Measuring the Bias of Technological Change*

Ulrich Doraszelski†
University of Pennsylvania and CEPR

Jordi Jaumandreu‡
Boston University and CEPR

This draft: November 20, 2014
First draft: December 13, 2008

Abstract

Technological change can increase the productivity of the various factors of production in equal terms or it can be biased towards a specific factor. We develop an estimator for production functions when productivity is multi-dimensional. We directly assess the bias of technological change by measuring, at the level of the individual firm, how much of it is factor neutral and how much is labor augmenting. Applying our estimator to panel data from Spain, we find that technological change is indeed biased, with both its factor-neutral and its labor-augmenting component causing output to grow by about 2% per year.

*We thank Pol Antràs, Matthias Doepke, Michaela Draganska, José Carlos Fariñas, Ivan Fernandez-Val, Paul Grieco, Chad Jones, Dale Jorgenson, Larry Katz, Pete Klenow, Jacques Mairesse, Ariel Pakes, Amil Petrin, Zhongjun Qu, Devesh Raval, Juan Sanchis, Matthias Schündeln, and John Van Reenen for helpful discussions and Sterling Horne, Mosha Huang, Thomas O'Malley, and Dan Sacks for research assistance. We gratefully acknowledge financial support from the National Science Foundation under Grants No. 0924380 and 0924282. An Online Appendix with additional results and technical details is available from the authors upon request.

†Wharton School, University of Pennsylvania, Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104, USA. E-mail: doraszelski@wharton.upenn.edu.

‡Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, USA. E-mail: jordij@bu.edu.

1 Introduction

When technological change occurs, it can increase the productivity of capital, labor, and the other factors of production in equal terms or it can be biased towards a specific factor. Whether technological change favors some factors of production over others is an empirical question that is central to economics. Yet, the empirical evidence is relatively sparse.

The literature on economic growth rests on the assumption that technological change increases the relative productivity of labor vis-à-vis other factors of production. It is well known that for a neoclassical growth model to exhibit steady-state growth either the production function must be Cobb-Douglas or technological change must be labor augmenting. Many models of endogenous growth (Romer 1986, Romer 1990, Lucas 1988) also assume labor-augmenting technological change, sometimes in the more specific form of human capital accumulation. A number of recent papers provide microfoundations for this extensive literature by theoretically establishing that profit-maximizing incentives can ensure that technological change is, at least in the long run, purely labor augmenting (Acemoglu 2003, Jones 2005). However, whether this is indeed the case is an empirical question that remains to be answered.

One reason for the scarcity of empirical assessments of the bias of technological change may be a lack of suitable data. Following early work by Brown & de Cani (1963) and David & van de Klundert (1965), economists have estimated aggregate production or cost functions that proxy for labor-augmenting technological change with a time trend (see, e.g., Lucas 1969, Kalt 1978, Antràs 2004, Binswanger 1974, Jin & Jorgenson 2010).[1] While this line of research has produced some evidence of labor-augmenting technological change, aggregation issues loom large in light of the staggering amount of heterogeneity across firms (see, e.g., Dunne, Roberts & Samuelson 1988, Davis & Haltiwanger 1992), as do the intricacies of constructing data series from national income and product accounts (see, e.g., Gordon 1990, Krueger 1999).

While traditionally using more disaggregated data, the productivity and industrial organization literatures assume that technological change is factor neutral. So-called Hicks-neutral technological change underlies, either explicitly or implicitly, most of the standard techniques for measuring productivity, ranging from the classic growth decompositions of Solow (1957) and Hall (1988) to the recent structural estimators for production functions (Olley & Pakes 1996, Levinsohn & Petrin 2003, Ackerberg, Caves & Frazer 2006, Doraszelski & Jaumandreu 2013, Gandhi, Navarro & Rivers 2013). In their present form these techniques therefore do not allow us to assess whether technological change is biased towards some factors of production.

In this paper, we combine firm-level panel data that is now widely available with ad-

[1] A much larger literature has estimated the elasticity of substitution using either aggregated or disaggregated data whilst maintaining the assumption of Hicks-neutral technological change, see Hammermesh (1993) for a survey.

vances in econometric techniques to directly assess the bias of technological change by measuring, at the level of the individual firm, how much of technological change is Hicks neutral and how much of it is labor augmenting. We account for firm-level heterogeneity in Hicks-neutral and labor-augmenting productivity by allowing their evolution to be subject to random shocks. Because these productivity innovations accumulate over time, they can cause persistent differences across firms. Hence, rather than assuming that a time trend can be interpreted as an average economy- or sector-wide measure of technological change, we obtain a much richer assessment of the impact of technological change at the level it takes place, namely the individual firm. Furthermore, we are able to relate the speed and direction of technological change to firms' R&D activities.

Our approach to separating Hicks-neutral from labor-augmenting technological change builds on recent advances in the structural estimation of production functions to tackle the endogeneity problem that arises because a firm's decisions depend on its productivity, and productivity is not observed by the econometrician (Marschak & Andrews 1944). We extend the insight of Olley & Pakes (1996) that the firm's decisions can be used to infer its productivity to a setting in which productivity is multi- instead of single-dimensional. In particular, we recover a firm's multi-dimensional productivity from its input usage. In doing so, we follow Doraszelski & Jaumandreu (2013) by exploiting the parameter restrictions between the production and input demand functions. This parametric inversion facilitates identification and estimation. We further contribute to the literature following Olley & Pakes (1996) by accounting for outsourcing and adjustment costs to permanent labor and highlighting the role they play for properly measuring the bias of technological change.

We apply our estimator to an unbalanced panel of 2375 Spanish manufacturing firms in ten industries from 1990 to 2006. Spain is an attractive setting for examining the speed and direction of technological change for two reasons. First, Spain became fully integrated into the European Union between the end of the 1980s and the beginning of the 1990s. Any trends in technological change that our analysis uncovers for Spain may thus be viewed as broadly representative for other continental European economies. Second, Spain inherited an industrial structure with few high-tech industries and mostly small and medium-sized firms. Traditionally, R&D is viewed as lacking and something to be boosted (OECD 2007).[2] Yet, Spain grew rapidly during the 1990s, and R&D became increasingly important (European Commission 2001). The accompanying changes in industrial structure are a useful source of variation for analyzing the role of R&D in stimulating different types of technological change.

The particular data set we use has several advantages. The broad coverage is unusual and allows us to assess the bias of technological change in industries that differ greatly in terms of firms' R&D activities. The data set also has an unusually long time dimension, enabling

[2] The Spanish government repeatedly attempted to stimulate R&D. Most recently, in 2005 it launched the ambitious Ingenio 2010 initiative targeted at funding large-size, high-risk research projects.

us to disentangle trends in technological change from short-term fluctuations. Finally, the data set has firm-level prices that we exploit heavily in the estimation.[3]

Our estimates provide clear evidence that technological change is biased. Ceteris paribus labor-augmenting technological change causes output to grow, on average, in the vicinity of 2% per year. While there is a shift from unskilled to skilled workers in our data, this skill upgrading explains some but not all of the growth of labor augmenting productivity. In many industries, labor-augmenting productivity grows because workers with a given set of skills become more productive over time.

At the same time, our estimates show that Hicks-neutral technological change plays an equally important role. In addition to labor-augmenting technological change, Hicks-neutral technological change causes output to grow, on average, in the vicinity of 2% per year.

Behind these averages lies a substantial amount of heterogeneity across industries and firms. At the level of the individual firm, the levels of labor-augmenting and Hicks-neutral productivity are positively correlated, as are their rates of growth. Moreover, for both components of productivity, the rate of growth is positively correlated with the level, indicating that differences in productivity between firms persist over time.

Finally, our estimates indicate that firms' R&D activities are associated with higher levels and rates of growth of labor-augmenting productivity and, perhaps to a lesser extent, with higher levels and rates of growth of Hicks-neutral productivity as well. Firms' R&D activities therefore are important for determining the differences in labor-augmenting and Hicks-neutral productivity across firms and the evolution of productivity over time.

Our paper is related to Van Biesebroeck (2003). Using plant-level panel data for the U.S. automobile industry, he estimates Hicks-neutral productivity as a fixed effect and recovers a plant's capital-biased (also called labor-saving) productivity from its input usage. Building on Doraszelski & Jaumandreu (2013), our approach is similar in that it uses a parametric inversion to recover unobserved productivity from observed inputs but differs in that productivity is multi-dimensional. Our model is also more general in that we allow both factor-neutral and factor-specific productivity to evolve over time and in response to firms' R&D activities.

Our paper is also related to Grieco, Li & Zhang (2014) who recover multiple unobservables from input usage. Because their data contains the materials bill rather than its split into price and quantity, they infer a firm's Hicks-neutral productivity and the price of materials that the firm faces by parametrically inverting the demand functions for labor and materials. In subsequent work in progress, Zhang (2014*a*, 2014*b*) applies the same idea to recover a firm's capital-augmenting productivity and its labor-augmenting productivity. We return to the related literature in Sections 5 and 6.

[3]There are other firm-level data sets such as the Colombian Annual Manufacturers Survey (Eslava, Haltiwanger, Kugler & Kugler 2004) and the Longitudinal Business Database at the U.S. Census Bureau that contain separate information on prices and quantities, at least for a subset of industries (Roberts & Supina 1996, Foster, Haltiwanger & Syverson 2008, Foster, Haltiwanger & Syverson 2013).

Finally, our paper touches—although more tangentially—on the literature on skill bias that studies the differential impact of technological change, especially in the form of computerization, on the various types of labor. Our approach is similar to some of the recent work on skill bias (see, e.g., Machin & Van Reenen 1998, Black & Lynch 2001, Bloom, Sadun & Van Reenen 2012, Abowd, Haltiwanger, Lane, McKinney & Sandusky 2007) in that it starts from a production function and focuses on the individual firm. While we focus on labor versus other factors of production, the techniques we develop may be applied to investigate the skill bias of technological change, although our particular data set is not ideal for this purpose. Our approach differs from the recent work on skill bias in that it explicitly models and estimates the differences in productivity across firms and the evolution of firm-level productivity over time. It is also more structural in tackling the endogeneity problem that arises in estimating production functions.

The remainder of this paper is organized as follows: Section 2 describes the data. Section 3 sets out a dynamic model of the firm. Productivity has a Hicks-neutral and a labor-augmenting component, and the evolution of both components is governed by stochastic processes that are controlled by the firm's R&D activities. Section 4 develops an estimator for production functions that allows us to retrieve Hicks-neutral and labor-augmenting productivity at the level of the individual firm. Sections 5 and 6 describe our results on labor-augmenting and Hicks-neutral technological change. Section 7 concludes and outlines a number of directions for future research.

2 Data

Our data comes from the Encuesta Sobre Estrategias Empresariales (ESEE) survey, a firm-level survey of the Spanish manufacturing sector sponsored by the Ministry of Industry. The unit of observation is the firm, not the plant or the establishment. Our data covers the 1990s and early 2000s. At the beginning of the survey in 1990, 5% of firms with up to 200 workers were sampled randomly by industry and size strata. All firms with more than 200 workers were asked to participate in the survey and 70% of them complied. Some firms vanish from the sample due to either exit (shutdown by death or abandonment of activity) or attrition. These reasons can be distinguished in the data and attrition remained within acceptable limits. To preserve representativeness, newly created firms were added to the sample every year. We provide details on industry and variable definitions in Appendix A.

Our sample covers a total of 2375 firms in ten industries when restricted to firms with at least three years of data. Columns (1) and (2) of Table 1 show the number of observations and firms by industry. Sample sizes are moderate. Newly created firms are a large share of the total number of firms, ranging from 26% to 50% in the different industries. There is a significant fraction of exiting firms, ranging from 6% to 15% and above in a few industries. Firms remain in the sample from a minimum of three years to a maximum of 16 years between 1990 and 2006.

The 1990s and early 2000s were a period of rapid output growth, coupled with stagnant or, at best, slightly increasing employment and intense investment in physical capital, see columns (3)–(6) of Table 1. Consistent with this rapid growth, firms on average report that their market is slightly more often expanding rather than contracting; hence, demand tends to shift out over time.

An attractive feature of our data is that it contains firm-specific price indices for output and inputs. The growth of prices, averaged from the growth of prices as reported individually by each firm, is moderate. The growth of the price of output in column (7) ranges from 0.8% to 2.1%. The growth of the wage ranges from 4.3% to 5.4% and the growth of the price of materials ranges from 2.8% to 4.1%.

Biased technological change. The evolution of the relative quantities and prices of the various factors of production already hint at an important role for labor-augmenting technological change. As columns (8) and (9) of Table 1 show, with the exception of industries 7, 8, and 9, the increase in materials M per unit of labor L is much larger than the decrease in the price of materials P_M relative to the wage W. One possible explanation is that the elasticity of substitution exceeds 1. To see this, recall that the elasticity of substitution is

$$\frac{d\ln\left(\frac{M}{L}\right)}{d\ln\left(\frac{MPR_M}{MPR_L}\right)} = -\frac{d\ln\left(\frac{M}{L}\right)}{d\ln\left(\frac{P_M}{W}\right)},$$

where the equality follows to the extent that the relative marginal products $\frac{MPR_M}{MPR_L}$ equal the relative prices $\frac{P_M}{W}$. However, because the estimates of the elasticity of substitution in the previous literature lie somewhere between 0 and 1 (see Chirinko (2008) and the references therein for the elasticity of substitution between capital and labor and Oberfield & Raval (2014) for the elasticity of substitution between an aggregate of capital and labor and materials), this explanation is implausible. Labor-augmenting technological change offers an alternative explanation. As it makes labor more efficient, it directly increases materials per unit of labor (see equation (15) in Section 5). Thus, labor-augmenting technological change may go a long way in rationalizing why the relative quantities $\frac{M}{L}$ change much more than the relative prices $\frac{P_M}{W}$.

