No 2018-21 – December

Working Paper

CEPII

Techies, Trade, and Skill-Biased Productivity

James Harrigan, Ariell Reshef & Farid Toubal

Highlights

- We develop a methodology to estimate both Hicks-neutral and Skill-biased productivity across firms.
- We apply this to administrative data on firms in the entire French private sector in 2009-2013.
- We estimate the effects of ICT, R&D, exporting and importing on the evolution of productivity.
- ICT, R&D and importing increase firm level productivity significantly.
- This helps explain relative demand shifts towards skilled labor at the firm level and at the aggregate level.

Abstract

We study the impact of firm level choices of ICT, R&D, exporting and importing on the evolution of productivity and its bias towards skilled occupations. We use a novel measure of the propensity of a firm to engage in technology investment and adoption: its employment of workers with STEM (science, technology, engineering and math) skills and experience who we call "techies". We develop a methodology for estimating firm level productivity that allows us to measure both Hicks-neutral and skill-augmenting technology differences, and apply this to administrative data on French firms in the entire private sector from 2009 to 2013. We find that techies and importing of intermediate inputs raise skill-biased productivity, while imports also raise Hicks-neutral productivity. We also find that higher firm-level skill biased productivity raises low-skill employment even as it raises the ratio of skilled to unskilled workers. This is because of the cost-reducing effect of higher productivity. The techie and trade effects are large, and can account for much of the aggregate increase in skilled employment from 2009 to 2013.

Keywords

Productivity, Skill Bias, Skill Augmenting, Labor Demand, Outsourcing, Globalization, R&D, ICT, Techies, STEM Skills.

JEL

D2, D24, F1, F16, F6, F66, J2, J23, J24, O52.

Working Paper

CEPII

CEPII (Centre d'Etudes Prospectives et d'Informations Internationales) is a French institute dedicated to producing independent, policy-oriented economic research helpful to understand the international economic environment and challenges in the areas of trade policy, competitiveness, macroeconomics, international finance and growth.

CEPII Working Paper
Contributing to research in international economics

Editorial Director:
Sébastien Jean

Production:
Laure Boivin

No ISSN: 1293-2574

CEPII
20, avenue de Ségur
TSA 10726
75334 Paris Cedex 07
+33 1 53 68 55 00
www.cepii.fr
Press contact: presse@cepii.fr

RESEARCH AND EXPERTISE
ON THE WORLD ECONOMY

Techies, Trade, and Skill-Biased Productivity*

James Harrigan[†], Ariell Reshef[‡] and Farid Toubal[§]

1 Introduction

Economists have been studying the nexus between labor demand, globalization and technology adoption for decades.[1] Theory and anecdote suggest that some combination of technology and globalization has raised the relative demand for more skilled workers, but quantifying the relative importance of these forces has proved challenging. While there is a consensus that skill-biased technological change (SBTC) has raised the relative demand for more skilled workers, direct micro evidence on the drivers of SBTC is remarkably sparse. One reason for this absence of evidence is that technological change is devilishly difficult to measure. In this paper we overcome these challenges, and estimate the separate firm-level effects of research and development (R&D), information and communication technology (ICT) adoption, exporting and importing on productivity and its bias towards skilled workers. Our paper is the first to estimate the effects of firm-level technology adoption and trade on productivity and labor demand in a unified framework. Using data from France between 2009 and 2013, we find large effects of importing, ICT and R&D on the relative demand for skill through their effects on skill-augmenting productivity.

Quantifying the importance of technology adoption and globalization for relative labor demand is hard for two reasons. First, while it is relatively easy to measure importing and exporting at the

*We thank Daron Acemoglu, Mary Amiti, Zsofia Barany, Eric Bartelsman, Flora Bellone, Andrew Bernard, Tibor Besedes, Esther Ann Bøler, Davin Chor, Chiara Criscuolo, Alon Eizenberg, Georg Graetz, Giammario Impullitti, Beata Javorcik, Guy Michaels and Jean-Marc Robin for helpful comments. The paper benefited from comments from seminar participants at Aarhus University, American University, AMSE, Bar Ilan University, Ben Gurion University, CREST, Dresden University, Gronigen University, France Stratégie, Hebrew University, University of Le Mans, LSE, Nottingham University, New York University and the Paris Trade Seminar, as well as workshop participants at DEGIT, ENEF, IFN (Stokholm), SAIS, SETC, and TRISTAN (Bayreuth). This research was supported by the Bankard Fund for Political Economy at the University of Virginia and by the Agence National de la Recherche under grant numbers ANR-16-CE60-0001-01 and ANR-10-EQPX-17 (*investissements d'avenir*).

[†]University of Virginia, Paris School of Economics and NBER, james.harrigan@virginia.edu
[‡]Paris School of Economics, CNRS and CEPII, ariell.reshef@psemail.eu
[§]ENS Paris-Saclay, CREST and CEPII, ftoubal@ens-cachan.fr

[1]Helpman (2018) and Acemoglu and Autor (2011) are insightful recent surveys of this literature.

level of the firm, it has proven very hard to measure technology adoption, except in case studies and in particular industries.[2] Second, it is difficult to identify causal effects since firms jointly choose whether to import, export and adopt technology.

In this paper, we solve both problems. We apply and extend new techniques from the structural production function estimation literature in order to consistently estimate both Hicks-neutral and skill-augmenting productivity shifters. We use a flexible specification of the firm's productivity process which permits us to make causal statements about the effects of firms' investment in ICT and R&D and of importing and exporting decisions on firm productivity. This then allows us to quantify the impact of these factors on the demand for skilled and unskilled labor. The focus on firms is important, because that is where decisions about technological change, globalization and employment are made.

In order to identify the effect of ICT and R&D on firm level productivity we use data on workers in technology-related occupations. We call these workers "techies". We view techies—engineers and technicians with skills and experience in science, technology, engineering and math (STEM)—as essential to productivity growth, by virtue of being the creators of new products and processes, and as mediators of technology adoption at the firm level. Techies are central in creating, planning, installing, and maintaining ICT, as well as in training and assisting other workers in the use of ICT. Techies are also involved in R&D. They design and lead R&D processes and ensure the transfer of know-how to other workers in the firm. Using data on R&D techies offers an alternative to R&D expenditure data.

Using administrative data on the entire French private sector from 2009 to 2013 we find that both technology and trade have large effects. Using our baseline estimates, we find that compared to firms that don't import and don't employ techies, firms with a lot of techies have skill-augmenting productivity which is 60 percent higher. Techies have no effect on Hicks-neutral productivity. Turning to the effects of trade, we find that firms that import have skill-augmenting productivity which is 120 percent higher, and Hicks-neutral productivity which is 50 percent higher.

Because our estimation procedure also delivers estimates of the elasticities of substitution and demand, we can translate the estimated effects on productivity into effects on firm-level labor demand. Again compared to firms that don't import and don't employ techies, firms with a lot of techies have employment of skilled labor that is 60 percent higher, employment of unskilled labor that is 15 percent higher, and skill intensity that is 40 percent higher. The effects of importing

[2] We discuss this literature below.

are even larger: firms that import have employment of skilled labor that is 115 percent higher, employment of unskilled labor that is 25 percent higher, and skill intensity that is 70 percent higher. When we aggregate across firms, we find that our model can match the overall increase in skill intensity in our sample from 2009 to 2013.

These results on the employment effects of techies and trade are crucial to public policy debates. They show that unskilled workers are right to be wary of technology and trade, which we find do indeed favor employment of skilled workers. But this is a relative effect: because of the powerful productivity effects of technology and trade, both skilled and unskilled workers see labor demand rise when the firms where they work hire techies and/or engage in offshoring.

All the results we find in this paper are within-firm effects of firm level decisions. We do not consider why firms choose to employ techies or import, nor do we consider the effects of these firm-level decisions on industry or economy-wide equilibrium labor demand and wages. These are limitations of the scope of our paper but do not impinge on the credibility of our research strategy. Furthermore, any credible analysis of the equilibrium effects of technology adoption and globalization on labor markets must be built on an understanding of what goes on inside firms. This is where our contribution lies.

The rest of the paper is organized as follows. In section 2, we discuss papers directly related to our research questions and methodology. After a brief discussion of why not all firms employ techies in section 3, we develop our econometric methodology in Section 4 and describe our data and construction of the estimation sample in Section 5. Estimation results and the quantitative implications for skill bias and labor demand are reported in Section 6.

2 Related research

SBTC and globalization have been of intense interest to economists for decades, but there are remarkably few papers that look for SBTC at the firm or plant level, and none that simultaneously estimate, as we do, the effect of ICT, R&D and globalization on SBTC. We discuss these few papers here to put our contribution in context. We also review research on the role of techies in production, and highlight recent developments in the estimation of production functions that we build on.

2.1 Firm-level biased technological change and globalization

In this subsection we discuss papers that study firm- or plant-level changes in the composition of employment as a result of technological change.[3] Several excellent papers study a single industry or firm. Bartel et al. (2007) look at the valve manufacturing industry between 1999 and 2003 to study how adoption of ICT caused reorganization within firms. They show that ICT adoption increased Hicks-neutral total factor productivity (TFP) (through faster setup times, greater customizability, and better quality control) and also raised the skill-requirements for machine operators. Autor et al. (2002) study how the introduction of digital check imaging affected reorganization and the allocation of tasks across workers within one large bank. Acemoglu and Finkelstein (2008) find that when a policy change in 1983 increased the relative price of labor for hospitals, hospitals increased both their capital-labor ratio and their skill/unskill ratio among nurses - a result that is suggestive of complementarity between capital and skilled labor. Our paper has a much broader scope, as we study all private sector firms in France.

Some of the most informative papers about firm-level SBTC are primarily descriptive. Dunne et al. (2004) uses the Census' Longitudinal Research Database and find that computer use within plants is not associated with higher overall labor productivity but is associated with greater non-production worker intensity, a common proxy for skilled workers. Helper and Kuan (2018) performed a 2011 survey of the auto parts industry, and find that most firms in this industry do not perform R&D, but innovate nonetheless through the efforts of their engineers and technicians. They also find that the tasks done by engineers overlap much more with skilled than with unskilled workers. Barth et al. (2017) study plant-level data on US manufacturing firms, with a focus on the role of scientists and engineers. In the private sector as a whole, they show that 80 percent of scientists and engineers worked outside R&D occupations in 2013. They estimate a simple gross revenue production function at the establishment level from 1992 to 2007 using fixed effects OLS, and find statistically significant effects of the science and engineer share of employment on revenue, a result suggestive of a positive effect of scientists and engineers on TFP. Bresnahan et al. (2002) argue persuasively for a complementarity between IT, decentralized firm organization, and skilled labor. They construct measures of "work organization", computer capital, and employee skill for a small number of very big publicly traded firms in the mid-1990s, and find robust correlations that are consistent with their view. Building on these insights, our paper moves beyond descriptive

[3] To keep this literature review manageable, we eschew discussion of the many important papers that analyze SBTC and related issues using industry-level data.

analysis to structural estimation.

Two recent papers on firm level skill-biased technology both use data from Norway to estimate causal effects. Akerman et al. (2015) exploit exogenous variation in the local availability of broadband internet in the 2000s, and find convincing intent-to-treat effects on both local skilled wages and firms' output elasticity for skilled workers. As they acknowledge, their evidence that firms who adopt broadband internet increase their skill intensity is weaker. Bøler (2015) uses a 2002 tax break for R&D expenditure to estimate the effects of R&D on firm-level skill intensity in manufacturing. Her reduced form evidence is supportive of a very strong effect of R&D on skill intensity, while her structural production function estimates find a smaller but still important effect. These Norwegian papers are importnat antecedents to our paper, but the greater size and diversity of the French economy, as well as our analysis of nonmanufacturing in addition to manufacturing firms, allows us to estimate broader and more nuanced effects.

Turning to the effects of globalization, Becker et al. (2013) use German micro data on employment and offshoring by multinational firms during 1998-2001. They find a positive association between offshoring and plant-level skill intensity. For Indonesia, Kasahara et al. (2016) find that plant-level use of imported materials raised the level of education within manufacturing plants between 1995 and 2007.[4] Bustos (2011) finds that Argentinian firms raised their productivity and skill intensity after a major trade reform in 1991. She also finds an association between spending on ICT and skill upgrading. Our results are consistent with the papers by Kasahara et al. (2016) and Bustos (2011), but we are not able to investigate the channel identified by Becker et al. (2013) since we don't have information on foreign affiliates of the French firms in our data.

Like us, Bender et al. (2018) use a framework that has both Hicks-neutral and management-augmenting technological differences across firms. Applying the methodology of Abowd et al. (1999), they construct individual-level measures of worker quality, which they match to firms. This very creative paper is hampered by matching problems that lead to a sample size of just 361 German firms across three years (2004, 2006, and 2009). As a consequence, their data analysis is mainly descriptive, while our much larger dataset allows us to do structural estimation.

Two recent papers estimate firm-level labor augmenting technology under the assumption that firms produce using a CES production function of capital, labor, and materials, which is very similar to what we do. Doraszelski and Jaumandreu (2018) develop a production function estimator that

[4] Amiti and Cameron (2012) also study firm-level data from Indonesia, but their primary focus is the skill premium rather than skill intensity.

is relevant to our methodology, so we defer discussion of this paper until section 2.3 below. Raval (2017) uses an equation implied by static cost minimization to estimate both the elasticity of substitution and the level of labor-augmenting technology differences across U.S. manufacturing plants between 1987 and 2007.[5] Though these two papers are not directly about skill bias, their findings of large labor augmenting technology differences across firms have a plausible interpretation as differences in the skill mix across firms. By contrast, our structural model directly estimates skilled labor augmenting technological differences across firms.

2.2 The role of techies

Though we are the first to analyze the impact of techies on SBTC, there is a small literature that has looked at the impact of techies on output, the structure of employment, and productivity at the firm level. The motivation for this literature is stated succinctly by Tambe and Hitt (2014): "the technical know-how required to implement new IT innovations is primarily embodied within the IT workforce".

Firm-level research on this proposition has been hampered by a lack of firm-occupation level data in most administrative and survey datasets. An exception is our Harrigan et al. (2016), which uses detailed occupational data (including data on techies) for the entire French private sector from 1994 to 2007. Harrigan et al. (2016) shows that French firms with more techies in 2002 grew more rapidly from 2002 to 2007, and also that increases in techies led to within-firm occupational polarization. Lichtenberg (1995) and Brynjolfsson and Hitt (1996), working with a small number of U.S. firms in the late 1980s and a simple Cobb-Douglas production function estimating equation, find that IT labor has a positive output elasticity. Tambe and Hitt (2012) use a newer data source and a more sophisticated estimation technique, and again find a positive output elasticity of IT labor. However, none of these three papers test the null that IT and non-IT labor have the same output elasticity. Using a remarkable dataset that tracks the movement of IT workers across firms, Tambe and Hitt (2014) find what they interpret as evidence for knowledge spillovers across firms through the channel of techie mobility. In the present paper, rather than treating IT labor as a simple input we estimate the effect of IT labor on productivity. In addition, our specification is not vulnerable to the endoegenity bias that plagues OLS estimation of Cobb-Douglas production functions.

