NBER WORKING PAPER SERIES

# TECHIES, TRADE, AND SKILL-BIASED PRODUCTIVITY

James Harrigan
Ariell Reshef
Farid Toubal

Working Paper 25295
http://www.nber.org/papers/w25295

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2018, Revised February 2021

We thank Daron Acemoglu, Mary Amiti, Zsoa Barany, Eric Bartelsman, Flora Bellone, Andrew Bernard, Tibor Besedes, Esther Ann Bøler, Davin Chor, Chiara Criscuolo, Jan De Loecker, Alon Eizenberg, Georg Graetz, Giammario Impullitti, Beata Javorcik, Guy Michaels, Jean-Marc Robin, and John Van Reenen 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 the CEPR, DEGIT, ENEF, ERWIT, 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). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Techies, Trade, and Skill-Biased Productivity
James Harrigan, Ariell Reshef, and Farid Toubal
NBER Working Paper No. 25295
November 2018, Revised
JEL No. D2,D24,F1,F16,F60,F66,J2,J23,J24,O52

**ABSTRACT**

We study the impact of firm level choices of ICT, R&D, exporting and importing on the evolution of productivity, its bias towards skilled workers, and the implications for labor demand. We use a novel measure of firm-level technology: firms' employment of workers in occupations related to R&D and ICT adoption, who we call “techies”. We develop a methodology for estimating nested CES production functions at the firm level, which allows us to measure both Hicks-neutral and skill-augmenting technology differences. Using administrative data on French firms we find that techies, exporting and importing raise skill-biased productivity. In contrast, only ICT techies raise Hicks-neutral productivity. On average, higher firm-level skill biased productivity does not affect low-skill employment even as it raises the ratio of skilled to unskilled workers, due to the cost-reducing effect of higher productivity. ICT techies account for large increases in aggregate demand for skill, mostly due to their effect on firm size, less so through within-firm changes. Exporting, importing, and R&D techies have smaller aggregate effects.

James Harrigan
Department of Economics
University of Virginia
P.O. Box 400182
Charlottesville, VA 22904-4182
and Paris School of Economics
and also NBER
harrigan@nber.org

Farid Toubal
University of Paris-Dauphine - PSL
Place du Maréchal de Lattre de Tassigny
75775 Paris Cedex 16
France
and CREST, CEPII
ftoubal@ens-cachan.fr

Ariell Reshef
Université Paris 1 - Panthéon Sorbonne
Maison des Sciences Economiques
106-112 Boulevard de l'Hopital
75647 Paris Cedex 13
France
and CNRS, Paris School of Economics, CEPII
ariell.reshef@psemail.eu

# 1 Introduction

Economists have been studying the nexus between labor demand, globalization and technology adoption for decades. Theory and anecdote suggest that some combination of technology and globalization has raised the relative demand for more skilled workers.[1] However, causal identification of these forces has proved challenging, as well as the channels through which they operate. 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.

We study how firm-level decisions on research and development (R&D), information and communication technology (ICT) adoption, as well as exporting and importing decisions affect firm-level productivity and its bias towards skilled workers. These shifts in productivity, in turn, affect firms' optimal demand for labor. We find large effects of global participation, R&D and, in particular, ICT on labor demand through their effects on SBTC.

We make several contributions. Our first is to develop new methodology to estimate a firm-level nested constant elasticity of substitution (CES) production function, where we nest skilled and unskilled labor in a labor composite. This allows us to compute both Hicks-neutral and skilled labor augmenting productivity shifters. We demonstrate that nesting skilled and unskilled labor has important implications for estimating the effects of participation in international trade and technology adoption on productivity.[2] Second, while most studies rely on small samples and/or focus on manufacturing, we apply our methodology to to most sectors of the French private sector, including both manufacturing and non-manufacturing industries.[3] This is particularly important for evaluating the effects of ICT, which is pervasive outside of manufacturing.

We then jointly estimate the separate causal effects of firm-level decisions on R&D and ICT investments, as well as decisions on importing and exporting on both dimensions of productivity, Hicks neutral and skill-augmenting. Ours is the first paper to do this. Finally, we evaluate the quantitive implications of our estimates, both at the firm-level and for aggregate relative demand for skilled labor. One important finding is that ICT has the largest effect on aggregate demand for

[1] Helpman (2018) and Acemoglu and Autor (2011) provide insightful reviews of the literature.

[2] This goes beyond Fox and Smeets (2011), who find that adjusting inputs for quality in several dimensions lowers Hicks-neutral productivity dispersion. Our point here is that inference on the importance of forces that affect productitivy is changed once we allow labor to be composed of two types. We elaborate on this in Section 6.5.

[3] We describe the sample in Section 5.

skill, mostly through its effect on firm sizes, rather than through within-firm adjustment.[4] Another important finding is that increases in demand for skilled labor due to higher skill augmenting productivity are not on average accompanied by lower demand for unskilled labor, due to the cost-reducing effect of skill augmenting productivity.

Quantifying the importance of technology adoption and globalization for relative labor demand at the micro level is hard for two reasons. First, while it is relatively easy to measure importing and exporting at the level of the firm, it has proven very hard to measure technology adoption, except in case studies and in particular industries.[5] The focus on firms is important, because that is where decisions about technological change, globalization and employment are made. Second, it is difficult to identify causal effects since firms jointly choose whether to import, export and adopt technology. In order to overcome these challeneges, 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.

Given the absence of information about real output or real intermediate input use in our data, we build on the methodology proposed by Grieco et al. (2016) [GLZ]. We extend GLZ's approach in several ways. First, we separate labor into three components: skilled and unskilled labor which contribute to output in the standard way, and "techies", who are assumed to affect production only through their lagged impact on productivity. We discuss the sensitivity of the results to this assumption in Section 6.5. Second, we allow the elasticity of substitution between skilled and unskilled labor to differ from the elasticity among capital, materials, and composite labor by estimating a nested-CES production function. Third, we allow these firm production functions to include—in addition to a Hicks neutral term that is already present in GLZ—a skilled-labor augmenting term.

Our estimator extends that of GLZ to the case of a nested CES production function, while applying insights from León-Ledesma et al. (2010) on using multiple equations to identify parameters of the production function. The estimator exploits the first order conditions implied by profit maximization and monopolistic competition to recover unobserved quantities of intermediate inputs and (augmented) skilled labor services, to identify the production function parameters, and to fully recover both Hicks-neutral and skill augmenting productivity shifters.[6] Like GLZ, our approach

[4]The greater importance of changes in firm composition versus within-firm adjustment is also found in De Loecker and Eeckhout (2018) for increases in average markups, Autor et al. (2020) for the decline in the labor share, and Harrigan et al. (2021) for job polarization.

[5]We discuss this literature in Section 2.

[6]Doraszelski and Jaumandreu (2013) and Doraszelski and Jaumandreu (2018) also use the first order conditions

does not rely on proxy variable methods.[7]

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. As in Doraszelski and Jaumandreu (2013), we assume that productivity follows a controlled Markov process and is endogenously determined by lagged productivity as well as other lagged firm-level decisions. Once this process is estimated, we use the parameters of the production functions to quantify the impact of these factors on the demand for skilled and unskilled labor. Our approach resembles that of Doraszelski and Jaumandreu (2018), which is the first paper in the production function estimation literature to estimate both neutral and non-neutral technology differences. Doraszelski and Jaumandreu (2018) include, in addition to a Hicks neutral productivity shifter, a labor augmenting term. However, they do not distinguish between skilled and unskilled labor and they identify only one elasticity of substitution. In addition, they rely on availbility of real input use and input prices at the firm level to identify the production function, which are not available in our data. While they observe only 2,375 firms in 10 manufacturing industries, our sample includes roughly 193,000 firms in both manufacturing and non-manufacturing.

We apply our methodology using matched employer-employee administrative data for most of the French private sector from 2009 to 2013. The dataset has information on exporting and importing by firm. In order to identify the effect of technology adoption and R&D on firm level productivity, we use workers in technology-related occupations, who we call "techies". These workers are engineers and technicians with skills and experience in science, technology, engineering and math (STEM). They are 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 (Tambe and Hitt (2012, 2014); Harrigan et al. (2021)).

The detailed description of occupations in the dataset allow us to identify the techies who are central in creating, planning, installing, and maintaining ICT, as well as in training and assisting other workers in the use of ICT. We are able to separately identify other techies who 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 at the firm level, for all firms in the French private sector. As a result we are able to estimate the separate firm-level effects of R&D

to identify the production function.

[7] Olley and Pakes (1996), Levinsohn and Petrin (2003) and Ackerberg et al. (2015) all rely on proxy methods. Gandhi et al. (2020) propose an alternative estimator that exploits firms' optimality conditions.

and ICT investment on productivity and its bias towards skilled workers.

Very few papers have investigated R&D and ICT investment jointly due to the lack of data (Hall et al. (2013), Mohnen et al. (2018)). Hall et al. (2010) argue that R&D is related to product and process innovation, whereas others argue that ICT investments foster organizational changes within firms such as business processes and work practices (Bresnahan et al. (2002)) and span of control (Bloom et al. (2014)), both of which may enhance productivity (Brynjolfsson and Hitt (2000)). As we discuss in the literature review below, and as the results demonstrate, the distinction between R&D and ICT proves to be important as both investments may have different influences on productivity and relative labor demand.

For each industry, we find an elasticity between capital, materials and labor aggregate which is greater than unity (1.5 on average). These results contrast with Doraszelski and Jaumandreu (2018) and Raval (2019) who assume a CES functional form with Hicks-neutral productivity differences across firms and labor augmenting technology. They do not distinguish skilled from unskilled labor, treating them as perfect substitutes. As the authors acknowledge, their findings of labor-augmenting technological progress may be conflating skill-composition differences with labor augmenting technological differences across firms. We find moreover an elasticity between skilled and unskilled labor that is greater than the upper nest elasticity (2.8 on average). This finding implies that skill augmenting technological progress necessarily raises the relative demand of skilled workers and the firm's skill intensity.

We find that both technology and trade have large effects on skill-augmenting productivity. Our baseline estimates imply that compared to firms that don't employ techies, firms with a lot of techies (at the $75^{th}$ percentile) have skill-augmenting productivity which is 15 percent higher. Both ICT and R&D techies have a similarly large marginal effect on skill-augmenting productivity, but because firms that employ R&D techies do so at greater intensity than firms that employ ICT techies, R&D techies are more important in explaining cross-firm differences in skill-augmenting productivity. In contrast, the effect of techies on Hicks neutral productivity is driven only by ICT techies. Turning to the effects of trade, we find that the effect of exporting and importing is to raise skill-augmenting productivity by 15 percent and 7 percent, respectively. We don't find an effect of trade on hicks neutral productivity shifters.

We use the estimates of the elasticities of substitution and demand to translate the estimated effects on productivity into effects on firm-level labor demand. Compared to firms that don't employ techies, firms with a lot of techies have employment of skilled labor that is 28 percent higher, while employment of unskilled labor is no lower. The effects of trade are comparably large: exporting and importing raise employment of skilled workers by 28 percent and 13 percent respectively while not reducing unskilled employment. 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. However, on avereage, this is only a relative effect: because of the powerful productivity effects of technology and trade, skilled workers see labor demand rise when the firms where they work hire techies and/or engage in international trade, but employment of unskilled workers barely drops. Nevertheless, there is cross-industry heterogeneity in this respect, where SBTC has sometimes positive and sometimes negative effects on demand for unskilled labor.

Aggregating all firm-level decisions and using our estimates we find that ICT techies have the largest effect on aggregate relative demand for skill, mostly through their effect on firm sizes, not through within-firm adjustment. The reason is that we estimate that only ICT affects directly Hicks-neutral productivity. In addition, ICT techies are more prevalent in larger and skill intensive firms, which amplifies their firm-level effect.

All the results we find in this paper are driven by firm-level decisions. Our methodology consistently estimates the effects of firms' choices to employ techies and trade, but does not model these choices themselves. We also do not consider the effects of these firm-level decisions on industry or economy-wide equilibrium employment and wages. These are limitations of the scope of our paper but do not impinge on the internal consistency 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

Skill-biased technological change (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 trade on SBTC. We discuss these few papers here to put our contribution in context. We also review research on the specific role of techies.

## 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.[8] Several 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 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) survey 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 that suggests a positive effect of scientists and engineers on TFP. Bresnahan et al. (2002) argue persuasively for a complementarity

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

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 analysis to structural estimation, and offers causal inference.

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 important 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.[9]

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.[10] 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

[9] Bøler (2015) sets up a similar nested-CES production function to ours, but does not fully estimate it. Instead, she only considers relative demand for skill, inferred from the lower nest.

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

firms across three years (2004, 2006, and 2009). As a consequence, their data analysis is mainly descriptive, while our dataset of nearly 200,000 French firms 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. The production function estimator of Doraszelski and Jaumandreu (2018) is related to our methodology, as we discuss above, but relies on the availablility of real output and inputs and their prices—information that is absent in our data. Raval (2019) 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.[11] While these two papers are not about SBTC, 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". Similarly, Deming and Noray (2018) show that, in their words, "STEM jobs are the leading edge of technology diffusion in the labor market".

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 Harrigan et al. (2021), which uses detailed occupational data (including data on techies) for the entire French private sector from 1994 to 2007. Harrigan et al. (2021) show that employment growth is higher in French firms with more techies, and also that more techies leads to within-firm skill upgrading. 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. Using a remarkable dataset that tracks the movement of IT workers across firms, Tambe and Hitt (2014) find what they

[11] The industry-level estimates from an earlier version of Raval (2019) are used to compute the aggregate elasticity of substitution in Oberfield and Raval (2020).

