenditures. Our paper is the first to measure the number of these workers at the firm level in the entire economy, which allows us to associate techies with firm outcomes, including internal occupational change and employment growth.

The econometric results suggest strong effects of techies on firm-level employment growth which are mostly driven by engineers rather than technicians. We also find a positive association between techies and the share of managers within firms and a negative association between techies and the share of skilled and services workers.

Our findings on the importance of techies are consistent with their hypothesized role as a conduit for technological change, whether through ICT adoption or more broadly. Since techie workers have STEM skills and education, an alternative interpretation of our results is that it is STEM skills themselves that explain our findings, rather than technology adoption.

Our paper is organized as follows. Section 2 discusses the relationship to the existing literature. We then discuss our theoretical framework for firm-level analysis in Section 5. In Section 4 we discuss data, document job polarization in the French labor market and show how polarization has evolved mostly due to firm composition. In section 6, we develop the econometric framework for estimating the impact of technology and trade between and within firms. In Section 7, we present econometric results showing how firm characteristics are associated both with firm employment growth and with within-firm occupational change. Section 8 concludes.

2 Relationship to existing literature

This is one of the first papers to describe and analyze polarization across and within firms. Since employment decisions are made at the firm level, firm-level data is ideal for studying polarization. Other papers have looked at polarization across industries or regions (see Beaudry et al. (2010), Autor and Dorn (2013), Michaels et al. (2014), and Goos et al. (2014). Goos et al. (2014) is the only previous paper to have found evidence of job polarization in France, although their discussion is limited to a single line in their

Table 2, and using survey data that are inferior to our data sources. Caliendo et al. (2015) analyze within-firm internal occupational change in France with the same data that we use, focusing on hierarchies, but restrict their attention to manufacturing firms (25% of private sector hours) and do not discuss ICT or polarization.

The main explanation for job polarization in the literature is the “routinization hypothesis” (Goos and Manning (2007)). As argued in Autor et al. (2003), technological progress in information and communications technology (ICT) allows machines to replace codifiable cognitive routine tasks that were once performed by humans. These tasks happen to be more common in occupations that are, on average, in the middle of the wage distribution. Thus, the diffusion of ICT lowers demand for these occupations. At the same time, ICT complements non-routine cognitive tasks, and demand for occupations that are characterized by these tasks—which are higher up in the wage distribution—rises. Occupations at the bottom of the wage distribution are less affected by ICT, and they absorb the residual supply of labor.⁴ Our results broadly support the importance of the “routinization hypothesis”, but with important nuances.

Our findings on the importance of techies for overall employment growth, while substituting for some occupations, are consistent with Graetz and Michaels (2018), who find that industries that invested more in robotization reduced relative demand for workers who are close substitutes to robots, but overall increased labor demand. As they do, we interpret this finding through a product demand mechanism: the main effect of firm-level technological change is to lower prices and increase competitiveness, which boosts labor demand. Similarly, Goos et al. (2014) and Gregory et al. (2016) also estimate increases in overall labor demand within economic units (industries or local labor markets) that undergo differentially more technological change through a similar competitive effect.

The effect of ICT investment on growth and internal occupational change is shown in several studies that typically focus on a particular industry or specific settings. In contrast, our paper considers the entire French private sector. Lichtenberg (1995) and Brynjolfsson and Hitt (1996), working with a small number of U.S. firms in the late 1980s find that IT labor has a positive output elasticity. Tambe and Hitt (2012) corroborate this finding in a newer data source and more sophisticated estimation technique. Autor et al. (2002) study organizational change due to the introduction of digital check imaging within one large bank. Bresnahan et al. (2002) discuss the complementarity between IT, decentralized firm organization, and skilled labor. Our results below are consistent with this; we find that techies (which we also associate with IT) cause an increase in the employment share of top managers and a decrease of middle management employment shares within firms. However, our descriptive results imply that within-firm re-organization may not be the most important margin of adjustment.

Several studies consider the manufacturing sector due to data availability. Our work is closely related to Maurin and Thesmar (2004), who investigate changes in organization within French manufacturing firms in 1984–1995, the period preceding ours, where they focus on skill upgrading and do not consider the role of techies. Dunne et al. (2004) find that computer use within U.S. manufacturing plants is associated with greater skill intensity, although not with greater labor productivity. Bartel et al. (2007) show that ICT

⁴Acemoglu and Autor (2011) provide an analytical framework that suggests how tasks are bundled across types of workers (differentiated by education level or skill), and how changes in demand for these tasks affect employment shares of these types.

adoption in the U.S. valve manufacturing industry caused reorganization within firms, raised the skill-requirements for machine operators and increased productivity (through faster setup times, greater customizability, and better quality control). Similar to our results on techies and employment growth, Barth et al. (2017) find a positive association between the science and engineer share of employment with revenue in U.S. manufacturing plants. In contrast, we study the entire private sector and make causal inferences.

The idea that engineers and other technically-trained workers are important for productivity growth has also found support in the economic history literature. Kelly et al. (2014) and Ben Zeev et al. (2017) highlight the importance of the British apprentice system during the British Industrial Revolution in supplying the basic skills needed for technology adoption (whether British technology or other). Maloney and Valencia-Caicedo (2017) construct a dataset of engineer intensity for the Americas and for U.S. counties around 1880, and show that this intensity helps predict income today. In addition, engineers are at the center of modern (endogenous) growth theory, e.g., Romer (1990).

A second force that could help explain job polarization and firm re-organization in general is offshoring, where domestic labor is replaced by imported inputs (see among many others Feenstra and Hanson (1996), Grossman and Rossi-Hansberg (2008), Rodriguez-Clare (2010), and Blinder and Krueger (2013)). Empirically, our results suggest a relatively modest role for offshoring in explaining polarization, as have a number of other studies including Feenstra and Hanson (1999), Michaels et al. (2014), Oesch (2013) and Biscourp and Kramarz (2007).

Most complementary to our paper are Böckerman et al. (2019) and Heyman (2016), who use samples of matched employee-employer data for Finland and Sweden, respectively, and focus on explaining within-firm organization. In contrast, our data cover every firm and employee in the French private sector, and we highlight the paramount importance of changes in firm composition and study what drives these changes—in addition to studying changes in within-firm organization.

A key objective of our paper is to identify causal effects of technology and trade on firms' occupational composition and size. Our identification strategy relies on initial conditions across firms to explain changes in occupational composition and size. This strategy is similar to that of Beaudry et al. (2010) and Autor and Dorn (2013), who exploit variation across space and use lagged initial conditions as instruments that help identify the propensity of local labor markets to respond to technological change. Michaels et al. (2014) estimate long differences specifications and exploit variation across industries (within countries), and instrument for differences in ICT intensity by using initial conditions in the United States.

Our results on the importance of firm composition is related to a recent literature that studies the relationship between the growth of so-called “superstar firms” and declines in the labor share. The role of firm composition is illustrated in Autor et al. (2020), who argue that when superstar firms, which are larger and more capital intensive to begin with, increase market and employment shares, then aggregate labor payments fall. We find a positive covariance between firm employment growth and techie intensity, suggesting that techie-intensive firms grow much faster than other firms. This is consistent with increasing competitiveness for these firms as they gain more from reductions in the cost and increases in availability of ICT.

3 Data

We use detailed panel data on firms in the French private sector economy in 1994–2007. Our empirical analysis proceeds in two steps. First, we describe the pattern of job polarization in France, then we provide an econometric analysis of the drivers of polarization. These steps require merging detailed firm-level information on employment by occupation to firm-level information on trade. The matching process is straightforward because firms in France are identified by the same identification number (called SIREN), which can be followed across years in both data sets.⁵ This section gives an overview of our data sources that will be used in subsequent analyses.