In contrast, columns (10) and (11) of Table 1 provide no evidence for capital-augmenting technological change. The investment boom in Spain in the 1990s and early 2000s was fueled by improved access to European and international capital markets. With the exception of industries 5, 6, and 8, the concomitant decrease in materials M per unit of capital K is much smaller than the increase in the price of materials P_M relative to the user cost of capital in our data, a notably rough measure of the price of capital P_K.[4] This pattern is

[4]In particular, the price of capital includes adjustment costs, and as a shadow price, it is unobservable. The user cost of capital, in contrast, is based solely on observables (see Appendix A).

consistent with an elasticity of substitution between 0 and 1. Indeed, capital-augmenting technological change can only directly contribute to the decline in materials per unit of capital in the unlikely scenario that it makes capital less efficient.

Temporary labor. Our empirical strategy treats temporary labor as a static (or "variable") input that is chosen each period to maximize short-run profits. This is appropriate because Spain greatly enhanced the possibilities for hiring and firing temporary workers during the 1980s and by the beginning of the 1990s had one the highest shares of temporary workers in Europe (Dolado, Garcia-Serrano & Jimeno 2002). Temporary workers are employed for fixed terms with no or very small severance pay. In our sample, between 72% and 84% of firms use temporary labor and among the firms that do its share of the labor force ranges from 16% in industry 10 to 32% in industry 9, see columns (1) and (2) of Table 2.

Rapid expansions and contractions of temporary labor are common: The difference between the maximum and the minimum share of temporary labor within a firm ranges on average from 20% to 33% across industries (column (3)). In addition to distinguishing temporary from permanent labor, we measure labor as hours worked (see Appendix A for details). At this margin, firms enjoy a high degree of flexibility: Within a firm, the difference between the maximum and the minimum hours worked ranges on average from 43% to 56% across industries, and the difference between the maximum and the minimum hours per worker ranges on average from 4% to 13% (columns (4) and (5)).

Outsourcing. Our empirical strategy accounts for outsourcing. Outsourcing may directly contribute to the shift from labor to materials that column (8) of Table 1 documents as firms procure parts and pieces from their suppliers rather than make them in house from scratch. As can be seen in columns (6) and (7) of Table 2, between 21% and 57% of firms in our sample engage in outsourcing. Among the firms that do, the share of outsourcing in the materials bill ranges from 14% in industry 7 to 29% in industry 4. While the share of outsourcing remains stable over our sample period, the standard deviation in column (7) indicates a substantial amount of heterogeneity across the firms within an industry, similar to the share of temporary labor in column (2).

Firms' R&D activities. The R&D intensity of Spanish manufacturing firms is low by European standards, but R&D became increasingly important during the 1990s (see, e.g., European Commission 2001).[5] Columns (8)–(10) of Table 2 show that the ten industries differ markedly in terms of firms' R&D activities and that there is again substantial heterogeneity across the firms within an industry. Industries 3, 4, 5, and 6 exhibit high innovative activity. More than two thirds of firms perform R&D during at least one year in the sample

[5] R&D intensities for manufacturing firms are 2.1% in France, 2.6% in Germany, and 2.2% in the UK as compared to 0.6% in Spain (European Commission 2004).

period, with at least 36% of stable performers engaging in R&D in all years (column (8)) and at least 28% of occasional performers engaging in R&D in some but not all years (column (9)). The R&D intensity among performers ranges on average from 2.2% to 2.9% (column (10)). Industries 1, 2, 7, and 8 are in an intermediate position. Less than half of firms perform R&D, and there are fewer stable than occasional performers. The R&D intensity is on average between 1.1% and 1.7% with a much lower value of 0.7% in industry 7. Finally, industries 9 and 10 exhibit low innovative activity. About a third of firms perform R&D, and the R&D intensity is on average between 1.0% and 1.5%.

3 A dynamic model of the firm

Our model relates to and builds on the previous literature on the structural estimation of production functions in a number of ways. First, Olley & Pakes (1996), Levinsohn & Petrin (2003), Ackerberg et al. (2006), Doraszelski & Jaumandreu (2013), and many others specify a Cobb-Douglas production function. Productivity is single-dimensional or, equivalently, technological change is Hicks neutral by construction. By generalizing the Cobb-Douglas production function and allowing productivity to be multi-dimensional, our model is able to capture the factor-specific nature or bias of technological change. Second, we follow Doraszelski & Jaumandreu (2013) and endogenize the evolution of productivity by incorporating R&D expenditures into the model. This accounts for nonlinearities and uncertainties in the link between R&D and productivity and generalizes the classic knowledge capital model (Griliches 1979). Third, we extend the parametric inversion in Doraszelski & Jaumandreu (2013) to infer a firm's multi-dimensional productivity from its input usage. Because it exploits the parameter restrictions between the production and input demand functions, this parametric inversion is less demanding on the data than the nonparametric inversion in Olley & Pakes (1996), Levinsohn & Petrin (2003), and Ackerberg et al. (2006), especially in settings like ours in which the input demand functions are high-dimensional (see Doraszelski & Jaumandreu (2013) for details on the pros and cons of the parametric inversion). Fourth, in recognition of the dual nature of the labor market in Spain, our model distinguishes between permanent and temporary workers and accounts for the substantial cost of hiring, training, and firing permanent workers. Fifth, to properly measure the bias of technological change our model distinguishes between in-house and outsourced materials.

Production function. The firm has the constant elasticity of substitution (CES) production function

$$Y_{jt} = \beta_0 \left[\beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \beta_L \left(\exp(\omega_{Ljt}) L_{jt}^*\right)^{-\frac{1-\sigma}{\sigma}} + \beta_M \left(M_{jt}^*\right)^{-\frac{1-\sigma}{\sigma}}\right]^{-\frac{\nu\sigma}{1-\sigma}} \exp(\omega_{Hjt}) \exp(e_{jt}), \tag{1}$$

where Y_{jt} is the output of firm j in period t, K_{jt} is capital, $L^*_{jt} = \Lambda(L_{Pjt}, L_{Tjt})$ is an aggregate of permanent labor L_{Pjt} and temporary labor L_{Tjt}, and $M^*_{jt} = \Gamma(M_{Ijt}, M_{Ojt})$ is an aggregate of in-house materials M_{Ijt} and outsourced materials (parts and pieces) M_{Ojt}. ω_{Ljt} and ω_{Hjt} are labor-augmenting and Hicks-neutral productivity, respectively, and e_{jt} is a mean zero random shock that is uncorrelated over time and across firms.

The parameters ν and σ are the elasticity of scale and substitution, respectively. Depending on the elasticity of substitution, the production function in equation (1) encompasses the special cases of a Leontieff ($\sigma \to 0$), Cobb-Douglas ($\sigma = 1$), and linear ($\sigma \to \infty$) production function. The remaining parameters are the constant of proportionality β_0 and the distributional parameters β_K, β_L, and β_M. Because β_0 cannot be separated from an additive constant in Hicks-neutral productivity ω_{Hjt}, we estimate them jointly. To simplify the notation and without loss of generality, we set $\beta_0 = 1$ in what follows.[6] We similarly set $\beta_L = 1$. Viewing technological change as operating by changing the efficiencies of the various factors of production is therefore equivalent to viewing it as changing these parameters of the production function.

The production function in equation (1) is the most parsimonious we can use to separate labor-augmenting from Hicks-neutral productivity.[7] It encompasses two restrictions. First, technological change does not affect the parameters ν and σ, as we are unaware of evidence suggesting that the elasticity of scale or the elasticity of substitution varies over our sample period. Second, the efficiencies of capital and materials are restricted to change at the same rate and in lockstep with Hicks-neutral technological change. Focusing on labor-augmenting productivity while abstracting from capital-augmenting productivity is justified by the patterns in the data described in Section 2. Treating capital and materials the same is also in line with the fact that both are, at least to a large extent, produced goods, whereas labor is traditionally viewed as unique among the various factors of production.[8] In Section 6, we explore further whether there is a potential role for capital-augmenting productivity in our data.

Laws of motion. The productivity of firm j in period t is the tuple $(\omega_{Ljt}, \omega_{Hjt})$. The components of productivity are presumably correlated with each other and over time and possibly also correlated across firms. Following Doraszelski & Jaumandreu (2013), we assume that they are governed by controlled first-order, time-inhomogeneous Markov processes with transition probabilities $P_{Lt+1}(\omega_{Ljt+1}|\omega_{Ljt}, R_{jt})$ and $P_{Ht+1}(\omega_{Hjt+1}|\omega_{Hjt}, R_{jt})$,

[6] We carefully ensure that the reported results depend only on the sum of β_0 and the additive constant in Hicks-neutral productivity ω_{Hjt}.

[7] As is well known, a Cobb-Douglas production function has an elasticity of substitution of one and therefore cannot be used to separate labor-augmenting from Hicks-neutral productivity. Our data rejects a Cobb-Douglas production function (see Section 5). Our empirical strategy generalizes to a translog production function. Other production functions may require numerically solving a system of equations to infer unobservables from observables.

[8] Marshall (1920), for example, writes in great detail about the variability of workers' efforts and its relationship to productivity.

where R_{jt} is R&D expenditures. Despite their parsimony, these stochastic processes accommodate correlation between the components of productivity.[9] Moreover, because they are time-inhomogeneous, they accommodate secular trends in productivity.

While the firm knows its productivity when it makes its decisions for period t, we follow Olley & Pakes (1996) and often refer to the tuple $(\omega_{Ljt}, \omega_{Hjt})$ as "unobserved productivity" since it is not observed by the econometrician. In contrast, the firm does not know the random shock e_{jt} when it makes its decisions.

The firm anticipates the effect of R&D on productivity when making its decisions for period t. The Markovian assumption implies

$$\omega_{Ljt+1} = E_t\left[\omega_{Ljt+1}|\omega_{Ljt}, R_{jt}\right] + \xi_{Ljt+1} = g_{Lt}(\omega_{Ljt}, R_{jt}) + \xi_{Ljt+1}, \tag{2}$$

$$\omega_{Hjt+1} = E_t\left[\omega_{Hjt+1}|\omega_{Hjt}, R_{jt}\right] + \xi_{Hjt+1} = g_{Ht}(\omega_{Hjt}, R_{jt}) + \xi_{Hjt+1}. \tag{3}$$

That is, *actual* labor-augmenting productivity ω_{Ljt+1} in period $t+1$ decomposes into *expected* labor-augmenting productivity $g_{Lt}(\omega_{Ljt}, R_{jt})$ and a random shock ξ_{Ljt+1}. This *productivity innovation* is by construction mean independent (although not necessarily fully independent) of ω_{Ljt} and R_{jt}. It captures the uncertainties that are naturally linked to productivity as well as those that are inherent in the R&D process such as chance of discovery, degree of applicability, and success in implementation. Actual Hicks-neutral productivity ω_{Hjt+1} decomposes similarly.

Capital accumulates according to $K_{jt+1} = (1-\delta)K_{jt} + I_{jt}$, where δ is the rate of depreciation. As in Olley & Pakes (1996), investment I_{jt} chosen in period t therefore becomes productive in period $t+1$. Choosing I_{jt} is therefore equivalent to choosing K_{jt+1}. Permanent labor is subject to convex adjustment costs $C_{L_P}(L_{Pjt}, L_{Pjt-1})$ that reflect the substantial cost of hiring and firing that the firm may incur (Hammermesh 1993, Hammermesh & Pfann 1996). The choice of permanent labor thus may have dynamic implications. In contrast, temporary labor is a static input.

Outsourcing is, to a large extent, based on contractual relationships between the firm and its suppliers (Grossman & Helpman 2002, Grossman & Helpman 2005). The ratio of outsourced to in-house materials $Q_{Mjt} = \frac{M_{Ojt}}{M_{Ijt}}$ is subject to (convex or not) adjustment costs $C_{Q_M}(Q_{Mjt+1}, Q_{Mjt})$ that stem from forming and dissolving these relationships. The firm must maintain Q_{Mjt} but may scale M_{Ijt} and M_{Ojt} up or down at will; in-house materials, in particular, is a static input.

[9] Our empirical strategy generalizes to a joint Markov process $P_{t+1}(\omega_{Ljt+1}, \omega_{Hjt+1}|\omega_{Ljt}, \omega_{Hjt}, r_{jt})$. While R&D is widely seen as a major source of productivity growth (see Griliches (1998, 2000) for surveys of the empirical literature), our empirical strategy extends to other sources such as technology adoption. Our data has investment in computer equipment and indicators of whether a firm has adopted digitally controlled machine tools, CAD, and robots. Both extensions are demanding on the data, however, as they increase the dimensionality of the functions that must be nonparametrically estimated.