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 methodology is not vulnerable to the endoegenity bias that plagues OLS estimation of production functions (the so-called "transmission bias") .

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

# 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-augmenting productivity, yet only few firms employ them. 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 $\omega_{ft-1}$ as given and has to choose optimal techie employment $T_{ft-1}$ to maximize profits. The relationship from techies to changes in productivity is

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

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

Although the elasticity of productivity with respect to techies is constant and equal to $\delta$, the level of techie employment required to attain a given growth in productivity $\Delta\omega_{ft}$ will differ across firms because of differences in $\gamma_{1f}$. Fixed costs of employing positive techies are $\gamma_{0f}$ and the wage of techies is $r$, so the cost of hiring techies is $rT_{ft-1} + \gamma_{0f}$. With heterogeneity in the costs $\gamma_{0f}$ and $\gamma_{1f}$ 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. All regressions include firm×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.

# 4 Econometric methodology

Most firm- or plant-level datasets (including ours) include information on revenue and the value of expenditures on materials but not data on the corresponding output and materials prices. 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.[13]

Many papers follow some variant of the Olley and Pakes (1996), Levinsohn and Petrin (2003) and

[13] 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.

Ackerberg et al. (2015) [OP/LP/ACF] proxy variable methodology to first estimate productivity and then study its determinants, often in the context of estimating the effects of exporting or importing on productivity.[14] 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.

As discussed above, Doraszelski and Jaumandreu (2013) take a different approach to estimating endogenous productivity, combining the firm's optimal demand with a controlled Markov specification for productivity, while Doraszelski and Jaumandreu (2018) apply this approach to estimating Hicks-neutral and separate labor-augmenting productivity shifters. We cannot apply their methodology because it relies on availbility of real input use and input prices at the firm level to identify the production function, which are not available in our data. Doraszelski and Jaumandreu (2018) estimate a 3-factor CES production function with Hicks-neutral $\omega_H$ and labor-augmenting $\omega_N$ technology differences across firms. Labor and materials are static inputs, which implies that the optimal labor to materials ratio depends on $\omega_N$ but not on $\omega_H$. This insight motivates a two stage procedure. In the first stage they recover $\widehat{\omega}_N$ 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 $\widehat{\omega}_N$ and $\widehat{\sigma}$ as data in a second stage which allows them to recover $\widehat{\omega}_H$.[15] In contrast, our approach identifies all parameters of the production function and productivity shifters in one step, as we describe below.

Our approach extends Grieco et al. (2016) [GLZ] in three 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 $\Omega_H$, already present in GLZ, we consider a skilled-labor augmenting term $\Omega_S$. Third, we allow the elasticity of substitution between $S$ and $L$ to differ from the elasticity among capital, materials, and composite labor.

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

[15] 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 $\omega_N$ growth of 1.5 percent per year.

### 4.1 Estimating the production function and productivity

We begin with a constant returns to scale nested CES production function, where physical output of firm $f$ in year $t$ $Y_{ft}$ is produced using composite labor $N_{ft}$, capital $K_{ft}$, and materials $M_{ft}$. The labor composite $N_{ft}$ is a CES function of skilled labor $S_{ft}$ and unskilled labor $L_{ft}$, both measured as hours worked. These functions are assumed to be the same for all firms in an industry up to the two productivity levels $\Omega_{Hft}$ and $\Omega_{Sft}$. For reasons discussed by GLZ, it is important for identification to normalize each data series by its geometric mean. We do this, so in what follows all variables should be understood as values relative to their geometric means, such that $\overline{L} = \overline{S} = \overline{K} = \overline{M} = \overline{Y} = 1$, where an overbar denotes the geometric mean of the respective variable.[16]

The normalized production function is

$$Y_{ft} = \Omega_{Hft} \left[ \alpha_N N_{ft}^{\gamma} + \alpha_K K_{ft}^{\gamma} + \alpha_M M_{ft}^{\gamma} \right]^{\frac{1}{\gamma}} , \ \gamma = \frac{\sigma - 1}{\sigma} \leq 1 \tag{1}$$

$$N_{ft} = \left[ \alpha_L L_{ft}^{\rho} + \alpha_S \left( \Omega_{Sft} S_{ft} \right)^{\rho} \right]^{\frac{1}{\rho}} , \ \rho = \frac{\varphi - 1}{\varphi} \leq 1 \, . \tag{2}$$

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

$$P_{ft} = A_t Y_{ft}^{\frac{1}{\eta}} , \tag{3}$$

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

[16] To understand the importance of normalizing the CES production function, see León-Ledesma et al. (2010) and the discussion and references on page 668 of Grieco et al. (2016). Below we illustrate one important outcome of normalization: it helps identify the distribution parameters of the production function.

levels have been realized.

#### 4.1.1 The estimating framework

Since $L$, $S$ and $M$ are static inputs, their first order conditions for expected profit maximization will always hold with equality.[17] As shown in Appendix A.2, these conditions imply

$$M_{ft} = \left( \frac{\alpha_N}{\alpha_M} \frac{E^M_{ft}}{E^N_{ft}} \right)^{1/\gamma} N_{ft} \tag{4}$$

$$\Omega_{Sft} = \left( \frac{S_{ft}}{L_{ft}} \right)^{\frac{1}{\varphi-1}} \left( \frac{\alpha_S W_{Lft}}{\alpha_L W_{Sft}} \right)^{\frac{\varphi}{1-\varphi}} , \tag{5}$$

where $E^X_{ft} = P_{Xft} X_{ft}$ is expenditures on production factor $X$ and $P_X$ is the price (or wage, $W$) of $X$. The derivation of (4) and (5) requires that $\sigma \neq 1$ and $\varphi \neq 1$, which rules out the Cobb-Douglas case that is assumed in most of the productivity estimation literature.[18] A few more steps yield an estimating equation,

$$lnR_{ft} = ln\left[ \frac{\eta}{1+\eta} \right] + ln\left[ E^M_{ft} + E^N_{ft}(1+\tau)\left( \frac{K_{ft}}{L_{ft}} \right)^{\gamma} \left( \frac{\lambda_{ft}}{\alpha_L} \right)^{\frac{\gamma}{\rho}} \right] + u_{ft} , \tag{6}$$

where $E^N_{ft} = E^L_{ft} + E^S_{ft}$ is the wage bill, $\lambda_{ft} = (E^L_{ft}/E^N_{ft})$ is unskilled labor's share of the wage bill, and $\tau = \alpha_K/\alpha_N$.[19]

Equation (6) has just four parameters ($\eta, \gamma, \rho$ and $\tau$) and no endogenous or unobservable variables. The key to the derivation is that there are three flexible ("static") inputs ($S$, $L$ and $M$), which gives us two ratios of static first order conditions, (4) and (5). These two equations allow us to eliminate the two unobservables, $M_{ft}$ and $\Omega_{Sft}$, and (A.23) in the Appendix allows us to eliminate $\Omega_{Hft}$. Because we have eliminated the unobserved productivity terms from our estimating equation, we do not need to use proxy variable methodology. Our timing assumptions are key: firms choose static inputs after observing both productivity shocks but before observing the revenue shock.

The five distribution parameters $\alpha_S, \alpha_L, \alpha_N, \alpha_K$ and $\alpha_M$ are identified by $\tau = \alpha_K/\alpha_N$ and the following equations,

$$\alpha_N + \alpha_K + \alpha_M = 1 \tag{7}$$

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

[18] This includes OP/LP/ACF. In our results below in Table 5, the point estimates for $\sigma$ and $\varphi$ exceed one for every industry, and in Table 6 we can always reject the null hypotheses $\sigma = 1$ and $\varphi = 1$.

[19] See Appendix A.2.1 for the derivation of (6).

$$\alpha_M \overline{E}^N = \alpha_N \overline{E}^M \tag{8}$$

$$\alpha_L = \overline{E}^L / \overline{E}^N \ , \quad \alpha_S = \overline{E}^S / \overline{E}^N \ . \tag{9}$$

Equations (7) is implied by constant returns to scale. Equation (8) follows from taking the geometric mean of (4), and using the normalization conditions. Equations (9) follow from taking the geometric mean of (5), using the normalization conditions (which imply $\overline{\Omega}_S = 1$), and constant returns. Equations (9) show that $\alpha_L$ and $\alpha_S$ are identified directly from geometric means of the data, and do not require estimation.

Equation (6) can be estimated consistently by nonlinear least squares. Following León-Ledesma et al. (2010), in order to increase efficiency we make use of the implication of (5) that the skill ratio can be written as

$$\ln\left(\frac{S}{L}\right)_{ft} = \beta_0 - \varphi \ln\left(\frac{W_S}{W_L}\right)_{ft} + v_{ft}\,, \tag{10}$$

where $v_{ft} = (\varphi - 1)\ln \Omega_{Sft}$. Equation (10) is an estimating equation, but since the unobservable $v_{ft}$ contains $\Omega_{Sft}$ it is likely that $Cov[\ln(W_S/W_L)_{ft}, v_{ft}] \neq 0$ . Therefore using (10) in our estimation framework requires an instrument for $\ln(W_S/W_L)_{ft}$. To form such an instrument, we exploit the rich occupational detail in our data on employment. As discussed in Section 5.1 below, we measure $S$ and $L$ as aggregates of employment in many detailed occupational categories. Denote the industry average wage in detailed occupation $o$ as $\overline{W}_o$, and the share of occupation $o$ in firm $f$'s employment of labor aggregate $j \in \{S, L\}$ as $\lambda_{jof}$. By definition, these shares sum to one within $S$ and $L$, $\sum_o \lambda_{Sof,t} = \sum_o \lambda_{Lof,t} = 1 \quad \forall f, t$. Our instrument $\ln Z_{ft}$ is then defined as

$$\ln Z_{ft} = \sum_o \lambda_{Sof,t-1} \ln \overline{W}_{ot-1} - \sum_o \lambda_{Lof,t-1} \ln \overline{W}_{ot-1}\,. \tag{11}$$

Equation (11) has a form similar to a Bartik or shift-share instrument. Goldsmith-Pinkham et al. (2020) show that a sufficient condition for exogeneity of instruments like (11) is that the shares $\lambda_{jof,t-1}$ are exogenous to the shock $v_{ft}$ in equation (10).[20] Exogeneity of shares follows in our case from our assumption that $\Omega_{Sft}$ affects the productivity of aggregate $S$, but not the individual occupations that make up $S$. An implication is that $\Omega_{Sft}$ will affect $S/L$ but not the composition of $S$ or $L$. Because there is substantial heterogeneity across firms in the detailed occupational

[20] Adao et al. (2019) propose methods for inference in single-equation linear shift-share designs, but their methods are not applicable to our GMM estimator developed below.

makeup of $S$ and $L$, there is ample cross-section variation in $\ln Z_{ft}$ to identify the parameter on $\ln(W_S/W_L)_{ft}$ in equation (10).

To estimate the parameters of interest we form a GMM estimator as follows. Write (6) as $lnR_{ft} = f(\eta, \gamma, \rho, \tau\,;\, E^M_{ft}, E^N_{ft}, K_{ft}, L_{ft}, \lambda_{ft}) + u_{ft}$. We compute the derivatives of $f$ with respect to ($\eta, \gamma, \rho$ and $\tau$) and set the product of these derivatives with the error $u_{ft}$ to zero, which gives four moment conditions. Then, using the instrument defined by 11, equation (10) gives us two addition moment conditions that identify $\rho$ and the constant in (10). We thus have six moment conditions to identify five parameters (recall that $\varphi = (1/(1-\rho))$.

Our GMM estimator weights each firm-year observation by total firm employment. 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 wish to avoid.

#### 4.1.2 Recovering productivity

We recover estimated skill augmenting productivity using (5). We show in Appendix A.2.1 that estimated Hicks neutral productivity is

$$\omega_{Hft} = \frac{\eta}{1+\eta} log \left\{ \frac{1}{A_t} \frac{\eta}{1+\eta} \left( \frac{E_{Nft}}{\alpha_N N^{\gamma}_{ft}} \right) \times \left[ \alpha_N \left( \frac{E_{Nft} + E_{Mft}}{E_{Nft}} \right) N^{\gamma}_{ft} + \alpha_K K^{\gamma}_{ft} \right]^{-\delta} \right\} \tag{12}$$

where $\delta = [\eta\,(1-\gamma) + 1]/\gamma\eta$. 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 firms' productivities 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 $\Omega_{Hft}$ is physical TFP. This follows from deflating revenue by price using equation (3). In practice, it is not plausible that the simple demand system given by (3) solves all the problems related to unobservable prices and quality that are required to distinguish revenue TFP from physical TFP, as Foster et al. (2008) are able to do. Our interpretation of $\Omega_{Hft}$ is revenue TFP, where some but not all of the variation in revenue has been controlled for by the demand system.

#### 4.1.3 Skill-augmenting productivity and skill bias

Hicks-neutral technology differences $\Omega_{Hft}$ have no implications for relative skill demand because they do not affect the relative marginal products of different inputs. Re-arranging equation (5) shows that the effect of skill augmenting technology differences $\Omega_{Sft}$ on relative skill demand depends crucially on the elasticity of substitution $\varphi$,

$$\frac{S_{ft}}{L_{ft}} = \Omega_{Sft}^{\varphi-1}\left(\frac{\alpha_S W_{Lft}}{\alpha_L W_{Sft}}\right)^{\varphi} \tag{13}$$

If $\varphi > 1$, a higher level of $\Omega_{Sft}$ raises relative skill demand, which is to say that skill-*augmenting* technology differences are skill-*biased*. In our empirical results below we estimate $\widehat{\varphi} > 1$ for all industries.[21] When $\varphi > 1$, the identification of $\Omega_{Sft}$ is transparent from equation (5): it is a residual that rationalizes greater skill intensity, conditional on parameters and factor prices.