3.1 Data on workers

Our source for information on workers is the *DADS Poste*, which is based on mandatory annual reports filed by all firms with employees, so our data includes all private sector French workers except the self-employed.⁶ Our unit of analysis is annual hours paid in a firm, by occupation. For each worker, the DADS reports gross and net wages, hours paid, and the two-digit occupation of each worker.

Every job is categorized by a two-digit PCS occupation code.⁷ Excluding agricultural and public sector categories, the PCS has 22 two-digit occupational categories. For our analysis, we aggregate these 22 codes into seven broader categories, which are the bold face headings in Table 1. Table 1 lists the constituent 2-digit codes of each heading, the share of hours and the relative wage of occupations in 1994, and growth in each occupations shares of hours from 1994 to 2007.

– Insert Table 1 about here –

In order to aggregate occupations, we base our judgment on the routine intensity of an occupation on the INSEE documentation for occupations coded in the DADS dataset, which provides a full description of each occupation. For example, an office worker “*performs a set of activities related to administrative tasks, copyediting, typesetting or information transcription, control of administrative operations or linked to the reception of clients*”. One needs to be more cautious when associating a type of task with an occupation. We do our best to aggregate occupations within broad categories that are (i) meaningful in terms of tasks and jobs that they unite and (ii) are in similar parts of the wage distribution in 1994. Some occupations are in different parts of the wage distribution in 1994. The deviation is however small enough so that we still find employment polarization with 22 occupations as we will see later on.

Techies are central to our research. Techie occupations combine two of the 2-digit occupations: “*technical managers and engineers*” and “*technicians*”. As is clear from the detailed descriptions in Table 2, many workers in these categories are closely connected with the installation, management, maintenance, and support of information and

⁵The data is reported at the level of establishments, which are identified by their SIRET. The first nine digits of each SIRET is the firm-level SIREN, which makes it easy to aggregate across establishments for each firm.

⁶The *DADS Poste* is an INSEE database compiled from the mandatory firm-level DADS (*Déclaration Annuelle des Données Sociales*) reports.

⁷PCS stands for *Professions et Catégories Socioprofessionnelles*.

communications technology (ICT), and even if they do not work with ICT these are jobs that require STEM and technical training, skill, and experience.⁸

– Insert Table 2 about here –

As we will outline in the theory section 5 below, techies mediate the effects of new technology within firms: they are the ones who plan, purchase, and install new ICT equipment, and who train and support other workers in the use of ICT. Inspection of Table 2 supports this argument, although the table also makes it clear that not all techie workers necessarily work primarily with ICT. In a nutshell, if a firm invests in ICT, it needs techies, and firms with more techies are probably more technologically sophisticated firms.

Each two-digit PCS category is an aggregate of as many as 40 four digit subcategories. Although hours data is not available by four-digit category, the descriptions of the four digit categories in Table 2 are helpful in understanding the kinds of tasks performed within two-digit categories. The subcategories listed in Table 1 also suggest differences in the susceptibility of jobs to automation and/or offshoring. For example, service workers such as restaurant servers, hair stylists, and child care providers do the sort of “*non-routine manual*” tasks (*c.f.* Autor et al. (2003)) that require both proximity and human interaction. The same can be said for both skilled and unskilled manual laborers, whose jobs include gardening, cooking, repair, building trades, and cleaning. In contrast, mid-level managers and professionals often do routine cognitive tasks that can be done more cheaply by computers or overseas workers. Skilled and unskilled industrial workers doing routine manual work are unquestionably directly in competition with both robots and imported intermediate goods.

One potential problem with our hypothesis that firm-level techies are an indicator of firm-level technological sophistication is that firms can purchase ICT consulting services. By hiring a consultant, firms can obtain and service new ICT without increasing their permanent staff of techies. However, the vast majority of techies work outside the IT consulting sector (95% in 1994 and 91% in 2007), which means that almost all techie services are provided by in-house staff rather than hired consultants.⁹

Techies are far from equally dispersed over all sectors. They are more prevalent in the IT consulting sector—but not exclusively concentrated there. To see this, we compute the following ratio which divides techie employment shares across two-digit sectors by overall employment shares:

$$\frac{\left(\frac{\text{techies in industry } i}{\text{total techies}}\right)}{\left(\frac{\text{employment in industry } i}{\text{total employment}}\right)}$$

This ratio is about 4.4 in the IT consulting sector in 2007, implying that techies are 4.4 times more prevalent in this sector compared to the overall representation of the sector

⁸The detailed data on the 4-digit occupations codes described in Table 2 are available in the DADS dataset from 2009 only. As our sample ends in 2007, we cannot include these numbers.

⁹The IT consulting sector is a subset of sectors 62 and 63 of the NAF Rev. 2 classification and comprises of “*Consulting in computer systems and software*” and “*Third party maintenance of computer systems and applications*”. These numbers are somewhat smaller if we use the corresponding sector 72 in the NAF Rev. 1 classification. The shares of engineers or technicians in the IT consulting sector are similar over the sample period.

in aggregate employment. The ratio is stable from 1994 to 2007. Outside of the IT consulting sector the same ratio in 2007 is on average 1.17, with a standard deviation of 1.19.¹⁰ This implies that overall, techies are widely dispersed across industries, in a way that is broadly commensurate with the size of the sector in which they are employed.¹¹

3.2 Data on firm-level trade.

The DADS Poste dataset has no information about the firm beyond the firm identifier and industry and, implicitly, firm-level aggregates related to employment such as total hours by occupation, average wages, etc. Our source for firm-level trade data is the French Customs. For each trade flow observation, we use the identity of the French importing or exporting firm and construct an indicator of the firm status into exporting and importing. We merge the customs data into the DADS Poste database, and keep all non-matched firms. Firms that are present in the DADS and that are not matched with customs data are assumed to have zero exports and imports. The customs data indeed covers tradable goods, and both exporting and importing firms are mostly in manufacturing and retailers. The share of internationally engaged firms, who either export, import or do both, in the sample is about 8.6% in 1994. These firms account for 49.9% of total hours paid.

From 1994 to 2007, roughly 2.9 million private sector firms appear in the DADS Poste data. Our descriptive analysis includes all 2.9 million firms, but in our econometric analysis we focus on the subset of firms that were in operation continuously from 1994 to 2007. There are 297,402 of these “*permanent*” firms, of whom 85% are in non-manufacturing. The permanent firms represent about a quarter of firms and half of hours paid in our sample in each year.

4 Facts

France has experienced rapid structural change across sectors over the last few decades, which can be seen in the sharp increase in the share of nonmanufacturing (mainly services) employment from 72 to 78 percent between 1994 and 2007. The share of hours employed in the non-manufacturing sector (mostly services) within the private sector grew from an already high share of 72 percent in 1994 to 78 percent in 2007. This transition is not the subject of our paper, but a potential explanation is that as productivity growth in the manufacturing sector is much faster than in non-manufacturing, while households’ demand for manufactured goods relative to services remains relatively constant, relatively fewer workers are needed in manufacturing (Baumol (1967)).¹² The evolution of employment within the non-manufacturing sector accounts for most of the economy and the employment composition at aggregate level.

¹⁰The average can be greater than 1 because it is an unweighted average across industries.

¹¹This does not address the issue that the techie “*treatment effect*” may be different in the IT consulting sector versus other sectors. In order to address this, we estimated all our regressions without the sectors 62-63. The estimates from these regressions are virtually the same as the main regressions. This is because we always control for industry fixed effects and weight our regressions by employment. Sectors 62-63 account for only 0.43% of firms in our regression sample, or 1.15% of employment.

¹²This explanation is consistent with findings in Buera and Kaboski (2009) and in Reshef (2013).

In this section, we show how the French job market polarized between 1994 and 2007. We show that polarization was accompanied by the *March of the Techies*: the growing importance of technologically-oriented occupations that require STEM skills. We also show that polarization occurred both between and within firms, and that the between component was much larger than the within component.