Input and output markets. The firm is a price-taker in input markets, where it faces W_{Pjt}, W_{Tjt}, P_{Ijt}, and P_{Ojt} as prices of permanent and temporary labor and in-house and outsourced materials, respectively. The firm has market power in the output market, e.g., because products are differentiated. Its inverse residual demand function $P(Y_{jt}, D_{jt})$ depends on its output Y_{jt} and the demand shifter D_{jt}.[10]

Bellman equation. The firm makes its decisions in a discrete time setting with the goal of maximizing the expected net present value of future cash flows. The Bellman equation for its dynamic programming problem is

$$\begin{aligned} V_t(\Omega_{jt}) \quad = \quad & \max_{K_{jt+1}, L_{Pjt}, L_{Tjt}, Q_{Mjt+1}, M_{Ijt}, R_{jt}} P\left(X_{jt}^{-\frac{\nu\sigma}{1-\sigma}} \exp(\omega_{Hjt}), D_{jt}\right) X_{jt}^{-\frac{\nu\sigma}{1-\sigma}} \exp(\omega_{Hjt})\mu \\ & - C_I(K_{jt+1} - (1-\delta)K_{jt}) - W_{Pjt}L_{Pjt} - C_{L_P}(L_{Pjt}, L_{Pjt-1}) - W_{Tjt}L_{Tjt} \\ & - \left(P_{Ijt} + P_{Ojt}Q_{Mjt}\right) M_{Ijt} - C_{Q_M}(Q_{Mjt+1}, Q_{Mjt}) - C_R(R_{jt}) \\ & + \frac{1}{1+\rho} E_t\left[V_{t+1}(\Omega_{jt+1}) | \Omega_{jt}, R_{jt}\right], \end{aligned}$$

where

$$X_{jt} = \beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \left(\exp(\omega_{Ljt}) L_{jt}^*\right)^{-\frac{1-\sigma}{\sigma}} + \beta_M \left(M_{jt}^*\right)^{-\frac{1-\sigma}{\sigma}}, \quad \mu = E_t\left[\exp(e_{jt})\right],$$

$\Omega_{jt} = (K_{jt}, L_{Pjt-1}, Q_{Mjt}, \omega_{Ljt}, \omega_{Hjt}, W_{Pjt}, W_{Tjt}, P_{Ijt}, P_{Ojt}, D_{jt})$ is the vector of state variables, and ρ is the discount rate. $C_I(I_{jt})$ and $C_R(R_{jt})$ are the cost of investment and R&D, respectively, and accommodate indivisibilities in investment and R&D projects.

Input usage. Our empirical strategy infers the firm's productivity $(\omega_{Ljt}, \omega_{Hjt})$ from its labor and materials decisions. The first-order conditions for permanent and temporary labor are

$$\nu\mu X_{jt}^{-\left(1+\frac{\nu\sigma}{1-\sigma}\right)} \exp\left(\omega_{Hjt}\right) \exp\left(-\frac{1-\sigma}{\sigma}\omega_{Ljt}\right) \left(L_{jt}^*\right)^{-\frac{1}{\sigma}} \frac{\partial L_{jt}^*}{\partial L_{Pjt}} = \frac{W_{Pjt}(1+\Delta_{jt})}{P_{jt}\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}, \quad (4)$$

$$\nu\mu X_{jt}^{-\left(1+\frac{\nu\sigma}{1-\sigma}\right)} \exp\left(\omega_{Hjt}\right) \exp\left(-\frac{1-\sigma}{\sigma}\omega_{Ljt}\right) \left(L_{jt}^*\right)^{-\frac{1}{\sigma}} \frac{\partial L_{jt}^*}{\partial L_{Tjt}} = \frac{W_{Tjt}}{P_{jt}\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}, \quad (5)$$

where $\eta(p_{jt}, D_{jt})$ is the absolute value of the price elasticity of the residual demand that the firm faces,[11] and by the envelope theorem, the gap between the wage of permanent workers

[10] In general, the residual demand that the firm faces depends on its rivals' prices. In taking the model to the data, one may replace rivals' prices by an aggregate price index or dummies, although this substantially increases the dimensionality of the functions that must be nonparametrically estimated.

[11] Throughout, we follow the convention that lower case letters denote logs and upper case letters denote levels.

W_{Pjt} and the shadow wage is

$$\begin{aligned}\Delta_{jt} &= \frac{1}{W_{Pjt}}\left(\frac{\partial C_{L_P}(L_{Pjt}, L_{Pjt-1})}{\partial L_{Pjt}} - \frac{1}{1+\rho}E_t\left[\frac{\partial V_{t+1}(\Omega_{jt+1})}{\partial L_{Pjt}}|\Omega_{jt}, R_{jt}\right]\right)\\ &= \frac{1}{W_{Pjt}}\left(\frac{\partial C_{L_P}(L_{Pjt}, L_{Pjt-1})}{\partial L_{Pjt}} + \frac{1}{1+\rho}E_t\left[\frac{\partial C_{L_P}(L_{Pjt+1}, L_{Pjt})}{\partial L_{Pjt}}|\Omega_{jt}, R_{jt}\right]\right).\end{aligned}$$

Our data combines the wages of permanent and temporary workers into $W_{jt} = W_{Pjt}(1 - S_{Tjt}) + W_{Tjt}S_{Tjt}$, where $S_{Tjt} = \frac{L_{Tjt}}{L_{jt}}$ is the (quantity) share of temporary labor and $L_{jt} = L_{Pjt} + L_{Tjt}$ is hours worked by permanent and temporary workers in our data. To make do, we assume that the aggregator $\Lambda(L_{Pjt}, L_{Tjt})$ is linearly homogenous. This implies $L^*_{jt} = L_{jt}\Lambda(1 - S_{Tjt}, S_{Tjt})$, $\frac{\partial L^*_{jt}}{\partial L_{Pjt}} = \Lambda_P(1 - S_{Tjt}, S_{Tjt})$, and $\frac{\partial L^*_{jt}}{\partial L_{Tjt}} = \Lambda_T(1 - S_{Tjt}, S_{Tjt})$. Using Euler's theorem to combine equations (4) and (5) yields

$$\begin{aligned}&\nu\mu X_{jt}^{-\left(1+\frac{\nu\sigma}{1-\sigma}\right)} \exp\left(\omega_{Hjt}\right) \exp\left(-\frac{1-\sigma}{\sigma}\omega_{Ljt}\right) L_{jt}^{-\frac{1}{\sigma}} \Lambda(1 - S_{Tjt}, S_{Tjt})^{-\frac{1-\sigma}{\sigma}}\\ &= \frac{W_{jt}\left(1 + \frac{\Delta_{jt}}{1+\frac{W_{Tjt}}{W_{Pjt}}\frac{S_{Tjt}}{1-S_{Tjt}}}\right)}{P_{jt}\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)} = \frac{W_{jt}\left(\frac{\frac{\Lambda_P(1-S_{Tjt}, S_{Tjt})}{\Lambda_T(1-S_{Tjt}, S_{Tjt})} + \frac{S_{Tjt}}{1-S_{Tjt}}}{\frac{W_{Pjt}}{W_{Tjt}} + \frac{S_{Tjt}}{1-S_{Tjt}}}\right)}{P_{jt}\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)},\end{aligned} \tag{6}$$

where the second equality follows from dividing equations (4) and (5) and solving for Δ_{jt}.

Because our data does not have the ratio $\frac{W_{Pjt}}{W_{Tjt}}$, we assume that $\frac{W_{Pjt}}{W_{Tjt}} = \lambda_0$ is an (unknown) constant[12] and treat $\frac{\frac{\Lambda_P(1-S_{Tjt}, S_{Tjt})}{\Lambda_T(1-S_{Tjt}, S_{Tjt})} + \frac{S_{Tjt}}{1-S_{Tjt}}}{\lambda_0 + \frac{S_{Tjt}}{1-S_{Tjt}}} = \lambda_1(S_{Tjt})$ as an (unknown) function of S_{Tjt} that must be estimated nonparametrically along with the parameters of the production function. Because equation (6) presumes interior solutions for permanent and temporary labor, we exclude observations with $S_{Tjt} = 0$ and thus $L_{Tjt} = 0$ from the subsequent analysis.

The first-order condition for in-house materials is

$$\nu\beta_M\mu X_{jt}^{-\left(1+\frac{\nu\sigma}{1-\sigma}\right)} \exp\left(\omega_{Hjt}\right) \left(M^*_{jt}\right)^{-\frac{1}{\sigma}} \frac{\partial M^*_{jt}}{\partial M_{Ijt}} = \frac{P_{Ijt} + P_{Ojt}Q_{Mjt}}{P_{jt}\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}, \tag{7}$$

where $P_{Ijt} + P_{Ojt}Q_{Mjt}$ is the effective cost of an additional unit of in-house materials if the firm must maintain the ratio of outsourced to in-house materials Q_{Mjt}.

Our data has the materials bill $P_{Mjt}M_{jt} = P_{Ijt}M_{Ijt} + P_{Ojt}M_{Ojt}$, the (value) share of outsourced materials $S_{Ojt} = \frac{P_{Ojt}M_{Ojt}}{P_{Mjt}M_{jt}}$, and the price of materials P_{Mjt}. We assume $P_{Mjt} = P_{Ijt} + P_{Ojt}Q_{Mjt}$ so that the price of materials is the effective cost of an additional unit of in-house materials. This implies $M_{jt} = M_{Ijt}$. To map the model to the data, we

[12] In Appendix D, we use a wage regression to estimate wage premia of various types of labor. In the Online Appendix, we extend the specification and demonstrate that the wage premia do not change much if at all over time in line with our assumption that the ratio $\frac{W_{Pjt}}{W_{Ljt}}$ is constant.

further assume that $\Gamma(M_{Ijt}, M_{Ojt})$ is linearly homogenous and normalize $\Gamma(M_{Ijt}, 0) = M_{Ijt}$. This implies $M^*_{jt} = M_{Ijt}\Gamma\left(1, \frac{P_{Ijt}}{P_{Ojt}} S_{Ojt}\right)$ and $\frac{\partial M^*_{jt}}{\partial M_{Ijt}} = \Gamma\left(1, \frac{P_{Ijt}}{P_{Ojt}} S_{Ojt}\right)$. Rewriting equation (7) yields

$$\nu\beta_M \mu X_{jt}^{-\left(1+\frac{\nu\sigma}{1-\sigma}\right)} \exp\left(\omega_{Hjt}\right) M_{jt}^{-\frac{1}{\sigma}} \Gamma\left(1, \frac{P_{Ijt}}{P_{Ojt}} S_{Ojt}\right)^{-\frac{1-\sigma}{\sigma}} = \frac{P_{Mjt}}{P_{jt}\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}. \tag{8}$$

Because our data does not have the ratio $\frac{P_{Ijt}}{P_{Ojt}}$, we assume that $\frac{P_{Ijt}}{P_{Ojt}} = \gamma_0$ is an (unknown) constant and treat $\ln \Gamma\left(1, \gamma_0 S_{Ojt}\right) = \gamma_1(S_{Ojt})$ as an (unknown) function of S_{Ojt}.[13] Equation (8) presumes an interior solution for in-house materials; it is consistent with a corner solution for outsourced materials. Indeed, absent outsourcing equation (8) reduces to the first-order condition for in-house materials.

Our primary interest is the bias of technological change. Thus, we think of $\lambda_1(S_{Tjt})$ and $\gamma_1(S_{Ojt})$ as "correction terms" on labor and, respectively, materials. While these correction terms help account for the substantial heterogeneity across the firms within an industry, they are not uniquely tied to particular theories about the Spanish labor market or the role of outsourcing. In the Online Appendix, we develop an alternative model of outsourcing that assumes that both in-house and outsourced materials are static inputs that the firm may mix-and-match at will. We show that the resulting correction term is indistinguishable from that in equation (8).

4 Empirical strategy

Perhaps the major obstacle in production function estimation is the endogeneity problem that arises because a firm's decisions depend on its productivity and productivity is not observed by the econometrician (Marschak & Andrews 1944).[14] Intuitively, if the firm adjusts to a change in its productivity by expanding or contracting its production, then unobserved productivity and input usage are correlated and biased estimates result.

The recent literature on the structural estimation of production functions resolves the endogeneity problem. The insight of Olley & Pakes (1996) is that if the decisions that the firm makes can be used to infer its productivity, then productivity can be controlled for in the estimation. Our empirical strategy extends their insight from a setting with single-dimensional productivity to one with multi-dimensional productivity.

[13] We have experimented with assuming that $\frac{P_{Ijt}}{P_{Ojt}} = \gamma_0(t)$ is an (unknown) function of time t and treating $\ln \Gamma\left(1, \gamma_0(t) S_{Ojt}\right) = \gamma_1(\gamma_0(t) S_{Ojt})$ as an (unknown) function of $\gamma_0(t) S_{Ojt}$. As we show in the Online Appendix, not much changes. Equation (11) tends to yield somewhat lower estimates of σ compared to our leading estimates in column (3) of Table 4. Compared to our leading estimates in columns (1) and (2) of Table 6 equation (13) tends to yield somewhat lower estimates of β_K and similar estimates of ν in the eight industries where we have been able to obtain estimates. Our conclusions about technological change remain the same.

[14] See Griliches & Mairesse (1998) and Ackerberg, Benkard, Berry & Pakes (2007) for reviews of this and other problems involved in the estimation of production functions.

From the labor and materials decisions in equations (6) and (8) we recover (conveniently rescaled) labor-augmenting productivity $\widetilde{\omega}_{Ljt} = (1-\sigma)\omega_{Ljt}$ and Hicks-neutral productivity ω_{Hjt} as

$$\begin{aligned}
\widetilde{\omega}_{Ljt} &= \widetilde{\gamma}_L + m_{jt} - l_{jt} + \sigma(p_{Mjt} - w_{jt}) - \sigma\lambda_2(S_{Tjt}) + (1-\sigma)\gamma_1(S_{Ojt}) \\
&\equiv \widetilde{h}_L(m_{jt} - l_{jt}, p_{Mjt} - w_{jt}, S_{Tjt}, S_{Ojt}), \qquad (9)\\
\omega_{Hjt} &= \gamma_H + \frac{1}{\sigma} m_{jt} + p_{Mjt} - p_{jt} - \ln\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right) \\
&\quad + \left(1 + \frac{\nu\sigma}{1-\sigma}\right) x_{jt} + \frac{1-\sigma}{\sigma}\gamma_1(S_{Ojt}) \\
&\equiv h_H(k_{jt}, m_{jt}, S_{Mjt}, p_{jt}, p_{Mjt}, D_{jt}, S_{Tjt}, S_{Ojt}), \qquad (10)
\end{aligned}$$

where $\widetilde{\gamma}_L = -\sigma \ln \beta_M$, $\lambda_2(S_{Tjt}) = \ln\left(\lambda_1(S_{Tjt})\Lambda\left(1 - S_{Tjt}, S_{Tjt}\right)^{\frac{1-\sigma}{\sigma}}\right)$, $\gamma_H = -\ln\left(\nu\beta_M\mu\right)$,

$$X_{jt} = \beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \beta_M \left(M_{jt} \exp\left(\gamma_1(S_{Ojt})\right)\right)^{-\frac{1-\sigma}{\sigma}} \left(\frac{1 - S_{Mjt}}{S_{Mjt}}\lambda_1(S_{Tjt}) + 1\right),$$

and $S_{Mjt} = \frac{P_{Mjt}M_{jt}}{W_{jt}L_{jt} + P_{Mjt}M_{jt}}$ is the share of materials in variable cost. Without loss of generality, we set $\beta_K + \beta_M = 1$. The functions $\widetilde{h}_L(\cdot)$ and $h_H(\cdot)$ allow us to recover unobservable productivity $(\widetilde{\omega}_{Ljt}, \omega_{Hjt})$ from observables, and we refer to them as inverse functions from hereon.