Estimating equations similar to (13) 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 (2019) 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 $\varphi$ in the data, through equation (10); as (6) illustrates, variation in levels also plays an important role.

### 4.2 Employment effects of productivity

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 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

[21] See Table 5.

A.2.4 that the effects of productivity differences across firms on employment are

$$\widehat{L} = (-\eta - 1)\,\widehat{\Omega}_H + (-\eta - \sigma)\,\theta_S\widehat{\Omega}_S + (\sigma - \varphi)\,\theta_{SN}\widehat{\Omega}_S \tag{14}$$

$$\widehat{S} = (-\eta - 1)\,\widehat{\Omega}_H + (-\eta - \sigma)\,\theta_S\widehat{\Omega}_S + (\sigma - \varphi)\,\theta_{SN}\widehat{\Omega}_S + (\varphi - 1)\,\widehat{\Omega}_S \tag{15}$$

where $\theta_{SN}$ is the share of $S$ in the unit cost of the labor composite $N$ and $\theta_S$ is the share of $S$ in total unit cost.[22] We show in Appendix A.2.4 that our normalization implies that $\theta_{SN} = \alpha_S$ and $\theta_S = \alpha_S\alpha_N$ at the geometric mean of the data. Each equation combines a labor saving effect and a cost reducing demand (or substitution) effect of technological progress:

- For both $S$ and $L$, the effect of Hicks-neutral technological progress $\widehat{\Omega}_H > 0$ is to reduce the employment that is required to produce a unit of ouput, and thus decrease employment with an elasticity of $-1$. But at the same time costs decrease with an elasticity of $-1$ and thus increase demand with elasticity $-\eta > 0$, so the net effect on employment of Hicks-neutral technological progress is $(-\eta - 1)\,\widehat{\Omega}_H$.

- The coefficient $(-\eta - \sigma)\,\theta_S$ that appears in both equations represents the effect of skill-augmenting technological progress $\widehat{\Omega}_S > 0$ through the labor composite $N$. This reduces the employment of both $S$ and $L$ that is required to produce a unit of output with an elasticity of $-\sigma\theta_S$. At the same time, production costs decrease with an elasticity of $-\theta_S$ and this increases overall demand with elasticity $-\eta\theta_S$. Thus, labor demand increases with an elasticity $(-\eta - \sigma)\,\theta_S$ through this channel.

- The coefficient $(\sigma - \varphi)\,\theta_{SN}$ that appears in both equations represents the effect of skill-augmenting technological progress $\widehat{\Omega}_S > 0$ through its effect on substitution towards the labor composite $N$. Skill-augmenting technological progress reduces the employment that is required to produce a unit of $N$ with an elasticity of $-\varphi\theta_{SN}$. At the same time, the cost of a unit of $N$ decreases with an elasticity of $-\theta_{SN}$ and this induces substitution towards $N$ within overal inputs with elasticity $\sigma\theta_{SN}$. Thus, labor demand increases with an elasticity $(\sigma - \varphi)\,\theta_{SN}$. This coefficient is negative if $\varphi$, the elasticity of substitution between $L$ and $S$ within $N$, exceeds $\sigma$, the elasticity of substitution between $N$ and the other factors (this is what we find empirically below, see Table 5).

- The term $(\varphi - 1)\,\widehat{\Omega}_S$ represents the familiar balance between the efficiency effect which re-

[22] Recall that we define the elasticity of demand $\eta$ to be less than $-1$, so $-\eta > 1$.

duces employment with an elasticity of $-1$ and the substitution between the other factors and $S$ which increases employment of $S$ with an elasticity $\varphi$.

Subtracting the first equation from the second shows that the elasticity of skill intensity with respect to skill-augmenting technological progress is $\varphi - 1$. When this elasticity is positive, skill-*augmenting* technological progress is said to be skill-*biased*. We apply equations (14) and (15) in our empirical analysis below.

### 4.3 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 estimator in Section 4.1.1, did not make any assumptions about stochastic processes that characterize the evolution of productivity shifters $\omega_{Hft}$ and $\omega_{Sft}$. Therefore, 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: lagged productivity, a $k \times 1$ vector of lagged characteristics of the firm $z_{ft-1}$, and a shock which is orthogonal to all other shocks and lagged variables 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 $\omega_{Hft}$ and $\omega_{Sft}$ to influence each other we specify the following two equations,

$$\omega_{Hft} = \mu_{Ht} + \beta_{HH}\omega_{Hft-1} + \beta_{HS}\omega_{Sft-1} + \beta_{HZ}z_{ft-1} + \xi_{Hft} \tag{16}$$

$$\omega_{Sft} = \mu_{St} + \beta_{SH}\omega_{Hft-1} + \beta_{SS}\omega_{Sft-1} + \beta_{SZ}z_{ft-1} + \xi_{Sft} \tag{17}$$

The shocks $\xi_{Hft}$ and $\xi_{Sft}$ are assumed to be serially uncorrelated. The industry$\times$time fixed effects $\mu_{Ht}$ and $\mu_{St}$ 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 (16) and (17). 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,

$$\begin{bmatrix} \omega_{Hf} \\ \omega_{Sf} \end{bmatrix} = (I-B)^{-1}\beta_Z z_f, \qquad B = \begin{bmatrix} \beta_{HH} & \beta_{HS} \\ \beta_{SH} & \beta_{SS} \end{bmatrix}, \ \beta_Z = \begin{bmatrix} \beta_{HZ} \\ \beta_{SZ} \end{bmatrix}. \tag{18}$$

#### 4.3.1 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 (16) and (17) 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 (16) 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 (16) and (17) 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. First, it is *lagged* exporting that enters the Markov process, which is to say that productivity (more precisely, the shock to productivity $\xi_{Hft}$) is realized after the exporting decision is made. Second, 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 $M_{ft}$. 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.

#### 4.3.2 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 capital in time $t$. Similarly, Beaudry et al. (2016) use a model with cognitive labor in $t$ affecting output only through its effect on organizational capital in *t+1*.

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 (6) does not go through. For this reason we drop from our analysis the Computer and electronics industry, where ICT techies almost surely contribute to output directly. 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.

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 $\varphi > 1$, firms with high techie shares will have measured $\omega_S$ which is biased down by more than for firms with low techie shares. The reason is that with $\varphi > 1$ greater $\omega_S$ leads to greater employment of skilled workers $S$, which implies that higher $S$ indicates higher true $\omega_S$. Incorrectly removing techies from $S$ will thus lead to an underestimate of $\omega_S$, 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 $\omega_H$ in equation (16) may be biased up, and our estimated effect of techies on $\omega_S$ in equation (17) 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 (16) and (17). Second, although mis-specification leads to a mechanical *intra*-temporal correlation between techies and $\omega_S$ and between techies and $\omega_H$, 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 (6), they should have no explanatory power in the second stage regressions (16) and (17). In section 6.5 below, we test this possibility by estimating the first stage with techies as part of labor inputs, 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.

### 4.4 Estimation details

We estimate equations (6) and (10) jointly by GMM, separately for 14 industries, which include both manufacturing and non-manufacturing sectors, and we weigh observations by employment

(Section 4.1.1 above discusses estimator in greater depth).[23] We report robust standard errors clustered by firm. Each industry sample 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 was made. This (often implicit) assumption is now standard in the literature, and we make it here. The estimated elasticities of substitution is given by the formulas $\widehat{\sigma} = (1 - \widehat{\gamma})^{-1}$ and $\widehat{\varphi} = (1 - \widehat{\rho})^{-1}$ and standard errors take this into account. 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 (5) and (12). After dropping the highest and lowest percentile of estimated productivity to trim outliers, we estimate the controlled Markov processes given by equations (16) and (17). 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 we compute bootstrap standard errors clustered by firm to take into account that the second stage uses estimated productivities as regressors.[24]

# 5 Data

We construct a detailed panel data on firms in the French private sector between 2009 and 2013. The panel is the result of merging three confidential, administrative firm-level datasets. Martching firms across these datasets is straightforward because firms 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

[23] 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 (6) failed to converge. We also drop the financial intermediation sector.

[24] Our bootstrap re-samples firms rather than individual firm-year observations, so the resulting bootstrap covariance matrices are effectively clustered by firms. All bootstrap results are computed using 800 replications. For more details on how we compute standard errors, see Appendix A.2.2.

on employees in each of their establishments, which are identified by their SIRET.[25] 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 4-digit occupational codes (there are almost 500 in total) 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 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 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. Table 3 also reports wide dispersion in these shares across industries.

The techie wage bill share 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.[26] 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.[27] 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 precise measure of firm-level effort devoted to tech-

[25]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 concern us given our focus on firm-level rather than individual outcomes.

[26]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.

[27]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.

nology 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.[28] Similarly, Schweitzer (2019) finds that in 2014, labor costs account for 60 percent of aggregate R&D spending in France.[29]

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.[30]

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. Overall, relatively few of the jobs in this category require a university degree, with the exception of the fairly large category 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.[31] Most French workers are on permanent labor contracts which make them 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.[32] This is all that is required for our

[28] 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.

[29] The remiainder 40 percent are split into 6 percent capital expenditures and 34 percent "other current expenses".

[30] 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.

[31] See equations (A.23) and (A.24) in the appendix.

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

estimating equation (6) to be appropriate for French firms.

## 5.2 Other firm level data

The DADS has information on the two-digit sector of activity of the firm. The estimation sample includes firms in 14 industries, which include both manufacturing and non-manufacturing sectors.[33] Firm balance sheet information comes from the FARE dataset for the years 2009-2013.[34] 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 dataset.

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. This is not a big concern for us, because most of the missing values are in the oil refining industry, which we drop from our sample.

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.[35] 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.

[33]Two sectors (coke and refined petroleum, and pharmaceutical products) are dropped because they have tiny shares of total hours worked, and three sectors (transport equipment, transportation and storage and publishing/broadcasting) are dropped because estimation of equation (6) failed to converge. We also drop the computers and electronic sector because of its intensity in techie workers.

[34]Fichier Approché des Résultats Ésane (FARE)

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

# 6 Estimation results

We start by reporting the results of production function estimation and the implied elasticities of employment with respect to skill augmenting productivity. We then turn to our estimates of endogenous productivity (equations 16 and 17). Using our regression results, we derive implications for labor demand. Finally, we estimate the aggregate effect of techies and trade on relative demand for skill.

## 6.1 Production function estimates

Table 5 reports industry-by-industry estimates of the production function and demand parameters. The derivation of equation (6) requires $\sigma \neq 1$ and $\varphi \neq 1$. All of our estimates of $\sigma$ and $\varphi$ are greater than one, and in all industries we can reject the nulls that $\sigma = 1$ and $\varphi = 1$ at conventional levels of statistical significance (see Table 6). The employment-weighted averages are roughly 1.5 for $\sigma$ and 2.8 for $\varphi$. Economic logic requires that the estimated elasticity of demand $\eta$ satisfies $-\widehat{\eta} > 1$, which holds and is statistically significant for all industries. The estimated $\eta$'s are plausible in magnitude, with an employment-weighted average value of -6.5. 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 $-\widehat{\eta} > \widehat{\sigma}$ and $\widehat{\varphi} > \widehat{\sigma} > 1$, and in all but one industry we can reject the null that $\varphi = \sigma$ (see column (3) of Table 6). From equations (14)-(15), this implies that skill-augmenting technological progress is skill-biased and thus necessarily raises the firm's skill intensity. In Table 7 we show that the elasticity of unskilled employment is on average slightly negative, although this conceals substantial heterogeneity, with substantial positive effects in many sectors and large negative effects for two industries, Wholesale and Retail, that together comprise more than a third of hours worked in our sample. By contrast, the elasticity of skilled employment is greater than or equal to one in every industry, with an employment-weighted average of 1.7. We return to these results when we calculate the employment effects of techies and trade in Table 12 below.

Our estimation procedure imposes that the distribution parameters $\alpha_N, \alpha_M$ and $\alpha_K$ 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, and the results are reasonable, given average shares in factor payments: $\alpha_M$ is the largest, followed by $\alpha_N$, while $\alpha_K$ is the smallest. Note that $\alpha_L$ and $\alpha_S$ are directly given by their corresponding average wage bill shares in (9), so they have no standard errors.

Our finding of $\sigma > 1$ contrasts with $\sigma < 1$ found by Doraszelski and Jaumandreu (2018) for Spanish manufacturing firms and Raval (2019) 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 (2019) 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 (2019) 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 $\varphi > 1$ 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. In the first row of Table 8 we find that the overall level of techies has a statistically significant effect on skill augmenting productivity $\omega_S$, but no significant effect on Hicks-neutral productivity $\omega_H$.[36] The second and third rows of the table report estimates when techies are broken down by their detailed job descriptions. ICT techies have a large and statistically significant effect on $\omega_H$, while R&D techies have a tiny and imprecisely estimated effect. Both ICT and R&D techies have large, precisely estimated and almost equal effects on $\omega_S$.

Turning to the effects of trade participation on productivity, we find that the effect of exporting on $\omega_H$ is very small and imprecisely estimated. 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. Slovenia in the 1990s was an emerging transition economy while France is a mature developed country, so our results are

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

consistent with the consensus in the literature that learning-by-exporting is found only in (some) developing countries. By contrast, exporting does have a positive and statistically significant effect on $\omega_S$.