4.1 Occupational polarization and the *March of the Techies*

The French occupational structure polarized between 1994 and 2007, with high-wage and low-wage occupation shares growing at the expense of middle-wage occupations. Polarization is shown in Figure 1, which illustrates the changes in each broad occupation's share of hours worked that are reported in Table 1. Figure ?? in Appendix visualizes the same pattern of polarization using 22 occupations confirming the results based on broad categories.

As in Table 1, occupations are ranked by average wage, and the width of the bars is proportional to each occupation's employment share in 1994.

– Insert Figure 1 about here –

Figure 1 shows that the share of hours by upper managers and especially engineers and technicians grew substantially, while the share worked by blue and white collar workers fell. In sharp contrast, low-wage personal service jobs grew. This pattern is driven by the nonmanufacturing sector (about 80 percent of employment), as shown in Figure 2. The pattern is somewhat different in manufacturing. As Figure 3 shows, employment in manufacturing was characterized by skill upgrading, with lower-skilled and office workers replaced by higher-skilled workers, along with a huge increase in the share of technicians and engineers.

– Insert Figures 2 and 3 about here –

Particularly striking in Figures 1 through 3 is the strong growth in the techie occupations. The share of techie hours increased from 9.04% to 12.21% of total hours from 1994 to 2007.¹³ We call this phenomenon *The March of the Techies* and illustrate it in Figure 4. Techie growth was steady over this period, with most of the growth since the late 1990s accounted for by the higher-paid and more-educated Technical Managers and Engineers category within techies.

4.2 Decompositions of aggregate trends

To better understand the aggregate patterns, we decompose changes in hours shares into changes within economic units and changes in composition. This allows us to evaluate, for example, whether aggregate changes occur due to within-firm adjustments, changes

¹³The share of techie hours in non-manufacturing increase from 58% to 64% between 1994 and 2007. The share of techie hours in manufacturing increase from roughly 12.7% to 20% while in non-manufacturing it increased from roughly 7% to 10%. The rate of increase of techie employment in manufacturing and non-manufacturing is very similar, but the level is higher in manufacturing. This implies that the relative expansion of non-manufacturing does not account for the increase in techie employment growth (in fact, it reduces it).

in firm size composition, or both. Thus, the decompositions help identify mechanisms that underlie the aggregate trends.

The change in occupation o 's share of aggregate hours over some period can be decomposed as follows

$$\Delta S_o = \underbrace{\sum_i \Delta \lambda_i \bar{s}_{io}}_{\text{between}} + \underbrace{\sum_i \bar{\lambda}_i \Delta s_{io}}_{\text{within}} , \quad (1)$$

where i indexes different partitions of the aggregate economy into distinct units: industries, firms, geographic units (departments or urban/rural distinctions) and gender. Here $\Delta \lambda_i$ is the change in the hours share of unit i in the aggregate, \bar{s}_{io} is the average share of occupation o within unit i , $\bar{\lambda}_i$ is the average hours share of unit i , and Δs_{io} is the change in the share of occupation o within unit i . The between component tells how much of the aggregate change is due to changes in composition, holding the within-unit occupational shares fixed. The within component tells how much of the aggregate change is due to changes in within-unit occupational shares, holding fixed the composition. At the level of firm for instance, the shift-share analyses examine whether firms that always had a high fraction of workers in an occupation grew more (between firm component), or whether firms started to hire more workers of that occupation over time (within firm component).

In the case of firms, we take into account net entry and exit in (1) as follows

$$\begin{aligned} \Delta S_o &= \underbrace{\sum_{i \in P} \Delta \lambda_i \bar{s}_{io}}_{\text{between, perm}} + \underbrace{\sum_{i \in P} \bar{\lambda}_i \Delta s_{io}}_{\text{within, perm}} + \underbrace{\sum_{i \in E} \lambda_{i2} s_{i2}}_{\text{entry}} - \underbrace{\sum_{i \in X} \lambda_{i1} s_{i1}}_{\text{exit}} \\ &= \underbrace{\sum_{i \in P} \Delta \lambda_i \bar{s}_{io}}_{\text{between, perm}} + \underbrace{\sum_{i \in P} \bar{\lambda}_i \Delta s_{io}}_{\text{within, perm}} + \text{net entry} , \end{aligned} \quad (2)$$

where P denotes the set of permanent firms, i.e., those that are present in both the first and last period of the decomposition, E denotes the set of firms that enter, i.e. are not present in the first period but are present in the last, and X denotes the set of firms that exit, i.e. are present in the first period but are not present in the last.¹⁴

We also use (1) in order to evaluate the contribution of gender and urban versus rural areas to overall change as follows

$$\text{Contribution}_i^o = \Delta \lambda_i \bar{s}_{io} + \bar{\lambda}_i \Delta s_{io} , \quad (3)$$

where $\sum_i \text{Contribution}_i^o = \Delta S_o$. For example, in the case that $\Delta S_o > 0$ urban areas (i.e., all i 's that are urban) can contribute to aggregate changes either by virtue of having a relatively high occupational share (\bar{s}_{io}) and growing disproportionately ($\Delta \lambda_i$), or by large increases of the share of occupation o within urban areas (Δs_{io}) given their relative size ($\bar{\lambda}_i$). While equation (3) is feasible with any number of units, considering two specific units is particularly informative.

¹⁴As it is defined, the set P may include firms that report positive hours in 1994 and in 2007, but no hours in one or more of the intervening years. We exclude a tiny number of firms of this type from our data. These firms account for a negligible share of aggregate hours.

4.3 Industries versus firms

We start by juxtaposing decompositions for industries with decompositions for firms. Table 3 reports the decompositions of ΔS_o for all seven occupational groups according to equation (1) in 1994–2007 for industries, firms, and then further breaking down firms according to equation (2). In the last row we report the weighted average contribution of each component to changes: $\sum_o \lambda_{o,1994} \times \text{between}_o / \Delta S_o$ and $\sum_o \lambda_{o,1994} \times \text{within}_o / \Delta S_o$, where $\lambda_{o,1994}$ are aggregate hours shares of occupation o in 1994.

The first message from Table 3 is that the lion’s share of the change in aggregate occupational shares occurs within industries, not due to industry composition. In all occupations but “Other white collar”, the within component is larger than the between component. The within component is more than all of the aggregate change for Managers, Techies and Service workers, implying that composition works in the opposite direction of the aggregate change. The weighted average contribution of the within component is 91%.

The second message from Table 3 is that in contrast to industries, most of the variation in aggregate occupational shares is due to firm composition, rather than changes within firms. For six out of seven occupations the between-firm component is much larger than the within component. For example, of the overall 3.17 percent point increase in the share of techies, four-fifths came from between-firm adjustments while only one-fifth came from increases in the within-firm share of techies. This finding suggests that firms with above-average techie shares grew much faster than average. The pattern is particularly stark for office workers: virtually all of the 3.15 percent point drop in office workers’ share of employment came from slower growth of firms with above-average shares of office workers. The only exception to this pattern is in the lowest paid and fast-growing services occupations, where the between and within split is roughly equal: firms with lots of service workers grew faster than average, while at the same time firms were on average increasing their employment of these low-wage occupations. The weighted average contribution of the between firm component is 125%, implying that, on average, within-firm adjustments work in the opposite direction of aggregate changes.

The third part of Table 3 reports within and between components for permanent firms (those in set P) in addition to their overall contribution (within, perm + between, perm) as in (2). Thus, the third message from Table 3 is that the permanent firm sample exhibits similar patterns for decompositions as the overall sample, with few significant differences. Overall, the pattern of polarization within permanent firms is similar to the aggregate, with one difference, which is the share of Unskilled workers increases, rather than decreases. However, this class of workers is second to last in terms of wage ranks, so polarization is still apparent. The weighted average contribution of permanent firms is 0.42 of the aggregate changes. Since the set of P firms account for almost half of employment in the sample, 0.5 is the relevant reference point for these contributions. The residual 0.58 is accounted for by net entry. As in the aggregate sample, the between firm component accounts for more than the entire change, while within firm changes work in the opposite direction.