[5] The industry-level estimates from an earlier version of Raval (2017) are used to compute the aggregate elasticity of substitution in Oberfield and Raval (2014).

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 predicting income today.[6] Indeed, engineers are at the center of modern (endogenous) growth theory, e.g., Romer (1990).

2.3 Production function estimation

The modern approach to firm-level productivity estimation is to specify output as a function of inputs, develop an estimation methodology that identifies the parameters of the production function, and then back out the implied estimated productivity. The foundational papers in this literature are Olley and Pakes (1996)[OP], Levinsohn and Petrin (2003)[LP], and Ackerberg et al. (2015)[ACF]. The OP/LP/ACF methodology takes it as given that data on real inputs and outputs are available. This methodology also treats productivity differences across firms as exogenous and Hicks-neutral. We build on recent methodological developments that move beyond these assumptions.

Most firm- or plant-level datasets (including ours) include information on revenue R_{ft} for firm f in year t and the value of expenditures on materials M_{ft} but not data on the corresponding output and materials prices p^Y_{ft} and p^M_{ft}. De Loecker and Goldberg (2014) give a clear exposition of the estimation and interpretation problems that arise when real input and output quantities are unavailable. Grieco et al. (2016) [GLZ] show how to estimate the parameters of a CES production function even in the absence of real output or input data, by exploiting the firm's first-order conditions for profit maximization. In section 4.1 below, we extend the GLZ estimator to allow for factor-augmenting productivity differences.

Economists have long been interested in the determinants of firm level productivity, and there are many papers that follow some variant of the OP/LP/ACF methodology to first estimate productivity and then study its determinants, most often in the context of estimating the effects of importing or exporting on productivity.[7] De Loecker (2013) points out the limitations of this approach. In particular, he shows that if productivity is an endogenous function of exporting then

[6] See also Murphy, Shleifer and Vishny (1991) for evidence on the relationship between engineers (versus lawyers) and income.

[7] Pavcnik (2002) may have been the first to do this, and Amiti and Konings (2007) is another well-known example.

a measure of exporting must be included in the moment conditions to get a consistent estimator for the production function. Doraszelski and Jaumandreu (2013) take a different approach to estimating endogenous productivity, combining the firm's parametric labor demand function with a nonparametric controlled Markov specification for productivity. Our estimation of endogenous productivity in section 4.4 below builds on these two papers.

Doraszelski and Jaumandreu (2018) is the first paper in the production function estimation literature to estimate both neutral and non-neutral technology differences. Their model has a 3-factor CES production function with Hicks-neutral ω_H and labor-augmenting ω_L technology differences across firms. Labor and materials are static inputs, which implies that the optimal labor to materials ratio depends on ω_L but not on ω_H. This insight motivates a two stage procedure. In the first stage they recover $\widehat{\omega}_L$ and the estimated elasticity of substitution $\widehat{\sigma}$ through estimation of a relative factor demand equation implied by the CES functional form. They then use $\widehat{\omega}_L$ and $\widehat{\sigma}$ as data in a second stage which allows them to recover ω_H. In section 4.1 below we show how a similar model can be estimated in a single step.

3 Why don't all firms employ techies?

Many of the papers in section 2.1 find that employment of techies enhances productivity, which raises a simple question: why don't all firms employ them? As we will show in section 6.2, in our sample of French firms techies are found to have strong positive effects on skill-biased productivity, yet most firms do not employ techies. A similar finding is well-known to trade economists: in some studies of developing countries, exporting is found to raise productivity, yet a minority of firms export. Following Melitz (2003), the consensus explanation for this phenomenon is fixed costs: firms choose to export only when the extra revenue from exporting exceeds the fixed costs of exporting. Alternatively, the variable costs of exporting may make it unprofitable for high-cost firms, as shown by Melitz and Ottaviano (2008). Here we sketch a simple model that makes a similar point about techies, and that gives a rationale for a constant elasticity relationship between techies and productivity. We do not estimate this model, rather we use it here to make a few simple theoretical points.

For maximum simplicity, suppose there are only two periods and one type of productivity. The firm takes demand, costs and initial period log productivity ω_{ft-1} as given and has to choose optimal techie employment T_{ft-1} to maximize profits. The relationship from techies to changes in productivity is

$$\omega_{ft} = \omega_{ft-1} + Max\left[\delta\, ln\left(\frac{T_{ft-1}}{\gamma_{1f}}\right), 0\right], \quad \delta \geq 0$$

Although the elasticity of productivity with respect to techies is constant and equal to δ, the level of techie employment required to attain a given growth in productivity $\Delta\omega_{ft}$ will differ across firms because of differences in γ_{1f}. Fixed costs of employing positive techies are γ_{0f} and the wage of techies is r, so the cost of hiring techies is $rT_{ft-1} + \gamma_{0f}$. With heterogeneity in the costs γ_{0f} and γ_{1f} not all firms will employ techies, and we derive the following very intuitive conclusions in Appendix A.2.6. First, the optimal amount of techies is more likely to be positive when demand and/or initial productivity are higher. Conversely, the optimal amount of techies is more likely to be zero when fixed costs of techies are high and/or when the efficiency of techies are low. Second, the optimal amount of techies may be zero even if the fixed cost of employing techies is zero. Finally, when the optimal amount of techies is positive, it is increasing in initial productivity and the efficiency of techies. A further implication of this framework is that since firms that export will have a higher demand level, they will also be more likely to employ techies.

The evidence in Table 1 is consistent with these simple predictions. The table reports regressions of firm-level outcomes on an indicator for positive techie employment (column 1) and the techie share of the wage bill when it is positive (column 2). All regressions include firm $\times$ industry fixed effects, so the reported results are identified by variation across firms within industry-years. Just 11 percent of firm-year observations have positive techies, as shown by the relative sample sizes in the two columns. The results show that techies are associated with greater revenue and a greater propensity to import and export and, in a preview of our structural analysis below, are also associated with a higher share of managers in the firm's wage bill. Column (1) shows that larger firms, importing and exporting firms, and firms that have a higher managerial intensity are far more likely to employ techies.[8] Column (2) shows that these relationships are somewhat weaker on the intensive margin: among firms who employ techies, the effect of size and export propensity are smaller, and the effect is indistinguishable from zero for importing. The exception is the large marginal effect of the share of the managerial wage bill share. The final row of Table 1 hints at our findings in Section 6.2, where we show that more techies lead to greater employment of managers.

[8] Since we exclude techies from the wage bill in these regressions, the techie correlation with manager's wage bill is not mechanical.

4 Econometric methodology

The modern approach to firm-level productivity estimation is to specify output as a function of inputs, develop an estimation methodology that identifies the parameters of the production function, and then back out the implied estimated productivity. The foundational papers in this literature are Olley and Pakes (1996)[OP], Levinsohn and Petrin (2003)[LP], and Ackerberg et al. (2015)[ACF]. The OP/LP/ACF methodology takes it as given that data on real inputs and outputs are available. But this is rarely the case. This methodology also treats productivity differences across firms as exogenous and Hicks-neutral. We build on recent methodological developments that move beyond these assumptions.

Most firm- or plant-level datasets (including ours) include information on revenue R_{ft} for firm f in year t and the value of expenditures on materials M_{ft} but not data on the corresponding output and materials prices p^Y_{ft} and p^M_{ft}. De Loecker and Goldberg (2014) give a clear exposition of the estimation and interpretation problems that arise when real input and output quantities are unavailable. Grieco et al. (2016) [GLZ] show how to estimate the parameters of a CES production function even in the absence of real output or input data, by exploiting the firm's first-order conditions for profit maximization.

Many papers follow some variant of the OP/LP/ACF methodology to first estimate productivity and then study its determinants, often in the context of estimating the effects of exporting or importing on productivity.[9] De Loecker (2013) points out the limitations of this approach. In particular, he shows that if productivity is an endogenous function of exporting then a measure of exporting must be included in the moment conditions to get a consistent estimator for the production function. Doraszelski and Jaumandreu (2013) take a different approach to estimating endogenous productivity, combining the firm's parametric labor demand function with a nonparametric controlled Markov specification for productivity.

Doraszelski and Jaumandreu (2018) is the first paper in the production function estimation literature to estimate both neutral and non-neutral technology differences. Their model has a 3-factor CES production function with Hicks-neutral ω_H and labor-augmenting ω_L technology differences across firms. Labor and materials are static inputs, which implies that the optimal labor to materials ratio depends on ω_L but not on ω_H. This insight motivates a two stage procedure. In the first stage they recover $\widehat{\omega}_L$ and the estimated elasticity of substitution $\widehat{\sigma}$ through estimation of

[9] For example, Pavcnik (2002) and Amiti and Konings (2007).

a relative factor demand equation implied by the CES functional form. They then use $\widehat{\omega}_L$ and $\widehat{\sigma}$ as data in a second stage which allows them to recover ω_H.[10]

Our approach extends Grieco et al. (2016) [GLZ] in two ways. First, we separate labor into three components: skilled and unskilled labor S and L, which contribute to output in the standard way, and workers T in technical occupations ("techies") who are assumed to affect production only through their lagged impact on productivity. Second, we allow firm production functions within an industry to differ in two dimensions: in addition to a Hicks neutral term $\Omega_{Hft} = e^{\omega_{Hft}}$, already present in GLZ, we add a skilled-labor augmenting term $\Omega_{Sft} = e^{\omega_{Sft}}$. Here we outline our approach, with details discussed in Appendix A.2.

4.1 Estimating productivity

We begin with a CES function where physical output Y_{ft} is produced using skilled labor S_{ft}, unskilled labor L_{ft}, capital K_{ft}, and materials M_{ft}. This function is assumed to be the same for all firms in an industry up to the two productivity levels Ω_{Hft} and Ω_{Sft}. For reasons discussed by GLZ, it is important for identification to normalize each data series by its geometric mean, and we choose units/minimize notation such that the geometric means $\overline{L} = \overline{S} = \overline{K} = \overline{M} = \overline{Y} = 1$.[11] Skilled labor services are the product of hours worked and skilled labor augmenting productivity Ω_{Sft}. The normalized production function is then

$$Y_{ft} = \Omega_{Hft} \left[\alpha_L L_{ft}^{\gamma} + \alpha_S \left(\Omega_{Sft} S_{ft} \right)^{\gamma} + \alpha_K K_{ft}^{\gamma} + \alpha_M M_{ft}^{\gamma} \right]^{\frac{1}{\gamma}}, \ \gamma = \frac{\sigma - 1}{\sigma} \leq 1 \tag{1}$$

Higher skill-augmenting technology Ω_{Sft} increases the effective supply of skilled labor services holding hours worked constant. Similarly, better Hicks-neutral technology Ω_{Hft} shifts physical output holding all physical inputs and skill-augmenting technology constant. Input and output prices may differ across firms, but the researcher only observes revenue R_{ft} and the value of materials purchases E_{ft}^M, along with physical L, S and K and wages. The labor and materials inputs are assumed to be chosen after Ω_{Hft} and Ω_{Sft} are observed. To go from revenue to output requires an assumption on demand, and we follow GLZ in assuming that firms face a common, industry-level constant elasticity of demand $\eta < -1$. The inverse demand function facing the firm is

[10] In their application to Spanish manufacturing data in 1990–2006, Doraszelski and Jaumandreu (2018) find great heterogeneity in the level and growth of labor-augmenting productivity differences, with weighted average annual ω_L growth of 1.5 percent per year..

[11] To fully understand the relevance and importance of normalizing the CES production function, see the discussion and references on page 668 of Grieco et al. (2016). Below we illustrate one important outcome of normalization: it helps identifying the distribution parameters of the production function.

$$P_{ft} = A_t Y_{ft}^{\frac{-1}{\eta}} \tag{2}$$

where A_t is an exogenous industry-level demand shifter.

A revenue shock u_{ft} is realized after all input choices have been made and both productivity levels have been realized.

4.1.1 The estimating equation

Our approach to developing an estimating equation uses economic theory. Since L, S and M are static inputs, their first order conditions for expected profit maximization will always hold with equality.[12] As shown in Appendix A.2, these imply

$$M_{ft} = \left(\frac{\alpha_L}{\alpha_M} \frac{E^M_{ft}}{E^L_{ft}} \right)^{1/\gamma} L_{ft} \tag{3}$$

$$\Omega_{Sft} = \left(\frac{S_{ft}}{L_{ft}} \right)^{\frac{1}{\sigma-1}} \left(\frac{\alpha_S W^L_{ft}}{\alpha_L W^S_{ft}} \right)^{\frac{\sigma}{1-\sigma}} \tag{4}$$

where $E^M_{ft} = P_{Mft} M_{ft}$ is expenditures on materials and $E^L_{ft} = W_{Lft} L_{ft}$ is the unskilled labor wage bill. The derivation of (3) and (4) requires that $\sigma \neq 1$, which ironically rules out the Cobb-Douglas case that is the assumption in most of the productivity estimation literature, including OP/LP/ACF.[13] A few more steps yields an estimating equation,

$$lnR_{ft} = ln \left[\frac{\eta}{1+\eta} \right] + ln \left[E^S_{ft} + E^M_{ft} + E^L_{ft} \left\{ 1 + \frac{\alpha_K}{\alpha_L} \left(\frac{K_{ft}}{L_{ft}} \right)^{\gamma} \right\} \right] + u_{ft} \tag{5}$$

where $E^S_{ft} = W_{Sft} S_{ft}$ is the skilled labor wage bill.[14] Equation (5) has just three parameters (η, γ and α_K/α_L), and as in GLZ it can be estimated by nonlinear least squares. The key to the derivation is that there are three static inputs (S, L and M), which gives us two ratios of static first order conditions, (3) and (4). These two equations allow us to eliminate the two unobservables, M_{ft} and Ω_{Sft}, and (A.19) in the Appendix allows us to eliminate Ω_{Hft}. Because we have eliminated the unobserved productivity terms from our estimating equation, we do not need to use the proxy methods of the OP/LP/ACF methodology. Our timing assumptions, which are the same as in

[12] That is, before the revenue shock u_{ft} is realized.

[13] In our results below in Table 5, the point estimate for σ exceeds one for every industry, and we can always reject the null hypothesis that $\sigma = 1$,

[14] See Appendix A.2.1 for the derivation of (5).