We find no statistically significant effects of importing on $\omega_H$, but we do find positive effects on $\omega_S$, though the effect is less than half as big as the effect of exporting. Our finding that both importing and exporting raise $\omega_S$, combined with our estimates of $\widehat{\varphi} > 1$, implies that global engagement causes French firms to increase their skill intensity.

Table 9 reports additional results about the effect of importing on productivity, and allows us to see whether the country of import sourcing and/or whether importing inputs affects our conclusions. In columns 1 and 4, we add an indicator variable for imports of intermediate goods, so that the effect of importing intermediates is the sum of the effect of importing in general and the incremental effect of importing inputs. Introducing this channel shows that the importing effect on $\omega_S$ is no different when we break out intermediate inputs. Columns (2) and (5) of Table 9 consider an alternative split, with separate indicators for importing from high income countries and all other countries, and we find that the effect of importing on $\omega_S$ is about twice as high when imports come from low and medium income countries.[37]

Finally, in columns (3) and (6) we report the interactions of the indicator variables for income class and intermediate imports. Consistent with what we found in column (4), whether imports of inputs come from high or medium/low income countries makes no difference. 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. Our baseline estimate of 0.0143 in column 4 of Table 9 means that the impact effect of offshoring is to raise $\omega_S$ by 1.4 percent compared to firms who source only from France. In column (3), the coefficient on importing from other country has a positive sign on on $\omega_H$ but a low degree of significance. However, the effects of importing from both high income and other countries which is reported in the bottom panel of the table is not significant.

The final rows of Table 8 report the effects of other controls. Productivity is very persistent, with a coefficient on lagged productivity of about 0.8 for both $\omega_H$ and $\omega_S$. By contrast, lagged $\omega_H$ has a very small and imprecisely estimated effect on $\omega_S$, and vice versa. Firm size has no discernible effect on $\omega_H$ but a positive effect on $\omega_S$, suggesting that bigger firms become more skill intenisve relative to smaller firms. Firm age has negative effects on both $\omega_H$ and $\omega_S$, suggesting that older

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

firms have slower productivity growth.[38]

Because of the persistence of productivity, the techie and trade effects reported in Table 8 understate the long-run impact of these variables on cross-sectional productivity differences. Table 10 reports the associated steady-state effects, using equation (18). In addition to reporting the long-run effects and their standard errors, Table 10 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 11. The number *0.148* in the third row of column 3 of Table 10 means that compared to firms with no techies, in the long run firms with a lot of techies have $\omega_S$ which is 15 percent higher. This is a big effect in economic terms. It is also large relative to the variation in $\omega_S$, as 0.148 is equal to 17 percent of the $75^{th} - 25^{th}$ percentile range of $\omega_S$, which is 0.86 (second row, last column of Table 11). In column 4 of Table 10 we find that the scaled long-run effect is about twice as big for R&D techies as for ICT techies, a difference due to the fact that the $75^{th}$ percentile of R&D techies is about 2.5 times as large as for ICT techies. In other words, both ICT and R&D techies have a large long-run effect on $\omega_S$, but the greater expenditure on R&D techies means that they have a larger economic effect on $\omega_S$ across firms.

ICT techies also have a large long-run effect on $\omega_H$, as seen in column (2) of Table 10. The steady state effect of 0.175 is comparable in size to the effect of techies on $\omega_S$, and is about 12 percent of the interquartile range of $\omega_H$.

The long-run effects of importing and exporting on $\omega_S$ are also substantial, both in economic terms and relative to the variation in productivity. Since the trade variables are indicators, they are simple to interpret: the number 0.147 in the "Exporting" row of column (3) and (4) of Table 10 means that compared to firms that don't export, exporting causes $\omega_S$ to be 15 percent higher in the long run. Importing causes an effect which is about half as big, with firms that import having 7 percent higher steady state $\omega_S$. Thus, the causal effects of exporting are about the same as the causal effects of employing techies.

### 6.3 The skill bias of techies and trade

Since the estimated effects of techies on $\omega_H$ are statistically insignificant in Tables 8 through 10, here we focus on quantifying the employment effects of techies and trade through their effects on $\omega_S$. Applying equations (14)-(15), our quantification uses the industry-level estimates of $\eta$ and $\sigma$

[38] As a robustness check, in unreported results we re-estimated equations (16) and (17) 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.

from Table 5 together with the long-run effects reported in Table 10. To arrive at an economy-wide number, we compute the elasticities defined in equations (14)-(15) for each industry. We then construct an employment-weighted average of the industry elasticities, which we report in the first line of Table 12. In Panel A of Table 12 we multiply the elasticities by the estimated effects of techies and trade from Table 10. 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 12 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 12 shows that skill augmenting productivity raises $S$ and $S/L$, and has a tiny negative effect on $L$. The elasticities are big: a one percent increase in $\Omega_S$ raises skilled employment by 1.7 percent, reduces unskilled employment by less than a tenth of a percent, and increases skill intensity by 1.8 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 0.25 log points (28 percent) higher than firms with no techies, employment of $L$ that is 0.01 log points lower, and a skill intensity that is 0.26 log points (30 percentage points) higher. This effect is driven more by R&D techies than by ICT techies.

The last two lines of Panel A show that exporting and importing are also strongly pro-employment and skill biased: firms that export have employment of $S$ that is 0.26 log points (28 percent) higher than firms that do not, and firms that import have employment of $S$ that is 0.12 log points (13 percent) higher than firms that do not. Putting these two effects together, global engagement causes firms to increase $S$ by a substantial 0.37 log points (45 percent).

Table 12 is one of the bottom lines of our paper. For the first time in the literature, we have jointly estimated the firm-level labor demand effects of ICT, R&D, importing and exporting in a unified framework, which allows us to compare their importance. Table 12 shows that techies and trading 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 trading raise the skill premium.

### 6.4 Aggregating the effects of techies and trade

In this section we ask: 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, exporting and importing decisions from 2009 through 2012 by iterating forward using the estimated parameters of equations (16) and (17). 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 equations (14)-(15). Finally, we sum over all firms to get predicted aggregate skill intensity. These calculations take into account both within-firm adjustment and as changes in firm sizes, but exclude firm entry and exit (85% of employment is accounted by "continuous" 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 $100 \times S/L$. This measure increases in our sample of continuous firms by 2.7 percentage points (pp), from 16.7 percent in 2009 to 19.4 percent in 2013 (the numbers for the entire French private sector are very similar). Our first computation uses only the statistically significant direct effects of techies and trade on $\omega_S$ from column 3 of Table 8. Using these estimates, we calculate that techie employment decisions in 2009–2013 imply an increase in relative demand for skill of 0.40pp. The same exercise using exporting and importing decisions amounts to 0.50pp and 0.23pp increase in aggregate relative demand for skill, respectively. Adding these three effects together gives an increase of 1.13pp, almost half the total increase in the data.

When we evaluate firm-level ICT techie employment decisions in 2009–2013, we use the estimates from columns 2 and 4 of Table 8. We find a much larger implied increase in relative demand for skill amounting to 3.5pp, which is more than the actual increase of 2.7pp. This is driven by the large direct effect of ICT techies on $\omega_H$. In fact, 87% of the 3.5pp implied increase in aggregate relative demand for skill is driven by changes in firm sizes, holding constant firm-level skill intensities. Skill-intensive firms are those who are also more likely to employ techies, who cause greater employment growth. The greater importance of changes in firm composition versus within-firm adjustment is also found in De Loecker and Eeckhout (2018) for increases in average markups, and Autor et al. (2020) for the decline in the labor share, and Harrigan et al. (2021) for job polarization.

The calculations in this section do not take account of equilibrium conditions in the markets for goods or labor. As such, they should be interpreted as rough estimates of the relative demand effects of techies and trade on labor demand, and in particular their relative importance. Computation of the general equilibrium effects on employment and wages is beyond the scope of this paper, but any general equilibrium model should be consistent with the aggregate relative demand effects that we compute here.

### 6.5 Robustness

There are three novel elements of the empirical approach we have implemented in the previous sections. The first is our specification of the nested CES production function in equations (1) and (2), which allows for both Hicks-neutral and skill-augmenting technology differences across firms as well as two elasticities of substitution. The second is the way we treat techies, assuming that they affect output only through their lagged effect on productivity. Finally, our controlled Markov specification allows productivity to be an endogenous outcome of firm decisions. In this section, we consider the sensitivity of our conclusions to each of these elements.

#### 6.5.1 Hicks-neutral productivity only

A key message of our paper is that technological differences across firms have a skill-augmenting component $\omega_{Sft}$ as well as a Hicks-neutral component $\omega_{Hft}$. If our model is correct, then computing $\omega_{Hft}$ while ignoring variation in $\omega_{Sft}$ will bias the estimates of $\omega_{Hft}$: firms that are highly productive bcause of high $\omega_{Sft}$ will incorrectly be measured as having high $\omega_{Hft}$. Additionally, if techies and trade affect productivity through their effect on $\omega_{Sft}$, then ignoring that channel will lead to an over-estimate of the effect of techies on $\omega_{Hft}$.

To investigate this bias, we estimate a simplified production function which aggregates all non-techie labor and compute the implied $\omega_{Hft}$, preccisely as in GLZ. We then estimate the controlled Markov equation (16) for $\omega_{Hft}$ excluding $\omega_{Sft-1}$.

Results are reported in Table A2, where we find that techies, R&D techies, and imports have a large and precisely estimated effect on $\omega_{Hft}$. This contrasts with our baseline results in Table 8, where the effect of techies and imports on firm productivity comes mainly through their effect on $\omega_{Sft}$ rather than through an effect on $\omega_{Hft}$. We conclude from comparing Tables 8 and A2 that an accurate accounting of technology differences across firms requires estimating $\omega_{Sft}$ as well as $\omega_{Hft}$.

#### 6.5.2 Including techies in the production function

In Section 4.3.2 above (with details in Appendix A.2.3) we discuss the biases in productivity estimation that would result from assuming (as we do) that techie labor services in $t$ have no effect on output in $t$ when in fact they do. Here we consider an alternative specification where we include techies in the definition of employment when estimating the production function, and then compute implied Hicks-neutral $\omega_{Hft}$ and skill-augmenting $\omega_{Sft}$ productivity. Since techies are part of production in this specification, they should not affect productivity when we re-estimate the controlled Markov specification for productivity as given by equations (16) and (17). Table A3 reports the results of this exercise. The estimated effects of techies on both $\omega_{Hft}$ and $\omega_{Sft}$ are somewhat smaller than in our baseline estimates in Table 8, but the null hypothesis that the effects are zero can be rejected. We thus conclude that the data reject the model that techies affect output only through a contemporaneous effect on output. This does not establish that techies have no contemporaneous effect on output, but, as discussed in Section 4.3.2, a model where techies belong in the production function and, with a lag, also in the productivity process is not identified.

#### 6.5.3 Relative labor demand

In Section 6.2 we found large effects of techies and trade on skill-augmenting technology $\omega^S_{ft}$. What this means is that, conditional on wages, techies and trade raise firm-level employment of more-skilled relative to less-skilled workers. This interpretation follows from the relative skill demand equation (13) when $\varphi > 1$ (which is what we find for all industries in Table 5).

We can estimate the relationship between relative skill demand and techies and trade more directly. Re-writing equation (13) gives

$$(\varphi - 1)\,\omega_{Sft} = \ln\left(\frac{S}{L}\right)_{ft} + \varphi \ln\left(\frac{W_L}{W_S}\right)_{ft} + constant \tag{19}$$

Substituting equation (19) into equation (17) and moving the relative wage term to the righthand side of the equation gives

$$\ln\left(\frac{S}{L}\right)_{ft} = -\varphi \ln\left(\frac{W_S}{W_L}\right)_{ft} + \beta_{SH}\omega_{Hft-1} + \beta_{SS}\omega_{Sft-1} + \beta_{SZ}z_{ft-1} + constant + \xi_{Sft} \tag{20}$$

Equation (20) does not impose all of the structure of our model, but is nonetheless intuitive: it states that relative skill demand depends on relative wages, lagged productivity, and the lagged determinants of current skill-augmenting productivity.

We estimate equation (20) by two-stage least squares, using $\ln Z_{ft}$ defined in equation (11) as an instrument for $\ln\left(W_S/W_L\right)_{ft}$. We report the results in Table A4. As in our baseline specifications for skill-augmenting productivity (columns 3 and 4 in Table 8 and columns 4, 5 and 6 in Table 9), we find large and statistically significant positive effects of techies, importing and exporting. The most interesting result from Table A4 is the estimate of 3.1 for the elasticity of substitution between skilled and less-skilled labor, only slightly larger than the 2.8 average of industry-specific estimates reported in Table 5. While these parameter estimates cannot be directly compared to the more model-based estimates discussed in Section 6, it is reassuring that the implications are broadly consistent: the elasticity of substitution between more-skilled and less-skilled labor is substantially greater than 1, and techies and trade raise the relative demand for skill.

# 7 Conclusion

We show how firm-level decisions on R&D, ICT and trade affect firm-level productivity and its bias towards skilled workers. We do this by estimating firm-level nested CES production functions, which allows us to infer both Hicks-neutral and skill-augmenting technology differences. We use matched employer-employee data in manufacturing and non-manufacturing industries in France, from 2009 to 2013. The data has information on exporting, importing and technology adoption at the firm level, as well as detailed information on each worker's occupation. Our measure of technology adoption is firm-level employment of workers in technology and research related occupations, who we call "techies".