We also correlate the overall between and within components with the corresponding contributions of the permanent sample firms. The correlation coefficient is .99 (both for the seven aggregate occupation classes, and for the more detailed 22 PCS codes). The

regression of the aggregate change on the change for permanent sample firms gives a coefficient of 1.087 with a t-statistic of 21.56. Thus, while there are specific differences in the evolutions in shares comparing overall and permanent firms, they are not very important on average.

The dominance of the between-firm channel in accounting for occupational change means that an explanation for polarization must explain why some firms grew faster than others. We return to this in Section 7 below after describing our estimation procedure.

4.4 Geography

We now turn to decompositions of aggregate occupational hours share by geography. We start by reporting in Table 4 decompositions of ΔS_o according to (1) in 1994–2007 for 94 French departments.¹⁵ Here the message is very clear: aggregate changes in occupational shares occur almost entirely within departments. Except for Managers, where composition accounts for -43% of the change in the aggregate hours share (i.e., works in the opposite direction), the contribution of composition is small. The weighted average contribution of the within departments components is 106%.

We then aggregate departments into two groups—22 urban departments and 62 rural departments – according to the definition of the French National Institute of Statistics (INSEE). Urban departments are those where the majority of the population lives in an urban area of 200,000 inhabitants or more. The 22 urban departments account for 55.1% of aggregate hours in 1994, and this share is remarkably stable, increasing slightly to 55.7 in 2007.¹⁶ Therefore, it is not surprising that, similar to the decomposition by departments, here too we find that virtually all of the aggregate changes come from within urban and rural areas, not from changes in urban-rural composition of employment.

We then turn to evaluate the separate contributions of urban versus rural areas to changes in aggregate occupational hours shares according to equation (3). On average, urban areas contribute virtually all of the aggregate changes, 95%. However, this average masks much heterogeneity across occupations. For example, urban areas contribute more than the entire increase in the hours share of Managers, whereas rural areas see declines in employment shares of managers.¹⁷ Rural areas contribute roughly three quarters of the aggregate changes for Techies, Other white collar, and for Office workers. In contrast, rural areas contribute two-thirds of the decline in Unskilled workers. Overall, job polarization is more of an urban phenomenon, as found by Autor (2019) for the U.S., whereas the pattern of change for occupational shares in rural areas is less clear.

4.5 Gender

Finally, we turn to decompositions of aggregate occupational hours share by gender. Table 5 reports decompositions of ΔS_o according to (1) and the separate contributions of

¹⁵There are 95 departments in “Metropolitan France”, i.e., excluding overseas departments. We drop Corsica because of reporting discrepancies over time.

¹⁶This stability masks some increase in urban density within departments. Appendix Table ?? lists the urban departments, together with their respective hours shares in 1994.

¹⁷The increased concentration of managers in cities is reminiscent of the finding of Duranton and Puga (2005).

women versus men according to (3). As with the urban-rural split, here too we find that virtually all of the changes occur within groups, with relatively little variation. In only one occupation, Skilled Workers, is the between-gender component of the change large (about a quarter of the total change). This is because men are much more prevalent in this type of occupation (notably in industrial and construction jobs, drivers, and transport jobs) compared to women, and the employment composition shifts slightly towards women, from 34.7% in 1994 to 35.6% in 2007 of total hours worked in France. The weighted average contribution of the within gender component is 96%.

The contribution of each sex is, however, relatively even, where women on average account for 47% of aggregate changes and the remaining 53% is accounted for by men. However, the average masks significant variation by occupation. Since the aggregate hours share of women increases only slightly, the relative contributions by gender for each occupation are driven by differential changes within each sex. Most notably, women alone account for more than the entire increase in Managers, increasing their employment share in managerial jobs by 0.9 percent points, whereas men see their share decline by 0.53 percent points. Women also increase their employment share in mid-level managerial jobs, where the overall decline is completely accounted for by men. Men account for most of the increase in Techies, Office workers and Skilled workers. Overall, the pattern of polarization is clearly apparent for women, but less so for men due to the large decline in employment in top managerial occupations.

5 Theoretical Framework

Our hypothesis is that techies are a channel through which falling ICT prices cause polarization. In the Online Appendix, we develop a simple model of firm-level outcomes that illustrates this. The model shows how a drop in the price of ICT leads to greater employment growth in ICT-intensive firms and to polarization of employment within a firm. These results help motivate our between- and within-firm descriptive and econometric analyses below. Here we describe the economic logic of the model, and refer interested readers to the details in the Online Appendix.

Assume that firms use three types of workers (high, medium and low skilled) along with ICT capital to produce output. In addition, the firm hires techies to manage and maintain the stock of ICT capital. Crucially, suppose that skilled workers are complements to ICT, while ICT substitutes for medium skilled workers, and that techies are a necessary input that is required for productively using ICT.

How does technological progress affect firms' employment? We model technological progress as a drop in the cost of ICT capital, or equivalently in the cost of the productive services provided by ICT. This price drop is common across firms, but the effects differ depending on how intensive firms are in their use of ICT. Since ICT and techies are strongly complementary, ICT-intensive firms are identified by their techie intensity.

There are two implications of such technological progress. First, firms' overall costs will fall, and the drop in costs will be larger for firms who use more ICT and employ more techies. This competitiveness effect implies that techie-intensive firms will gain market share, boosting their employment relative to less techie-intensive firms. The second implication is that firms will increase their employment of skilled relative to both

medium and unskilled workers. The negative effect is stronger for medium skilled workers, for whom ICT is a substitute. Thus, cheaper ICT capital causes polarization within the firm: employment of both the highest and lowest skilled workers increases relative to medium-skilled workers. This effect will be stronger in firms who are more reliant on ICT and have greater techie intensity.

Putting these two effects together implies that technological progress will lead to polarization of overall labor demand through both a between-firm "competitiveness" channel and a within-firm "substitution" channel. We investigate the importance of both channels in our empirical analysis in the following sections.

Turning to the effects of globalization, Feenstra and Hanson (1996) were the first to show how purchases of imported intermediates (offshoring) can affect the skill composition of employment. More recently, Acemoglu and Autor (2011) show how offshoring can contribute to firm-level polarization. These and related analyses make the point that offshoring has competing effects on total firm-level employment: a direct substitution effect of imported intermediates for workers within the firm, and a cost-reducing effect that can raise demand for workers whose jobs are not offshored.¹⁸ We estimate these polarization and net employment effects in our econometric analysis below.

A large empirical literature, reviewed by Bernard et al. (2007), finds that exporting is associated with higher skill intensity in cross-sections of firms. We know of no theoretical or empirical analysis of the link between exporting and firm-level employment, perhaps because the first-order effect is too obvious: firms that also sell abroad will tend to have greater labor demand than those who sell only at home. In our econometric analysis below we look for such firm-level employment effects both within and between French firms.

6 Econometric Framework

The theoretical framework implies that firms' responses to economy wide trends in reductions in barriers to trade and to reductions in ICT prices will depend on their characteristics. In this section, we develop an econometric framework for estimating how firms' global engagement and technology intensity cause changes in employment growth and within-firm changes in occupational composition. In particular, how do predetermined firm differences in the propensity to trade and adopt technology lead firms to change their size and employment mix over time? We answer these questions by regressing *changes* in firm-level employment and occupational shares from on *initial levels* of trade and techies. In some specifications we address possible endogeneity of the initial levels by using long lags as instrumental variables. We always control for 2-digit industry specific effects, which absorb industry-wide trends such as import competition, as well as industry-wide average levels of control variables. Thus, our identification comes from cross-firm variation within industries in initial conditions. We report separate regressions for manufacturing and for non-manufacturing, since we expect firm level responses to be different across these two sectors of the economy.