OP/LP/ACF, are key: firms choose static inputs after observing both productivity shocks but before observing the revenue shock.

The model has six parameters of interest ($\eta, \gamma, \alpha_S, \alpha_L, \alpha_K$ and α_M). Given the estimates of η, γ and α_K/α_L, the four distribution parameters are identified by the following equations,

$$\alpha_L + \alpha_S + \alpha_K + \alpha_M = 1 \tag{6}$$

$$\alpha_M \overline{E}^L = \alpha_L \overline{E}^M \tag{7}$$

$$\alpha_S \bar{E}^L = \alpha_L \bar{E}^S \tag{8}$$

Equation (6) is implied by constant returns to scale. Equations (7) and (8) follow by taking the geometric means of (3) and (4) respectively, and using the normalization conditions.

Our estimator for (5) is weighted nonlinear least squares, with weights given by total employment in firm *f*. This is appropriate given that we want to estimate population average partial effects, where the population of interest is industry employment (see Solon et al. (2015) for a discussion of this rationale for weighting). In the absence of employment weights, firms with few workers would have the same influence on the estimates as firms with very many workers, which we want to avoid.

4.1.2 Recovering productivity

Once this estimator is implemented, we can recover estimated Hicks neutral and skill augmenting productivity using (A.20) in the Appendix and (4), respectively. Fully recovering Hicks neutral productivity would also require an estimate of the unobservable aggregate A_t. This doesn't matter for the cross sectional distribution at a point in time, but it does imply that our Hicks neutral productivity estimates are comparable over time only in relative terms. That is, we can compare two firm's productivity in a given year, and we can say how this comparison changes over time, but we cannot compare Hicks-neutral productivity shifters for a given firm over time.

Formally, the Hicks neutral parameter Ω_{Hft} is physical TFP. This follows from deflating revenue by price using equation (2). In practice, it is not plausible that the simple demand system given by (2) solves all the problems related to unobservable prices and quality that are required to distinguish revenue TFP from physical TFP, the way that Foster et al. (2008) do in their classic paper. Our interpretation of Ω_{Hft} will be as revenue TFP, where some but not all of the variation in revenue has been controlled for by the demand system.

4.2 Skill-augmenting productivity and skill bias

Hicks-neutral technology differences Ω_{Hft} have no implications for relative skill demand because they do not affect the relative marginal products of different inputs. Re-arranging equation (4) shows that the effect of skill augmenting technology differences Ω_{Sft} on relative skill demand depends crucially on the elasticity of substitution σ,

$$\frac{S_{ft}}{L_{ft}} = \Omega_{Sft}^{\sigma-1} \left(\frac{\alpha_S W_{ft}^L}{\alpha_L W_{ft}^S} \right)^{\sigma} \tag{9}$$

If $\sigma > 1$, a higher level of Ω_{Sft} raises relative skill demand, which is to say that skill *augmenting* technology differences are skill *biased*. In our empirical results below we estimate $\widehat{\sigma} > 1$ for all industries.[15] When $\sigma > 1$, the identification of Ω_{Sft} is transparent from equation (4): it is a residual that rationalizes greater skill intensity, conditional on parameters and factor prices.

Estimating equations similar to (9) in log-linear form and identifying the elasticity of substitution have a long history in the macro-labor literature on SBTC; see Acemoglu and Autor (2011) for discussion and references. More recently, Raval (2017) and Doraszelski and Jaumandreu (2018) estimate the elasticity of substitution between capital and labor and labor-augmenting technology shifters using similar equations. In our methodology this source of variation is only part of what identifies σ in the data; as (5) illustrates, variation in levels also plays an important role. See Grieco et al. (2016) and references therein on the advantages of our approach.

4.3 Output and employment effects of productivity

When productivity differences are solely Hicks-neutral across firms, the output effect of productivity is obvious: the elasticity of output with respect to productivity is one. Because skill-augmenting productivity Ω_S enters the production function by multiplying skilled labor hours S, the output elasticity with respect to Ω_S is the same as the output elasticity with respect to S.

There are two effects of productivity improvements on factor demand, holding product demand curves and factor prices constant. First, greater productivity lowers costs which increases final demand and thus the demand for inputs. Second, greater productivity means fewer inputs are

[15] See Table 5.

required per unit output, which reduces the demand for inputs. The net effect on factor demand depends on the balance between these two effects. When technological change is pervasive, there will be general equilibrium effects on both factor prices and market demand which are beyond the scope of this paper. But holding factor prices and market demand constant, we show in Appendix A.2.4 that the effects of productivity differences across firms on employment are

$$\begin{aligned} \frac{dS_{ft}}{S_{ft}} &= (\eta - 1)\, d\,\omega_{Hft} + [\sigma - 1 + (\eta - \sigma)\,\lambda_{Sft}]\, d\,\omega_{Sft} \\ \frac{dL_{ft}}{L_{ft}} &= (\eta - 1)\, d\,\omega_{Hft} + \qquad\qquad (\eta - \sigma)\,\lambda_{Sft}\, d\,\omega_{Sft}\,, \end{aligned} \tag{10}$$

where λ_{Sft} is the share of S in cost. For Hicks-neutral productivity ω_H, the intuition for the elasticity of $\eta - 1$ is that the labor-saving effect has an elasticity of -1, while the demand effect through lower costs has an elasticity η. The elasticities with respect to skill-augmenting productivity ω_S incorporate both the demand effect through η, which tends to raise employment, and the substitution effect through σ. When $\sigma > 1$, the substitution effect raises S/L when $d\,\omega_S > 0$, which is the SBTC case. We apply equations (10) in our empirical analysis below.

4.4 Endogenous productivity

In the OP/LP/ACF methodology, productivity is treated as completely exogenous. But one reason to do firm-level productivity estimation (and one of our main research questions) is to be able to study what causes the estimated productivity differences. In the trade literature, this has been done repeatedly in the context of explaining the fact that exporters have higher productivity: is this fact due to selection à la Melitz (2003), or is there an additional causal "learning-by-exporting" effect?

In developing our estimating equation (5), we made no assumptions about the stochastic processes that characterize productivity. Because of this, we are free to study the determinants of productivity in a flexible way, using firm-level explanatory variables. Following Doraszelski and Jaumandreu (2013), we now assume that productivity is given by a "controlled Markov" process, where productivity depends on three factors: (1) lagged productivity, (2) a $k \times 1$ vector of lagged characteristics of the firm z_{ft-1}, and (3) a shock which is orthogonal to all the other shocks in the model.

The lagged firm characteristics z_{ft-1} include choice variables of the firm such as exporting, importing and employment of techies as well as predetermined firm characteristics such as age and size which are known to help predict productivity. To allow ω_{Hft} and ω_{Sft} to influence each other we specify the following two equations,

$$\omega_{Hft} = \mu_{Ht} + \beta_{HH}\omega_{Hft-1} + \beta_{HS}\omega_{Sft-1} + \beta_{HZ}z_{ft-1} + \xi_{Hft} \tag{11}$$

$$\omega_{Sft} = \mu_{St} + \beta_{SH}\omega_{Hft-1} + \beta_{SS}\omega_{Sft-1} + \beta_{SZ}z_{ft-1} + \xi_{Sft} \tag{12}$$

The shocks ξ_{Hft} and ξ_{Sft} are assumed to be serially uncorrelated. The industry×time fixed effects μ_{Ht} and μ_{St} control for among other things the demand shifter A_t. These equations can be consistently estimated by OLS. De Loecker (2013) and Doraszelski and Jaumandreu (2013) estimate more general non- or semi-parametric versions of (11) and (12). A virtue of our parametric specification is that it is straightforward to calculate the steady-state cross-sectional effects of persistent differences in firm characteristics,

$$\begin{bmatrix} \omega_{Hf} \\ \omega_{Sf} \end{bmatrix} = (I-B)^{-1}\beta_Z z_f, \qquad B = \begin{bmatrix} \beta_{HH} & \beta_{HS} \\ \beta_{SH} & \beta_{SS} \end{bmatrix}, \; \beta_Z = \begin{bmatrix} \beta_{HZ} \\ \beta_{SZ} \end{bmatrix} \tag{13}$$

4.4.1 Interpretation and identification

It is important to be clear about what is meant by a "controlled Markov process". The key is that the Markov assumption breaks realized productivity into expected and unexpected components. Thus, statistical exogeneity of lagged productivity and firm characteristics in (11) and (12) is assured, but can we interpret the estimated effects of (say) techies as causal in the cross section? For example, if the estimated effect of techies in (11) is positive, can we say "techies cause higher Hicks-neutral productivity"? If the answer is yes, that raises the question, what determines the choice of techies and trade status, and why don't all firms make the same choices? In the trade context, underlying differences in firm-specific trade costs have been used to explain why not all firms trade, and similar reasoning can be applied in the case of techies: some products/processes are simply harder to improve using ICT, and/or firms have unobservable heterogeneity in their aptitude for applying IT and thus employing techies. In Section 3 above, we presented a simple model of optimal techie choice to clarify the insight that equations (11) and (12) can consistently estimate the effect of techies on productivity even in the absence of a structural model of techie choice.

De Loecker (2013) has a persuasive discussion of how to interpret the learning-by-exporting effect in his version of the controlled Markov process (page 8). He emphasizes two things. One, it is *lagged* exporting that enters the Markov process, which is to say that productivity (more precisely, the shock to productivity ξ_{Hft}) is realized after the exporting decision is made. Two, the persistence of the exporting decision is controlled for by having lagged *realized* productivity in the equation for current productivity. These arguments extend directly to our setting.

The way that Doraszelski and Jaumandreu (2013) discuss their estimated effects of R&D on productivity is to remain silent on the issue of how R&D decisions are decided. That is, they answer the question: given that a firm has decided to do R&D, what is the estimated effect on productivity? We will take the same approach, and will interpret our estimates as answering the question: given that a firm has decided to trade and/or employ techies, what is the estimated effect on productivity?

As in De Loecker (2013) and Doraszelski and Jaumandreu (2013), identification of the effects of firm choices on productivity is based on cross-sectional differences in productivity growth between firms that do or do not make a given choice. For example, consider two firms with the same lagged productivity and all other explanatory variables except that one firm chooses to employ techies and the other does not. If the firm with techies has higher productivity in the next period, the estimator attributes that to the firm's employment of techies.

In our application we measure techies by the share of techies in the firms' wage bill, which has the virtue of capturing both the extensive and intensive margin of techie employment. Imported inputs are already included in a firm's purchases of materials M_{ft}, though our data do not allow us to distinguish between domestic and imported inputs. To allow for a productivity effect of importing while avoiding double counting of imported inputs, we measure importing by an indicator variable. As discussed in Grieco et al. (2017), the importing indicator can be thought of as measuring the firm's access to a broader and/or cheaper range of inputs than are available domestically. For symmetry with how we treat imports, we also measure a firm's exporting activity by an indicator variable.

4.4.2 Do techies belong in the production function?

A central element of our methodology is that we assume that techies affect output only through their effect on future productivity, and not through any contemporaneous contribution to factor services that affect current output. This assumption is analogous to the standard assumption that investment in *t-1* has no effect on output in *t-1*, but raises output in t through its contribution to K_t. Our reasons for specifying the role of techies in this way are both theoretical and empirical. Theoretically, if techie employment in t affects output in t as part of labor input in t as well as productivity in *t+1,* then the static first order conditions for optimal employment would not hold and the derivation of our estimating equation (5) does not go through. Empirically, if techies enter the production function (1) as a separate factor, an implication is that employment of techies would be strictly positive for all firms in all periods, which is emphatically not the case in our data, where only 11 percent of firms employ techies. As discussed above, a number of empirical studies provide support for our specification of the effect of techies on output.

While our assumption that techies affect output only through their effect on future productivity is well-grounded, it is important to consider how our measurement of productivity could go awry if techies do in fact increase current output directly, which we will call the "orthodox case". If the orthodox case is the right specification, then leaving techies out of the first-stage production function (1) will understate labor inputs in the first stage. In Appendix A.2.3 we show that the greater is the underestimate of true inputs the greater will be the overestimate of Hicks-neutral productivity. Less obviously, we also show that when $\sigma > 1$, firms with high techie shares will have measured ω_S which is biased down by more than for firms with low techie shares. The reason is that with $\sigma > 1$ greater ω_S leads to greater employment of skilled workers S, which implies that higher S indicates higher true ω_S. Incorrectly removing techies from S will thus lead to an underestimate of ω_S, and the underestimate will be larger the greater is the share of techies in S. The key implication is that if the orthodox model is correct, our estimated effect of techies on ω_H in equation (11) may be biased up, and our estimated effect of techies on ω_S in equation (12) may be biased down. However, these biases are largely mitigated due to two related reasons. First, any biases in the estimated levels of productivity appear on both sides of equations (11) and (12). Second, although mis-specification leads to a mechanical *intra-* temporal correlation between techies and ω_S and between techies and ω_H, since we estimate the effect of techies on *future* productivity controlling for current productivity, i.e. an *inter*-temporal relationship, this direct effect washes out.

A further implication of the orthodox model is that if we do include techies as part of labor input in the first-stage estimating equation (5), they should have no explanatory power in the second stage regressions (11) and (12). We test this possibility be estimating the first stage with techies as part of labor input, and then testing the null that techies have no effect in the second stage. The null is rejected at the 0.01 percent significance level.

4.5 Estimation details

We estimate equation (5) by nonlinear least squares, separately for 16 industries including both manufacturing and non-manufacturing sectors, with observations weighted by employment (Section 4.1.1 above discusses estimation in greater depth).[16] Standard errors are clustered by firm. Each industry NLLS regression is an unbalanced panel, which raises the issue of selection bias due to endogenous exit. As pointed out by Ackerberg et al. (2007), endogenous exit will not bias production function estimation as long as the firm exits in the period after the exit decision has been made. This (often implicit) assumption is now standard in the literature, and we make it here. The estimated elasticity of substitution is given by the formula $\widehat{\sigma} = (1 - \widehat{\gamma})^{-1}$, with the standard error of $\widehat{\sigma}$ computed by bootstrapping (see discussion below). Our estimation sample is summarized in Table 4.