We find that techies, exporting and importing raise skill-biased productivity. In contrast, only ICT techies raise Hicks-neutral productivity. We show that both trade and employment of techies lead to greater employment of skilled workers without reducing employment of less skilled workers. The result for unskilled workers is surprising but easy to explain: the direct substitution effect away from unskilled labor is offset by the powerful employment-enhancing effect of greater productivity. When aggregating our firm-level estimates, we find that ICT has the largest effect on aggregate demand for skill, mostly through its effect on firm sizes, not through within-firm adjustment. These conclusions are based on firm-level effects, calculated holding product 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.

We develop a new methodology that allows us to identify the causal effects on productivity and employment of firms' decisions on trade and employment of techies through their effects on both Hicks-neutral and skill-augmenting technology differences. 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 international trade is an important question that is beyond the scope of this paper.

## 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.

Adao, R., M. Kolesár, and E. Morales (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics 134*(4), 1949–2010.

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., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics 135*(2), 645–709.

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.

Beaudry, P., D. A. Green, and B. M. Sand (2016). The great reversal in the demand for skill and cognitive tasks. *Journal of Labor Economics 34*(S1), S199–S247.

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

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.

Bloom, N., L. Garicano, R. Sadun, and J. Van Reenen (2014). The distinct effects of information technology and communication technology on firm organization. *Management Science 60*(12), 2859–2885.

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.

Brynjolfsson, E. and L. M. Hitt (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic perspectives 14*(4), 23–48.

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 J. Eeckhout (2018). Global market power. NBER Working Paper No. 24768.

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

Deming, D. J. and K. L. Noray (2018, September). Stem careers and technological change. Working Paper 25065, National Bureau of Economic Research.

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., L. Foster, J. Haltiwanger, and K. Troske (2004). Wage and productivity dispersion in United States manufacturing: The role of computer investment. *Journal of Labor Economics 22*(2), 397–429.

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

Fox, J. T. and V. Smeets (2011). Does input quality drive measured differences in firm productivity? *International Economic Review 52*(4), 961–989.

Gandhi, A., S. Navarro, and D. A. Rivers (2020). On the identification of gross output production functions. *Journal of Political Economy 128*(8), 2973–3016.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review 110*(8), 2586–2624.

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.

Hall, B. H., F. Lotti, and J. Mairesse (2013). Evidence on the impact of r&d and ict investments on innovation and productivity in italian firms. *Economics of Innovation and New Technology 22*(3), 300–328.

Hall, B. H., J. Mairesse, and P. Mohnen (2010). Measuring the returns to r&d. In *Handbook of the Economics of Innovation*, Volume 2, pp. 1033–1082. Elsevier.

Harrigan, J., A. Reshef, and F. Toubal (2021). The March of the Techies: Job Polarization Within and Between Firms. *Research Policy*.

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*. 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.

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.

León-Ledesma, M. A., P. McAdam, and A. Willman (2010). Identifying the elasticity of substitution with biased technical change. *The American Economic Review 100*(4), 1330–1357.

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.

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.

Mohnen, P., M. Polder, and G. van Leeuwen (2018). Ict, r&d and organizational innovation: Exploring complementarities in investment and production. Technical report, National Bureau of Economic Research.

Oberfield, E. and D. Raval (2020, 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. R. (2019). The micro elasticity of substitution and non-neutral technology. *The RAND Journal of Economics 50*(1), 147–167.

Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy 98*(5, Part 2), S71–S102.

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

Schweitzer, C. (2019). L'enquête R&D: mesurer l'effort de R&D des entreprises, au-delà du crédit d'impôt recherche. INSEE Références, édition 2019 – Dossier.

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.

Table 1: Covariates of techies

| | $I\,(techies > 0)$ |
|---|---|
| Log revenue | 0.120 |
| | (0.001) |
| Exporter dummy | 0.272 |
| | (0.005) |
| Importer dummy | 0.275 |
| | (0.004) |
| Manager wage bill share | 0.142 |
| | (0.011) |
| Obs. | 568,650 |

Notes to Table 1: Each entry is the weighted least squares coefficient in a regression of the row variable on an indicator equal to 1 if the firm employs any techies. All regressions include industry × year fixed effects. Regressions weighted by firm employment. Standard errors clustered by firm in parentheses. All estimated coefficients are statistically significant at the 0.01 level.

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 |
| 2010 | 38 | 62 |
| 2013 | 39 | 61 |
| | Variation across industries in 2009 | |
| std. dev. | 31.0 | 31.0 |
| Min | 17.4 | 4.2 |
| Max | 95.8 | 82.6 |

Notes to Table 3: Wage bill shares within techies in the estimation sample. Wage bill shares sum to 100 across the two categories ICT and R&D.

Table 4: Estimation sample, 2010-2013

| Industry | Obs. | Obs. (%) | Firms | Firms (%) | Revenue (%) | Hours (%) |
|---|---|---|---|---|---|---|
| Food, beverage, tobacco | 18,224 | 3.2 | 6,072 | 3.1 | 8.8 | 6.8 |
| Textiles, wearing apparel | 6,771 | 1.2 | 2,201 | 1.1 | 1.3 | 1.9 |
| Wood, paper products | 13,810 | 2.4 | 4,491 | 2.3 | 2.5 | 3.6 |
| Chemical products | 4,075 | 0.7 | 1,202 | 0.6 | 5.0 | 3.2 |
| Rubber and plastic | 11,257 | 2.0 | 3,469 | 1.8 | 4.3 | 5.0 |
| Basic metal and fabricated metal | 23,908 | 4.2 | 7,457 | 3.9 | 4.7 | 7.2 |
| Electrical equipment | 3,242 | 0.6 | 1,001 | 0.5 | 1.9 | 2.2 |
| Machinery and equipment | 8,111 | 1.4 | 2,447 | 1.3 | 3.1 | 3.9 |
| Other manufacturing | 19.210 | 3.4 | 6,275 | 3.3 | 2.6 | 4.5 |
| Construction | 108,919 | 19.2 | 38,389 | 19.9 | 9.0 | 14.1 |
| Wholesale | 122,317 | 21.5 | 40,032 | 20.7 | 37.6 | 21.8 |
| Retail | 145,271 | 25.5 | 48,680 | 25.2 | 14.0 | 12.6 |
| Accommodation and food services | 64,172 | 11.3 | 24,212 | 12.5 | 2.2 | 6.4 |
| Administrative and support activities | 19,363 | 3.4 | 7,033 | 3.6 | 3.1 | 6.8 |
| *Total* | *568,650* | *100* | *192,961* | *100* | *100* | *100* |

Notes to Table 4: We lose 9.8% of total revenue and 9.2% of total hours due to dropping the sectors of coke and refined petroleum, pharmaceutical products, computer electornic, transport equipment, publishing and broadcasting and transportation and storage.

Table 5: Production function estimates

| Industry | $\alpha_L$ | $\alpha_S$ | $\alpha_N$ | $\alpha_M$ | $\alpha_K$ | $\sigma$ | $\varphi$ | $\eta$ | obs | firms |
|---|---|---|---|---|---|---|---|---|---|---|
| Food, beverage, tobacco | 0.787 | 0.213 | 0.251 | 0.537 | 0.212 | 1.875 | 2.776 | -5.91 | 19,307 | 6,289 |
| | | | *0.003* | *0.006* | *0.009* | *0.098* | *0.196* | *0.27* | | |
| Textiles, apparel | 0.731 | 0.269 | 0.419 | 0.504 | 0.077 | 1.483 | 1.986 | -3.14 | 7,164 | 2,272 |
| | | | *0.007* | *0.008* | *0.015* | *0.156* | *0.196* | *0.10* | | |
| Wood, paper products | 0.743 | 0.257 | 0.386 | 0.390 | 0.224 | 1.088 | 1.585 | -3.83 | 14,616 | 4,638 |
| | | | *0.008* | *0.008* | *0.015* | *0.044* | *0.156* | *0.17* | | |
| Chemical products | 0.664 | 0.336 | 0.213 | 0.510 | 0.277 | 1.344 | 1.949 | -5.84 | 4,316 | 1,250 |
| | | | *0.005* | *0.013* | *0.019* | *0.069* | *0.298* | *0.67* | | |
| Rubber & plastic | 0.764 | 0.236 | 0.300 | 0.521 | 0.179 | 1.676 | 2.102 | -4.02 | 11,883 | 3,579 |
| | | | *0.005* | *0.008* | *0.013* | *0.124* | *0.204* | *0.16* | | |
| Basic & fabricated metal | 0.773 | 0.227 | 0.373 | 0.321 | 0.307 | 1.263 | 1.493 | -4.25 | 25,347 | 7,717 |
| | | | *0.005* | *0.004* | *0.010* | *0.031* | *0.062* | *0.16* | | |
| Electrical equipment | 0.704 | 0.296 | 0.275 | 0.491 | 0.234 | 1.238 | 2.111 | -5.22 | 3,425 | 1,026 |
| | | | *0.009* | *0.016* | *0.024* | *0.063* | *0.393* | *0.66* | | |
| Machinery & equipment | 0.694 | 0.306 | 0.281 | 0.446 | 0.274 | 1.075 | 2.397 | -5.64 | 8,625 | 2,532 |
| | | | *0.007* | *0.012* | *0.019* | *0.024* | *0.191* | *0.56* | | |
| Other manufacturing | 0.718 | 0.282 | 0.361 | 0.310 | 0.329 | 1.170 | 2.522 | -5.04 | 20,337 | 6,515 |
| | | | *0.008* | *0.007* | *0.016* | *0.031* | *0.172* | *0.39* | | |
| Construction | 0.743 | 0.257 | 0.414 | 0.369 | 0.217 | 1.505 | 2.433 | -3.09 | 115,537 | 39,884 |
| | | | *0.003* | *0.003* | *0.007* | *0.030* | *0.056* | *0.05* | | |
| Wholesale | 0.634 | 0.366 | 0.158 | 0.708 | 0.134 | 1.370 | 3.677 | -10.94 | 129,894 | 41,557 |
| | | | *0.001* | *0.003* | *0.003* | *0.027* | *0.092* | *0.36* | | |
| Retail | 0.671 | 0.329 | 0.145 | 0.749 | 0.106 | 1.382 | 3.316 | -8.94 | 154,133 | 50,483 |
| | | | *0.000* | *0.002* | *0.003* | *0.030* | *0.059* | *0.21* | | |
| Accommodation and food | 0.777 | 0.223 | 0.449 | 0.286 | 0.265 | 1.956 | 1.982 | -3.99 | 67,903 | 25,012 |
| | | | *0.007* | *0.004* | *0.011* | *0.080* | *0.065* | *0.16* | | |
| Admin & support | 0.761 | 0.239 | 0.531 | 0.070 | 0.399 | 1.989 | 4.056 | -5.95 | 20,937 | 7,326 |
| | | | *0.012* | *0.002* | *0.014* | *0.057* | *0.178* | *0.56* | | |
| hours weighted average | 0.71 | 0.29 | 0.30 | 0.49 | 0.21 | 1.48 | 2.78 | -6.46 | | |

Notes to Table 5: Each row reports GMM estimates of equations (6) and (10) for a given industry. Observations weighted by firm employment. Reported parameters are functions of the parameters of (6). Parameters $\alpha_L$ and $\alpha_S$ are computed directly from the data, see text for details. Standard errors clustered by firm in italics. All estimates are significantly different from zero at the 0.01 level. The hours weighted averages in the final final row are for descriptive purposes only.

Table 6: Test statistics for production function estimates

| Industry | (1) $H_0 : \sigma = 1$ | (2) $H_0 : \varphi = 1$ | (3) $\varphi - \sigma$ | (4) *J(1)* |
|---|---|---|---|---|
| Food, beverage, tobacco | *9.0* | *9.1* | 0.901 | *9.9* |
| | | | *4.4* | |
| Textiles, apparel | *3.1* | *5.0* | 0.503 | *2.8* |
| | | | *2.2* | |
| Wood, paper products | *2.0* | *3.8* | 0.497 | *0.5* |
| | | | *3.2* | |
| Chemical products | *5.0* | *3.2* | 0.604 | *6.1* |
| | | | *2.2* | |
| Rubber & plastic | *5.5* | *5.4* | 0.426 | *3.7* |
| | | | *2.2* | |
| Basic & fabricated metal | *8.4* | *7.9* | 0.230 | *19.9* |
| | | | *4.1* | |
| Electrical equipment | *3.8* | *2.8* | 0.873 | *2.6* |
| | | | *2.3* | |
| Machinery & equipment | *3.2* | *7.3* | 1.322 | *7.3* |
| | | | *6.9* | |
| Other manufacturing | *5.4* | *8.9* | 1.352 | *3.5* |
| | | | *8.0* | |
| Construction | *16.7* | *25.6* | 0.928 | *64.6* |
| | | | *14.8* | |
| Wholesale | *13.9* | *29.1* | 2.307 | *3.7* |
| | | | *24.4* | |
| Retail | *12.8* | *39.5* | 1.934 | *33.4* |
| | | | *30.4* | |
| Accommodation and food | *12.0* | *15.2* | 0.026 | *37.0* |
| | | | *0.3* | |
| Admin & support | *17.3* | *17.2* | 2.067 | *43.1* |
| | | | *10.7* | |

Notes to Table 6: Columns (1) and (2) report $t$-statistics for the null hypotheses that the elasticities of substitution $\sigma$ and $\varphi$ in the production function are equal to one. Column (3) reports the point estimate of $\varphi - \sigma$ from Table 5 and, in italics, the $t$-statistic for $H_0 : \varphi - \sigma = 0$. Column (4) reports Hansen's $\chi^2$ $J$ statistic test of overidentification.