¹⁸This is related to the wage effects studied in Grossman and Rossi-Hansberg (2008).

6.1 Firm employment growth equation

Optimal firm-level employment depends on both demand and cost conditions. We specify optimal log employment for firm f in year t , $\ln h_{ft}$, as

$$\ln h_{ft} = \beta_f + D(X_f) \cdot t + W_{ft}\gamma + \varepsilon_{ft} . \quad (4)$$

Here β_f a firm-specific intercept. The firm specific trend $D(X_f) \cdot t$ is determined by firm-specific values X_f . This trend absorbs how different firms respond to common, aggregate trends, for example, the drop in the price of ICT. The corresponding element of X_f is our proxy for θ , the firm-specific techie share. In this case $D(X_f)$ captures how firms with different θ (techie intensity) respond differently to the drop in ICT prices. The effect of time-varying, firm-specific characteristics is given by $W_{ft}\gamma$. Here W_{ft} includes firm characteristics which we cannot measure, such as capital, intermediate inputs, and demand shocks. First-differencing (4) gives

$$\Delta \ln h_f = D(X_f) + u_f , \quad (5)$$

where $u_f = \Delta W_{ft}\gamma + \Delta \varepsilon_f$ is a composite error term that includes changes in the firm characteristics and changes in the error term ε_f .

We model the firm-specific time trend $D(X_f)$ as a function of the *initial level* of the techie share and trade status at the beginning of the interval over which first differences are taken. Firms that do not trade at all, and/or that have no techies at all, are likely to be distinctly different from firms that do trade and/or have techies, so to accommodate this we include in X_f indicators for positive values of techies and trade. We also entertain gradual and non-linear effects of techies, by constructing indicators for three terciles of techie intensity over the positive support (where no techies is the omitted category). In some specifications we also separate techies into engineers (PCS code 38) and technicians (PCS code 47). To control for the well-established fact that firm growth rates decline with size, we also include the initial level of employment h_{f0} as a regressor. Finally, we include industry i fixed effects β_i to absorb industry trends.

Let $techies_{ft}$ be the share of techies in firm f 's total hours worked in period t and $techpos_{ft}$ be an indicator equal to one if $techies_{ft} > 0$. Let $exppos_{ft}$ and $imppos_{ft}$ be indicators equal to one if firm f exports or imports in period t , respectively. The basic equation to be estimated is the empirical implementation of (5)

$$\Delta \ln h_f = \beta_i + \beta_1 techpos_{f0} + \beta_2 exppos_{f0} + \beta_3 imppos_{f0} + \beta_4 h_{f0} + u_f , \quad (6)$$

where $\beta_i + X_{f0}\beta$ summarizes the $D(X_f)$ function, and $X_f = \beta_1 techpos_{f0} + \beta_2 exppos_{f0} + \beta_3 imppos_{f0}$. This specification captures the theoretical prediction that in the face of falling ICT prices, employment will grow more in firms that have higher initial levels of techies, as shown in section 5 above. Similarly, a firm that exports final goods or purchases imported inputs will be more affected by the increased integration of Eastern Europe, China, and India into the world economy than will a firm that does not trade. Thus, equation (6) allows us to estimate the heterogeneous effect of aggregate trends on firm outcomes, where the heterogeneity is captured by firm characteristics in the initial period. With industry fixed effects β_i , the parameters of interest are identified by variation across

firms within industries in the levels of techies, trade, and employment.¹⁹

We estimate equation (6) by weighted least squares (WLS) on a single 14-year difference from 1994 to 2007, so that for example, $techpos_{f0} = techpos_{f,1994}$. Because (6) is a long differences specification, it is not informative about the speed of adjustment of employment or possible lags in the response to drops in economywide ICT prices or the growth of globalization. We also estimate versions of (6) on changes in 2002–2007 by weighted two stage least squares (W2SLS), using 1994 values to instrument for initial levels. Accordingly, when estimating by W2SLS $techpos_{f0} = techpos_{f,2002}$ in (6), and the instrument is $techpos_{f,1994}$, and similarly for $exppos_{f0}$ and $imppos_{f0}$. The choice to base the changes for W2SLS in 2002–2007 is done in order to allow a sufficiently long lag for the instruments to be considered plausibly exogenous and excludable from (6). We weight all regressions using 1994 firm hours worked. This ensures that the estimated regressions can be interpreted as conditional expectations across the distribution of hours worked, which is what is relevant for our research question. In particular, this allows for stronger influence of larger firms on the estimator.

6.2 Within-firm organization equation

In order to analyze the changes in the occupational structure within French firms, we replace the dependent variable in (6) by changes in the share of hours employed in each occupation, where the denominator of the share excludes techies, in order to avoid a completely mechanical relationship with initial levels of techies. The approach here is very similar to the approach in Section 6.1. The firm-occupation outcome measure of interest is Δs_{fo} , the change in the *ex-techie* share of hours of the six large non-techie occupations listed in Table 1. For each occupation o our estimating equation is thus

$$\Delta s_{fo} = \beta_i^o + \beta_1^o techpos_{f0} + \beta_2^o exppos_{f0} + \beta_3^o imppos_{f0} + u_f^o. \quad (7)$$

As with (6), we estimate (7) by WLS in 1994–2007 and by W2SLS in 2002–2007 using 1994 levels as instruments for levels in 2002.

7 Econometric Results

7.1 Firm-level employment growth regressions

The theoretical framework of Section 5 provides predictions on the effect of firms' initial level of technology intensity and global engagement on their employment growth rate. We investigate these predictions by estimating the empirical model (6). The estimator uses variation across firms within the same industry. Any industry-level characteristic, including demand, import competition, sector-specific trends in ICT prices, etc., is absorbed by the inclusion of industry fixed effects.

¹⁹We reported above that the share of techies working in the IT consulting sector out of all techies rises from 4.7% in 1994 to 9% in 2007. This implies that we may mis-measure the techie labor services indicator for a non-negligible number of firms that do not employ techies, but outsource these services. This mis-measurement biases the estimator of the effects of techies towards zero, and makes it harder to infer large impacts (either positive or negative), and our estimates can be thought of as lower bounds.

We report the results of the WLS specifications in Table 6. We regress annualized growth rates in 1994-2007 on initial levels in 1994. In those regressions, we use either indicator variables for techie employment (overall techies, engineers and technicians), or indicators for terciles of techie employment intensity on the positive part of the support (the omitted category remains no techie employment). We also estimate specifications that split techies into engineers and technicians.

– Insert Table 6 about here –

Columns (1) to (3) of Table 6 report the effect of techies, exports and imports on the annualized growth rate of firm-level hours in the manufacturing sector. Columns (4) to (6) report the results for the non-manufacturing sector.²⁰ Columns (1) and (4) show that firms with a positive techie share in 1994 saw significantly faster employment growth than firms without techies. The presence of techies involves a sizeable acceleration of firm-level employment growth, on the order of 3.5 percentage points in manufacturing industries and 1.1 percentage points in non-manufacturing sectors per year. This result is consistent with the theoretical prediction in Section 5, where falling ICT prices raise the competitiveness of firms that employ more techies. For manufacturing firms, exporters grew more rapidly while importers grew more slowly, though these effects are not statistically significant.

Columns (2) and (5) display results where we allow for differential effects of techie intensity within firms. The specifications include three indicators for terciles of techie intensity over the positive support. The coefficients on the second and third tercile indicators are statistically significant in manufacturing, and the third tercile indicator is somewhat larger. In non-manufacturing industries all tercile coefficients are statistically significant, and the third is clearly larger than the first and second. In both manufacturing and non-manufacturing samples, the joint tests of equality between terciles reject the hypothesis that they are equal. These results suggest that firms that are the most intensive in techies have a faster employment growth than firms that have smaller techie shares.