Industry-level production function estimation generates estimated Hicks neutral and skill augmenting productivity for each firm-year, computed using equations (4) and (A.20). After dropping the highest and lowest percentile of estimated productivity to trim outliers, we estimate the controlled Markov processes given by equations (11) and (12). In these regressions, we measure techies by the lagged share of techies in the firm's wage bill, and import and export participation are measured by indicator variables. We also include lagged firm size (measured by lagged revenue) and firm age as additional controls. Estimation is by weighted least squares with industry $\times$ year fixed effects and bootstrap standard errors clustered by firm.[17]

5 Data

We use detailed panel data on firms in the French private sector economy between 2009 and 2013. Estimating productivity requires information on firms' revenue, labor force composition, capital, and materials expenditures. As our second step hypothesizes that productivities are determined by techies, trade and other firm characteristics, we need additional information at the level of the firm. We therefore merge a set of three confidential firm-level datasets. This matching process is

[16] GLZ show that this is equivalent to the GMM estimator. Two sectors (coke and refined petroleum, and pharmaceutical products) are dropped because they have tiny shares of total hours worked and very few firms, and two sectors (transport equipment and publishing/broadcasting) are dropped because estimation of equation (5) failed to converge. We also drop the financial intermediation sector.

[17] Since the regressors in (11) and (12) are generated by first-stage estimation of (5) for each industry, we bootstrap the full two-stage procedure to calculate standard errors. Our bootstrap re-samples firms rather than individual firm-year observations, so the resulting bootstrap co variance matrices are effectively clustered by firms. All bootstrap results are computed using 400 replications. For more details on how we compute standard errors, see Appendix A.2.2.

straightforward because firms in France are identified by the same identification number (SIREN), which can be followed across years in each of the three datasets. We highlight key features of the data here, and relegate other details to Appendix A.1.

5.1 The composition of labor within firms

Our first source of information is taken from the annual declaration of social data (DADS) dataset. The DADS is a requirement for all businesses with employees. Employers provide information on employees in each of their establishments, which are identified by their SIRET.[18] The first nine digits of each SIRET is the firm-level identification number, SIREN, which makes it easy to aggregate across establishments for each firm. For each worker, the DADS reports gross and net wages, hours paid, tenure, gender, age and occupation. It also reports the sector of activity of the firm. There is no information about workers' education.

We use the French occupational classification PCS-ESE to allocate all workers to one of three broad categories (Appendix Table A1 lists the two-digit PCS codes). Detailed four-digit occupational codes are reported in the DADS beginning in 2009, which determines the first year of our sample.

Table 2 lists the 4-digit occupations that we classify as techies, based on the occupational descriptions. Techie occupations are a subset of the two digit occupations "technical managers and engineers" (38) and "technicians" (47), and are closely related to the installation, management, maintenance, and support of ICT, as well as product and process design and longer-term R&D activities. In our empirical anaylysis, we will exploit the distinctions in our data between the more senior and highly-educated technical managers engineers and the lower-ranked technicians. Similarly, we will look separately at the effect of techies whose job descriptions mention ICT and those who work in R&D occupations. Table 3 reports the shares of ICT and R&D workers, as well as of engineers and technicians, within the overall techie wage bill. R&D workers are a somewhat larger share of the techie wage bill than the share of ICT workers, and the R&D share increases slightly from 2009 to 2013. The share of engineers in the techie wage bill is more than twice as large as the share of technicians, and also increases from 2009 to 2013. Table 3 also reports wide dispersion in these shares across industries.

[18]The declaration file serves both fiscal and social administrative purposes. All employers and their employees are covered by the DADS declaration with the exception of self-employed and government bodies, domestic services (section 97-98 of NAF rev. 2) and employees in businesses outside French territory (section 99 of NAF rev. 2). The data do not include worker identifiers, so we cannot track workers over time, but this does not concerns us given our focus on firm-level rather than individual outcomes.

The techie share of hours as a measure of firm-level technological sophistication can be compared to R&D expenditures, a common metric for technology adoption in the literature. Firm-level R&D is a useful measure, but it excludes much of the ongoing expenditure and managerial attention that firms devote to technology adoption and ICT use.[19] In fact, reported R&D is not even a necessary condition for technology adoption and innovation, and firms employ many scientists and engineers in non-R&D occupations.[20]. Conversely, R&D is likely to be impossible without the employment of techies, who are needed to install, maintain and manage the ICT used in R&D departments. Thus, the techie share is a more comprehensive measure of firm-level effort devoted to technology adoption than R&D expenditures. The wage bill of techies is a big chunk of overall ICT spending: Saunders and Brynjolfsson (2016) found that for a sample of US firms, more than half of all spending on IT was on techies[21].

One potential threat to our approach that treats firm-level techies as an indicator of firm-level technological sophistication is that firms can purchase ICT and R&D consulting services. By hiring a consultant, firms can obtain and service ICT without increasing its permanent staff of techies. However, less than 4% of techie hours are in the IT and R&D consulting sectors, which implies that over 96% of the hourly services supplied by techies are obtained in-house rather than purchased from consultants.[22]

In order to construct the broad managerial occupation category S we aggregate the number of hours worked by firm owners (proprietors, CEOs or directors of firms), workers in top management positions, and professionals and engineers whose tasks are not related to either ICT or R&D. These are non-technical management and professional occupations that are dominated by workers with a university education. We allocate to the broad non-managerial occupation category category L clerical employees, blue-collar workers, services workers, and technicians who do not work in ICT or R&D related occupations. Though our mnemonic for these workers is "unskilled", the category L includes a wide variety of occupations, some of which are highly skilled. Overall, relatively few of the jobs in this category require a university degree, with the exception of the fairly large category

[19] Firm-level R&D expenditures also include expenditures on R&D capital goods, which are a component of the firm's investment. Thus using R&D expenditures in the context of production function estimation raises the potential for double-counting of inputs.

[20] As noted above, Barth et al. (2017) find that 80 percent of U.S. private sector scientists and engineers worked outside R&D occupations in 2013.

[21] Saunders and Brynjolfsson (2016) find that for a sample of 127 large publicly traded US firms from 2003 to 2006, half of all spending on IT is for "Internal IT Services (e.g., custom software, design, maintenance, administration)". Including training services brings the share to 0.54.

[22] We refer to the IT and R&D consulting sectors as industry codes 62 (Computer Programming, consultancy and related activities), 631 (Data Processing, Hosting and related activities ; web portals), and 72 (Scientific R&D) in the NAF classification.

of "middle managers", most of whom probably have university degrees.

A key feature of our methodology is that firms are assumed to be able to choose their labor inputs to satisfy the static first order conditions for profit maximization after observing productivity.[23] Most French workers are on permanent labor contracts which make them very difficult and expensive to lay off, which at first glance makes it implausible that firms can choose employment to satisfy their static first order conditions. However, many French workers are on temporary labor contracts which make adjustment of labor input *at the margin* cheap and easy.[24] This is all that is required for our estimating equation (5) to be appropriate for French firms.

5.2 Other firm level data

Firm balance sheet information comes from the FARE dataset for the years 2009-2013.[25] The source of the information is firms' tax declarations. We use information on total revenues, material expenditures and the necessary series that we need to construct the capital stock at the level of the firm. Appendix A.1 describes the source data and explains how we construct firm-level capital stocks using this data.

Data on bilateral exports and imports of firms located in France are provided by French Customs. For each observation, we know the importing or exporting firm, trading partner country, the product traded, and the value of trade. We use the firm-level SIREN identifier to match the trade data to our two other data sources. This match is not perfect: we fail to match about 11 percent of imports and exports to firms. The reason for the imperfect match is that there are SIRENs in the trade data for which there is no corresponding SIREN in our other data sources. This is likely to lead to a particular type of measurement error: for some firms, we will observe zero trade even when true trade is positive.

In some of our specifications, we classify trade by country and/or product category. Countries are classified as High Income based on the 2011 World Bank classification.[26] In order to identify intermediate inputs, we use the Broad Economic Categories (BEC rev. 4) classification from the United Nations that classifies HS6 products into final, intermediate and capital goods. See Appendix A.1 for details.

[23] See equations (A.15) and (A.16) in the appendix.

[24] In our sample, the share of hours worked on temporary contracts is 3 percent for S, 11 percent for L, and 4 percent for T.

[25] Fichier Approché des Résultats Ésane (FARE)

[26] https://datahelpdesk.worldbank.org/knowledgebase/articles/906519-world-bank-country-and-lending-groups

6 Estimation results

In this section we start by reporting the results of production function estimation (equation 5), in Table 5. We then turn to the central results of our paper, estimation of endogenous productivity (equations 11 and 12). Tables 6, 7 and 8 show the effects of techies and trade on productivity. Using our regression results, Table 10 shows the implications for skill-bias and employment. Finally, we estimate the aggregate effect of techies and imports on relative demand for skill.

6.1 Production function estimates

Table 5 reports industry-by-industry estimates of the parameters identified by equations (5)-(8). While estimation of these parameters is central to our research strategy, most of the results reported in Table 5 are not that interesting in and of themselves, so we discuss this Table briefly.

The derivation of the estimating equation (5) requires $\sigma \neq 1$. All of our estimates of σ are greater than one, and in all but one industry (Textiles and apparel) we can reject the null that $\sigma = 1$ at conventional levels of statistical significance (t-statistics for this null are reported in the final column of the Table). The fact that $\widehat{\sigma} > 1$ implies that skill augmenting technological progress is skill-biased (see equation 9), an implication that we will return to below.

Economic logic requires that the estimated elasticity of demand η satisfies $\widehat{\eta} > 1$, which holds for all industries and is statistically statistically significant in all but one industry (Transport and storage). The estimated η's are plausible in magnitude, with a weighted average value of 6.3. For example, we find particularly large elasticities in Wholesale and Retail, which makes sense based on the nature of these industries.

For each industry, we find that $\widehat{\eta} > \widehat{\sigma}$. Equations (10) show that this is a necessary condition for skill-augmenting technological progress to raise employment of both skilled and unskilled workers. This is an important finding, since it implies that skill-biased technological change at the level of the firm is good news for less skilled workers despite the fact that they are easily substitutable ($\widehat{\sigma} > 1$) for skilled workers. The finding that $\widehat{\eta} > \widehat{\sigma}$ is also important when we calculate the employment effects of techies and trade in Table 10 below.

Our estimation procedure imposes that the distribution parameters $\alpha_L, \alpha_S, \alpha_M$ and α_K are strictly positive and sum to one, which helps explain the generally small standard errors on these parameters. The final row of the table reports the weighted average of the parameters (purely for descriptive purposes), and the results are reasonable given average shares in factor payments:

the materials and capital parameters are the largest while the managerial labor parameter is the smallest.

Our finding of $\sigma > 1$ contrasts with $\sigma < 1$ found by Doraszelski and Jaumandreu (2018) for Spanish manufacturing firms and Raval (2017) for U.S. manufacturing plants. Like us, these two papers assume a CES functional form with Hicks-neutral productivity differences across firms, where capital and materials are two of the inputs. The key difference from our paper is that Doraszelski and Jaumandreu (2018) and Raval (2017) combine all workers into a single labor aggregate and allow for labor-augmenting technological differences, while we divide labor into two categories and allow for skill-augmenting technological differences. Both Doraszelski and Jaumandreu (2018) and Raval (2017) acknowledge that their findings of labor-augmenting technological progress may be conflating skill-composition differences with labor augmenting technological differences across firms. Our finding of $\sigma > 1$ is consistent with the findings of Bøler (2015), who estimates a model much like ours on Norwegian manufacturing firms. Our result is also consistent with most of the industry- and macro-level labor literature on substitution between skilled and unskilled labor (see the results and discussion in Acemoglu and Autor (2011)).

6.2 Endogenous productivity

We begin by discussing the impact effects of techies and trade on productivity before reporting the steady state effects.

The first five rows of Table 6 report different specifications of the impact effect of techies on productivity. Whatever the specification, techies have a statistically significant effect on skill augmenting productivity ω_S, but no significant effect on Hicks-neutral productivity ω_H.[27].

The second and third rows of the table show the separate techie effects that come from higher-paid and more-educated technical managers and engineers (about 70 percent of the techie wage bill, as shown in Table 3) versus less senior technicians. The overall effect on ω_S is driven by engineers rather than technicians, with the point estimate for technicians actually small and negative (though not statistically significant). The next two rows of the Table report estimates when techies are broken down by their detailed job descriptions. ICT techies have twice as big an effect on ω_S as do R&D techies.[28]

[27] Similarly, Doraszelski and Jaumandreu (2018) find that firm level R&D expenditure has a larger effect on labor-augmenting than Hicks-neutral technological progress.

[28] What about the separate effects of ICT engineers, ICT technicians, R&D engineers and R&D technicians? When we estimated this specification the results were very imprecise, so to reduce table clutter we do not report them. We conclude that the data do not allow separately estimating these effects.

Turning to the effects of trade participation on productivity, we find that the impact effect of exporting is very small and (except in the final two specifications for ω_S) statistically insignificant. This does not imply that exporting firms are not more productive and/or skill intensive, as has been shown in countless studies. Rather, our estimates show that conditional on lagged productivity, exporting does not cause higher productivity. Thus our results for France contrast with the results of De Loecker (2013), who finds that exporting leads to productivity increases for Slovenian firms during the 1990s using a very similar specification. Slovenia in the 1990s was an emerging transition economy while France is a mature developed country, so our results are consistent with the consensus in the literature that learning-by-exporting is found only in (some) developing countries.

By contrast, we find statistically significant effects of importing on productivity: firms that import have about 2.5 percent higher Hicks-neutral and 6.6 percent higher skill-augmenting productivity. Our result for ω_H is consistent with studies of developing and transition economies that find that access to foreign intermediate inputs leads to productivity increases, but we are the first to find such a result for a developed country.[29] Our results for ω_S are more novel. When combined with our Table 5 results that $\sigma > 1$, our estimated effect of importing on ω_S offers direct support for the hypothesis that offshoring leads to skill upgrading within firms.[30]

Table 7 reports additional results about the effect of importing on productivity. In columns (1) and (4), we add an indicator variable for imports of intermediate goods. Introducing this channel shows that the importing effect on ω_S is primarily through imported inputs. Columns (2) and (5) of Table 7 consider an alternative split, with separate indicators for importing from high income countries and all other countries.[31] Although not very precisely estimated, this split implies that most of the effect of importing on ω_H comes from high income countries, which is consistent with higher quality imports sourced from such countries. In contrast, the effect of importing on ω_S is somewhat larger when coming from lower income countries, though the difference is not statistically significant. Finally, in columns (3) and (6) we report the interactions of the indicator variables for income class and intermediate imports. The effect of importing inputs from both high income and other countries is the sum of the two coefficients, and these linear combinations are reported in the bottom panel of the Table. We find that importing inputs from both high and lower income

[29] See Halpern et al. (2015) for Hungary, Amiti and Konings (2007) for Indonesia, Kasahara and Rodrigue (2008) for Chile, and Topalova and Khandelwal (2011) for India.