Table 7: Employment elasticities of skill-augmenting productivity

| | $S$ | $L$ | $S/L$ |
|---|---|---|---|
| Food, beverage, tobacco | 1.80 | 0.02 | 1.78 |
| | *0.16* | *0.04* | *0.20* |
| Textiles, apparel | 1.04 | 0.05 | 0.99 |
| | *0.15* | *0.05* | *0.20* |
| Wood, paper products | 0.73 | 0.14 | 0.59 |
| | *0.12* | *0.04* | *0.16* |
| Chemical products | 1.07 | 0.12 | 0.95 |
| | *0.22* | *0.09* | *0.30* |
| Rubber & plastic | 1.17 | 0.07 | 1.10 |
| | *0.17* | *0.04* | *0.20* |
| Basic & fabricated metal | 0.69 | 0.20 | 0.49 |
| | *0.05* | *0.02* | *0.06* |
| Electrical equipment | 1.18 | 0.07 | 1.11 |
| | *0.29* | *0.11* | *0.39* |
| Machinery & equipment | 1.38 | -0.01 | 1.40 |
| | *0.14* | *0.07* | *0.19* |
| Other manufacturing | 1.53 | 0.01 | 1.52 |
| | *0.13* | *0.06* | *0.17* |
| Construction | 1.36 | -0.07 | 1.43 |
| | *0.04* | *0.01* | *0.06* |
| Wholesale | 2.39 | -0.29 | 2.68 |
| | *0.06* | *0.04* | *0.09* |
| Retail | 2.04 | -0.27 | 2.32 |
| | *0.04* | *0.02* | *0.06* |
| Accommodation and food | 1.18 | 0.20 | 0.98 |
| | *0.05* | *0.02* | *0.06* |
| Admin & support | 3.07 | 0.01 | 3.06 |
| | *0.15* | *0.07* | *0.18* |
| Weighted average | 1.71 | -0.07 | 1.78 |

Notes to Table 7: Elasticities computed using equations (14)-(15) and estimates from Table 5. Standard errors in italics. Weighted average elasticities computed using industry hours shares from Table 4.

Table 8: Techie and trade effects on productivity

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Hicks neutral $\omega_{ft}^H$ | | Skill augmenting $\omega_{ft}^S$ | |
| Techies | 0.4533 | | 0.3501*** | |
| | (0.397) | | (0.043) | |
| Techies : ICT | | 0.9137** | | 0.3519*** |
| | | (0.4558) | | (0.051) |
| Techies : R&D | | 0.0465 | | 0.3485*** |
| | | (0.6165) | | (0.061) |
| Exporting | -0.0130 | -0.0124 | 0.0305*** | 0.0305*** |
| | (0.027) | (0.0269) | (0.0032) | (0.0032) |
| Importing | 0.0236 | 0.0237 | 0.0139*** | 0.0139*** |
| | (0.035) | (0.0349) | (0.0032) | (0.0032) |
| lagged $\omega_{ft}^H$ | 0.8108*** | 0.8107*** | 0.0049* | 0.0049* |
| | (0.025) | (0.0250) | (0.0026) | (0.0026) |
| lagged $\omega_{ft}^S$ | 0.0486 | 0.0479 | 0.7939*** | 0.7939*** |
| | (0.034) | (0.0337) | (0.0054) | (0.0054) |
| firm size | 0.0005 | 0.0005 | 0.0004*** | 0.0004*** |
| | (0.0004) | (0.0004) | (0.0001) | (0.0001) |
| firm age | -0.0383** | -0.0382** | -0.0064*** | -0.0064*** |
| | (0.0166) | (0.0165) | (0.0013) | (0.0013) |

Notes to Table 8: Weighted least squares estimation of equations (16) and (17), pooled across industries, with industry $\times$ year fixed effects. Observations weighted by firm employment. Bootstrapped standard errors clustered by firm in parentheses. Sample: 568,650 observations on 192,961 firms during 2009-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table 9: Techie and trade effects on productivity, importing detail

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | Hicks neutral $\omega_{ft}^{H}$ | | | Skill augmenting $\omega_{ft}^{S}$ | | |
| Techies | 0.4544 | 0.4515 | 0.4596 | 0.3499*** | 0.3399*** | 0.3396*** |
| | (0.398) | (0.399) | (0.399) | (0.043) | (0.043) | (0.043) |
| Exporting | -0.0123 | -0.0092 | -0.0083 | 0.0303*** | 0.0265*** | 0.0265*** |
| | (0.026) | (0.024) | (0.023) | (0.0032) | (0.0032) | (0.0032) |
| Importing | 0.0370** | | | 0.0112** | | |
| | (0.017) | | | (0.0048) | | |
| Importing inputs | -0.0158 | | | 0.0031 | | |
| | (0.036) | | | (0.0047) | | |
| Importing high income | | 0.0085 | 0.0232 | | 0.0087*** | 0.0096* |
| | | (0.031) | (0.020) | | (0.0032) | (0.0049) |
| Importing other income | | 0.0080 | 0.0319* | | 0.0178*** | 0.0169*** |
| | | (0.029) | (0.019) | | (0.0036) | (0.0051) |
| Imp. inputs high income | | | -0.0180 | | | -0.0010 |
| | | | (0.029) | | | (0.0049) |
| Imp. inputs other income | | | -0.0318 | | | 0.0014 |
| | | | (0.033) | | | (0.0059) |
| Linear combinations of estimates | | | | | | |
| Importing + imp. inputs | 0.0212 | | | 0.0143*** | | |
| | (0.040) | | | (0.0033) | | |
| High inc., imp. + imp. inputs | | | 0.00526 | | | 0.0086** |
| | | | (0.035) | | | (0.0034) |
| Other inc., imp. + imp. inputs | | | 0.0001 | | | 0.0183*** |
| | | | (0.036) | | | (0.0042) |
| Other controls | Yes | Yes | Yes | Yes | Yes | Yes |

Notes to Table 9: Weighted least squares estimation of equations (16) and (17), pooled across industries, with industry × 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. Observations weighted by firm employment. Bootstrapped standard errors clustered by firm in parentheses. Sample: 568,650 observations on 192,961 firms during 2009-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table 10: Long run techie and trade effects on productivity

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Hicks neutral $\omega_{ft}^H$ | | Skill augmenting $\omega_{ft}^S$ | |
| Techies | 2.851 | | 1.767*** | |
| | (2.30) | | (0.201) | |
| | *0.239* | | *0.148* | |
| Techies : ICT | | 5.291** | | 1.834*** |
| | | (2.65) | | (0.235) |
| | | *0.175* | | *0.061* |
| Techies : R&D | | 0.678 | | 1.707*** |
| | | (3.09) | | (0.294) |
| | | *0.054* | | *0.137* |
| Exporting | -0.0308 | -0.0282 | 0.147*** | 0.147*** |
| | (0.138) | (0.137) | (0.015) | (0.015) |
| Importing | 0.143 | 0.143 | 0.0707*** | 0.0707*** |
| | (0.192) | (0.192) | (0.015) | (0.015) |

Notes to Table 10: Effects are long-run steady state effects on productivity given by equation (18), based on results in Table 8. 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 11, are reported in *italics*. Sample: 568,650 observations on 192,961 firms during 2010-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table 11: Summary statistics for second stage regressions

| | Mean | Std Dev | *p25* | *p50* | *p75* | *p75-p25* |
|---|---|---|---|---|---|---|
| Hicks neutral productivity $\omega_H$ | 0 | 1.95 | -0.72 | 0.08 | .79 | 1.514 |
| Skill augmenting productivity $\omega_S$ | 0 | 0.83 | -0.42 | .01 | 0.43 | 0.86 |
| Techie share of wage bill | 0.063 | 0.069 | 0.019 | 0.041 | 0.084 | 0.065 |
| Techies: ICT | 0.031 | 0.052 | 0.094 | 0.018 | 0.033 | 0.024 |
| Techies: R&D | 0.058 | 0.060 | 0.016 | 0.039 | 0.080 | 0.062 |
| Exporting | 0.44 | 0.50 | 0 | 0 | 1 | 1 |
| Importing | 0.48 | 0.50 | 0 | 0 | 1 | 1 |

Notes to Table 11: Summary statistics for variables used to estimate equations (16) and (17), weighted by firm employment. Statistics for the three 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: 568,650 observations on 192,961 firms during 2010-2013.

Table 12: Employment effects of skill augmenting technology differences

| | $S$ | $L$ | $S/L$ |
|---|---|---|---|
| Elasticities | 1.71 | -0.07 | 1.78 |
| A. Elasticities $\times$ second stage estimates | | | |
| Techies | 3.02 | -0.12 | 3.14 |
| Techies : ICT | 3.13 | -0.12 | 3.26 |
| Techies : R&D | 2.92 | -0.12 | 3.03 |
| Exporting | 0.25 | -0.01 | 0.26 |
| Importing | 0.12 | 0.00 | 0.13 |
| B. Elasticities $\times$ second stage estimates $\times\, p75$ | | | |
| Techies | 0.25 | -0.01 | 0.26 |
| Techies : ICT | 0.10 | 0.00 | 0.11 |
| Techies : R&D | 0.23 | -0.01 | 0.24 |

Notes to Table 12: 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 equations (14)-(15) and the estimates from Table 5. Panel A multiplies the elasticities by the corresponding estimates from columns (3) and (4) of Table 10. Panel B multiplies the panel A numbers by the 75th percentile of the distribution of techies, from Table 11.

# 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.[39] 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.[40]

### A.1.2 Balance sheet data: FARE

Firm-level balance sheet information is reported in an INSEE dataset called FARE.[41] 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

[39]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.

[40]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.

[41]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 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 ($K_{i,s,t}^{BV}$) 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 $p_{s,t+1}^{K} = p_{s,t-T}^{I}$. The average age of capital is computed using the share of depreciated capital, $DK_{i,s,t}^{BV}$ in the capital stock at historical cost.

$$T = \frac{DK_{i,s,t}^{BV}}{K_{i,s,t}^{BV}} \times \widetilde{A}$$

where

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

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[42]. We use the Broad Economic Categories (BEC rev. 4) classification

[42] 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 Estimating equation

Using equations (1) and (3), revenue is given by

$$R_{ft} = e^{u_{ft}} P_{ft} Y_{ft} = e^{u_{ft}} A_t \Omega_{Hft}^{\frac{\eta+1}{\eta}} \left[ \alpha_N N_{ft}^{\gamma} + \alpha_K K_{ft}^{\gamma} + \alpha_M M_{ft}^{\gamma} \right]^{\frac{\eta+1}{\eta\gamma}}, \tag{A.21}$$

Equation (A.21) contains three unobservable shocks ($u_{ft}$, $\omega_{Hft}$ and $\omega_{Sft}$) and one unobservable variable $M_{ft}$. $N_{ft}$ can be constructed from observables and parameters. Before the revenue shock is realized, firms choose inputs optimally to maximize *ex ante* profit. Using (2) to substitute for $N_{ft}$ in (A.21) yields the following first-order conditions for this problem,

$$\alpha_N N_{ft}^{\gamma-1} X_{ft} = P_{Nft} \tag{A.22}$$

$$\alpha_L L_{ft}^{\rho-1} \left[ \alpha_L L_{ft}^{\rho} + \alpha_S \Omega_{Sft}^{\rho} (S_{ft})^{\rho} \right]^{\frac{1-\rho}{\rho}} \alpha_N N_{ft}^{\gamma-1} X_{ft} = W_{Lft} \tag{A.23}$$

$$\alpha_S \Omega_{Sft}^{\rho} (S_{ft})^{\rho-1} \left[ \alpha_L L_{ft}^{\rho} + \alpha_S \Omega_{Sft}^{\rho} (S_{ft})^{\rho} \right]^{\frac{1-\rho}{\rho}} \alpha_N N_{ft}^{\gamma-1} X_{ft} = W_{Sft} \tag{A.24}$$

$$\alpha_M M_{ft}^{\gamma-1} X_{ft} = P_{Mft} \tag{A.25}$$

where $X_{ft} = \frac{\eta+1}{\eta} A_t \Omega_{Hft}^{\frac{\eta+1}{\eta}} \left[ \alpha_N N_{ft}^{\gamma} + \alpha_K K_{ft}^{\gamma} + \alpha_M M_{ft}^{\gamma} \right]^{\delta}$, $\delta = \frac{\eta(1-\gamma)+1}{\gamma\eta}$. Dividing (A.22) by (A.25) and solving for $M_{ft}$ gives equation (4), and dividing (A.23) by (A.24) and solving for $\Omega_{Sft}$ gives (5). Using (5) in (2) gives

$$N_{ft} = L_{ft} \left( \frac{\alpha_L}{\lambda_{ft}} \right)^{\frac{1}{\rho}} \tag{A.26}$$

where $\lambda_{ft} = \frac{E_{ft}^L}{E_{ft}^L + E_{ft}^S} = \frac{E_{ft}^L}{E_{ft}^N}$ is the share of unskilled labor in the wage bill.

If $N_{ft}$ were data, using (4) to eliminate $M_{ft}$ and (A.23) to eliminate $\Omega_{Hft}$ from (A.21) would yield the following estimating equation from GLZ (their equation 8, see their Appendix A.1 for derivation),

$$\ln R_{ft} = \ln \frac{\eta}{\eta+1} + \ln \left[ E_{ft}^{M} + E_{ft}^{N} \left(1 + \frac{\alpha_K}{\alpha_N}\right) \left(\frac{K_{ft}}{N_{ft}}\right)^{\gamma} \right] + u_{ft} \tag{A.27}$$

$N_{ft}$ depends on parameters and is thus not data, so we substitute (A.26) into (A.27) which yields our estimating equation (6).