The results so far suggest that firm-level employment growth is faster in techie-intensive firms. The effects of techies may be different across different group of techie workers, highly-educated technical managers and engineers versus lower-ranked technicians. In columns (3) and (6), we exploit the distinction between engineers and technicians in our data to contrast their effects on firm-level employment growth. We find a strong and positive effect of engineers on non-manufacturing firms' employment growth. Non-manufacturing firms with a positive presence of engineers increased their employment growth by about 1.3 percentage point compared to firms without techies, but there is virtually no effect of employment of technicians.²¹ Given the size of the non-manufacturing sector, the effect is economically large. In manufacturing, the data do not allow us to separately estimate the effect of engineers from that of technicians because their employment

²⁰To address the issue that the techie treatment effect may be different in digital sectors versus others, we estimated all our regressions without these sectors (NAF sectors 62 and 63, Rev. 2 classification). The estimates from these regressions are virtually the same as the main regressions. This is because we always control for industry fixed effects and weight our regressions by employment. Moreover, sectors 62-63 account for only 0.43% of firms in our regression sample, or 1.15% of employment.

²¹This somewhat surprising result is not because there are no technicians in non-manufacturing: their share of hours is roughly the same as those of engineers in 1994.

is too strongly correlated across firms, causing the standard errors to be very big. Both engineers and technician have a positive effect on employment growth but the effects are statistically insignificant at conventional levels.

In Table 7 below, we report results using just-identified W2SLS regressions of changes in 2007-2007 on levels in 2002, using 1994 lags as instruments. We report the Kleibergen-Paap F tests which yield values larger than 20. This suggest that our estimations are unlikely to suffer from weak instruments bias. In these regressions we do not estimate models with nonlinear effects, and focus only on indicators for techies and trade. We present the W2SLS regressions alongside WLS estimate of the same specification to compare the results.

– Insert Table 7 about here –

Columns (1-4) report the results for the manufacturing sample while Columns (5-8) concern the non-manufacturing sample. The OLS estimates suggest a strong and positive impact of techies on employment growth over the period 2002-2007. We find that the presence of techies involves an increase of firm-level employment growth of 1.8 percentage points in manufacturing industries and 1.2 percentage points in non-manufacturing sectors per year. The impact of techies on employment growth is smaller for manufacturing during the shorter period from 2002 to 2007 compared to the longer period starting in 1994, but it is very similar for non-manufacturing sectors. Columns (3) and (7) reproduce the specifications in Columns (3) and (6) of Table 6 on the shorter sample. We now find positive and statistically significant effects of both engineers and technicians on employment growth of manufacturing firms. The results for non-manufacturing firms are similar to the longer sample, with a positive and statistically significant effects of engineers on non-manufacturing firms employment growth. Overall, our findings are broadly consistent across periods.

The techie estimates from just-identified W2SLS models are presented in Columns (2), (4), (6) and (8). The results confirm that techies have a strong effect on manufacturing employment growth of 4.6 percent per year. The W2SLS regressions indicate that engineers are associated with manufacturing firms’ employment growth much more than technicians. Concerning the non-manufacturing sector, we find that the main effect of techies drops slightly, from 1.1 to 0.9 percent per year, but we lose precision because standard errors increase so much that these effects on employment growth are not statistically significant.

7.2 Changes in the occupational structure of firms

We next examine the changes in the occupational structure within French firms, and in particular the job polarization that was documented in Section 4. Our hypothesis is that both the techie intensity and the global engagement of firms are important factors, and that firms responses to economy wide changes will depend on these characteristics. We measure changes in a firm’s occupational structure by changes in the share of hours in each of six PCS occupations, excluding the share of techies. As with the employment growth regression, we compute “ex-techie” shares of employment and ask whether these are associated with the exporter and importer status of the firm as well as the presence of techies in the firm. Our estimation approach here is very similar to the approach

described above. Each regression at occupational level includes a set of 2-digit sector fixed effects. The regression coefficients are identified by cross-firm variation in changes within 2-digit sectors.

We discuss our results in detail in Appendix D, and summarize them briefly here. We find that the presence of techies in 1994 is followed by within-firm skill upgrading, with an increase in the share of top and middle managers and a decrease in the share of the three lowest paid occupations. Turning to the effects of international trade, both importing and exporting in 1994 are associated with a subsequent increase in the top manager share.

8 Conclusions

In this paper we use administrative employee-firm-level data to show that the labor market in France polarized between 1994 and 2007: employment shares of high and low wage occupations have grown, while middle wage occupations have shrunk. This has profound implications for inequality. We show that job polarization occurs mainly within industries, through changes in firm sizes rather than through within-firm adjustment. The importance of firm size reallocations for polarization implies that simple theories of substitution across workers within firms miss an important margin of adjustment. This margin is changes in competitiveness, where more tech-savvy firms gain market share at the expense of less tech-savvy firms.

We develop a stylized model which illustrates that variation in initial levels of ICT intensity predict differential responses of firms to the drop in the price of ICT capital. Motivated by the fact that technology adoption is mediated by technically qualified managers and technicians, we develop a novel measure of the propensity to adopt new technology: the firm-level employment share of techies. Firm-level techie intensity helps predict firm-level outcomes that are consistent with the effects of falling ICT prices and the growing importance of STEM skills generally, in particular through the competitiveness margin. As techies become more numerous in the economy, we predict that these effects will be commensurately more widespread.

We develop an empirical framework that allows us to study the effect of techie intensity on both between-firm and within-firm employment changes. We show that firms with more techies in 1994 saw substantially faster employment growth in 1994–2007. The effect is pervasive across sectors and is mostly driven by engineers rather than technicians. Our empirical analysis is consistent with technological change improving the competitiveness of these firms relative to other firms with fewer techies. Concerning the impact of techies on within-firm occupational changes, we find that techies are associated with skill upgrading, and particularly an increase in the employment share of managers.

References

- Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 4b. North Holland.
- Autor, D. H. (2019). Aea papers and proceedings. *AEA Papers and Proceedings* 109, 1–32.
- Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *The American Economic Review* 103(5), 1553–1597.
- Autor, D. H., D. Dorn, and G. H. Hanson (2013). The china syndrome: Local labor market effects of import competition in the united states. *The American Economic Review* 103(6), 2121–2168.
- Autor, D. H., D. Dorn, and G. H. Hanson (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal* 125(584), 621–646.
- Autor, D. H., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, forthcoming.
- Autor, D. H., L. F. Katz, and M. S. Kearney (2006). The polarization of the u.s. labor market. *The American Economic Review*, *AEA Papers and Proceedings* 96(2), 189–194.
- Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in u.s. wage inequality: Revising the revisionists. *The Review of Economics and Statistics* 90(2), 300–323.
- Autor, D. H., F. Levy, and R. J. Murnane (2002). Upstairs, downstairs: Computers and skills on two floors of a large bank. *Industrial and Labor Relations Review* 55(3), 432–447.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118(4), 1279–1333.
- Bartel, A., C. Ichniowski, and K. Shaw (2007). How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *The quarterly journal of Economics* 122(4), 1721–1758.
- Barth, E., J. C. Davis, R. B. Freeman, and A. J. Wang (2017). The effects of scientists and engineers on productivity and earnings at the establishment where they work. NBER Working Paper No. 23484.
- Baumol, W. J. (1967). Macroeconomics of unbalanced growth: The anatomy of urban crisis. *The American Economic Review* 57(3), 415–426.
- Beaudry, P., M. Doms, and E. Lewis (2010). Should the personal computer be considered a technological revolution? evidence from u.s. metropolitan areas. *Journal of Political Economy* 118(5), 988–1036.
- Ben Zeev, N., J. Mokyr, and K. Van Der Beek (2017). Flexible supply of apprenticeship in the british industrial revolution. *The Journal of Economic History* 77(1), 208–250.
- Bernard, A., B. Jensen, S. Redding, and P. Schott (2007). Firms in international trade. *Journal of Economic Perspectives* 21(3), 105–130.