[30] Other studies that have found evidence for within firm/plant skill upgrading due to importing include Kasahara et al. (2016) for Indonesia and Martins and Opromolla (2011) for Portugal. Becker et al. (2013) do not study firm-level importing, but they do find that increases in foreign affiliate employment raises skill intensity at home for German multinationals.

[31] Countries are defined as high income on the basis of the 2002 World Bank classification.

countries raises ω_S, with a small and statistically insignificant difference across country groups. Overall, Table 7 shows that offshoring (that is, imports of intermediate inputs) increases within-firm skill intensity, as suggested by Feenstra and Hanson (1996), Grossman and Rossi-Hansberg (2008) (in very different models) and many others. Our baseline estimate of 0.075 in column 4 of Table 7 means that the impact effect of offshoring is to raise ω_S by 7.5 percent compared to firms who source only from France. The effect on ω_H is one-third as large.

The final rows of Table 6 report the effects of other controls. Productivity is very persistent, with a coefficient on lagged productivity of nearly 0.9 for both ω_H and ω_S. By contrast, lagged ω_H has virtually no effect on ω_S, and vice versa. Firm size has a small negative effect on ω_H , which is consistent with the well-known result that larger firms have slower productivity growth. By contrast, firm size has a positive effect on ω_S. Firm age has no independent effect on productivity.[32]

Because of the persistence of productivity, the techie and trade effects reported in Table 6 understate the long-run impact of these variables on cross-sectional productivity differences. Table 8 reports the associated steady-state effects, using equation (13). In addition to reporting the long-run effects and their standard errors, Table 8 also reports scaled long-run techie effects in italics. These are computed by multiplying the long-run coefficients times the 75^{th} percentile of the corresponding variable, as reported in Table 9. The number *0.48* in the third row of column (4) means that compared to firms with no techies, firms with a lot of techies have ω_S which is 0.48 higher, or 62 percent. This is a very big effect in economic terms. It is also large relative to the variation in ω_S, as 0.48 is equal to about a quarter of the $75^{th} - 25^{th}$ percentile range of ω_S, which is 1.95 (second row, last column of Table 9). This scaled effect is entirely due to engineers, and is split roughly 50–50 between the effect of ICT and R&D techies.

The long-run trade effects of importing are also very big, both in economic terms and relative to the variation in productivity. Since the trade variables are indicators, they are simple to interpret: the number 0.78 at the bottom of column 4 of Table 8 means that compared to firms that don't import, importing firms have ω_S which is 0.78 higher, or 118 percent. The effect is smaller but still big for ω_H: looking at the bottom of columns 1 to 3, firms that import have ω_H which is higher by 0.42 (52 percent) compared to non-importers.

[32] As a robustness check, in unreported results we re-estimated equations (11) and (12) with two lags of all variables. The sum of the coefficients is not appreciably different than the corresponding baseline coefficients with only one lag. This implies that the long-run effects with one or two lags are essentially the same.

6.3 The skill bias of techies and trade

Since the estimated effects of techies on ω_H are statistically insignificant in Tables 6 through 8, here we focus on quantifying the employment effects of techies and trade through their effects on ω_S. Applying equations (10), our quantification uses the industry-level estimates of η and σ from Table 5 together with the long-run effects reported in Table 8. To arrive at an economy-wide number, we compute the elasticities defined in equations (10) for each industry, setting the cost share of skilled labor λ_{Sft} equal to the CES distribution parameter α_S.[33] We then construct an employment-weighted average of the industry elasticities, which we report in the first line of Table 10. In Panel A of Table 10 we multiply the elasticities times the estimated effects of techies and trade from Table 8. Finally, to give a sense of the magnitude of the techie effects, in Panel B we multiply the Panel A estimates by the 75^{th} percentile of the employment-weighted distribution of techies. Thus, the numbers in Panel B of Table 10 answer the question: how does employment differ between a firm with no techies and a firm with a lot of techies?

The first line of Table 10 shows that skill augmenting productivity raises S, L and S/L, and the elasticities are big: a one percent increase in Ω_S raises skilled employment by almost one percent, unskilled employment by 0.3 percent, and skill intensity by 0.7 percent. Techies are an important driver of these effects: as shown in the first row of Panel B, high techie firms have employment of S that is 0.47 log points higher than firms with no techies, employment of L that is 0.14 log points higher, and a skill intensity that is 0.33 log points higher. This effect is driven entirely by engineers, and is about equally driven by ICT and R&D techies.

The last line of Panel A shows that importing is also strongly pro-employment and skill biased: firms that import have employment of S that is 0.76 log points higher than firms that do not import, employment of L that is 0.22 log points higher, and a skill intensity that is 0.54 log points higher.

Table 10 is one of the bottom lines of our paper. For the first time in the literature, we have estimated the effects of technology and globalization in a unified firm-level framework, which allows us to compare their importance. Table 10 shows that techies and importing raise both skill intensity and employment, and the effects are big. These are firm-level effects, calculated holding market demand and factor prices fixed. In a general equilibrium full employment model, the partial equilibrium effects found here would have clear implications for relative wages: techies and importing raise the skill premium.

[33] This holds identically at the geometric mean of the sample as a consequence of our data normalization.

6.4 Aggregating the effects of techies and trade

Our final empirical exercise asks: how much of the change in the aggregate skill intensity in our sample period can be explained by firms' choices on techies and trade? To do this we proceed in three steps.

First, we construct predicted changes in productivity shifters across firms. Starting from estimated Hicks-neutral and skill augmenting productivity levels (in logs) in 2009, we predict productivity levels in 2013 using actual techie and importing decisions from 2009 through 2012 by iterating forward using the estimated parameters of equations (11) and (12). The log differences between the predicted productivity in 2013 and the actual values in 2009 are due to techies and trade choices made by firms between 2009 and 2012. In the second step we use the predicted changes in productivity from the first step to calculate the predicted change in employment of S and L for each firm between 2009 and 2013 using equation (10). Finally, we sum over all firms to get predicted aggregate skill intensity. These calculations take into account both within-firm adjustment as well as changes in firm sizes, but exclude firm entry and exit (85% of employment is accounted by firms who exist throughout our sample). Details of how we perform these calculations are given in Appendix A.2.5.

We measure aggregate skill intensity by the employment share of S $(= S/(L+S))$. This measure increases in our sample by 0.0186, from 0.141 in 2009 to 0.160 in 2013 (the numbers for the entire French private sector are similar). Our aggregation of firm level techie employment decisions in 2009–2013 implies an increase in relative demand for skill amounting to 0.011. The same exercise using importing decisions amounts to a somewhat larger increase in relative demand for skill of 0.016. Adding these two effects together gives an increase of 0.027, which is more than the actual increase of 0.019. This is not surprising, since the calculations do not take into account equilibrium conditions in the markets for goods and of labor. Nonetheless, these calculations illustrate the quantitative importance of the firm-level techie and offshoring channels that we estimate.

7 Conclusion

In this paper we have shown, for the first time in the literature, that technology adoption and importing are both skill biased at the firm level. The focus on firms is important, because firms are where decisions about technological change, globalization and employment are made. Because we identify these effects in a common framework it is simple to compare their magnitudes, and we find that they are big and of the same order of magnitude.

Our findings do not challenge the common wisdom that R&D is the key to expansion of the set of available technologies. However, in order for these technologies to diffuse in the economy, firms need to adopt them. We identify the adoption of ICT through the employment of ICT techies. Under this interpretation of our results, it is not surprising that R&D techies increase firm productivity less than ICT techies on the margin. This interpretation also offers an important policy implication, which is that investments in technology adoption may be at least as important as investment in basic research. This conclusion has implications for a whole range of public policies related to education, research, and investment.

We reach these conclusions using administrative data that cover the whole French private sector from 2009 to 2013. Our methodology applies new techniques from the structural production function literature to consistently estimate both Hicks-neutral and skill-augmenting technology. This methodology also allows us to identify the causal effects on productivity and employment of firms' decisions on trade and employment of techies. This identification is well-founded even though we remain silent on what drives firm decisions on techies and trade. Understanding why some, but not all, firms employ techies and engage in offshoring is an important research question that is beyond the scope of this paper.

Our bottom line is that both employment of techies and offshoring lead to greater employment of both skilled and less skilled workers. The result for unskilled workers is surprising but easy to explain: the direct substitution effect away from unskilled labor is outweighed by the powerful employment-enhancing effect of greater productivity. These are firm-level conclusions, calculated holding market demand and economy-wide aggregates constant. Analyzing market and general equilibrium effects that are consistent with our firm-level findings is an important task for future research.

8 References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica 67*(2), 251–333.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter and D. Card (Eds.), *Hanbook of Labor Economics*, Volume 4b. North Holland.

Acemoglu, D. and A. Finkelstein (2008). Input and technology choices in regulated industries: Evidence from the health care sector. *Journal of Political Economy 116*(5), 837–880.

Ackerberg, D., C. L. Benkard, S. Berry, and A. Pakes (2007). *Econometric tools for analyzing market outcomes*, Volume 6, Chapter James J. Heckman and Edward E. Leamer, pp. 4171–4276. Elsevier.

Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica 83*(6), 2411–2451.

Akerman, A., I. Gaarder, and M. Mogstad (2015). The skill complementarity of broadband internet. *The Quarterly Journal of Economics 130*(4), 1781–1824.

Amiti, M. and L. Cameron (2012). Trade liberalization and the wage skill premium: Evidence from Indonesia. *Journal of International Economics 87*, 277–287.

Amiti, M. and J. Konings (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from Indonesia. *The American Economic Review*, 1611–1638.

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.

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, June). The effects of scientists and engineers on productivity and earnings at the establishment where they work. Working Paper 23484, National Bureau of Economic Research.

Becker, S. O., K. Ekholm, and M.-A. Muendler (2013). Offshoring and the onshore composition of tasks and skills. *Journal of International Economics 90*(1), 91–106.

Ben Zeev, N., M. Joel, and K. Van Der Beek (2017). Flexible supply of apprenticeship in the british industrial revolution. *The Journal of Economic History 77*(1), 208–250.

Bender, S., N. Bloom, D. Card, J. Van Reenen, and S. Wolter (2018). Management practices, workforce selection, and productivity. *Journal of Labor Economics 36*(S1), S371–S409.

Bøler, E. A. (2015, November). Technology-skill complementarity in a globalized world. *University of Oslo mimeo*.

Bonleu, A., G. Cette, and G. Horny (2013). Capital utilization and retirement. *Applied Economics 45*(24), 3483–3494.

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.

Bustos, P. (2011). The impact of trade on technology and skill upgrading: Evidence from Argentina. Working Papers, Universitat Pompeu Fabra. Departamento de Economía y Empresa.

Cette, G., N. Dromel, R. Lecat, and A.-C. Paret (2015). Production factor returns: the role of factor utilization. *Review of Economics and Statistics 97*(1), 134–143.

De Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics 5*(3).

De Loecker, J. and P. K. Goldberg (2014). Firm performance in a global market. *Annual Review of Economics 6*, 201–27.

Doraszelski, U. and J. Jaumandreu (2013). R&D and productivity: Estimating endogenous productivity. *The Review of Economic Studies 80*(4), 1338–1383.

Doraszelski, U. and J. Jaumandreu (2018). Measuring the bias of technological change. *Journal of Political Economy 126*(3), 1027–1084.

Dunne, T., LuciaÂ Foster, JohnÂ Haltiwanger, and KennethÂ Troske (2004). Wage and productivity dispersion in United States manufacturing: The role of computer investment. *Journal of Labor Economics 22*(2), 397–429.

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.

Foster, L., JohnÂ Haltiwanger, and S. Chad (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review 98*(1), 394–425.

Grieco, P. L., S. Li, and H. Zhang (2016). Production function estimation with unobserved input price dispersion. *International Economic Review 57*(2), 665–690.

Grieco, P. L., S. Li, and H. Zhang (2017). Imported inputs, productivity, and the impact of import liberalization: Evidence from chinese paint manufacturers.

Grossman, G. M. and E. Rossi-Hansberg (2008). Trading tasks: A simple theory of offshoring. *The American Economic Review 98*(5), 1978–97.

Halpern, L., M. Koren, and A. Szeidl (2015). Imported inputs and productivity. *The American Economic Review 105*(12), 3660–3703.

Harrigan, J., A. Reshef, and F. Toubal (2016). The March of the Techies: Technology, trade, and job polarization in France, 1994-2007. Working Paper 22110, National Bureau of Economic Research.

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.

Helpman, E. (2018). *Globalization and Inequality*. Cambridge: Harvard University Press.

Kasahara, H., Y. Liang, and J. Rodrigue (2016). Does importing intermediates increase the demand for skilled workers? plant-level evidence from Indonesia. *Journal of International Economics 102*, 242 – 261.

Kasahara, H. and J. Rodrigue (2008). Does the use of imported intermediates increase productivity? plant-level evidence. *Journal of development economics 87*(1), 106–118.

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.

Levinsohn, J. and A. Petrin (2003). Estimating production functions using inputs to control for unobservables. *The Review of Economic Studies 70*(2), 317–341.

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.

Maloney, W. F. and F. Valencia-Caicedo (2017). Engineering Growth: Innovative Capacity and Development in the Americas. Working Paper 6339, CESIfo Working Paper.

Martins, P. S. and L. D. Opromolla (2011). Why ex (im) porters pay more: Evidence from matched firm-worker panels.

Melitz, M. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica 71*, 1695–1725.

Melitz, M. J. and G. I. Ottaviano (2008). Market size, trade, and productivity. *The Review of Economic Studies 75*(1), 295–316.

Oberfield, E. and D. Raval (2014, August). Micro data and macro technology. Technical report.

Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 1263–1297.

Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants. *The Review of Economic Studies 69*(1), 245–276.

Raval, D. (2017, February). The micro elasticity of substitution and non-neutral technology. *RAND Journal of Economics*, forthcoming.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy 98*(5), S71–S102.

Saunders, A. and E. Brynjolfsson (2016). Valuing information technology related intangible assets. *MIS Quarterly 40*(1).

Solon, G., S. J. Haider, and J. M. Wooldridge (2015). What are we weighting for? *Journal of Human resources 50*(2), 301–316.

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.

Topalova, P. and A. Khandelwal (2011). Trade liberalization and firm productivity: The case of India. *Review of economics and statistics 93*(3), 995–1009.