Labor-augmenting productivity. Substituting the inverse function in equation (9) into the law of motion for labor-augmenting productivity ω_{Ljt} in equation (2), we form our first estimation equation

$$\begin{aligned}
m_{jt} - l_{jt} = -\sigma(p_{Mjt} - w_{jt}) + \sigma\lambda_2(S_{Tjt}) - (1-\sigma)\gamma_1(S_{Ojt}) \\
+ \widetilde{g}_{Lt-1}(\widetilde{h}_L(m_{jt-1} - l_{jt-1}, p_{Mjt-1} - w_{jt-1}, S_{Tjt-1}, S_{Ojt-1}), R_{jt-1}) + \widetilde{\xi}_{Ljt}, \qquad (11)
\end{aligned}$$

where the (conveniently rescaled) conditional expectation function is

$$\widetilde{g}_{Lt-1}(\widetilde{h}_L(\cdot), R_{jt-1}) = (1-\sigma) g_{Lt-1}\left(\frac{\widetilde{h}_L(\cdot)}{1-\sigma}, R_{jt-1}\right)$$

and $\widetilde{\xi}_{Ljt} = (1-\sigma)\xi_{Ljt}$.

In estimating equation (11), we allow $\widetilde{g}_{Lt-1}(\widetilde{h}_L(\cdot), R_{jt-1})$ to differ between zero and positive R&D expenditures and specify

$$\begin{aligned}
\widetilde{g}_{Lt-1}(\widetilde{h}_L(\cdot), R_{t-1}) &= \widetilde{g}_{L0}(t-1) + 1(R_{jt-1} = 0)\widetilde{g}_{L1}(\widetilde{h}_L(\cdot)) \\
&\quad + 1(R_{jt-1} > 0)\widetilde{g}_{L2}(\widetilde{h}_L(\cdot), r_{jt-1}), \qquad (12)
\end{aligned}$$

where $1(\cdot)$ is the indicator function and the functions $\widetilde{g}_{L1}(\widetilde{h}_L(\cdot))$ and $\widetilde{g}_{L2}(\widetilde{h}_L(\cdot), r_{jt-1})$ are modeled as described in Appendix B. Because the Markov processes governing productivity is time-inhomogeneous, we allow the conditional expectation function $\widetilde{g}_{Lt-1}(\widetilde{h}_L(\cdot), R_{jt-1})$ to shift over time by $\widetilde{g}_{L0}(t-1)$. In practice, we model this shift with time dummies.

Equation (11) is a semiparametric, partially linear, model with the additional restriction that the inverse function $\widetilde{h}_L(\cdot)$ is of known form. Identification in the sense of the ability to separate the parametric and nonparametric parts of the model follows from standard arguments (Robinson 1988, Newey, Powell & Vella 1999).

Labor l_{jt}, materials m_{jt}, the share of temporary labor S_{Tjt}, and the share of outsourced materials S_{Ojt} are correlated with $\widetilde{\xi}_{Ljt}$ (by virtue of the first-order conditions (4), (5), and (7) since $\widetilde{\xi}_{Ljt}$ is part of $\widetilde{\omega}_{Ljt}$). We therefore base estimation on the moment conditions

$$E\left[A_{Ljt}(z_{jt})\widetilde{\xi}_{Ljt}\right] = 0,$$

where $A_{Ljt}(z_{jt})$ is a vector of functions of the exogenous variables z_{jt} as described in Appendix B.

In considering instruments it is important to appreciate that equation (11) models the evolution of labor-augmenting productivity $\widetilde{\omega}_{Ljt}$. As a consequence, instruments have to be uncorrelated with the productivity innovation $\widetilde{\xi}_{Ljt}$ but not necessarily with productivity itself. Because $\widetilde{\xi}_{Ljt}$ is the innovation to productivity $\widetilde{\omega}_{Ljt}$ in period t, it is not known to the firm when it makes its decisions in period $t-1$. All past decisions are therefore uncorrelated with $\widetilde{\xi}_{Ljt}$.[15] Time t and the demand shifter D_{jt} are exogenous by construction.

Our data has firm-specific prices w_{jt} and p_{Mjt} for labor and materials that vary substantially both across firms and across periods. The larger part of the variation in the wage can be attributed to geographic and temporal differences in the supply of labor and the fact that firms operate in different product submarkets (see Appendix D for details). It is thus exogenous and useful for estimating equation (11).

The smaller part of the variation in the wage can be attributed to differences in the quality of labor and may therefore be correlated with the productivity innovation $\widetilde{\xi}_{Ljt}$ in equation (11). Extending the reasoning underlying Olley & Pakes (1996), Levinsohn & Petrin (2003), and Ackerberg et al. (2006) that lagged values are less susceptible to endogeneity than current values, we restrict ourselves to the lagged wage w_{jt-1} and the lagged price of materials p_{Mjt-1} for instruments. A test for overidentifying restrictions in Section 5 cannot reject that w_{jt-1} and p_{Mjt-1} are uncorrelated with the productivity innovation $\widetilde{\xi}_{Ljt}$ in our data. In Section 5, we also show that our estimates are robust to purging the variation due to differences in the quality of labor from the lagged wage w_{jt-1}.

Hicks-neutral productivity. Substituting the inverse functions in equations (10) and (10) into the production function in equation (1) and the law of motion for Hicks-neutral

[15]By contrast, l_{jt-1} is correlated with $\widetilde{\omega}_{Ljt}$ as long as productivity is correlated over time.

productivity ω_{Hjt} in equation (3), we form our second estimation equation[16]

$$y_{jt} = -\frac{\nu\sigma}{1-\sigma}x_{jt}$$
$$+g_{Ht-1}(h_H(k_{jt-1}, m_{jt-1}, S_{Mjt-1}, p_{jt-1}, p_{Mjt-1}, D_{jt-1}, S_{Tjt-1}, S_{Ojt-1}), R_{jt-1}) + \xi_{Hjt} + e_{jt}. \tag{13}$$

We specify $g_{Ht-1}(h_H(\cdot), R_{jt-1})$ analogously to $\widetilde{g}_{Lt-1}(\widetilde{h}_L(\cdot), R_{jt-1})$ in equation (12). Equation (13) is again a semiparametric model with the additional restriction that the inverse function $h_H(\cdot)$ is of known form.

Because output y_{jt}, materials m_{jt}, the share of materials in variable cost S_{Mjt}, the share of temporary labor S_{Tjt}, and the share of outsourced materials S_{Ojt} are correlated with ξ_{Hjt}, we base estimation on the moment conditions

$$E\left[A_{Hjt}(z_{jt})(\xi_{Hjt} + e_{jt})\right] = 0,$$

where $A_{Hjt}(z_{jt})$ is a vector of function of the exogenous variables z_{jt}. As before, we exploit the timing of decisions to rely on lags for instruments. In addition, $k_{jt} = \ln((1-\delta)K_{jt-1} + I_{jt-1})$ is determined in period $t-1$ and therefore uncorrelated with ξ_{Hjt}.

Estimation. We use the two-step GMM estimator of Hansen (1982). Let $\nu_{Ljt}(\theta_L) = \widetilde{\xi}_{Ljt}$ be the residual of estimation equation (11) as a function of the parameters θ_L to be estimated and $\nu_{Hjt}(\theta_H) = \xi_{Hjt} + e_{jt}$ the residual of estimation equation (13) as a function of θ_H. The GMM problem corresponding to equation (11) is

$$\min_{\theta_L} \left[\frac{1}{N}\sum_j A_{Lj}(z_j)\nu_{Lj}(\theta_L)\right]' \widehat{W}_L \left[\frac{1}{N}\sum_j A_{Lj}(z_j)\nu_{Lj}(\theta_L)\right], \tag{14}$$

where $A_{Lj}(z_j)$ is a $Q_L \times T_j$ matrix of functions of the exogenous variables z_j, $\nu_{Lj}(\theta_L)$ is a $T_j \times 1$ vector, $\widehat{W}_L$ is a $Q_L \times Q_L$ weighting matrix, Q_L is the number of instruments, T_j is the number of observations of firm j, and N is the number of firms. We provide further details in Appendix B.

The GMM problem corresponding to equation (13) is analogous. Equation (13) is considerably more nonlinear than equation (11). To facilitate its estimation, we impose the estimated values of those parameters in θ_L that also appear in θ_H. We correct the standard errors as described in the Online Appendix. Because they tend to be more stable, we report first-step estimates for equation (13) and use them in the subsequent analysis; however, we

[16]There are other possible estimation equations. In particular, one can use the labor and materials decisions in equations (6) and (8) together with the production function in equation (1) to recover $\widetilde{\omega}_{Ljt}$, ω_{Hjt}, and e_{jt} and then set up separate moment conditions in $\widetilde{\xi}_{Ljt}$, ξ_{Hjt}, and e_{jt}. This may yield efficiency gains. Our estimation equation (13) has the advantage that it is similar to a CES production function that has been widely estimated in the literature.

use second-step estimates for testing.

5 Labor-augmenting technological change

To provide insight and relate our empirical strategy to the literature, it is helpful to abstract from the distinction between permanent and temporary labor and in-house and outsourced materials. To this end, we follow Levinsohn & Petrin (2003) and assume that labor l_{jt} and materials m_{jt} are homogenous inputs that are chosen each period to maximize short-run profits.[17] This implies $\lambda_1(S_{Tjt}) = 1$, $\lambda_2(S_{Tjt}) = 0$, and $\gamma_1(S_{Ojt}) = 0$, so that the simplified model emerges as a special case as the correction terms on labor and materials vanish.

In the simplified model, equation (9) can be rewritten as

$$m_{jt} - l_{jt} = -\widetilde{\gamma}_L - \sigma(p_{Mjt} - w_{jt}) + \widetilde{\omega}_{Ljt}. \tag{15}$$

Equation (15) shows that in the presence of labor-augmenting technological change, materials per unit of labor varies over time and across firms for two reasons. First, it varies according to the price of materials p_{Mjt} relative to the price of labor w_{jt}. For example, if the relative price of materials falls, then materials per unit of labor rises. Second, labor-augmenting technological change increases materials per unit of labor. A rise in $\widetilde{\omega}_{Ljt}$ ceteris paribus causes a rise in materials per unit of labor. This reflects the displacement effect of labor-augmenting technological change.

Related literature. While we use equation (9) to recover labor-augmenting productivity $\widetilde{\omega}_{Ljt}$, much empirical work directly estimates relationships like it. Equation (15) with skilled and unskilled workers in place of materials and labor is at the heart of the literature on skill bias (see Card & DiNardo (2002) and Violante (2008) and the references therein); with capital in place of materials, equation (15) serves to estimate the elasticity of substitution σ in an aggregate value-added production function (see Antràs 2004).

Equation (15) is often estimated by OLS. The problem is that labor-augmenting productivity, which is not observed by the econometrician, is correlated over time and also with the wage. The wage is likely to be higher when labor is more productive, even if it adjusts slowly with some lag. This positive correlation induces an upward bias in the estimate of the elasticity of substitution. This is a variant of the endogeneity problem in estimating production functions.

It is widely recognized that the estimate of the elasticity of substitution may be biased as a result. Proxying for unobserved productivity by a time trend, time dummies, or a measure of innovation is unlikely to completely remove the bias. Antràs (2004) shows that

[17] Levinsohn & Petrin (2003) invoke this assumption to establish in their equation (9) a sufficient condition for the invertibility of the intermediate input: On p. 320, just below equation (1), they assume that labor is "freely variable," on p. 322, just above equation (6), they assume that the intermediate input is also "freely variable," and they invoke short-run profit maximization at the start of the proof on p. 339.

the estimate of the elasticity of substitution improves by including a time trend and allowing for serial correlation in the remaining error term. However, less than fully accounting for the evolution of productivity leaves an error term that likely remains correlated with the ratio of prices. Using firm-level panel data, Van Reenen (1997) proxies for unobserved productivity by the number of innovations commercialized in a given year. His approach assumes that the remaining error term is white noise and is thus unlikely to succeed if productivity is governed by a more general stochastic process.[18] Also using firm-level panel data, Raval (2013) estimates the elasticity of substitution in a variant of equation (15) obtained from a value-added production function with capital- and labor-augmenting productivity.[19] This rests on the assumption that capital and labor are both static inputs that are chosen each period to maximize short-run profits. Proxying for the firm-specific wage by a regional wage index and for the price of capital by a dummy, Raval (2013) runs OLS by year and sometimes by industry. While not using time-series variation may alleviate the endogeneity problem, relying on proxies introduces measurement error as a source of bias.

Modeling the evolution of labor-augmenting productivity as we do in our estimation equation (11) is a natural way to resolve the endogeneity problem. Breaking out the part of $\widetilde{\omega}_{Ljt}$ that is observable via the conditional expectation function $\widetilde{g}_{Lt-1}(\cdot)$ leaves "less" in the error term and intuitively diminishes the endogeneity problem. Moreover, to the extent that a correlation between the included variables and the error term remains, it is straightforward to instrument: As discussed in Section 4, the key advantage of equation (11) over equation (15) is that instruments have to be uncorrelated with the productivity innovation $\widetilde{\xi}_{Ljt}$ but not necessarily with productivity itself.