- Biscourp, P. and F. Kramarz (2007). Employment, skill structure and international trade: Firm-level evidence for france. *Journal of International Economics* 72(1), 22–51.
- Blinder, A. S. (2006). Offshoring: The next industrial revolution? *Foreign Affairs* 85(2), 113–28.
- Blinder, A. S. and A. B. Krueger (2013). Alternative measures of offshorability: A survey approach. *Journal of Labor Economics* 31(2), S97–S128.
- Böckerman, P., S. Laaksonen, and J. Vainiomäki (2019). Does ict usage erode routine occupations at the firm level? *LABOUR* 33(1), 26–47.
- Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics* 117(1), 339–376.
- Brynjolfsson, E. and L. Hitt (1996). Paradox lost? firm-level evidence on the returns to information systems spending. *Management science* 42(4), 541–558.
- Brynjolfsson, E. and L. M. Hitt (2003). Computing productivity: Firm-level evidence. *Review of economics and statistics* 85(4), 793–808.
- Buera, F. J. and J. P. Kaboski (2009). Can traditional theories of structural change fit the data? *Journal of the European Economic Association* 7(2-3), 469–477.
- Caliendo, L., F. Monte, and E. Rossi-Hansberg (2015). The anatomy of french production hierarchies. *Journal of Political Economy* 123(4).
- Cerina, F., A. Moro, and M. Rendall (2017). The role of gender in employment polarization.
- Dunne, T., L. Foster, J. Haltiwanger, and K. Troske (2004). Wage and productivity dispersion in United States manufacturing: The role of computer investment. *Journal of Labor Economics* 22(2), 397–429.
- Duranton, G. and D. Puga (2005). From sectoral to functional urban specialisation. *Journal of urban Economics* 57(2), 343–370.
- Dustmann, C., J. Ludsteck, and U. Schönberg (2007). Revisiting the german wage structure. *The Quarterly Journal of Economics* 124(2), 843–881.
- Feenstra, R. C. and G. H. Hanson (1996). Foreign investment, outsourcing, and relative wages. In G. M. Grossman, R. C. Feenstra, and D. R. Irwin (Eds.), *The Political Economy of Trade Policy*, pp. 89–127. Oxford: MIT Press.
- Feenstra, R. C. and G. H. Hanson (1999). The impact of outsourcing and high-technology capital on wages: Estimates for the united states, 1979-1990. *The Quarterly Journal of Economics* 114(3), 907–940.
- Firpo, S., N. M. Fortin, and T. Lemieux (2011). Occupational tasks and changes in the wage structure. Working Paper, University of British Columbia.
- Goos, M. and A. Manning (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics* 89(1), 118–133.
- Goos, M., A. Manning, and A. Salomons (2009). Job polarization in europe. *The American Economic Review, AEA Papers and Proceedings* 99(2), 58–63.

- Goos, M., A. Manning, and A. Salomons (2014). Explaining job polarization: Routine-biased technological change and offshoring. *The American Economic Review* 104(8), 2509–2526.
- Graetz, G. and G. Michaels (2018). Robots at work. *Review of Economics and Statistics* 100(5), 753–768.
- Gregory, T., A. Salomons, and U. Zierahn (2016, July). Racing with or against the machine? evidence from europe. Technical report, unpublished.
- Grossman, G. M. and E. Rossi-Hansberg (2008). Trading tasks: A simple theory of offshoring. *The American Economic Review* 98(5), 1978–97.
- Helper, S. and J. Kuan (2018). What goes on under the hood? how engineers innovate in the automotive supply chain. In *US Engineering in the Global Economy*. University of Chicago Press.
- Heyman, F. (2016). Job polarization, job tasks and the role of firms. *Economics letters* 145, 246–251.
- Kelly, M., J. Mokyr, and C. Ó Gráda (2014). Precocious albion: a new interpretation of the british industrial revolution. *Annu. Rev. Econ.* 6(1), 363–389.
- Lichtenberg, F. R. (1995). The output contributions of computer equipment and personnel: A firm-level analysis. *Economics of innovation and new technology* 3(3-4), 201–218.
- Malgouyres, C. (2017). The impact of chinese import competition on the local structure of employment and wages: Evidence from france. *Journal of Regional Science* 57(3), 411–441.
- Maloney, W. F. and F. Valencia-Caicedo (2017). Engineering Growth: Innovative Capacity and Development in the Americas. CESifo Working Paper No. 6339.
- Maurin, E. and D. Thesmar (2004). Changes in the functional structure of firms and the demand for skill. *Journal of labor economics* 22(3), 639–664.
- Michaels, G., A. Natraj, and J. Van Reenen (2014). Has ict polarized skill demand? evidence from eleven countries over twenty-five years. *Review of Economics and Statistics* 96(1), 60–77.
- Oesch, D. (2013). *Occupational Change in Europe*. United Kingdom: Oxford University Press.
- Reshef, A. (2013). Is technological change biased towards the unskilled in services? an empirical investigation. *Review of Economic Dynamics* 16, 312–331.
- Rodriguez-Clare, A. (2010). Offshoring in a ricardian world. *American Economic Journal: Macroeconomics* 2(2), 227–58.
- Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy* 98(5), S71–S102.
- Spitz-Oener, A. (2006). Technical change, job tasks and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics* 24(2), 235–270.
- Tambe, P. and L. M. Hitt (2012). The productivity of information technology investments: New evidence from it labor data. *Information Systems Research* 23(3-part-1), 599–617.
- Tambe, P. and L. M. Hitt (2014). Job hopping, information technology spillovers, and productivity growth. *Management science* 60(2), 338–355.

Table 1: Occupational definitions and employment shares

Occupation	Employment Shares		
	level 1994	change 1994-2007	relative wage 1994
Business owners, professionals, managers	9.24	0.37	1.96
Craftsmen	0.70	-0.65	1.32
Shopkeepers	0.61	-0.55	1.39
Heads of small business	0.75	0.00	2.70
Scientific professionals	0.48	-0.04	1.54
Creative professionals	0.65	-0.07	1.48
Managers	6.04	1.68	2.04
Techies	9.04	3.17	1.59
Engineers & technical managers	4.51	2.29	2.04
Technicians	4.53	0.87	1.13
Other white collar	4.85	-0.26	1.15
Teachers	0.32	0.04	1.05
Health and social workers	1.23	0.15	0.95
Foremen, supervisors	3.30	-0.46	1.19
Office workers	25.35	-3.15	1.00
Office workers, middle-level	12.85	-0.68	1.12
Office workers, lower-level	12.50	-2.47	0.84
Skilled workers	27.36	-0.62	0.82
Skilled industrial	11.35	-1.14	0.87
Skilled manual	8.96	-0.41	0.73
Drivers	4.58	0.61	0.74
Skilled transport & wholesale	2.48	0.32	0.78
Unskilled workers	13.81	-2.56	0.70
Unskilled industrial	9.71	-2.89	0.71
Unskilled manual	3.96	0.33	0.61
Services workers	10.49	3.04	0.66
Private security	0.69	0.36	0.70
Retail	6.30	1.47	0.65
Personal services	3.51	1.22	0.63

Note to table: “*re*”
wage in 1994.