Table 1: Covariates of techies

	(1) $I\,(techies > 0)$	(2) Techie wage bill share for techies > 0
Log revenue	0.131***	0.003***
	(0.003)	(0.001)
Exporter dummy	0.548***	0.004**
	(0.013)	(0.002)
Importer dummy	0.591***	-0.002
	(0.012)	(0.002)
Manager wage bill share	0.525***	0.134***
	(0.049)	(0.013)
Industry × year	Yes	Yes
Obs.	612,769	69,175

Notes to Table 1: Each entry is the weighted least squares coefficient in a regression of the column dependent variable on the row explanatory variable. All regressions include industry × year fixed effects. Techie measures are: column (1) indicator equal to 1 if the firm employs any techies, column (2) the share of techies in the firm's wage bill for firms that employ techies. Regressions weighted by firm employment. Standard errors clustered by firm in parenthesis. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table 2: Techie occupations

Technical managers & engineers (*Ingénieurs et cadres techniques d'entreprise*)		
383a	R&D	Engineers and R&D managers, electricity and electronics
384a	R&D	Mechanical engineers and R&D managers
385a	R&D	Materials and chemical engineers and R&D managers
386a	R&D	Engineers and R&D managers, intermediate goods
388a	ICT	Information technology R&D engineers and managers
388b	ICT	Information technology support engineers and managers
388c	ICT	Information technology project managers
388e	ICT	Telecommunications engineers and specialists
Technicians (*Techniciens*)		
473b	R&D	R&D technicians, electrical and electronic equipment
474b	R&D	R&D technicians, mechanical and metalworking equipment
475a	R&D	R&D technicians, processing industries
478a	ICT	R&D technicians, information technology
478b	ICT	Computer production and operation technicians
478c	ICT	Computer installation and maintenance technicians
478d	ICT	Telecommunications and computer network technicians

Notes to Table 2: First column is the occupational code of the PCS classification, and the third column is our translation of the official descriptions. The second column is our categorization based on the descriptions.

Table 3: Wage bill shares within Techies (%)

	Whole economy			
	ICT	R&D	Engineers	Technicians
2009	43	57	71	29
2013	41	59	73	27
	Variation across industries in 2009			
std. dev.	30.1	30.1	10.0	10.0
Min	17.5	4.8	49.9	17.8
Max	95.2	82.5	82.2	50.1

Notes to Table 3: Wage bill shares sum to 100 across the two categories ICT, R&D, and also across the two categories Engineers and Technicians.

Table 4: Estimation sample, 2009-2013

Industry	Obs.	Obs. (%)	Firms	Firms (%)	Revenue (%)	Hours (%)
Food, beverage, tobacco	26,407	3.1	6,516	3.0	7.08	5.48
Textiles, wearing apparel	9,732	1.1	2,337	1.1	0.94	1.39
Wood, paper products	19,892	2.3	4,808	2.2	1.72	2.42
Chemical products	5,748	0.7	1,292	0.6	3.68	2.29
Rubber and plastic	15,960	1.9	3,702	1.7	3.04	4.06
Basic metal and fabricated metal	34,074	4.0	7,960	3.7	3.96	5.29
Computer, electronic	5,462	0.6	1,267	0.6	1.64	2.09
Electrical equipment	4,597	0.5	1,063	0.5	1.54	1.79
Machinery and equipment	11,475	1.3	2,605	1.2	2.33	2.72
Other manufacturing	27,648	3.2	6,722	3.1	1.98	3.47
Construction	159,641	18.6	41,175	19.2	7.18	12.14
Wholesale	176,180	20.5	42,846	20.0	37.6	17.6
Retail	210,097	24.5	52,078	24.3	17.1	15.3
Transportation and storage	27,683	3.2	6,899	3.2	5.94	8.76
Accommodation and food services	95,571	11.1	25,817	12.0	1.93	6.29
Administrative and support activities	28,674	3.3	7,557	3.5	2.42	6.98
Total	*858,841*	*100*	*214,644*	*100*	*100*	*100*

Notes to Table 4: We lose 10.2% of total revenue and 7.5% of total hours due to dropping the sectors of coke and refined petroleum, pharmaceutical products, transport equipment, and publishing and broadcasting.

Table 5: Production function estimates

Industry	α_L	α_S	α_M	α_K	σ	η	Obs.	Firms	t, $H_0 : \sigma = 1$
Food, beverage, tobacco	0.190***	0.051***	0.574***	0.185***	2.856***	5.085***	26,407	6,516	4.78
	(0.003)	(0.001)	(0.010)	(0.014)	(0.388)	(0.366)			
Textiles, apparel	0.295***	0.109***	0.549***	0.047**	2.240***	2.947***	9,732	2,337	1.59
	(0.006)	(0.002)	(0.011)	(0.019)	(0.782)	(0.136)			
Wood, paper products	0.285***	0.099***	0.433***	0.182***	1.156***	3.481***	19,892	4,808	1.79
	(0.006)	(0.002)	(0.010)	(0.018)	(0.087)	(0.169)			
Chemical products	0.137***	0.070***	0.553***	0.239***	1.316***	4.729***	5,748	1,292	2.34
	(0.004)	(0.002)	(0.014)	(0.020)	(0.135)	(0.477)			
Rubber & plastic	0.230***	0.072***	0.567***	0.131***	2.919***	3.443***	15,960	3,702	2.95
	(0.005)	(0.002)	(0.013)	(0.019)	(0.651)	(0.187)			
Basic & fabricated metal	0.249***	0.073***	0.307***	0.371***	1.484***	5.866***	34,074	7,960	4.52
	(0.018)	(0.005)	(0.022)	(0.045)	(0.107)	(1.283)			
Computer, electronic	0.186***	0.124***	0.463***	0.228**	1.536***	2.985***	5,462	1,267	2.22
	(0.021)	(0.014)	(0.053)	(0.089)	(0.242)	(0.539)			
Electrical equipment	0.184***	0.079***	0.533***	0.205***	1.382***	5.068***	4,597	1,063	2.85
	(0.008)	(0.003)	(0.022)	(0.033)	(0.134)	(0.937)			
Machinery & equipment	0.190***	0.085***	0.485***	0.241***	1.404***	4.852***	11,475	2,605	2.35
	(0.009)	(0.004)	(0.023)	(0.035)	(0.172)	(0.738)			
Other manufacturing	0.253***	0.101***	0.347***	0.299***	1.334***	4.076***	27,648	6,722	4.70
	(0.009)	(0.004)	(0.012)	(0.024)	(0.071)	(0.378)			
Construction	0.280***	0.098***	0.377***	0.245***	1.238***	2.870***	159,641	41,175	4.26
	(0.011)	(0.004)	(0.015)	(0.030)	(0.056)	(0.234)			
Wholesale	0.092***	0.054***	0.764***	0.091***	1.475***	7.966***	176,180	42,846	2.27
	(0.001)	(0.000)	(0.005)	(0.006)	(0.209)	(0.432)			
Retail	0.087***	0.042***	0.758***	0.114***	1.774***	11.991***	210,097	52,078	3.97
	(0.001)	(0.000)	(0.006)	(0.007)	(0.195)	(1.705)			
Transport & storage	0.303***	0.051***	0.081***	0.565***	1.362***	4.225	27,683	6,899	4.95
	(0.082)	(0.014)	(0.022)	(0.117)	(0.073)	(3.032)			
Accommodation and food	0.299***	0.085***	0.273***	0.342***	1.668***	6.729***	95,571	25,817	5.02
	(0.020)	(0.006)	(0.019)	(0.045)	(0.133)	(2.252)			
Admin & support	0.427***	0.133***	0.087***	0.352***	2.515***	5.606***	28,674	7,557	8.51
	(0.020)	(0.006)	(0.004)	(0.030)	(0.178)	(0.983)			
weighted average	0.209	0.073	0.485	0.233	1.690	6.262			

Notes to Table 5: Each row reports weighted nonlinear least squares estimates of equation (5) for a given industry. Reported parameters are functions of the parameters of (5). Regressions weighted by firm employment. Bootstrapped standard errors clustered by firm in parentheses. The final column reports t-statistics for the null hypothesis that $\sigma = 1$. The final row reports hours-weighted averages of the industry level parameters. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table 6: Techie and trade effects on productivity

	(1)	(2)	(3)	(4)	(5)	(6)
	Hicks neutral ω_{ft}^{H}			Skill augmenting ω_{ft}^{S}		
Techies	0.042			0.700***		
	(0.115)			(0.215)		
Techies : engineers		-0.008			1.015***	
		(0.152)			(0.300)	
Techies : technicians		0.188			-0.233	
		(0.166)			(0.249)	
Techies : ICT			0.086			0.908***
			(0.163)			(0.276)
Techies : R&D			-0.012			0.450*
			(0.173)			(0.252)
Exporting	-0.011	-0.011	-0.011	0.033	0.034*	0.033*
	(0.012)	(0.012)	(0.013)	(0.020)	(0.021)	(0.019)
Importing	0.026**	0.025**	0.026***	0.066***	0.068***	0.066***
	(0.011)	(0.011)	(0.009)	(0.016)	(0.017)	(0.017)
lagged ω_{ft}^{H}	0.904***	0.904***	0.905***	0.025***	0.025***	0.025***
	(0.010)	(0.010)	(0.011)	(0.007)	(0.007)	(0.007)
lagged ω_{ft}^{S}	0.018***	0.018***	0.018***	0.901***	0.901***	0.901***
	(0.004)	(0.004)	(0.004)	(0.007)	(0.007)	(0.009)
firm size	-0.031***	-0.030***	-0.030***	0.027***	0.026***	0.028***
	(0.005)	(0.005)	(0.005)	(0.007)	(0.007)	(0.007)
firm age	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

Notes to Table 6: Weighted least squares estimation of equations (11) and (12), pooled across industries, with industry × year fixed effects. Observations weighted by firm employment. Bootstrapped standard errors clustered by firm in parentheses. Sample: 612,769 observations on 205,472 firms during 2009-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table 7: Techie and trade effects on productivity, importing detail

	(1)	(2)	(3)	(4)	(5)	(6)
	Hicks neutral ω_{ft}^{H}			Skill augmenting ω_{ft}^{S}		
Techies	0.042	0.040	0.042	0.700***	0.684***	0.680***
	(0.115)	(0.115)	(0.114)	(0.214)	(0.211)	(0.213)
Exporting	-0.010	-0.007	-0.006	0.030	0.024	0.022
	(0.012)	(0.011)	(0.011)	(0.020)	(0.019)	(0.020)
Importing	0.034			0.023		
	(0.022)			(0.020)		
Importing inputs	-0.009			0.052***		
	(0.026)			(0.020)		
Importing high income		0.017	0.020		0.040***	0.017
		(0.011)	(0.026)		(0.014)	(0.022)
Importing other income		-0.001	0.011		0.054***	0.032*
		(0.014)	(0.020)		(0.014)	(0.018)
Importing inputs high income			-0.005			0.029
			(0.025)			(0.022)
Importing inputs other income			-0.016			0.028
			(0.020)			(0.021)
Linear combinations of estimates						
Importing + inputs	0.024*			0.075***		
	(0.013)			(0.006)		
Imp. high income + inputs			0.015			0.046***
			(0.01)			(0.009)
Imp. other inc. + inputs			-0.004			0.059***
			(0.012)			(0.019)
Other controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes to Table 7: Weighted least squares estimation of equations (11) and (12), pooled across industries, with industry $\times$ year fixed effects. Estimates of other controls (lagged productivity, firm size and age) omitted. Effects reported in bottom panel are linear combinations of WLS estimates. Regressions weighted by firm employment. Bootstrapped standard errors clustered by firm in parentheses. Sample: 612,769 observations on 205,472 firms during 2009-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table 8: Long run techie and trade effects on productivity

	(1)	(2)	(3)	(4)	(5)	(6)
	Hicks neutral ω_{ft}^{H}			Skill augmenting ω_{ft}^{S}		
Techies	1.87			7.57***		
	(1.41)			(2.23)		
	0.12			*0.48*		
Techies : engineers		1.96			10.7***	
		(1.91)			(3.21)	
		0.09			*0.49*	
Techies : technicians		1.59			-1.95	
		(1.93)			(2.61)	
		0.04			*-0.05*	
Techies : ICT			2.78			9.90***
			(2.04)			(2.95)
			0.07			*0.25*
Techies : R&D			0.78			4.76*
			(2.06)			(2.64)
			0.05			*0.28*
Exporting	-0.06	-0.05	-0.06	0.32	0.33	0.32)
	(0.15)	(0.14)	(0.15)	(0.22)	(0.22)	(0.21)
Importing	0.42***	0.42***	0.42***	0.78***	0.79***	0.77***
	(0.14)	(0.14)	(0.12)	(0.19)	(0.19)	(0.20)

Notes to Table 8: Effects are long-run steady state effects on productivity given by equation (13), based on results in Table 6. Bootstrapped standard errors clustered by firm in parentheses. Scaled techie coefficients, defined as the estimate times the 75th percentile of the techie distribution reported in Table 9, are reported in *italics*. Sample: 612,769 observations on 205,472 firms during 2009-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table 9: Summary statistics for second stage regressions

	Mean	Std Dev	*p25*	*p50*	*p75*	*p75-p25*
Hicks neutral productivity ω_H	0	1.80	-0.97	0.03	1.10	2.07
Skill augmenting productivity ω_S	0	2.25	-1.00	-0.05	0.95	1.95
Techie share of wage bill	0.054	0.086	0.008	0.024	0.064	0.055
Techies: engineers	0.043	0.076	0.006	0.017	0.045	0.039
Techies: technician	0.019	0.032	0.002	0.009	0.025	0.022
Techies: ICT	0.026	0.055	0.005	0.013	0.025	0.020
Techies: R&D	0.046	0.069	0.003	0.021	0.058	0.055
Exporting	0.59	0.49	0	1	1	1
Importing	0.65	0.48	0	1	1	1

Notes to Table 9: Summary statistics for variables used to estimate equations (11) and (12), weighted by firm employment. Statistics for the five Techie variables are for firms with positive employment of each techie variable separately. *p25* is the value of the variable at the 25th percentile of it's distribution, etc. Sample: 612,769 observations on 205,472 firms during 2009-2013.

A Appendix

A.1 data definitions and construction

Here we discuss in detail the three administrative datasets used in our paper, as well as details on supplementary publicly available data.

A key feature of the French statistical system is that establishments are identified by a unique number, the SIRET, which is used by all data sources. The first 9 digits of an establishment's SIRET comprise the SIREN of the firm to which the establishment belongs. This makes it easy to aggregate from establishments to firms.