Elasticity of substitution. Tables 3 and 4 summarize different estimates of the elasticity of substitution. To facilitate the comparison with the existing literature, we proxy for $\widetilde{\omega}_{Ljt}$ in equation (15) by a time trend $\widetilde{\delta}_L t$ and estimate by OLS. As can be seen from columns (3) and (4) of Table 3, with the exception of industry 9, the estimates of the elasticity of substitution are in excess of one, whereas the estimates in the previous literature lie somewhere between 0 and 1 (Chirinko 2008, Oberfield & Raval 2014). This reflects, first, that a time trend is a poor proxy for labor-augmenting technological change at the firm level and, second, that the estimates are upward biased as a result of the endogeneity problem. Nevertheless, the significant positive time trend once again previews the importance of labor-augmenting technological change.

We resolve the endogeneity problem by modeling the evolution of labor-augmenting productivity and estimating equation (11) by GMM. Columns (5)–(10) of Table 3 refer to the simplified model with $\lambda_1(S_{Tjt}) = 1$, $\lambda_2(S_{Tjt}) = 0$, and $\gamma_1(S_{Ojt}) = 0$. As expected the

[18] Indeed, Van Reenen (1997) obtains a positive direct effect of innovation on employment, contrary to the displacement effect of labor-augmenting technological change.

[19] See Gandhi et al. (2013) for a recent discussion of the drawbacks of estimating a value-added instead of a gross-output production function.

estimates of the elasticity of substitution are much lower and range from 0.45 to 0.64 as can be seen from column (5). With the exception of industries 6 and 8 in which σ is either implausibly high or low, we clearly reject the special cases of both a Leontieff ($\sigma \to 0$) and a Cobb-Douglas ($\sigma = 1$) production function.

Testing for overidentifying restrictions, we reject the validity of the moment conditions at a 5% level in five industries and we are close to rejecting in two more industries (columns (6) and (7)). To pinpoint the source of this problem, we exclude the subset of moments involving lagged materials m_{jt-1} from the estimation. As can be seen from columns (8)–(10), the estimates of the elasticity of substitution lie between 0.46 and 0.84 in all industries and at a 5% level we can no longer reject the validity of the moment conditions in any industry.

To see why the exogeneity of lagged materials m_{jt-1} is violated contrary to the timing of decisions in our model, recall that a firm engages in outsourcing if it can procure parts and pieces from its suppliers that are cheaper or better than what the firm can make in house from scratch. These quality differences are pushed into the error term by lumping in-house and outsourced materials together. As outsourcing often relies on contractual relationships between the firm and its suppliers, the resulting error term is likely correlated over time and thus with lagged materials m_{jt-1} as well.

The correction terms in equation (11) account for outsourcing as well as adjustment costs on permanent labor that may drive a wedge between the relative quantities and prices of materials and labor. As can be seen in columns (3)–(5) of Table 4, the correction terms duly restore the exogeneity of lagged materials m_{jt-1} as we cannot reject the validity of the moment conditions at a 5% level in any industry except for industry 7 in which we (barely) reject.[20] Our leading estimates of σ in column (3) of Table 4 lie between 0.44 and 0.80. Compared to the estimates in column (8) of Table 3, there are no systematic changes and our leading estimates are somewhat lower in five industries and somewhat higher in five industries. In sum, accounting for outsourcing and adjustment costs on permanent labor is an improvement over the assumption in Levinsohn & Petrin (2003) and many others that labor and materials are homogenous and static inputs and a key step in estimating the elasticity of substitution.

Our estimates of the elasticity of substitution are robust to purging the variation due to differences in the quality of labor from the lagged wage w_{jt-1}. In Appendix D, we use a wage regression to isolate the part of the wage that depends on the available data on the skill mix of a firm's labor force. Using $\widehat{w}_{Qjt-1}$ to denote this part, we replace w_{jt-1} as an instrument by $w_{jt-1} - \widehat{w}_{Qjt-1}$. Compared to column (3) of Table 4, the estimates of the elasticity of substitution in column (6) decrease somewhat in three industries, remain essentially unchanged in two industries, and increase somewhat in five industries.[21] The absence of

[20] As noted in Section 3, we exclude observations with $S_{Tjt} = 0$ and thus $L_{Tjt} = 0$ because equation (6) presumes interior solutions for permanent and temporary labor.

[21] As we show in the Online Appendix, not much changes if we isolate the part of the wage that additionally depends on firm size to try and account for the quality of labor beyond our rather coarse data on the skill mix

substantial and systematic changes confirms that the variation in w_{jt-1} is exogenous with respect to $\widetilde{\xi}_{Ljt}$ and therefore useful in estimating equation (11).

Labor-augmenting technological change. With equation (11) estimated, we recover the labor-augmenting productivity $\omega_{Ljt} = \frac{\widetilde{\omega}_{Ljt}}{1-\sigma}$ of firm j in period t up to an additive constant from equation (9). We take the growth of labor-augmenting productivity at firm j in period t to be $\Delta\omega_{Ljt} = \omega_{Ljt} - \omega_{Ljt-1} \approx \frac{\exp(\omega_{Ljt})-\exp(\omega_{Ljt-1})}{\exp(\omega_{Ljt-1})}$.[22] To obtain aggregate measures representing an industry, we account for the survey design by replicating the subsample of small firms $\frac{70\%}{5\%} = 14$ times before pooling it with the subsample of large firms. We report weighed averages of individual measures in Table 5, where the weight $\mu_{jt} = Y_{jt-2}/\sum_j Y_{jt-2}$ is the share of output of firm j in period $t-2$.

In line with the patterns in the data described in Section 2, our estimates imply an important role for labor-augmenting technological change. As can be seen from column (1), labor-augmenting productivity grows quickly, on average, with rates of growth ranging from 1.0% per year in industry 7 to 18.3% in industry 6 and above in industry 5. The rate of growth is, on average, slightly negative in industry 9. Hidden behind these averages is a tremendous amount of heterogeneity across firms. The rate of growth is positively correlated with the level of labor-augmenting productivity (column (2)), indicating that differences in labor-augmenting productivity between firms persist over time.

Ceteris paribus $\Delta\omega_{Ljt} \approx \frac{\exp(\omega_{Ljt})L^*_{jt-1}-\exp(\omega_{Ljt-1})L^*_{jt-1}}{\exp(\omega_{Ljt-1})L^*_{jt-1}}$ approximates the rate of growth of a firm's effective labor force $\exp(\omega_{Ljt-1})L^*_{jt-1}$. We approximate the rate of growth of the firm's output Y_{jt-1} by $\epsilon_{Ljt-2}\Delta\omega_{Ljt}$, where ϵ_{Ljt-2} is the elasticity of output with respect to the firm's effective labor force in period $t-2$ (see Appendix C for details). This output effect, while close to zero in industry 9, ranges on average from 0.7% per year in industries 7 and 8 to 3.6% in industry 6, see column (3) of Table 5. Overall, labor-augmenting technological change causes output to grow in the vicinity of 2% per year.

Figure 1 illustrates the magnitude of labor-augmenting technological change and the heterogeneity in its impact across industries. The depicted index cumulates the year-to-year changes in labor-augmenting productivity in terms of output effects and is normalized to one in 1991.

Firms' R&D activities. While there is practically no difference in two industries, in eight industries firms that perform R&D have higher levels of labor-augmenting productivity than firms that do not perform R&D as can be seen from column (4) of Table 5. The rate of

of a firm's labor force (Oi & Idson 1999). Compared to column (3) of Table 4, the estimates of the elasticity of substitution decrease somewhat in three industries, remain essentially unchanged in three industries, and increase somewhat in four industries.

[22]Given the specification of $\widetilde{g}_{Lt-1}(\widetilde{h}_L(\cdot), R_{jt-1})$ in equation (12), we exclude observations where a firm switches from performing to not performing R&D or vice versa between periods $t-1$ and t from the subsequent analysis. We further exclude observations where a firm switches from zero to positive outsourcing or vice versa.

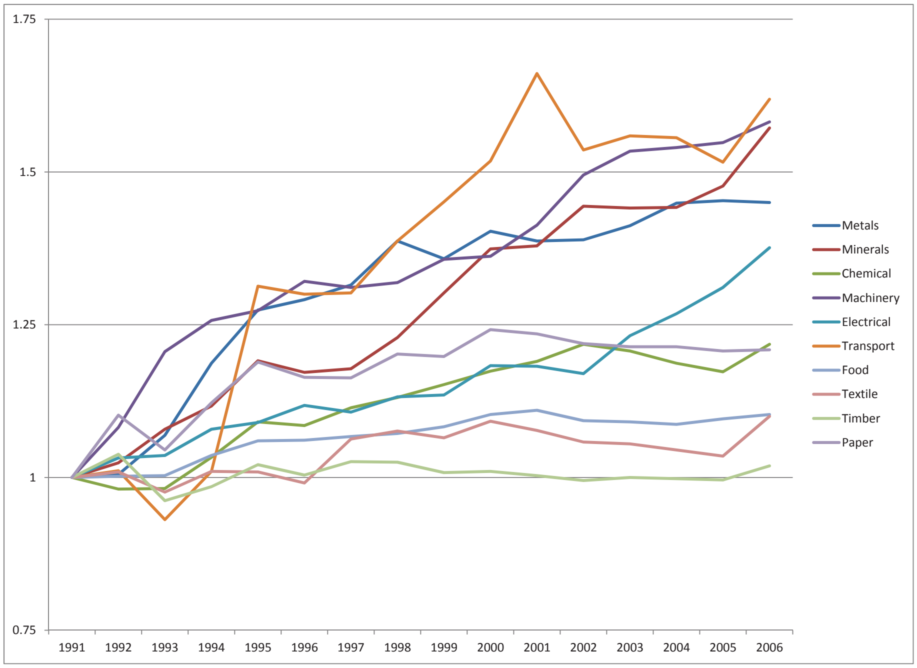

Figure 1: Labor-augmenting technological change. Output effect. Index normalized to one in 1991.

growth of labor-augmenting productivity for firms that perform R&D, on average, exceeds that of firms that do not perform R&D in eight industries. As can be seen from columns (5) and (6) of Table 5, the output effect for firms that perform R&D exceeds that of firms that do not perform R&D in six industries. Overall, our estimates indicate that firms' R&D activities are associated not only with higher levels of labor-augmenting productivity but by and large also with higher rates of growth of labor-augmenting productivity. Firms' R&D activities play a key role in determining the differences in labor-augmenting productivity across firms and the evolution of this component of productivity over time.

Skill upgrading. In our data, there is a shift from unskilled to skilled workers. For example, the share of engineers and technicians in the labor force increases from 7.2% in 1991 to 12.3% in 2006. While this shift has to be seen against the backdrop of a general increase of university graduates in Spain during the 1990s and 2000s, it begs the question how much skill upgrading contributes to the growth of labor-augmenting productivity.

To answer this question, we leverage our rather coarse data on the skill mix of a firm's labor force. Besides the share of temporary labor S_{Tjt}, our data has the share of white collar workers and the shares of engineers and technicians, respectively.[23]

We assume that there are Q types of permanent labor with qualities $1, \theta_2, \ldots, \theta_Q$ and corresponding wages $W_{P1jt}, W_{P2jt}, \ldots, W_{PQjt}$. The firm, facing this menu of qualities and wages, behaves as a price-taker in the labor market. In recognition of their different qualities, $L^*_{Pjt} = L_{P1jt} + \sum_{q=2}^{Q} \theta_q L_{Pqjt}$ is an aggregate of the Q types of permanent labor, with L_{Pqjt} being the quantity of permanent labor of type q at firm j in period t. $L^*_{jt} = \Lambda(L^*_{Pjt}, L_{Tjt})$ is the aggregate of permanent labor L^*_{Pjt} (instead of $L_{Pjt} = \sum_{q=1}^{Q} L_{Pqjt}$) and temporary labor L_{Tjt} in the production function in equation (1). Permanent labor is subject to convex adjustment costs $C_{B_P}(B_{Pjt}, B_{Pjt-1})$, where $B_{Pjt} = \sum_{q=1}^{Q} W_{Pqjt} L_{Pqjt}$ is the wage bill for permanent labor. The state vector Ω_{jt} therefore includes B_{Pjt-1}, $W_{P1jt}, W_{P2jt}, \ldots, W_{PQjt}$ instead of L_{Pjt-1} and W_{Pjt}.

The first-order condition for permanent labor of type q is

$$\nu \mu X_{jt}^{-\left(1+\frac{\nu\sigma}{1-\sigma}\right)} \exp(\omega_{Hjt}) \exp\left(-\frac{1-\sigma}{\sigma}\omega_{Ljt}\right) \left(L^*_{jt}\right)^{-\frac{1}{\sigma}} \frac{\partial L^*_{jt}}{\partial L^*_{Pjt}} \theta_q = \frac{W_{Pqjt}(1+\Delta_{jt})}{P_{jt}\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}, \quad (16)$$

where $\theta_1 = 1$ and the gap between the wage W_{Pqjt} and the shadow wage is

$$\begin{aligned} \Delta_{jt} &= \frac{\partial C_{B_P}(B_{Pjt}, B_{Pjt-1})}{\partial B_{Pjt}} - \frac{1}{W_{Pqjt}} \frac{1}{1+\rho} E_t\left[\frac{\partial V_{t+1}(\Omega_{jt+1})}{\partial L_{Pqjt}} | \Omega_{jt}, R_{jt}\right] \\ &= \frac{\partial C_{B_P}(B_{Pjt}, B_{Pjt-1})}{\partial B_{Pjt}} + \frac{1}{1+\rho} E_t\left[\frac{\partial C_{B_P}(B_{Pjt+1}, B_{Pjt})}{\partial B_{Pjt}} | \Omega_{jt}, R_{jt}\right]. \end{aligned}$$

Equation (16) implies that $\theta_q = \frac{W_{Pqjt}}{W_{P1jt}}$ at an interior solution. While our data does not have $W_{P1jt}, W_{P2jt}, \ldots, W_{PQjt}$, the wage regression in Appendix D enables us to recover θ_q by estimating the wage premium $\left(\frac{W_{Pqjt}}{W_{P1jt}} - 1\right)$ of permanent labor of type q over type 1.