Table 2: Techies, representative suboccupations

Technical managers and engineers
Technical managers for large companies
Engineers and R&D managers
Electrical, mechanical, materials and chemical engineers
Purchasing, planning, quality control, and production managers
Information technology R&D engineers and managers
Information technology support engineers and managers
Telecommunications engineers and specialists
Technicians
Designers of electrical, electronic, and mechanical equipment
R&D technicians, general and IT
Installation and maintenance of non-IT equipment
Installation and maintenance of IT equipment
Telecommunications and computer network technicians
Computer operation, installation and maintenance technicians

Figure 1: CHANGE IN ECONOMYWIDE HOURS SHARES (1994-2007)




Figure 2: CHANGE IN HOURS SHARES (NON-MANUFACTURING: 1994-2007)




Figure 3: CHANGE IN HOURS SHARES (MANUFACTURING: 1994-2007)




Figure 4: THE MARCH OF THE TECHIES




Table 3: DECOMPOSITION OF AGGREGATE TRENDS – INDUSTRIES VERSUS FIRMS (1994-2007)

	Level in 1994	Change	Decomposition				Contribution share of		
			Industries		Firms		Change	Within	Between
			Within	Between	Within	Between			
Business owners, professionals, managers	9.24	0.37	0.95	-0.58	-0.95	1.32	0.41	-0.50	0.91
Techies	9.04	3.17	3.28	-0.11	0.6	2.57	1.08	-0.18	1.26
Other white collar	4.85	-0.26	-0.07	-0.19	-0.05	-0.21	-0.07	0.02	-0.08
Office workers	25.35	-3.15	-1.92	-1.23	0	-3.15	-1.01	-0.06	-0.95
Skilled workers	27.36	-0.62	-0.38	-0.24	0.13	-0.75	-0.40	0.01	-0.41
Unskilled workers	13.81	-2.56	-1.82	-0.74	0.56	-3.12	0.46	0.40	0.06
Services workers	10.49	3.04	4.16	-1.12	1.55	1.49	1.31	0.67	0.64
Wgt. Avg. share in Change			0.91	0.09	-0.25	1.25	0.42	-0.13	0.55

Table 4: DECOMPOSITION OF AGGREGATE TRENDS – GEOGRAPHY (1994-2007)

	Level in 1994	Change	Departments		Level in 1994		Urban/Rural		Decomposition	
			Within	Between	Urban	Rural	Within	Between	Urban	Rural
Business owners, professionals, managers	9.24	0.37	0.53	-0.16	11.47	6.48	0.32	0.05	1.06	-0.69
Techies	9.04	3.17	3.15	0.02	10.98	6.64	3.14	0.03	2.25	0.92
Other white collar	4.85	-0.26	-0.28	0.02	4.89	4.80	-0.26	0.00	-0.19	-0.07
Office workers	25.35	-3.15	-3.08	-0.07	28.97	20.87	-3.19	0.04	-2.37	-0.78
Skilled workers	27.36	-0.62	-0.61	-0.01	22.81	32.98	-0.57	-0.05	-0.65	0.03
Unskilled workers	13.81	-2.56	-2.83	0.27	10.88	17.43	-2.53	-0.03	-0.82	-1.74
Services workers	10.49	3.04	2.99	0.05	9.99	11.11	3.04	0.00	1.61	1.43
Wgt. Avg. share in Change			1.06	-0.06			0.97	0.03	0.95	0.05

Table 5: DECOMPOSITION OF AGGREGATE TRENDS – GENDER (1994-2007)

	Level in 1994	Change	Decomposition					Contribution	
	Women	Men	Level in 1994	Within	Between	Women	Men		
Business owners, professionals, managers	9.24	0.37	6.70	10.59	0.39	-0.02	0.90	-0.53	
Techies	9.04	3.17	3.06	12.22	3.25	-0.08	0.72	2.45	
Other white collar	4.85	-0.26	4.09	5.25	-0.25	-0.01	0.07	-0.33	
Office workers	25.35	-3.15	45.06	14.86	-3.45	0.30	-1.18	-1.97	
Skilled workers	27.36	-0.62	7.06	38.16	-0.47	-0.15	-0.01	-0.61	
Unskilled workers	13.81	-2.56	14.98	13.19	-2.56	0.00	-1.46	-1.10	
Services workers	10.49	3.04	19.04	5.94	2.92	0.12	1.65	1.39	
Wgt. Avg. share in Change					0.96	0.04	0.47	0.53	

Table 6: BASELINE RESULTS – WLS REGRESSIONS (1994-2007)

Dep. Variable	Annualized employment growth rate, 1994-2007					
	(1)	(2)	(3)	(4)	(5)	(6)
	Manufacturing			Non-Manufacturing		
Techies	0.0349*			0.0109***		
	(0.0208)			(0.00280)		
Techies (1 st terc.)		0.0404			0.0117***	
		(0.0285)			(0.00306)	
Techies (2 nd terc.)		0.0266**			0.00584*	
		(0.0110)			(0.00354)	
Techies (3 rd terc.)		0.0322***			0.0202***	
		(0.00571)			(0.00625)	
Engineers			0.0269			0.0128***
			(0.0174)			(0.00346)
Technicians			0.0259			0.00121
			(0.0163)			(0.00316)
Exporters	0.0281	0.0285	0.0246	0.00109	0.00115	0.000532
	(0.0378)	(0.0377)	(0.0354)	(0.00478)	(0.00480)	(0.00478)
Importers	-0.0448	-0.0452	-0.0488	-0.00202	-0.00190	-0.00259
	(0.0514)	(0.0535)	(0.0542)	(0.00504)	(0.00503)	(0.00500)
Hours	-0.0156***	-0.0153***	-0.0176***	-0.00880***	-0.00889***	-0.00915***
	(0.00420)	(0.00264)	(0.00545)	(0.00131)	(0.00131)	(0.00144)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	43,476	43,476	43,476	253,926	253,926	253,926
R ²	0.005	0.005	0.005	0.001	0.001	0.001
<i>F</i> -test (<i>p</i> -value)						

Equality of terciles 0.052 0.036 $\beta^{Engineers} = \beta^{Technicians}$ 0.793 0.030

Dependent variable is firm-level annualized employment growth rate. Each specification includes 2-digit sector fixed effects. WLS estimates with robust standard errors. Standard errors are in parentheses. ***, **, * significantly different from 0 at 1%, 5%, and 10% levels, respectively.

Table 7: IV RESULTS – WLS AND W2SLS REGRESSIONS (2002-2007)

Dep. Variable	Annualized employment growth rate, 2002-2007							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Manufacturing				Non-Manufacturing			
	WLS	W2SLS	WLS	W2SLS	WLS	W2SLS	WLS	W2SLS
Techies	0.0185*** (0.00271)	0.0464*** (0.0134)			0.0115*** (0.00377)	0.00895 (0.0183)		
Engineers			0.0119*** (0.00325)	0.0601** (0.0292)			0.0122*** (0.00333)	-0.00849 (0.0464)
Technicians			0.00874*** (0.00261)	-0.0201 (0.0319)			-0.00106 (0.00421)	-0.0112 (0.0398)
Exporters	-0.00517* (0.00288)	-0.0181 (0.0152)	-0.00576** (0.00293)	-0.0225 (0.0154)	-0.0144*** (0.00549)	-0.0275 (0.0225)	-0.0146*** (0.00548)	-0.0325 (0.0228)
Importers	0.00543* (0.00295)	-0.00762 (0.0182)	0.00421 (0.00304)	-0.00926 (0.0186)	-0.00229 (0.00561)	-0.0143 (0.0214)	-0.00238 (0.00561)	-0.00707 (0.0218)
Hours	-0.00617*** (0.00106)	-0.00561*** (0.00121)	-0.00666*** (0.00111)	-0.00579*** (0.00138)	0.000354 (0.00130)	0.00342 (0.00245)	0.000276 (0.00137)	0.00651** (0.00319)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	43,476	43,476	43,476	43,476	253,926	253,926	253,926	253,926
R ²	0.019		0.019		0.009		0.009	
F-Stat (1 ^{stage})		23.94		36.92		47.64		39.21

Dependent variable is firm-level annualized employment growth rate. Each specification includes 2-digit sector fixed effects. WLS and W2SLS estimates with robust standard errors. The F-Stat (1^{stage}) is the Kleibergen-Paap Wald statistics. Standard errors are in parentheses. ***, **, * significantly different from 0 at 1%, 5%, and 10% levels, respectively.