A.1.1 Workers: DADS Poste

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.[34] The DADS Poste is an INSEE database compiled from the mandatory firm-level DADS ("Déclaration Annuelle de Données Sociales") reports. Our unit of analysis is annual hours paid in a firm, by occupation. The 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. For each worker, the DADS Poste reports gross and net wages, hours paid, occupation, tenure, gender and age. There is no information about workers' education or overall labor market experience. The data do not include worker identifiers, so we can not track workers over time, but this is of no concern to us given our focus on firm-level rather than individual outcomes.[35]

A.1.2 Balance sheet data: FARE

Firm-level balance sheet information is reported in an INSEE dataset called FARE.[36] The balance sheet variables used in our empirical analysis include revenue, expenditure on materials, and the book value of capital. We do not use balance sheet data on employment or the wage bill, because

[34] All employers and their employees are covered by the DADS declaration with the exception of self-employed and government bodies, domestic services (section 97-98 of NAF rev. 2) and employees in businesses outside French territory (section 99 of NAF rev. 2). However, local authorities and public-employed hospital staff are included since 1992. Public institutions of industrial and commercial nature are also included.

[35] A related dataset, made famous by Abowd et al. (1999), is the DADS Panel. This sample from of the DADS data does include worker identifiers.

[36] FICUS (*Fichier complet unifié de SUSE*) reports balance sheet data through 2007, while FARE (*Fichier approché des résultats Ésane*) starts in 2008. The underlying data sources are identical.

Table 10: Employment effects of skill augmenting technology differences

	S	L	S/L
Elasticities	0.972	0.282	0.690
A. Elasticities × second stage estimates			
Techies	7.36***	2.135***	5.224***
Techies : engineers	10.40***	3.018***	7.385***
Techies : technicians	-1.896	-0.550	-1.346
Techies : ICT	9.625***	2.793***	6.833***
Techies : R&D	4.628*	1.343*	3.285*
Exporting	0.311	0.09	0.221
Importing	0.758***	0.220***	0.538***
B. Elasticities × second stage estimates $\times p75$			
Techies	0.468***	0.136***	0.333***
Techies : engineers	0.473***	0.137***	0.336***
Techies : technicians	-0.047	-0.014	-0.033
Techies : ICT	0.240***	0.070***	0.170***
Techies : R&D	0.269*	0.078*	0.191*

Notes to Table 10: This table reports estimated steady-state effects of skill augmenting productivity differences on employment of managers S, other workers L, and their ratio. The first row reports employment-weighted averages of industry level elasticities, computed using equation (10) and the estimates from Table 5. Panel A multiplies the elasticities by the corresponding estimates from Table 8. Panel B multiplies the panel A numbers by the 75th percentile of the distribution of techies, from Table 9. Asterisks denote the statistical significance of the corresponding estimates from Table 8.

Table A1: Occupational codes

PCS code	official description	unofficial English translation
21	Artisans	Small business owners and workers
22	Commerçants et assimilés	Shopkeepers
23	Chefs d'entreprise de 10 salariés ou plus	Heads of businesses
34	Professeurs, professions scientifiques	Scientific and educational professionals
35	Professions de l'information, des arts et des spectacles	Creative professionals
37	Cadres administratifs et commerciaux d'entreprise	Top managers and professionals
38	Ingénieurs et cadres techniques d'entreprise	Engineers and Technical managers
42	Professeurs des écoles, instituteurs et assimilés	Teachers
43	Professions intermédiaires de la santé et du travail social	Mid-level health professionals
46	Professions intermédiaires administratives et commerciales des entreprises	Mid-level managers and professionals
47	Techniciens	Technicians
48	Contremaîtres, agents de maîtrise	Supervisors and foremen
53	Policiers et militaires	Security workers
54	Employés administratifs d'entreprise	Office workers
55	Employés de commerce	Retail workers
56	Personnels des services directs aux particuliers	Personal service workers
62	Ouvriers qualifiés de type industriel	Skilled industrial workers
63	Ouvriers qualifiés de type artisanal	Skilled manual laborers
64	Chauffeurs	Drivers
65	Ouvriers qualifiés de la manutention, du magasinage et du transport	Skilled transport and wholesale workers
67	Ouvriers non qualifiés de type industriel	Unskilled industrial workers
68	Ouvriers non qualifiés de type artisanal	Unskilled manual laborers

Notes to Table A1: The PCS (*Professions et Catégories Socioprofessionnelles*) system of occupational codes is used to classify all workers in France. PCS codes are assigned by employers. This Table omits public sector and agricultural occupations that are not in our data.

the DADS Poste data is more detailed, but the FARE wage bill and employment data are extremely highly correlated with the corresponding DADS Poste data.

To construct capital stocks, we begin with the book value of capital recorded in FARE. We follow the methodology proposed by Bonleu et al. (2013) and Cette et al. (2015). Since the stocks are recorded at historical cost, i.e. at their value at the time of entry into the firm i's balance sheet, an adjustment, has to be made to move from stocks valued at historic cost ($K^{BV}_{i,s,t}$) to stocks valued at current prices ($K_{i,s,t}$). We deflate K^{BV} by a price by assuming that the sectoral price of capital is equal to the sectoral price of investment T years before the date when the first book value was available, where T is the corrected average age of capital, hence $p^K_{s,t+1} = p^I_{s,t-T}$. The average age of capital is computed using the share of depreciated capital, $DK^{BV}_{i,s,t}$ in the capital stock at historical cost.

$$T = \frac{DK^{BV}_{i,s,t}}{K^{BV}_{i,s,t}} \times \widetilde{A}$$

where

$$\widetilde{A} = median_{i \in S} \left(\frac{K^{BV}_{i,s,t}}{\Delta DK^{BV}_{i,s,t}} \right)$$

with S the set of firms in a sector. We use the median value $\widetilde{A}$ to reduce the volatility in the data, as investments within firms are discrete events.

A.1.3 Trade data: Douanes

Our source for firm-level trade data is the French Customs (*Douanes*). For each trade observation, we know the importing or exporting firm, trading partner country, the product traded, and the value of trade. We use the firm-level SIREN identifier to match the trade data to our two other data sources. This match is not perfect: we fail to match about 11 percent of imports and exports to firms. The reason for the imperfect match is that there are SIRENs in the trade data for which there is no corresponding SIREN in our other data sources. This is likely to lead to a particular type of measurement error: for some firms, we will observe zero trade even when true trade is positive.

In some of our specifications we classify trade by exporter per capita income and/or whether they are imports of intermediate goods. Countries are classified as High Income based on the 2011 World Bank classification[37]. We use the Broad Economic Categories (BEC rev. 4) classification

[37] https://datahelpdesk.worldbank.org/knowledgebase/articles/906519-world-bank-country-and-lending-groups

from the United Nations that classifies HS6 products into final, intermediate and capital goods.

A.2 Methodology

This section gives details on our two-step estimation procedure and related calculations.

A.2.1 Deriving the estimating equation

Using equations (1) and (2), revenue is given by

$$R_{ft} = e^{u_{ft}} P_{ft} Y_{ft} = e^{u_{ft} + \frac{\eta+1}{\eta} \omega_{Hft}} A_t \left[\alpha_L L_{ft}^{\gamma} + \alpha_S \left(\Omega_{Sft} S_{ft} \right)^{\gamma} + \alpha_K K_{ft}^{\gamma} + \alpha_M M_{ft}^{\gamma} \right]^{\frac{\eta+1}{\eta\gamma}}, \tag{A.14}$$

Equation (A.14) contains three unobservable shocks (u_{ft}, ω_{Hft} and ω_{Sft}) and one unobservable variable M_{ft}. The first-order conditions for profit maximization are

$$\alpha_L L_{ft}^{-1/\sigma} X_{ft} = W_{Lft} \tag{A.15}$$

$$\alpha_S \Omega_{Sft}^{\gamma} \left(S_{ft} \right)^{-1/\sigma} X_{ft} = W_{Sft} \tag{A.16}$$

$$\alpha_M M_{ft}^{-1/\sigma} X_{ft} = P_{Mft} \tag{A.17}$$

where $X_{ft} = \left[\frac{1+\eta}{\eta} \right] A_t \Omega_{Hft} \left[\alpha_L L_{ft}^{\gamma} + \alpha_S \left(\Omega_{Sft} S_{ft} \right)^{\gamma} + \alpha_K K_{ft}^{\gamma} + \alpha_M M_{ft}^{\gamma} \right]^{\frac{\eta(1-\gamma)+1}{\gamma\eta}}$. Dividing (A.15) by (A.17) and solving for M_{ft} gives equation (3), and dividing (A.15) by (A.16) and solving for Ω_{Sft} gives (4). Next, substitute for M_{ft} and Ω_{Sft} into the revenue function using (3) and (4) respectively,

$$R_{ft} = e^{u_{ft} + \frac{\eta+1}{\eta} \omega_{Hft}} A_t \left[\left(\frac{E_{ft}^L + E_{ft}^S + E_{ft}^M}{E_{ft}^L} \right) \alpha_L L_{ft}^{\gamma} + \alpha_K K_{ft}^{\gamma} \right]^{\frac{\eta+1}{\eta\gamma}} \tag{A.18}$$

Next, substitute (3) and (4) into (A.15), multiply both sides by L_{ft} and solve for $e^{\frac{\eta+1}{\eta}\omega_{ft}}$ to get

$$e^{\frac{\eta+1}{\eta} \omega_{Hft}} = \frac{E_{ft}^L}{A_t \alpha_L L_{ft}^{\gamma}} \left[\frac{\eta}{1+\eta} \right] \left[\alpha_L L_{ft}^{\gamma} \left(\frac{E_{ft}^L + E_{ft}^S + E_{ft}^M}{E_{ft}^L} \right) + \alpha_K K_{ft}^{\gamma} \right]^{-\delta} \tag{A.19}$$

which can be solved for Hicks-neutral productivity,

$$\omega_{Hft} = \frac{\eta}{1+\eta} log \left\{ \frac{1}{A_t \alpha_L} \frac{\eta}{1+\eta} L_{ft}^{-\gamma} E_{Lft} \times \left[\alpha_L \left(\frac{E_{Lft} + E_{ft}^S + E_{Mft}}{E_{Lft}} \right) L_{ft}^{\gamma} + \alpha_K K_{ft}^{\gamma} \right]^{\frac{-1}{\gamma}\left(\frac{\eta+1}{\eta}\right)} \right\} \tag{A.20}$$

Plugging (A.19) into (A.18) and taking logs gives the estimating equation (5).

A.2.2 Calculation of standard errors

All co variance matrices are computed by the bootstrap, with 400 replications (increasing the number of replications to 1000 has no effect on inference). For the first stage estimates reported in Table 5, the bootstrap is computed by drawing firms without replacement, which amounts to clustering standard errors by firm. For the second stage estimates reported in Tables 6 and 7, the bootstrap is run over the full two-stage procedure. For the long-run estimates reported in 8, we sample directly from the estimated second stage co variance matrix.

The empirical bootstrap distributions have thicker tails than the normal distribution. To accurately reflect this while maintaining standard reporting conventions, the reported standard errors in Tables 6 through 8 are computed so that the bootstrap 90% confidence intervals have the same width as 90% confidence intervals computed using the normal distribution with the reported standard error.

A.2.3 Effect of techies on output

A central element of our methodology is that we assume that techies affect output only through their effect on future productivity, and not through any contemporaneous contribution to factor services that affect current output. This assumption is analogous to the standard assumption that investment in *t-1* has no effect on output in *t-1*, but raises output in t through its contribution to K_t. While our assumption that techies affect output only through their effect on future productivity is well-grounded, it is important to consider how our measurement of productivity could go awry if techies do in fact increase current output directly, a case that we will call the "orthodox case". For concreteness, we suppose that in the orthodox case techies are a component of skilled labor S, so that techies T and managers B (for "bosses") together make up skilled labor S, and that the ratio of techies to managers, $\delta_{ft} = T_{ft}/B_{ft}$, varies across firms and time. In levels, this amounts to

$$S_{ft} = B_{ft} + T_{ft} = (1 + \delta_{ft})\, B_{ft}$$

Using the approximation $log\,(1 + \delta_{ft}) \simeq \delta_{ft}$ and the notation that lower case letters are the log of upper case variables gives $s_{ft} = \delta_{ft} + b_{ft}$. Similarly, define λ_{ft} as the ratio of the techie to manager wage bill in S,

$$E^{S}_{ft} = E^{B}_{ft} + E^{T}_{ft} = (1 + \lambda_{ft})\, E^{B}_{ft}$$

Under the assumption that our model is correct, we can write true Hicks-neutral productivity as

$$\omega^{1}_{Hft} = \theta_{ft} - \frac{1}{\gamma} log \left\{ \alpha_L \left(\frac{E^{L}_{ft} + E^{B}_{ft} + E^{M}_{ft}}{E^{L}_{ft}} \right) L^{\gamma}_{ft} + \alpha_K K^{\gamma}_{ft} \right\}$$

where $\theta_{ft} = \frac{\eta}{1+\eta} log \left\{ \frac{1}{A_t \alpha_L} \frac{\eta}{1+\eta} L^{-\gamma}_{ft} E_{Lft} \right\}$. Under the assumption that the orthodox model is correct, the term E^{B}_{ft} in this expression is multiplied by $(1 + \lambda_{ft})$, giving

$$\omega^{2}_{Hft} = \theta_{ft} - \frac{1}{\gamma} log \left\{ \alpha_L \left(\frac{E^{L}_{ft} + (1 + \lambda_{ft})\, E^{B}_{ft} + E^{M}_{ft}}{E^{L}_{ft}} \right) L^{\gamma}_{ft} + \alpha_K K^{\gamma}_{ft} \right\}$$

If the orthodox model is correct, but we incorrectly estimate Hicks-neutral productivity using ω^{1}_{Hft}, then the error is

$$\begin{aligned} \omega^{1}_{Hft} - \omega^{2}_{Hft} = & \frac{1}{\gamma} log \left\{ \alpha_L \left(\frac{E^{L}_{ft} + (1 + \lambda_{ft})\, E^{B}_{ft} + E^{M}_{ft}}{E^{L}_{ft}} \right) L^{\gamma}_{ft} + \alpha_K K^{\gamma}_{ft} \right\} \\ & - \frac{1}{\gamma} log \left\{ \alpha_L \left(\frac{E^{L}_{ft} + E^{B}_{ft} + E^{M}_{ft}}{E^{L}_{ft}} \right) L^{\gamma}_{ft} + \alpha_K K^{\gamma}_{ft} \right\} \end{aligned}$$

This expression is strictly positive and increasing in the techie share λ_{ft}. The intuition is simple: the larger is λ_{ft}, the greater is the underestimate of true inputs under the wrong model and thus the greater the overestimate of Hicks-neutral productivity.