Multiplying equation (16) by the share S_{Pqjt} of permanent workers of type q and summing yields

$$\nu \mu X_{jt}^{-\left(1+\frac{\nu\sigma}{1-\sigma}\right)} \exp(\omega_{Hjt}) \exp\left(-\frac{1-\sigma}{\sigma}\omega_{Ljt}\right) \left(L^*_{jt}\right)^{-\frac{1}{\sigma}} \frac{\partial L^*_{jt}}{\partial L^*_{Pjt}} \Theta_{jt} = \frac{W_{Pjt}(1+\Delta_{jt})}{P_{jt}\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}, \quad (17)$$

where $\Theta_{jt} = S_{P1jt} + \sum_{q=2}^{Q} \theta_q S_{Pqjt} = 1 + \sum_{q=2}^{Q} \left(\frac{W_{Pqjt}}{W_{P1jt}} - 1\right) S_{Pqjt}$ is a quality index and

[23] We have these latter measures in the year a firm enters the sample and every subsequent four years. We take the skill mix to be unchanging in the interim.

$W_{Pjt} = \sum_{q=1}^{Q} W_{Pqjt} S_{Pqjt}$. Using Euler's theorem to combine equations (5) and (17) yields

$$\nu\mu X_{jt}^{-\left(1+\frac{\nu\sigma}{1-\sigma}\right)} \exp\left(\omega_{Hjt}\right) \exp\left(-\frac{1-\sigma}{\sigma}\omega_{Ljt}\right) L_{jt}^{-\frac{1}{\sigma}} \Lambda((1-S_{Tjt})\Theta_{jt}, S_{Tjt})^{-\frac{1-\sigma}{\sigma}}$$

$$= \frac{W_{jt}\left(1+\frac{\Delta_{jt}}{1+\frac{W_{Tjt}}{W_{Pjt}}\frac{S_{Tjt}}{1-S_{Tjt}}}\right)}{P_{jt}\left(1-\frac{1}{\eta(p_{jt},D_{jt})}\right)} = \frac{W_{jt}\left(\frac{\frac{\Lambda_P((1-S_{Tjt})\Theta_{jt},S_{Tjt})\Theta_{jt}}{\Lambda_T((1-S_{Tjt})\Theta_{jt},S_{Tjt})}+\frac{S_{Tjt}}{1-S_{Tjt}}}{\frac{W_{Pjt}}{W_{Tjt}}+\frac{S_{Tjt}}{1-S_{Tjt}}}\right)}{P_{jt}\left(1-\frac{1}{\eta(p_{jt},D_{jt})}\right)}, \tag{18}$$

where the second equality follows from dividing equations (5) and (17) and solving for Δ_{jt}. We proceed as before by assuming that $\frac{W_{Pjt}}{W_{Ljt}} = \lambda_0$ is an (unknown) constant and treating $\frac{\frac{\Lambda_P((1-S_{Tjt})\Theta_{jt},S_{Tjt})\Theta_{jt}}{\Lambda_T((1-S_{Tjt})\Theta_{jt},S_{Tjt})}+\frac{S_{Tjt}}{1-S_{Tjt}}}{\lambda_0+\frac{S_{Tjt}}{1-S_{Tjt}}} = \lambda_1(S_{Tjt}, \Theta_{jt})$ as an (unknown) function of S_{Tjt} and Θ_{jt} that must be estimated nonparametrically. Replacing $\lambda_2(S_{Tjt}) = \ln\left(\lambda_1(S_{Tjt})\Lambda(1-S_{Tjt}, S_{Tjt})^{\frac{1-\sigma}{\sigma}}\right)$ by $\lambda_2(S_{Tjt}, \Theta_{jt}) = \ln\left(\lambda_1(S_{Tjt}, \Theta_{jt})\Lambda((1-S_{Tjt})\Theta_{jt}, S_{Tjt})^{\frac{1-\sigma}{\sigma}}\right)$ in our estimation equation (11) therefore accounts for types of permanent labor that differ in their qualities and wages.

The estimates of the elasticity of substitution in column (7) of Table 5 continue to hover around 0.6 across industries, with the exception of industries 4 and 8 in which they are implausibly low. Compared to column (3) of Table 4, they decrease somewhat in three industries, remain essentially unchanged in two industries, and increase somewhat in five industries. We develop the quality index Θ_{jt} to "chip away at the productivity residual" by improving the measurement of inputs in the spirit of Caselli (2005) and the earlier productivity literature (Jorgenson 1995*a*, Jorgenson 1995*b*). As can be seen from column (10) of Table 5, skill upgrading indeed explains some, but by no means all of the growth of labor augmenting productivity. Compared to column (1), the rates of growth stay the same or go down in all industries. In industries 7, 8, 9, and 10 labor-augmenting productivity is stagnant or declining after accounting for skill upgrading, indicating that improvements in the skill mix over time are responsible for most of the growth of labor-augmenting productivity. In contrast, in industries 1, 2, 3, 4, 5, and 6, labor-augmenting productivity continues to grow after accounting for skill upgrading, albeit often at a much slower rate. In these industries, labor-augmenting productivity grows also because workers with a given set of skills become more productive over time.

6 Hicks-neutral technological change

From equation (11) we obtain an estimate of the elasticity of substitution and recover labor-augmenting productivity at the firm level. To recover Hicks-neutral productivity and the remaining parameters of the production function, we have to estimate equation (13).

Distributional parameters and elasticity of scale. Table 6 reports the distributional parameters β_K and $\beta_M = 1 - \beta_K$ and the elasticity of scale ν. Our estimates of β_K range from 0.07 in industry 8 to 0.31 in industry 6 (column (1)). Although the estimates of the elasticity of scale are rarely significantly different from one, taken together they suggest slightly decreasing returns to scale (columns (2)). We cannot reject the validity of the moment conditions in any industry by a wide margin (columns (3) and (4)).

Hicks-neutral technological change. With equation (13) estimated, we recover Hicks-neutral productivity ω_{Hjt} up to an additive constant from equation (10). We take the growth of Hicks-neutral productivity at firm j in period t to be $\Delta\omega_{Hjt} = \omega_{Hjt} - \omega_{Hjt-1} \approx \frac{\exp(\omega_{Hjt})-\exp(\omega_{Hjt-1})}{\exp(\omega_{Hjt-1})}$. Ceteris paribus $\Delta\omega_{Hjt} \approx \frac{X_{jt-1}^{-\frac{\nu\sigma}{1-\sigma}}\exp(\omega_{Hjt})\exp(e_{jt-1})-X_{jt-1}^{-\frac{\nu\sigma}{1-\sigma}}\exp(\omega_{Hjt-1})\exp(e_{jt-1})}{X_{jt-1}^{-\frac{\nu\sigma}{1-\sigma}}\exp(\omega_{Hjt-1})\exp(e_{jt-1})}$ approximates the rate of growth of a firm's output Y_{jt-1}. The rate of growth of Hicks-neutral productivity is therefore directly comparable to the output effect of labor-augmenting technological change. We proceed as before to obtain aggregate measures representing an industry.

As can be seen from column (1) of Table 7, Hicks-neutral productivity grows quickly in five industries, with rates of growth ranging, on average, from 1.2% per year in industry 8 to 4.4% in industry 1. It grows much more slowly or barely at all in three industries, with rates of growth below 0.5% per year. While there is considerable heterogeneity in the rate of growth of Hicks-neutral productivity across industries, overall Hicks-neutral technological change causes output to grow in the vicinity of 2% per year. Once again, the rate of growth is positively correlated with the level of Hicks-neutral productivity (column (2)), indicating that differences in Hicks-neutral productivity between firms persist over time.

Figure 2 illustrates the magnitude of Hicks-neutral technological change. The depicted index cumulates the year-to-year changes in Hicks-neutral productivity and is normalized to one in 1991.[24] The heterogeneity in the impact of Hicks-neutral technological change across industries clearly exceeds that of labor-augmenting technological change (see again Figure 1).

Taken together labor-augmenting and Hicks-neutral technological change cause output to grow by, on average, between 0.7% in industry 7 and 7.8% in industry 6, as can be seen in column (3) of Table 7. The components of productivity are positively correlated. This correlation is slightly stronger in the rates of growth (column (4)) as it is in levels.

Firms' R&D activities. As can be seen from column (5) of Table 7, firms that perform R&D have higher levels of Hicks-neutral productivity than firms that do not perform R&D in six industries but lower levels of Hicks-neutral productivity in four industries. While there is practically no difference in industry 10, the rate of growth of Hicks-neutral productivity for firms that perform R&D, on average, exceeds that of firms that do not perform R&D

[24] In industry 9, in line with column (1) of Table 7, we trim values of $\Delta\omega_H$ below -0.25 and above 0.5.

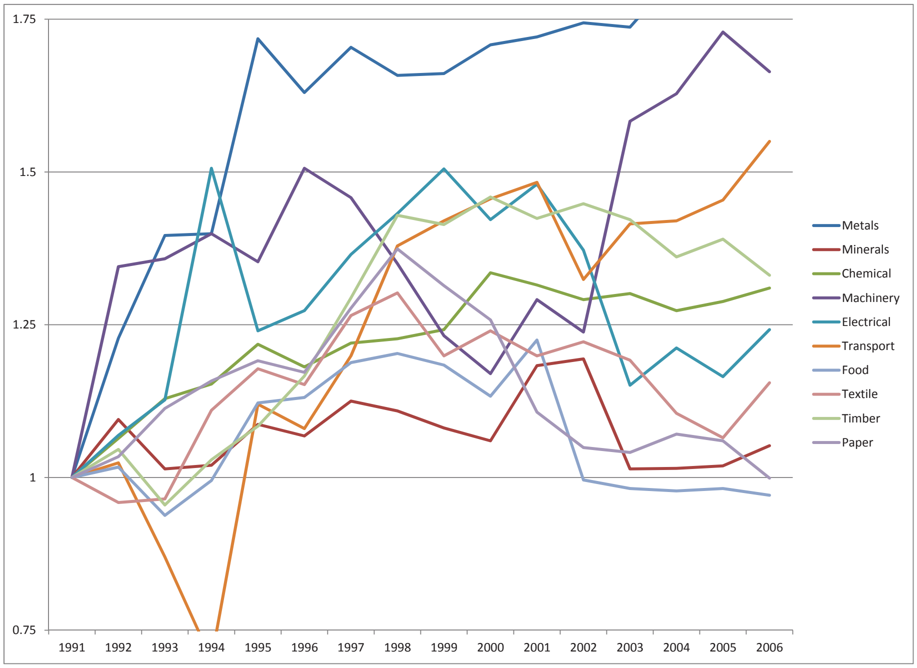

Figure 2: Hicks-neutral technological change. Index normalized to one in 1991.

in five industries, as can be seen from columns (6) and (7). Overall, our estimates indicate that firms' R&D activities are associated with higher levels and rates of growth of Hicks-neutral productivity, although firms' R&D activities seem less closely tied to Hicks-neutral than to labor-augmenting productivity. This is broadly consistent with the large literature on induced innovation that dates back at least to Hicks (1966) and argues that firms direct their R&D activities to conserve the relatively more expensive factors of production, in particular labor.

Capital-augmenting productivity. As discussed in Section 2, the evolution of the relative quantities and prices of the various factors of production provides no evidence for capital-augmenting technological change. Our leading specification therefore restricts the efficiencies of capital and materials to change at the same rate and in lockstep with Hicks-neutral technological change. A more general specification allows for capital-augmenting productivity ω_{Kjt} so that equation (1) (with $\beta_0 = \beta_L = 1$) becomes

$$Y_{jt} = \left[\beta_K \left(\exp(\omega_{Kjt}) K_{jt}\right)^{-\frac{1-\sigma}{\sigma}} + \left(\exp(\omega_{Ljt}) L^*_{jt}\right)^{-\frac{1-\sigma}{\sigma}} + \beta_M \left(M^*_{jt}\right)^{-\frac{1-\sigma}{\sigma}}\right]^{-\frac{\nu\sigma}{1-\sigma}} \exp(\omega_{Hjt}) \exp(e_{jt}). \tag{19}$$

We explore the potential role of capital-augmenting productivity in our data in two ways.

First, we follow Raval (2013) and parts of the previous literature on estimating aggregate production functions (see Antràs (2004) and the references therein) and assume that capital is a static input that is chosen each period to maximize short-run profits. In analogy to equation (9), we recover (conveniently rescaled) capital-augmenting productivity $\widetilde{\omega}_{Kjt} = (1-\sigma)\omega_{Kjt}$ as

$$\begin{aligned} \widetilde{\omega}_{Kjt} &= \widetilde{\gamma}_K + m_{jt} - k_{jt} + \sigma(p_{Mjt} - p_{Kjt}) + (1-\sigma)\gamma_1(S_{Ojt}) \\ &\equiv \widetilde{h}_K(m_{jt} - k_{jt}, p_{Mjt} - p_{Kjt}, S_{Ojt}), \end{aligned} \tag{20}$$

where $\widetilde{\gamma}_K = -\sigma \ln\left(\frac{\beta_M}{\beta_K}\right)$ and we use the user cost of capital in our data as a rough measure of the price of capital p_{Kjt}. Using our leading estimates from Section 5, we recover the capital-augmenting productivity $\omega_{Kjt} = \frac{\widetilde{\omega}_{Kjt}}{1-\sigma}$ of firm j in period t.[25] $\Delta\omega_{Kjt}$ in column (1) of Table 8 approximates the rate of growth of a firm's effective capital stock $\exp(\omega_{Kjt-1})K_{jt-1}$ and $\epsilon_{Kjt-2}\Delta\omega_{Kjt}$ in column (2) the rate of growth of the firm's output Y_{jt-1}, where ϵ_{Kjt-2} is the elasticity of output with respect to the firm's effective capital stock (see Appendix C). As can be seen from column (1), capital-augmenting productivity grows slowly, on average, with rates of growth of 0.8% per year in industry 6, 2.2% in industry 10, and 5.6% in industry 1. The rate of growth is negative in the remaining seven industries. The growth of capital-augmenting productivity is especially underwhelming in comparison to the growth of labor-augmenting productivity (see again column (1) of Table 5). The output effect in column (2) is also close to zero in all industries, although this likely reflects the fact that capital is not a static input. As the user cost of capital excludes adjustment costs, it falls short of the shadow price of capital, and using it drives down the elasticity of output with respect to the firm's effective capital stock.