To solve for Hicks-neutral productivity $\omega_{Hft}$, substitute (4) and (5) into (A.22), multiply both sides by $N_{ft}$, and take logs to obtain equation (12). An alternative expression for $\omega_{Hft}$ is obtained using (A.23) rather than (A.22),

$$\begin{aligned} \omega_{Hft} = \frac{\eta}{\eta+1} &\left[ \ln \left[ \left( \frac{\eta}{1+\eta} \right) \left( A_t \alpha_L^{\frac{\gamma}{\rho}} \alpha_N \right)^{-1} \right] + \frac{\gamma}{\rho} \ln E_{Lft} - \gamma \ln L_{ft} + \frac{\rho - \gamma}{\rho} \ln E_{ft}^{N} \right] \\ &+ \frac{\eta(\gamma-1)-1}{(1+\eta)\gamma} \ln \left[ \alpha_L^{\frac{\gamma}{\rho}} \alpha_N \left( \frac{E_{ft}^{N} + E_{ft}^{M}}{E_{ft}^{N}} \right) L_{ft}^{\gamma} \left( \frac{E_{ft}^{N}}{E_{ft}^{L}} \right)^{\frac{\gamma}{\rho}} + \alpha_K K_{ft}^{\gamma} \right] \end{aligned} \tag{A.28}$$

### A.2.2 Calculation of standard errors

For the first stage GMM estimates reported in Tables 5 through 7, standard errors are clustered by firm. For the second stage estimates reported in Tables 8 through 10 and A3 through A4, covariance matrices are computed by bootstrapping, with firms drawn without replacement, which is equivalent to clustering by firm. In all bootstraps we use 800 replications (decreasing the number of replications to 400 had virtually no effect on inference).

As explained in Section 4.1.1, the derivation of our production function estimator requires that that $\sigma \neq 1$ and $\varphi \neq 1$. Our point estimates satisfy these restrictions (see Table 6), but in a small number of our bootstrap draws the value of $\sigma$ is so close to one, which is equivalent to $\gamma$ close to zero, that some of the values of $\widehat{\omega}_{Hft}$ take on very extreme values (to see how this happens, note that equation (12) which defines $\omega_{Hft}$ includes a term with an exponent $\delta$ which goes to infinity as $\gamma$ goes to zero). To prevent our covariance estimates from being distorted by this small number of bootstrap draws, we drop all bootstraop draws $b$ where $|\widehat{\gamma}_b| < 0.01$.

#### 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". In general this is an intractable problem to analyze, so here we make two empirically relevant simplifications. First, we suppose that $\varphi = \sigma$ (which is not far from what we find in Table 5). Second, 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 $\tau_{ft}$ as the ratio of the techie to manager wage bill in $S$,

$$E^S_{ft} = E^B_{ft} + E^T_{ft} = (1 + \tau_{ft})\, E^B_{ft}$$

When $\varphi = \sigma$, the expression for Hicks-neutral productivity given by equation (A.28) simplifies to

$$\omega_{Hft} = \frac{\eta}{\eta + 1}\left[\ln\left(\frac{\eta}{1+\eta}\frac{E_{Lft}}{A_t\alpha_L\alpha_N L^{\gamma}_{ft}}\right)\right] + \beta\ln\left[\alpha_L\alpha_N\left(\frac{E^L_{ft} + E^S_{ft} + E^M_{ft}}{E^L_{ft}}\right)L^{\gamma}_{ft} + \alpha_K K^{\gamma}_{ft}\right]$$

where $\beta = \frac{\eta(\gamma-1)-1}{(1+\eta)\gamma}$ is negative as long as $|\eta| > \sigma$, a condition which holds in our estimates (see Table 5). Under the assumption that our model is correct, we can write true Hicks-neutral productivity as

$$\omega^1_{Hft} = \theta_{ft} + \beta \ln \left\{ \alpha_L \alpha_N \left( \frac{E^L_{ft} + E^B_{ft} + E^M_{ft}}{E^L_{ft}} \right) L^\gamma_{ft} + \alpha_K K^\gamma_{ft} \right\}$$

where $\theta_{ft} = \frac{\eta}{\eta+1} \left[ \ln \left( \frac{\eta}{1+\eta} \frac{E_{Lft}}{A_t \alpha_L \alpha_N L^\gamma_{ft}} \right) \right]$. Under the assumption that the orthodox model is correct, the term $E^B_{ft}$ in this expression is multiplied by $(1 + \tau_{ft})$, giving

$$\omega^2_{Hft} = \theta_{ft} + \beta \ln \left\{ \alpha_L \alpha_N \left( \frac{E^L_{ft} + (1 + \tau_{ft})\, E^B_{ft} + E^M_{ft}}{E^L_{ft}} \right) L^\gamma_{ft} + \alpha_K K^\gamma_{ft} \right\}$$

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

$$\begin{aligned} \omega^1_{Hft} - \omega^2_{Hft} &= \beta \ln \left\{ \alpha_L \alpha_N \left( \frac{E^L_{ft} + E^B_{ft} + E^M_{ft}}{E^L_{ft}} \right) L^\gamma_{ft} + \alpha_K K^\gamma_{ft} \right\} \\ &\quad - \beta \ln \left\{ \alpha_L \alpha_N \left( \frac{E^L_{ft} + (1 + \tau_{ft})\, E^B_{ft} + E^M_{ft}}{E^L_{ft}} \right) L^\gamma_{ft} + \alpha_K K^\gamma_{ft} \right\} \end{aligned}$$

This expression is strictly positive and increasing in the techie share $\tau_{ft}$. The intuition is simple: the larger is $\tau_{ft}$, 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, from (A.23) and (A.24) we can write true skill-augmenting productivity as

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

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

$$\omega^2_{Sft} = l_{ft} - b_{ft} - \delta_{ft} + \frac{\tau_{ft}}{\rho} + \frac{1}{\rho} log \left( \frac{\alpha_L E^B_{ft}}{\alpha_S E^L_{ft}} \right)$$

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

$$\omega^1_{Sft} - \omega^2_{Sft} = \delta_{ft} - \frac{\tau_{ft}}{\rho} + \frac{1}{\rho} log \left( \frac{\alpha_S}{\alpha_B} \right)$$

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

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

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

### A.2.4 Output and employment elasticities

The unit cost function corresponding to the nested CES production function given by equations (1) and (2) is

$$\begin{aligned} C &= \frac{1}{\Omega_H}\left[\alpha_N^\sigma\left[\alpha_L^\varphi w_L^{1-\varphi} + \alpha_S^\varphi\left(\frac{w_S}{\Omega_S}\right)^{1-\varphi}\right]^{\frac{1-\sigma}{1-\varphi}} + \alpha_K^\sigma r_K^{1-\sigma} + \alpha_M^\sigma p_M^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \\ &= \frac{1}{\Omega_H}\left[\alpha_N^\sigma w_N^{1-\sigma} + \alpha_K^\sigma r_K^{1-\sigma} + \alpha_M^\sigma p_M^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \\ &= \frac{1}{\Omega_H} X^{\frac{1}{1-\sigma}} \end{aligned} \tag{A.29}$$

where $w_N = \left[\alpha_L^\varphi w_L^{1-\varphi} + \alpha_S^\varphi\left(\frac{w_S}{\Omega_S}\right)^{1-\varphi}\right]^{\frac{1}{1-\varphi}}$ and $X = \left[\alpha_N^\sigma w_N^{1-\sigma} + \alpha_K^\sigma r_K^{1-\sigma} + \alpha_M^\sigma p_M^{1-\sigma}\right]$. The unit factor demands can be obtained using Shepard's lemma,

$$L_1 = \frac{\partial C}{\partial w_N} \times \frac{\partial w_N}{\partial w_L} = \frac{\alpha_N^\sigma \alpha_L^\varphi w_L^{-\varphi} w_N^{\varphi-\sigma}}{\Omega_H} X^{\frac{\sigma}{1-\sigma}} \tag{A.30}$$

$$S_1 = \frac{\partial C}{\partial w_N} \times \frac{\partial w_N}{\partial w_S} = \frac{\alpha_N^\sigma \alpha_S^\varphi \Omega_S^{\varphi-1} w_S^{-\varphi} w_N^{\varphi-\sigma}}{\Omega_H} X^{\frac{\sigma}{1-\sigma}} \tag{A.31}$$

$$K_1 = \frac{\partial C}{\partial r_K} = \frac{\alpha_K^\sigma r_K^{-\sigma}}{\Omega_H} X^{\frac{\sigma}{1-\sigma}} \tag{A.32}$$

$$M_1 = \frac{\partial C}{\partial p_M} = \frac{\alpha_M^\sigma p_M^{-\sigma}}{\Omega_H} X^{\frac{\sigma}{1-\sigma}} \tag{A.33}$$

Unit skill intensity is

$$\left(\frac{S}{L}\right)_1 = \Omega_S^{\varphi-1}\left(\frac{\alpha_L w_S}{\alpha_S w_L}\right)^{-\varphi}$$

For an inverse demand curve of the form $P = AQ^{\frac{1}{\eta}}$, revenue is $AQ^{\frac{\eta+1}{\eta}}$ and the marginal revenue = marginal cost condition is $\frac{\eta+1}{\eta}AQ^{\frac{1}{\eta}} = C$. Solving for profit maximizing output $Q$ and revenue $R$ gives

$$Q = BC^{-\eta} = B\Omega_H^{-\eta} X^{\frac{\eta}{1-\sigma 1}}$$

$$R = DC^{1-\eta} = D\Omega_H^{-(\eta+1)} X^{\frac{\eta+1}{\sigma-1}}$$

where $B = \left(\frac{\eta+1}{\eta}A\right)^{-\eta}$, $\quad D = AB^{\frac{\eta+1}{\eta}}$. Increasing marginal costs leads to proportionately increasing prices (because of the constant markup) and thus a decline in sales with elasticity $\eta < 0$ and a decline in revenue with elasticity $\eta+1 < 0$. Hicks-neutral TFP improvements raise sales and revenue with elasticities of $-\eta > 0$ and $-(\eta+1) > 0$ respectively. Multiplying the unit factor demands from (A.30) and (A.31) by the quantity demanded gives the full Marshallian factor demands,

$$L = L_1 Q = \frac{\alpha_N^\sigma \alpha_L^\varphi w_L^{-\varphi} w_N^{\varphi-\sigma}}{\Omega_H} X^{\frac{\sigma}{1-\sigma}} \times B\Omega_H^{-\eta} X^{\frac{\eta}{1-\sigma}}$$

$$S = S_1 Q = \frac{\alpha_N^\sigma \alpha_S^\varphi \Omega_S^{\varphi-1} w_S^{-\varphi} w_N^{\varphi-\sigma}}{\Omega_H} X^{\frac{\sigma}{1-\sigma}} \times B\Omega_H^{-\eta} X^{\frac{\eta}{1-\sigma}}$$

Taking logs of these factor demands and collecting constant parameters in $k$ gives

$$\begin{aligned}
\ln L &= k_L + (-\eta-1)\ln\Omega_H - \varphi\ln w_L + (\varphi-\sigma)\ln w_N + \left(\frac{\eta+\sigma}{1-\sigma}\right)\ln X \\
\ln S &= k_S + (-\eta-1)\ln\Omega_H - \varphi\ln w_S + (\varphi-\sigma)\ln w_N + \left(\frac{\eta+\sigma}{1-\sigma}\right)\ln X + (\varphi-1)\ln\Omega_S
\end{aligned} \tag{A.34}$$

To understand how these Marshallian factor demands respond to changes in technology, it is first necessary to calculate cost shares and elasticities. First we calculate $\theta_N, \theta_K$, and $\theta_M$ the shares of composite labor and capital in unit cost,

$$\theta_N = \frac{w_N N_1}{C} = \frac{\alpha_N^\sigma w_N^{1-\sigma}}{X}$$

$$\theta_K = \frac{r_K K_1}{C} = \frac{\alpha_K^\sigma r_K^{1-\sigma}}{X}$$

$$\theta_M = \frac{p_M M_1}{C} = \frac{\alpha_K^\sigma r_K^{1-\sigma}}{X}$$

Next, we calculate $\theta_{LN}$ and $\theta_{SN}$, the shares of skilled and unskilled labor in the cost of a unit of composite labor

$$\theta_{LN} = \frac{w_L L_N}{w_N} = \alpha_L^\varphi w_L^{1-\varphi} \left[ \alpha_L^\varphi w_L^{1-\varphi} + \alpha_S^\varphi \left( \frac{w_S}{\Omega_S} \right)^{1-\varphi} \right]^{-1}$$

$$\theta_{SN} = \frac{w_S L_N}{w_N} = \alpha_S^\varphi \left( \frac{w_S}{\Omega_S} \right)^{1-\varphi} \left[ \alpha_L^\varphi w_L^{1-\varphi} + \alpha_S^\varphi \left( \frac{w_S}{\Omega_S} \right)^{1-\varphi} \right]^{-1}$$

Then we calculate the shares of skilled and unskilled labor in the unit cost of output. We do this in two steps. First, we divide (A.30) by (A.29) to get

$$\frac{L_1}{C} = \frac{\alpha_N^\sigma w_N^{-\sigma}}{X} \times \left[ \alpha_L^\varphi w_L^{1-\varphi} + \alpha_S^\varphi \left( \frac{w_S}{\Omega_S} \right)^{1-\varphi} \right]^{\frac{\varphi}{1-\varphi}}$$

Next multiply by $\frac{w_N w_L}{w_N}$,

$$\begin{aligned} \frac{w_L L_1}{C} &= \frac{\alpha_N^\sigma w_N^{1-\sigma}}{X} \times \alpha_L^\varphi w_L^{1-\varphi} \left[ \alpha_L^\varphi w_L^{1-\varphi} + \alpha_S^\varphi \left( \frac{w_S}{\Omega_S} \right)^{1-\varphi} \right]^{-1} \\ &= \theta_N \times \theta_{LN} \end{aligned}$$

So we conclude as expected that

$$\theta_L = \theta_N \times \theta_{LN}$$

Similarly,

$$\theta_S = \theta_N \times \theta_{SN}$$

Notice that $\theta_{LN} + \theta_{SN} = 1$ and $\theta_L + \theta_S = \theta_N$ and $\theta_L + \theta_S + \theta_K + \theta_M = 1$.