Under the assumption that our model is correct, we can write true skill-augmenting productivity as

$$\omega^{1}_{Sft} = l_{ft} - b_{ft} + \frac{1}{\gamma} log \left(\frac{\alpha_L E^{B}_{ft}}{\alpha_B E^{L}_{ft}} \right)$$

Under the assumption that the orthodox model is correct, and using $log\,(1+\lambda_{ft}) \simeq \lambda_{ft}$, we can write true skill-augmenting productivity as

$$\omega_{Sft}^{2} = l_{ft} - b_{ft} - \delta_{ft} + \frac{\lambda_{ft}}{\gamma} + \frac{1}{\gamma} log\left(\frac{\alpha_L E_{ft}^{B}}{\alpha_S E_{ft}^{L}}\right)$$

If the orthodox model is correct, but we incorrectly estimate skill-augmenting productivity using ω_{Sft}^{1}, then the error is

$$\omega_{Sft}^{1} - \omega_{Sft}^{2} = \delta_{ft} - \frac{\lambda_{ft}}{\gamma} + \frac{1}{\gamma} log\left(\frac{\alpha_S}{\alpha_B}\right)$$

The third term in this expression is a constant, while the first is positive. In our application we always estimate $1 > \gamma > 0$, so the second term is negative. If techies are paid on average the same as managers, then $\delta_{ft} = \lambda_{ft}$ and we have

$$\omega_{Sft}^{1} - \omega_{Sft}^{2} = \delta_{ft}\left(\frac{\gamma - 1}{\gamma}\right) + \frac{1}{\gamma} log\left(\frac{\alpha_S}{\alpha_B}\right)$$

Since $\left(\frac{\gamma-1}{\gamma}\right) < 0$, we conclude that the error is negatively correlated with the techie share in the cross section: firms with high techie shares δ_{ft} will have measured ω_{Sft} which is biased down by more than for firms with low techie shares. The intuition is as follows. When $1 > \gamma > 0$, greater ω_{Sft} leads to greater employment of skilled workers S, which implies that higher S indicates higher true ω_{Sft}. Incorrectly removing techies from S will thus lead to an underestimate of ω_{Sft}, and the underestimate will be larger the greater is the share of techies in S.

A.2.4 Output and employment elasticities

For the normalized CES production function (1), the output elasticity of Ω_{Sft} is

$$\frac{\partial \ln Y_{ft}}{\partial \ln \Omega_{Sft}} = \epsilon_{YSft} = \frac{\alpha_S\left(\Omega_{Sft}S_{ft}\right)^{\gamma}}{\left[\alpha_L L_{ft}^{\gamma}\left(1 + \frac{E_{ft}^{M}}{E_{ft}^{L}}\right) + \alpha_S\left(\Omega_{Sft}S_{ft}\right)^{\gamma} + \alpha_K K_{ft}^{\gamma}\right]}$$

where we use (3) to eliminate the unobservable M_{ft}. At the point of normalization where both productivity and factor inputs equal unity, the elasticity simplifies to $\epsilon_{YSft} = \alpha_S$. The output effect of Hicks-neutral and skill augmenting technology differences together is

$$d\,lnY_{ft} = d\,\omega_{Hft} + \epsilon_{YSft} d\,\omega_{Sft} \tag{A.21}$$

The partial equilibrium effects of productivity on employment depend on elasticities of substitution and demand. Given factor prices and market demand, begin with the unit cost function corresponding to (1),

$$c = \frac{1}{\Omega_H}\left[\alpha_L^\sigma\left(\frac{w_L}{\Omega_L}\right)^{1-\sigma} + \alpha_S^\sigma\left(\frac{w_S}{\Omega_S}\right)^{1-\sigma} + \alpha_M^\sigma w_M^{1-\sigma} + \alpha_K^\sigma w_K^{1-\sigma}\right]^{\frac{1}{1-\sigma}} = \frac{1}{\Omega_H}X^{\frac{1}{1-\sigma}}$$

apply Shepard's Lemma, and set unit cost equal to marginal revenue from (2) to obtain

$$S_{ft} = \frac{\alpha_S^\sigma}{\Omega_{Sft}} B_t \Omega_{Hft}^{\eta-1}\left(\frac{w_{Sft}}{\Omega_{Sft}}\right)^{-\sigma} X_{ft}^{\frac{\sigma-\eta}{1-\sigma}}$$

$$L_{ft} = \frac{\alpha_L^\sigma}{\Omega_{Lft}} B_t \Omega_{Hft}^{\eta-1}\left(\frac{w_{Lft}}{\Omega_{Lft}}\right)^{-\sigma} X_{ft}^{\frac{\sigma-\eta}{1-\sigma}}$$

where $B_t = \left(\frac{\eta-1}{\eta}A_t\right)^\eta$. Log differentiating these two equations gives the elasticities of employment with respect to productivity differences in equation (10).

A.2.5 Quantifying the aggregate effect of techies and trade on relative demand for skill

Here we describe in detail how we quantify the effect of techies and trade on changes in the aggregate skill ratio between 2009 and 2013.

To get predicted values for productivity, we set all error terms to zero and iterate forward. Starting from $t-1 = 2009$, and setting the time fixed effects equal to zero to eliminate notational clutter, we have

$$\widehat{\omega}_{f2010} = \beta_Z z_{f2009} + B\omega_{f2009}\,,$$

where β_Z and B are estimated parameters from equations (11) and (12), ω_{f2009} is estimated productivity in 2009, and $\widehat{\omega}_{f2010}$ is predicted productivity in 2010. All ω_{ft} (whether predicted or not) are 2×1 vectors with elements ω_{Hft} and ω_{Sft}. Iterating forward gives

$$\widehat{\omega}_{f2013} = \beta_Z z_{f2012} + B\beta_Z z_{f2011} + B^2\beta_Z z_{f2010} + B^3\beta_Z z_{f2009} + B^4\omega_{f2009}\,. \tag{A.22}$$

Suppose that beginning in 2009, $z_{ft} = 0$. Plugging this into (A.22) we can define

$$\widehat{\omega}^0_{f2013} = B^4 \omega_{f2009}$$

and thus the effect of the actual path of firm decisions (captured in the sequence of z_{ft}s) on predicted productivity is

$$\widehat{\omega}_{f2013} - \widehat{\omega}^0_{f2013} = x = \beta_Z z_{f2012} + B\beta_Z z_{f2011} + B^2 \beta_Z z_{f2010} + B^3 \beta_Z z_{f2009}\,. \tag{A.23}$$

The 2×1 vector x defined by equation (A.23) is key in what follows. Note that (A.23) is not affected if we explicitly account for the estimated time fixed effects: since the terms involving the fixed effects are unchanged when we set $z_{ft} = 0$, they would show up in the definition of both $\widehat{\omega}_{f2013}$ and $\widehat{\omega}^0_{f2013}$ and thus would cancel out in the definition of (A.23). The same holds for any elements of the vector z_{ft} that we are not interested in (such as size and age): since these are identical in the definitions of both $\widehat{\omega}_{f2013}$ and $\widehat{\omega}^0_{f2013}$, their effects cancel out in (A.23).

When we consider the effect of firm decisions on techies, all elements of z_{ft} in (A.23) are set to zero, except for firm expenditures on techies. Similarly, when we consider the effect of firm importing decisions, the only non-zero element of z_{ft} in (A.23) are the firm indicators for importing.

To get the effect of predicted productivity on predicted S_{f2013} and L_{f2013}, we substitute the elements of x (under each experiment: techies or importing) for $d\,\omega_{Hft}$ and $d\,\omega_{Sft}$ in equations (10). In doing so, we replaceλ_{Sft} by the industry-specific α_S, because the data do not permit us to gauge the expenditure share on skilled labor. The reason is that we do not know the costs of firms, since we do not observe the cost of capital (we observe all other components of costs). Since we normalize all of our input data by the geometric mean for each industry, at the point of normalization $\overline{\lambda}_S = \alpha_S$ (see Grieco et al. (2016) for more details on the importance of normalization). We then use these predicted percent changes between 2009 and 2013 to get predicted levels in 2013, based on the actual 2009 levels of S_{f2009} and L_{f2009}. Finally, we sum across the predicted S_{f2013} and L_{f2013} to get predicted aggregate levels S_{2013} and L_{2013} in 2013. These are used to compute predicted $S_{2013}/(L_{2013} + S_{2013})$.

We compare the predicted change $S_{2013}/(L_{2013} + S_{2013}) - S_{2009}/(L_{2009} + S_{2009})$ to the actual change in the data, as described in the main text.

A.2.6 Firm choice of techies

In this section we describe a very simple model of a firm's choice of how many techies to employ. The purpose is to give intuition about why some but not all firms choose to hire techies, and to motivate the correlations that are reported in Table 1. We describe the firm's optimal choice of techies, given a simple function from current techies to future productivity. A simple two-period model is sufficient to illustrate the mechanisms at work. We also assume that there is just one dimension to productivity, and that the firm faces the demand curve given by equation (2).

The relationship from techies to changes in log productivity is

$$\omega_{ft} = \omega_{ft-1} + Max\left[\delta\, ln\left(\frac{T_{ft-1}}{\gamma_{1f}}\right), 0\right], \quad \delta \geq 0 \tag{A.24}$$

Here, effective techie services per unit of techies employed is $\frac{1}{\gamma_{1f}} \leq 1$. Fixed costs of employing positive techies are γ_{0f}. Although the elasticity of productivity with respect to techies is constant and equal to δ, the level of techie employment required to attain a given $\Delta\omega_{ft}$ will differ across firms because of differences in γ_{1f}.

The production function is

$$Y_{ft} = \Omega_{ft} L_{ft}$$

where $L_f = F(X_f)$, F is the CES aggregator, X_f is vector of inputs, and $\Omega_{ft} = e^{\omega_{ft}}$. Let w be the cost of the input bundle. By equation (2), revenue is

$$R_{ft} = A\left[\Omega_{ft} L_{ft}\right]^{\frac{\eta-1}{\eta}}$$

The static profit-maximizing input choice is

$$L_{ft} = \Omega_{ft}^{\eta-1}\left[\frac{\eta-1}{\eta}\frac{A}{w}\right]^{\eta}$$

Plugging this back into the expression for revenue gives optimized revenue for given productivity,

$$R_{ft} = B\Omega_{ft}^{\eta-1}, \qquad B = A^{\eta}\left(\frac{\eta-1}{\eta}\right)^{\eta-1} w^{1-\eta}$$

With no discounting, the firm chooses T_{ft-1} to maximize two-period profits,

$$\Pi_f = B\Omega_{ft-1}^{\eta-1} + B\Omega_{ft}^{\eta-1} - rT_{ft-1} - \delta_{0f} I\left(T_{ft-1} > 0\right)$$

where $I()$ is the indicator function. There will be two solutions, one the corner solution with $T_{ft-1} = 0$ and the other an interior optimum with $T_{ft-1} > 0$. When $T_{ft-1} > 0$, equation (A.24) implies $\Omega_{ft} = \left[\frac{T_{ft-1}}{\gamma_{1f}}\right]^{\delta} \Omega_{ft-1}$. Substituting this into the expression for profits gives

$$\Pi_f^T = B\Omega_{ft-1}^{\eta-1} - rT_{ft-1} - \gamma_{0f} + B\left(\left[\frac{T_{ft-1}}{\gamma_{1f}}\right]^{\delta} \Omega_{ft-1}\right)^{\eta-1} \tag{A.25}$$

At the interior solution, the firm chooses T_{ft-1} to maximize Π_f^T. The solution of this problem is

$$T_{ft} = (\delta\eta - \delta)^{\frac{1}{1-\delta(\eta-1)}} r^{\frac{1}{1-\delta(\eta-1)}} \gamma_{1f}^{\frac{\delta(\eta-1)}{\delta(\eta-1)-1}} \Omega_{f1}^{\frac{1-\eta}{\delta(\eta-1)-1}} \tag{A.26}$$

For high enough values of δ, the second order condition of the profit maximization problem doesn't hold and optimal techie employment is infinite. To rule this out we assume $\delta < \frac{1}{\eta-1} < 1$. This restriction implies that the elasticities of techies with respect to r and γ_{1f} are negative, and that the elasticity of techies with respect to ω_{f1} is positive.

Plugging the solution (A.26) back into the expression for Ω_{ft} gives

$$\Omega_{ft} = \left[\frac{\gamma_{1f} r}{\delta(\eta-1)}\right]^{\frac{-\delta}{1-\delta(\eta-1)}} \Omega_{ft-1}^{\frac{1}{1-\delta(\eta-1)}} \tag{A.27}$$

This equation establishes the intuitive result that optimized Ω_{ft} is decreasing in r and γ_{1f}, and increasing in Ω_{ft-1}.

To figure out whether $T_{f1} = 0$ or $T_{f1} > 0$ yields higher profits, the firm simply computes maximized profits in each case. Profits at the corner solution are

$$\Pi_f^C = 2B\Omega_{f1}^{\eta-1}$$

To compute profits at the interior solution, substitute (A.26) and (A.27) into (A.25) to obtain

$$\Pi_f^T = B\Omega_{f1}^{\eta-1} - r\gamma_{0f} + + \left(\frac{\Omega_{f1}}{\gamma_{1f}^{\delta}}\right)^{\frac{\eta-1}{1-\delta(\eta-1)}} \left[B\left[\frac{r}{\delta(\eta-1)}\right]^{\frac{\delta(\eta-1)}{\delta(\eta-1)-1}} - r\delta(\eta-1)^{\frac{1}{1-\delta(\eta-1)}}\right]$$

Thus the difference between the two profit levels is

$$\Pi_f^T - \Pi_f^C = -r\gamma_{0f} + \left(\frac{\Omega_{f1}}{\gamma_{1f}^{\delta}}\right)^{\frac{\eta-1}{1-\delta(\eta-1)}} \left[B\left[\frac{r}{\delta(\eta-1)}\right]^{\frac{\delta(\eta-1)}{\delta(\eta-1)-1}} - r\delta(\eta-1)^{\frac{1}{1-\delta(\eta-1)}}\right]$$

A necessary condition for this to be positive is that the term in brackets is positive. This will be more likely when demand (captured by B) is higher, and less likely when r is higher. If the term in brackets is positive, the whole expression is more likely to be positive the smaller is γ_{1f} and γ_{0f} and the larger is Ω_{f1}. If the term in brackets is negative, then $\Pi_f^T - \Pi_f^C < 0$ even if $\gamma_{0f} = 0$, which shows that fixed costs are not a necessary condition for zero techies to be optimal.

The lessons from this exercise are quite simple and intuitive:

- The optimal amount of techies is more likely to be positive when demand and/or initial productivity are higher.
- The optimal amount of techies is more likely to be zero when fixed costs of techies are high and/or when the efficiency of techies are low.
- The optimal amount of techies may be zero even if the fixed cost of employing techies is zero.
- When the optimal amount of techies is positive, it is increasing in initial productivity and the efficiency of techies.

Firms that export will have a higher demand level A, and thus will be more likely to employ techies. The predictions from this simple model are consistent with the correlations in Table 1.