Second, we return to the usual setting in the literature following Olley & Pakes (1996) and allow the choice of capital to have dynamic implications. We follow parts of the previous literature on estimating aggregate production functions and proxy for ω_{Kjt} by a time trend $\delta_K t$. Our second estimation equation (13) remains unchanged except that

$$X_{jt} = \beta_K \left(\exp(\delta_K t)K_{jt}\right)^{-\frac{1-\sigma}{\sigma}} + \beta_M \left(M_{jt} \exp\left(\gamma_1(S_{Ojt})\right)\right)^{-\frac{1-\sigma}{\sigma}} \left(\frac{1-S_{Mjt}}{S_{Mjt}}\lambda_1(S_{Tjt}) + 1\right).$$

Columns (3)–(7) of Table 8 summarize the resulting estimates of β_K, ν, and δ_K. The estimates of β_K and ν are very comparable to those in Table 5. Moreover, the insignificant time trend leaves little room for capital-augmenting technological change in our data.

[25] As an alternative to plugging our leading estimates from Section 5 into equation (20), in the Online Appendix we use equation (20) to form the analog to our first estimation equation (11):

$$\begin{aligned} m_{jt} - k_{jt} = -\sigma(p_{Mjt} - p_{Kjt}) - (1-\sigma)\gamma_1(S_{Ojt}) \\ + \widetilde{g}_{Kt-1}(\widetilde{h}_K(m_{jt-1} - k_{jt-1}, p_{Mjt-1} - p_{Kjt-1}, S_{Ojt-1}), R_{jt-1}) + \widetilde{\xi}_{Kjt}. \end{aligned}$$

Consistent with measurement error in p_{Kjt}, the resulting estimates of σ are very noisy and severely biased toward zero.

In sum, in line with the patterns in the data described in Section 2, there is little, if any, evidence for capital-augmenting technological change in our data. Of course, our ways of exploring the potential role of capital-augmenting productivity are less than ideal in that they either rest on the assumption that capital is a static input or abstract from firm-level heterogeneity in capital-augmenting productivity. An important question is therefore whether our approach can be extended to treat capital-augmenting productivity on par with labor-augmenting and Hicks-neutral productivity. Recovering a third component of productivity, at a bare minimum, requires a third decision to invert besides labor and materials. Investment is a natural candidate. Unlike the demand for labor and materials, however, investment depends on the details of the firm's dynamic programming problem. Hence, it may have to be inverted nonparametrically as in Olley & Pakes (1996). There are two principal difficulties. First, one has to prove that the observed demands for labor and materials along with investment are jointly invertible for unobserved capital-augmenting, labor-augmenting, and Hicks-neutral productivity. This is not an easy task given the difficulties Buettner (2005) encountered in a much simpler dynamic programming problem. Second, the inverse functions $\widetilde{h}_K(\cdot)$, $\widetilde{h}_L(\cdot)$, and $h_H(\cdot)$ are high-dimensional. Thus, estimating these functions nonparametrically is demanding on the data.

Related literature. As mentioned in Section 1, our paper is related to Grieco et al. (2014) and subsequent work in progress by Zhang (2014*a*, 2014*b*). These papers build on Doraszelski & Jaumandreu (2013) by exploiting the parameter restrictions between the production function and input demand functions to infer unobservables from observables. Because their data contains the materials bill rather than its split into price and quantity, Grieco et al. (2014) assume that labor and materials are both static inputs that are chosen each period to maximize short-run profits and solve the implied first-order conditions for the firm's Hicks-neutral productivity and the price of materials that the firm faces. Zhang (2014*a*, 2014*b*) proxies for the price of materials by a regional price index (similar to Raval 2013) and instead solves the first-order conditions for the firm's capital- and labor-augmenting productivity. One difference to our approach is that Grieco et al. (2014) and Zhang (2014*a*, 2014*b*) plug the recovered unobservables back into the production function. While this avoids assumptions on the law of motion for productivity, parameters of interest may cancel depending on the specification of the production function (see Example 3.1 of Grieco et al. (2014) and Section 4 of Ackerberg et al. (2006)).

Using firm-level panel data for the Chinese steel industry, Zhang (2014*b*) adds Hicks-neutral productivity to the model in Zhang (2014*a*) and specifies an uncontrolled first-order, time-homogenous Markov process for it. He infers this additional unobservable from investment (though without proving invertibility). The empirical strategy draws on Ackerberg et al. (2006) in that Hicks-neutral productivity ω_{Hjt} is separated from the random shock e_{jt} in a first nonparametric step (though without accounting for prices and correcting for the endogeneity of the revenue shares of labor and materials with respect to the random shock). In a second step, the parameters of the production function are estimated off the law of motion for ω_{Hjt}. By comparing the means and standard deviations of ω_{Kjt}, ω_{Ljt}, and ω_{Hjt}, Zhang (2014*b*) concludes that firm-level heterogeneity is largest in labor-augmenting productivity (though this conclusion can be questioned by recalling that ω_{Kjt}, ω_{Ljt}, and ω_{Hjt} can only be recovered up to additive constants and are measured in non-comparable units anyway).

7 Conclusions

Technological change can increase the productivity of capital, labor, and the other factors of production in equal terms or it can be biased towards a specific factor. In this paper, we directly assess the bias of technological change by measuring, at the level of the individual firm, how much of technological change is Hicks neutral and how much of it is labor augmenting.

To this end, we develop a dynamic model of the firm in which productivity is multidimensional. At the center of the model is a CES production function that parsimoniously relates the relative quantities of materials and labor to their relative prices and labor-augmenting productivity. To properly isolate and measure labor-augmenting productivity, we account for other factors that impact this relationship, in particular, outsourcing and adjustment costs on permanent labor.

We apply our estimator to an unbalanced panel of 2375 Spanish manufacturing firms in ten industries from 1990 to 2006. Our estimates indicate limited substitutability between the various factors of production. This calls into question whether the widely-used Cobb-Douglas production function with its unitary elasticity of substitution adequately represents firm-level production processes.

Our estimates provide clear evidence that technological change is biased. Ceteris paribus labor-augmenting technological change causes output to grow, on average, in the vicinity of 2% per year. While skill upgrading explains some of the growth of labor augmenting productivity, in many industries labor-augmenting productivity grows because workers with a given set of skills become more productive over time. In short, our estimates cast doubt on the assumption of Hicks-neutral technological change that underlies many of the standard techniques for measuring productivity and estimating production functions.

At the same time, our estimates do not validate the assumption that technological change is purely labor augmenting that plays a central role in the literature on economic growth. In addition to labor-augmenting technological change, our estimates show that Hicks-neutral technological change causes output to grow, on average, in the vicinity of 2% per year.

Behind these averages lies a substantial amount of heterogeneity across industries and firms. Our estimates point to substantial and persistent differences in labor-augmenting and Hicks-neutral productivity between firms. Firms' R&D activities play a key role in determining these differences and their evolution over time. Interestingly, our estimates indicate that labor-augmenting productivity is slightly more closely tied to firms' R&D activities than to Hicks-neutral productivity. Through the lens of the literature on induced innovation this may be viewed as supporting the argument that firms direct their R&D activities to conserve on labor.

An interesting avenue for future research is to investigate the implications of the different types of technological change for employment. Recent research points to biased technological change as a key driver of the diverging experiences of the continental European, U.S., and U.K. economies during the 1980s and 1990s (Blanchard 1997, Caballero & Hammour 1998, Bentolila & Saint-Paul 2004, McAdam & Willman 2013). Our estimates lend themselves to decomposing firm-level changes in employment into displacement, substitution, and output effects and to compare these effects between labor-augmenting and Hicks-neutral technological change. This may be helpful for better understanding and predicting the evolution of employment as well as for designing labor market and innovation policies in the presence of biased technological change.

Appendix A Data

We observe firms for a maximum of 17 years between 1990 and 2006. We restrict the sample to firms with at least three years of data on all variables required for estimation. The number of firms with 3, 4,..., 17 years of data is 313, 240, 218, 215, 207, 171, 116, 189, 130, 89, 104, 57, 72, 94, and 160, respectively. Table A1 gives the industry definitions along with their equivalent definitions in terms of the ESEE, National Accounts, and ISIC classifications (columns (1)–(3)). Based on the National Accounts in 2000, we further report the shares of the various industries in the total value added of the manufacturing sector (column (4)).

In what follows we define the variables we use. We begin with the variables that are relevant for our main analysis.

- *Investment.* Value of current investments in equipment goods (excluding buildings, land, and financial assets) deflated by the price index of investment. The price of investment is the equipment goods component of the index of industry prices computed and published by the Spanish Ministry of Industry. By measuring investment in operative capital we avoid some of the more severe measurement issues of the other assets.

- *Capital.* Capital at current replacement values $\widetilde{K}_{jt}$ is computed recursively from an initial estimate and the data on current investments in equipment goods $\widetilde{I}_{jt}$. We update the value of the past stock of capital by means of the price index of investment P_{It} as $\widetilde{K}_{jt} = (1-\delta)\frac{P_{It}}{P_{It-1}}\widetilde{K}_{jt-1} + \widetilde{I}_{jt-1}$, where δ is an industry-specific estimate of the rate of depreciation. Capital in real terms is obtained by deflating capital at current replacement values by the price index of investment as $K_{jt} = \frac{\widetilde{K}_{jt}}{P_{It}}$.

- *Labor.* Total hours worked computed as the number of workers times the average hours per worker, where the latter is computed as normal hours plus average overtime minus average working time lost at the workplace.

- *Materials.* Value of intermediate goods consumption (including raw materials, components, energy, and services) deflated by a firm-specific price index of materials.

- *Output.* Value of produced goods and services computed as sales plus the variation of inventories deflated by a firm-specific price index of output.

- *Wage.* Hourly wage cost computed as total labor cost including social security payments divided by total hours worked.

- *Price of materials.* Firm-specific price index for intermediate consumption. Firms are asked about the price changes that occurred during the year for raw materials, components, energy, and services. The price index is computed as a Paasche-type index of the responses.

- *Price of output.* Firm-specific price index for output. Firms are asked about the price changes they made during the year in up to 5 separate markets in which they operate. The price index is computed as a Paasche-type index of the responses.

- *Demand shifter.* Firms are asked to assess the current and future situation of the main market in which they operate. The demand shifter codes the responses as 0, 0.5, and 1 for slump, stability, and expansion, respectively.

- *Share of temporary labor.* Fraction of workers with fixed-term contracts and no or small severance pay.

- *Share of outsourcing.* Fraction of customized parts and pieces that are manufactured by other firms in the value of the firm's intermediate goods purchases.

- *R&D expenditures.* R&D expenditures include the cost of intramural R&D activities, payments for outside R&D contracts with laboratories and research centers, and payments for imported technology in the form of patent licensing or technical assistance, with the various expenditures defined according to the OECD Oslo and Frascati manuals.

We next turn to additional variables that we use for descriptive purposes, extensions, and robustness checks.

- *User cost of capital.* Computed as $P_{It}(r_{jt} + \delta - CPI_t)$, where P_{It} is the price index of investment, r_{jt} is a firm-specific interest rate, δ is an industry-specific estimate of the rate of depreciation, and CPI_t is the rate of inflation as measured by the consumer price index.

- *Skill mix.* Fraction of non-production employees (white collar workers), workers with an engineering degree (engineers), and workers with an intermediate degree (technicians).

- *Region.* Dummy variables corresponding to the 19 Spanish autonomous communities and cities where employment is located if it is located in a unique region and another dummy variable indicating that employment is spread over several regions.

- *Product submarket.* Dummy variables corresponding to a finer breakdown of the 10 industries into subindustries (restricted to subindustries with at least 5 firms, see column (5) of Table A1).

- *Technological sophistication.* Dummy variable that takes the value one if the firm uses digitally controlled machines, robots, CAD/CAM, or some combination of these procedures.

- *Identification between ownership and control.* Dummy variable that takes the value one if the owner of the firm or the family of the owner hold management positions.

- *Age.* Years elapsed since the foundation of the firm with a maximum of 40 years.

- *Firm size.* Number of workers in the year the firm enters the sample.

Appendix B Estimation

Unknown functions. The functions $\widetilde{g}_{L1}(\widetilde{h}_L(\cdot))$, $\widetilde{g}_{L2}(\widetilde{h}_L(\cdot), r_{jt-1})$, $g_{H1}(h_H(\cdot))$, and $g_{H2}(h_H(\cdot), r_{jt-1})$ that are part of the conditional expectation functions $\widetilde{g}_{Lt-1}(\widetilde{h}_L(\cdot), R_{jt-1})$ and $g_{Ht-1}(h_H(\cdot), R_{jt-1})$ are unknown and must be estimated nonparametrically, as must be the absolute value of the price elasticity $\eta(p_{jt}, D_{jt})$ and the correction terms $\lambda_1(S_{Tjt})$, $\lambda_2(S_{Tjt})$, and $\gamma_1(S_{Ojt})$. Following Wooldridge (2004), we model an unknown function $q(v)$ of one variable v by a univariate polynomial of degree Q. We model an unknown function $q(u, v)$ of two variables u and v by a complete set of polynomials of degree Q (see Judd 1998). Unless otherwise noted, we omit the constant in $q(\cdot)$ and set $Q = 3$ in the remainder of this paper.