Next, compute the following elasticities,

$$\frac{\partial \ln X}{\partial \ln \Omega_H} = \frac{\partial \ln w_N}{\partial \ln \Omega_H} = 0$$

$$\frac{\partial \ln w_N}{\partial \ln \Omega_S} = -\theta_{SN}$$

$$\frac{\partial \ln X}{\partial \ln \Omega_S} = -(1-\sigma)\, w_N^{\varphi-\sigma} \alpha_N^{\sigma} \alpha_S^{\varphi} \left(\frac{w_S}{\Omega_S}\right)^{1-\varphi} X^{-1} = (\sigma - 1)\,\theta_S$$

Substituting these expressions into equations (A.34) yields equations (14) and (15). These equations have straightforward intuition:

- For both $S$ and $L$, the effect of Hicks-neutral technological progress is to reduce the employment required to produce a unit of ouput, and thus decrease employment with an elasticity of -1. But at the same time, costs decrease with an elasticity of -1 and thus increase demand with elasticity $\eta$, so the net effect on employment of Hicks-neutral technological progress is $(\eta - 1)\,\widehat{\Omega}_H$.

- The effect of skill-augmenting technological progress has multiple channels. First, suppose that $\sigma = \varphi$, which means we have a simple 3-factor CES function and the second term disappears. The term $(\eta - \sigma)\,\theta_S \widehat{\Omega}_S$ that appears in both equations represents the cost-reducing effect of skill-augmenting technological progress, which reduces costs with an elasticity of $\theta_S$ and thus increases labor demand with an elasticity $(\eta - \sigma)\,\theta_S$. The term $(\varphi - 1)\,\widehat{\Omega}_S$ in equation (15) represents the usual balance between the efficiency effect which reduces employment with an elasticity of -1 and the substitution between the other factors and $S$ which increases employment of S with an elasticity $\sigma = \varphi$.

- The coefficient $(\sigma - \varphi)\,\theta_{SN}$ is negative if $\varphi$, the elasticity of substitution between $L$ and $S$ within $N$, exceeds $\sigma$, the elasticity of substitution between $N$ and the other factors. When $S$ becomes more productive, there is both substitution within $N$ towards $S$ and substitution towards $N$ from the other factors.

It is instructive to compute output elasticities directly from the primal production function (1) and (2). We have

$$
\begin{aligned}
\ln Y_{ft} &= \ln \Omega_{Hft} + \frac{1}{\gamma} \ln \widetilde{Y}_{ft}, \qquad \widetilde{Y}_{ft} = \left[\alpha_N N_{ft}^{\gamma} + \alpha_K K_{ft}^{\gamma} + \alpha_M M_{ft}^{\gamma}\right] \\
\ln N_{ft} &= \frac{1}{\rho} \ln \widetilde{N}_{ft}, \qquad \widetilde{N}_{ft} \left[\alpha_L L_{ft}^{\rho} + \alpha_S \left(\Omega_{Sft} S_{ft}\right)^{\rho}\right]
\end{aligned}
$$

Output elasticities $\epsilon$ are

$$
\epsilon_{YN} = \alpha_N \frac{N_{ft}^{\gamma-1}}{\widetilde{Y}_{ft}}, \quad \epsilon_{YK} = \alpha_K \frac{K_{ft}^{\gamma-1}}{\widetilde{Y}_{ft}}, \quad \epsilon_{YM} = \alpha_M \frac{M_{ft}^{\gamma-1}}{\widetilde{Y}_{ft}}
$$

$$
\epsilon_{NL} = \alpha_L \frac{L_{ft}^{\rho-1}}{\widetilde{N}_{ft}}, \quad \epsilon_{NS} = \alpha_S \frac{\Omega_{Sft}^{\rho} S_{ft}^{\rho-1}}{\widetilde{N}_{ft}}
$$

$$
\epsilon_{YL} = \epsilon_{YN}\epsilon_{NL}, \quad \epsilon_{YS} = \epsilon_{YN}\epsilon_{NS}
$$

Evidently, these output elasticities are complex functions of parameters and the data, and will differ depending on where in the sample they are evaluated. But recall that we normalize all data series, including productivity, by their geometric means. Thus at the geometric mean of the normalized data, which we denote by superscript $g$, we have

$$
Y^g = \Omega_H^g = \Omega_S^g = N^g = K^g = M^g = L^g = S^g = 1
$$

which along with the constant returns to scale assumptions $\alpha_N + \alpha_K + \alpha_M = \alpha_L + \alpha_S = 1$ in turn implies that the distribution parameters $\alpha_j$ have a very simple interpretation as output elasticities at the geometric mean of the data,

$$
\epsilon_{YN}^g = \alpha_N, \quad \epsilon_{YK}^g = \alpha_K, \quad \epsilon_{YM}^g = \alpha_M, \quad \epsilon_{NL}^g = \alpha_L, \quad \epsilon_{NS}^g = \alpha_S, \quad \epsilon_{YL}^g = \alpha_N\alpha_L, \quad \epsilon_{YS}^g = \alpha_N\alpha_S
$$

### 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 $\beta_Z$ and $B$ are estimated parameters from equations (16) and (17), $\omega_{f2009}$ is estimated productivity in 2009, and $\widehat{\omega}_{f2010}$ is predicted productivity in 2010. All $\omega_{ft}$ (whether predicted or not) are $2\times 1$ vectors with elements $\omega_{Hft}$ and $\omega_{Sft}$. Iterating forward gives

$$\widehat{\omega}_{f2013} = \beta_Z z_{f2012} + B\beta_Z z_{f2011} + B^2\beta_Z z_{f2010} + B^3\beta_Z z_{f2009} + B^4\omega_{f2009}\,. \tag{A.35}$$

Suppose that beginning in 2009, $z_{ft}=0$. Plugging this into (A.35) 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

$$\widehat{\omega}_{f2013} - \widehat{\omega}^0_{f2013} = x = \beta_Z z_{f2012} + B\beta_Z z_{f2011} + B^2\beta_Z z_{f2010} + B^3\beta_Z z_{f2009}\,. \tag{A.36}$$

The $2\times 1$ vector $x$ defined by equation (A.36) is key in what follows. Note that (A.36) 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.36). 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.36).

When we consider the effect of firm decisions on techies, all elements of $z_{ft}$ in (A.36) are set to zero, except for firm expenditures on techies. Similarly, when we consider the effect of firm trade decisions, the only non-zero element of $z_{ft}$ in (A.36) are the firm indicators for exporting or 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 (14)-(15). In doing so, we use industry-specific estimates of the elasticities $\eta$, $\sigma$ and $\varphi$. We replace $\theta_{SN}$ by the industry-specific estimate of $\alpha_S$, and $\theta_S$ by the industry-specific estimate of $\alpha_S \alpha_N$, because the data do not permit us to gauge the expenditure share on skilled labor in total costs. 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 $\overline{\theta}_{SN} = \alpha_S$ and $\overline{\theta}_S = \alpha_S \alpha_N$. 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 $S_{2013}/L_{2013}$).

We compare the predicted change $S_{2013}/L_{2013} - S_{2009}/L_{2009}$ 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 (3).

The relationship from techies to changes in log productivity is

$$\omega_{ft} = \omega_{ft-1} + Max\left[\delta\, ln\left(\frac{T_{ft-1}}{\gamma_{1f}}\right), 0\right], \quad \delta \geq 0 \qquad \text{(A.37)}$$

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

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 (3), 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.37) implies $\Omega_{ft} = \left[\frac{T_{ft-1}}{\gamma_{1f}}\right]^{\delta}\Omega_{ft-1}$. Substituting this into the expression for profits gives

$$\Pi_f^T = B\Omega_{ft-1}^{\eta-1} - rT_{ft-1} - \gamma_{0f} + B\left(\left[\frac{T_{ft-1}}{\gamma_{1f}}\right]^{\delta}\Omega_{ft-1}\right)^{\eta-1} \tag{A.38}$$

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

$$T_{ft} = (\delta\eta - \delta)^{\frac{1}{1-\delta(\eta-1)}} r^{\frac{1}{1-\delta(\eta-1)}} \gamma_{1f}^{\frac{\delta(\eta-1)}{\delta(\eta-1)-1}} \Omega_{f1}^{\frac{1-\eta}{\delta(\eta-1)-1}} \tag{A.39}$$

For high enough values of $\delta$, 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 $\gamma_{1f}$ are negative, and that the elasticity of techies with respect to $\omega_{f1}$ is positive.

Plugging the solution (A.39) back into the expression for $\Omega_{ft}$ gives

$$\Omega_{ft} = \left[\frac{\gamma_{1f} r}{\delta\left(\eta-1\right)}\right]^{\frac{-\delta}{1-\delta(\eta-1)}} \Omega_{ft-1}^{\frac{1}{1-\delta(\eta-1)}} \tag{A.40}$$

This equation establishes the intuitive result that optimized $\Omega_{ft}$ is decreasing in $r$ and $\gamma_{1f}$, and increasing in $\Omega_{ft-1}$.

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.39) and (A.40) into (A.38) to obtain

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

Thus the difference between the two profit levels is

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

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 $\gamma_{1f}$ and $\gamma_{0f}$ and the larger is $\Omega_{f1}$. 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.

### A.3 Additional tables

Table A2: Techie and trade effects on Hicks neutral productivity, simplified production function

| | (1) | (2) |
|---|---|---|
| Techies | 0.4872*** | |
| | (0.1068) | |
| Exporting | 0.0074 | 0.0077 |
| | (0.0049) | (0.0049) |
| Importing | 0.0214*** | 0.0214*** |
| | (0.0063) | (0.0063) |
| Techies : ICT | | 0.7129*** |
| | | (0.1253) |
| Techies : R&D | | 0.2846** |
| | | (0.1227) |
| lagged $\omega_{ft}^H$ | 0.8545*** | 0.8544*** |
| | (0.0102) | (0.0102) |
| Other controls | Yes | Yes |

Notes to Table A2: Dependent variable for all regressions is estimated Hicks neutral productivity $\omega_{ft}^H$ computed from simplified production function estimation, see text for details. Weighted least squares estimation of equation (16), pooled across industries, with industry $\times$ year fixed effects. Other controls include firm size and age. Observations weighted by firm employment. Bootstrapped standard errors clustered by firm in parentheses. Sample: 568,650 observations on 192,961 firms during 2009-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table A3: Techie and trade effects on productivity, techies in production

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Hicks neutral $\omega_{ft}^H$ | | Skill augmenting $\omega_{ft}^S$ | |
| Techies | 0.2828 | | 0.2735*** | |
| | (0.3136) | | (0.0465) | |
| Techies : ICT | | 0.7938** | | 0.2377*** |
| | | (0.3999) | | (0.0496) |
| Techies : R&D | | -0.2018 | | 0.3074*** |
| | | (0.4683) | | (0.0638) |
| Exporting | -0.0089 | -0.0084 | 0.0300*** | 0.0300*** |
| | (0.0308) | (0.0307) | (0.0030) | (0.0030) |
| Importing | 0.0241 | 0.0243 | 0.0136*** | 0.0136*** |
| | (0.0422) | (0.0422) | (0.0031) | (0.0031) |
| other controls | yes | yes | yes | yes |

Notes to Table A3: Weighted least squares estimation of equations (16) and (17), pooled across industries, with industry $\times$ year fixed effects. Techies included in $S$ and $L$ in production function estimation. Other controls included lagged $\omega_{ft}^H$ and $\omega_{ft}^S$, lagged firm size and age. Observations weighted by firm employment. Bootstrapped standard errors clustered by firm in parentheses. Sample: 568,650 observations on 192,961 firms during 2009-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.

Table A4: Techie and trade effects on skill demand

| | (1) | (2) |
|---|---|---|
| log Skill premium | -3.1312*** | -3.1299*** |
| | (0.0517) | (0.0516) |
| Techies | 0.9328*** | |
| | (0.1390) | |
| Exporting | 0.1490*** | 0.1502*** |
| | (0.0156) | (0.0156) |
| Importing | 0.1386*** | 0.1388*** |
| | (0.0149) | (0.0148) |
| Techies : ICT | | 1.8250*** |
| | | (0.1920) |
| Techies : R&D | | 0.1460 |
| | | (0.1365) |
| lagged $\omega_{ft}^{H}$ | 0.0067* | 0.0064* |
| | (0.0038) | (0.0038) |
| lagged $\omega_{ft}^{S}$ | 0.9795*** | 0.9780*** |
| | (0.0853) | (0.0852) |
| Other controls | Yes | Yes |

Notes to Table A4: Dependent variable for all regressions is log hours of skilled/unskilled hours worked. Weighted two stage least squares estimation of equation (20), pooled across industries, with industry × year fixed effects. Other controls include firm size and age. Observations weighted by firm employment. Bootstrapped standard errors clustered by firm in parentheses. Sample: 568,650 observations on 192,961 firms during 2009-2013. Asterisks indicate statistical significance, * = 0.10, ** = 0.05, *** = 0.01.