Starting with the conditional expectation functions, we specify $\widetilde{g}_{L1}(\widetilde{h}_L(\cdot)) = q(\widetilde{h}_L(\cdot) - \widetilde{\gamma}_L)$, $\widetilde{g}_{L2}(\widetilde{h}_L(\cdot), r_{jt}) = q_0 + q(\widetilde{h}_L(\cdot) - \widetilde{\gamma}_L, r_{jt})$, $g_{H1}(h_H(\cdot)) = q(h_H(\cdot) - \gamma_H)$, and $g_{H2}(h_H(\cdot), r_{jt}) = q_0 + q(h_H(\cdot) - \gamma_H, r_{jt})$, where q_0 is a constant and the function $q(\cdot)$ is modeled as described above. Without loss of generality, we absorb $\widetilde{\gamma}_L$ and γ_H into the overall constants of our estimation equations. Turning to the absolute value of the price elasticity, to impose the theoretical restriction $\eta(p_{jt}, D_{jt}) > 1$, we specify $\eta(p_{jt}, D_{jt}) = 1 + \exp(q(p_{jt}, D_{jt}))$, where the function $q(\cdot)$ is modeled as described above except that we suppress terms involving D_{jt}^2 and D_{jt}^3. Turning to the correction terms, we specify $\lambda_1(S_{Tjt}) = q(\ln S_{Tjt})$ and $\lambda_2(S_{Tjt}) = q(\ln S_{Tjt})$ in industries 2, 3, and 10 and $\lambda_1(S_{Tjt}) = q(\ln(1 - S_{Tjt}))$ and $\lambda_2(S_{Tjt}) = q(\ln(1 - S_{Tjt}))$ in the remaining industries.[26] Finally, we specify $\gamma_1(S_{Ojt}) = q(S_{Ojt})$; this ensures that $\gamma_1(S_{Ojt}) = 0$ if $S_{Ojt} = 0$ in line with the normalization $\Gamma(M_{Ijt}, 0) = M_{Ijt}$.

Parameters and instruments. Our first estimation equation (11) has 36 parameters: constant, σ, 15 parameters in $\widetilde{g}_{L0}(t-1)$ (time dummies), 3 parameters in $\widetilde{g}_{L1}(\widetilde{h}_L(\cdot))$, 10 parameters in $\widetilde{g}_{L2}(\widetilde{h}_L(\cdot), r_{jt-1})$, 3 parameters in $\lambda_2(S_{Tjt})$, and 3 parameters in $\gamma_1(S_{Ojt})$.

[26]To incorporate skill upgrading, we instead specify $\lambda_1(S_{Tjt}, \Theta_{jt}) = q(\ln S_{Tjt}, \ln \Theta_{jt})$ and $\lambda_2(S_{Tjt}, \Theta_{jt}) = q(\ln S_{Tjt}, \ln \Theta_{jt})$ in industries 2, 3, and 10 and $\lambda_1(S_{Tjt}, \Theta_{jt}) = q(\ln(1 - S_{Tjt}), \ln \Theta_{jt})$ and $\lambda_2(S_{Tjt}, \Theta_{jt}) = q(\ln(1 - S_{Tjt}), \ln \Theta_{jt})$ in the remaining industries, where the function $q(\cdot)$ is modeled as described above except that we suppress terms involving $(\ln \Theta_{jt})^2$ and $(\ln \Theta_{jt})^3$.

Our instrumenting strategy is adapted from Doraszelski & Jaumandreu (2013) and we refer the reader to Doraszelski & Jaumandreu (2013) and the references therein for a discussion of the use of polynomials for instruments. We use the constant, 15 time dummies, the dummy for performers $1(R_{jt-1} > 0)$, the demand shifter D_{jt}, and a univariate polynomial in $\ln S_{Ojt-1} + m_{jt-1}$ interacted with $1(S_{Ojt-1} > 0)$ (3 instruments). We further use a complete set of polynomials in l_{jt-1}, m_{jt-1}, and $p_{Mjt-1} - w_{jt-1}$ interacted with the dummy for nonperformers $1(R_{jt-1} = 0)$ (19 instruments). In industries 5 and 8 we replace $p_{Mjt-1} - w_{jt-1}$ by p_{Mjt-1} in the complete set of polynomials. Finally, we use a complete set of polynomials in l_{jt-1}, m_{jt-1}, and $p_{Mjt-1} - w_{jt-1}$ and r_{jt-1} interacted with the dummy for performers $1(R_{jt-1} > 0)$ (34 instruments). This yields a total of 74 instruments and $74 - 36 = 38$ degrees of freedom (see column (4) of Table 4).

After imposing the estimated values from equation (11), our second estimation equation (13) has 40 parameters: constant, β_K, ν, 15 parameters in $g_{H0}(t-1)$ (time dummies), 3 parameters in $g_{H1}(h_H(\cdot))$, 10 parameters in $g_{H2}(h_H(\cdot), r_{jt-1})$, 3 parameters in $\lambda_1(S_{Tjt})$, and 6 parameters in $\eta(p_{jt}, D_{jt})$.

As before, we use polynomials for instruments. We use the constant, 15 time dummies, the dummy for performers $1(R_{jt-1} > 0)$, the demand shifter D_{jt}, a univariate polynomial in p_{jt-1} (3 instruments), a univariate polynomial in $p_{Mjt-1} - p_{jt-1}$ (3 instruments), and a univariate polynomial in k_{jt} (3 instruments). We also use a complete set of polynomials in $M_{jt-1}\frac{1-S_{Mjt-1}}{S_{Mjt-1}}$ and K_{jt-1} interacted with the dummy for nonperformers $1(R_{jt-1} = 0)$ (9 instruments). Finally, we use a complete set of polynomials in $M_{jt-1}\frac{1-S_{Mjt-1}}{S_{Mjt-1}}$ and K_{jt-1} (9 instruments) and a univariate polynomial in r_{jt-1} interacted with the dummy for performers $1(R_{jt-1} > 0)$ (3 instruments). This yields a total of 48 instruments and $48 - 40 = 8$ degrees of freedom in industries 1, 2, 3, 6, 7, 9, and 10 (see column (3) of Table 6). In industries 4, 5, and 8, we add a univariate polynomial in $\ln(1 - S_{Tjt-1})$ (3 instruments). We replace the univariate polynomial in k_{jt} by k_{jt} in industries 4 and 8 and we drop D_{jt} in industry 5.

Estimation. From the GMM problem in equation (14) with weighting matrix $\widehat{W}_L = \left[\frac{1}{N}\sum_j A_{Lj}(z_j)A_{Lj}(z_j)'\right]^{-1}$ we first obtain a consistent estimate $\widehat{\theta}_L$ of θ_L. This first step is the NL2SLS estimator of Amemiya (1974). In the second step, we compute the optimal estimate with weighting matrix $\widehat{W}_L = \left[\frac{1}{N}\sum_j A_{Lj}(z_j)\nu_{Lj}(\widehat{\theta}_L)\nu_{Lj}(\widehat{\theta}_L)' A_{Lj}(z_j)'\right]^{-1}$. Throughout the paper, we report standard errors that are robust to heteroskedasticity and autocorrelation.

Implementation. Gauss code for our estimator is available from the authors upon request along with instructions for obtaining the data. To the reduce the number of parameters to search over in the GMM problem in equation (14), we "concentrate out" the parameters that enter it linearly (Wooldridge 2010, p. 435). To guard against local minima, we have extensively searched over the remaining parameters, often using preliminary estimates to narrow down the range of these parameters.

Testing. The value of the GMM objective function for the optimal estimator, multiplied by N, has a limiting χ^2 distribution with $Q - P$ degrees of freedom, where Q is the number of instruments and P the number of parameters to be estimated. We use it as a test for overidentifying restrictions or validity of the moment conditions.

Appendix C Output effect

Direct calculation starting from equation (1) yields the elasticity of output with respect to a firm's effective labor force:

$$\begin{aligned}\epsilon_{Ljt} &= \frac{\partial Y_{jt}}{\partial \exp(\omega_{Ljt})L^*_{jt}} \frac{\exp(\omega_{Ljt})L^*_{jt}}{Y_{jt}} \\ &= \frac{\nu \left(\exp(\omega_{Ljt})L^*_{jt}\right)^{-\frac{1-\sigma}{\sigma}}}{\beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \left(\exp(\omega_{Ljt})L^*_{jt}\right)^{-\frac{1-\sigma}{\sigma}} + \beta_M \left(M^*_{jt}\right)^{-\frac{1-\sigma}{\sigma}}}.\end{aligned}$$

Using equation (9) to substitute for ω_{Ljt} and simplifying we obtain

$$\epsilon_{Ljt} = \frac{\nu \frac{1-S_{Mjt}}{S_{Mjt}} \lambda_1(S_{Tjt})}{\frac{\beta_K}{\beta_M}\left(\frac{K_{jt}}{M_{jt}\exp\left(\gamma_1(\gamma_0(t)S_{Ojt})\right)}\right)^{-\frac{1-\sigma}{\sigma}} + \frac{1-S_{Mjt}}{S_{Mjt}}\lambda_1(S_{Tjt}) + 1}. \tag{21}$$

Recall from equation (6) that $\lambda_1(S_{Tjt}) = 1 + \frac{\Delta_{jt}}{1+\frac{W_{Tjt}}{W_{Pjt}}\frac{S_{Tjt}}{1-S_{Tjt}}}$, where Δ_{jt} is the gap between the wage of permanent workers W_{Pjt} and the shadow wage. To facilitate evaluating equation (21), we abstract from adjustment costs and set $\lambda_1(S_{Tjt}) = 1$.

Direct calculation starting from equation (19) also yields the elasticity of output with respect to a firm's effective capital stock:

$$\begin{aligned}\epsilon_{Kjt} &= \frac{\partial Y_{jt}}{\partial \exp(\omega_{Kjt})K_{jt}} \frac{\exp(\omega_{Kjt})K_{jt}}{Y_{jt}} \\ &= \frac{\nu \left(\exp(\omega_{Kjt})K_{jt}\right)^{-\frac{1-\sigma}{\sigma}}}{\left(\exp(\omega_{Kjt})K_{jt}\right)^{-\frac{1-\sigma}{\sigma}} + \left(\exp(\omega_{Ljt})L^*_{jt}\right)^{-\frac{1-\sigma}{\sigma}} + \beta_M \left(M^*_{jt}\right)^{-\frac{1-\sigma}{\sigma}}} \\ &= \frac{\nu}{1 + \frac{P_{Mjt}M_{jt}}{P_{Kjt}K_{jt}}\left(\frac{1-S_{Mjt}}{S_{Mjt}}\lambda_1(S_{Tjt}) + 1\right)},\end{aligned} \tag{22}$$

where we use equations (9) and (20) to substitute for ω_{Ljt} and ω_{Kjt}, respectively. As with equation (21), we set $\lambda_1(S_{Tjt}) = 1$ to evaluate equation (22).

Appendix D Wage regression

As column (1) of Table A2 shows, the coefficient of variation for the (level of the) wage W_{jt} ranges from 0.35 to 0.50 across industries.[27] The variance decomposition in columns (2)–(4) shows that around one quarter of the overall variation is within firms across periods. The larger part of this variation is across firms.

To explore the source of this variation, we regress the (log of the) wage w_{jt} on the skill mix of a firm's labor force as given by the share of temporary (as opposed to permanent) labor, the share of white (as opposed to blue) collar workers, and the shares of engineers and technicians (as opposed to unskilled workers), time dummies, region dummies, product

[27] The coefficient of variation for the price of materials ranges from 0.12 to 0.19 across industries.

submarket dummies, the demand shifter, and an array of other firm characteristics, namely dummies for technological sophistication and identification of ownership and control as well as univariate polynomials of degree 3 in age and firm size.

To motivate this regression, assume that there are Q types of labor with wages W_{1jt}, W_{2jt}, ..., W_{Qjt} and write the wage as

$$W_{jt} = \sum_{q=1}^{Q} W_{qjt} S_{qjt} = W_{1jt} \left(1 + \sum_{q=2}^{Q} \left(\frac{W_{qjt}}{W_{1jt}} - 1 \right) S_{qjt} \right),$$

where S_{qjt} is the share of labor of type q and $\sum_{q=1}^{Q} S_{qjt} = 1$. Because

$$w_{jt} \approx w_{1jt} + \sum_{q=2}^{Q} \left(\frac{W_{qjt}}{W_{1jt}} - 1 \right) S_{qjt},$$

the coefficient on S_{qjt} in the wage regression is an estimate of the wage premium $\left(\frac{W_{qjt}}{W_{1jt}} - 1 \right)$ of labor of type q over type 1. Because we do not have the joint distribution of skills (e.g., temporary white collar technician) in our data, we approximate it by the marginal distributions (e.g., share of temporary labor) and ignore higher-order terms. As columns (5)–(8) of Table A2 show, the estimated coefficients on the skill mix of a firm's labor force are often significant, have the expected signs and are quite similar across industries. On average across industries, temporary workers earn 36% less than permanent workers, white collar workers earn 26% more than blue collar workers, engineers earn 85% more than unskilled workers, and technicians earn 23% more than unskilled workers.

The wage regression also shows that some, but by no means all variation in the wage is due to worker quality. To isolate the part of the wage that depends on the skill mix of a firm's labor force, we decompose the predicted wage $\widehat{w}_{jt}$ into a prediction $\widehat{w}_{Qjt}$ based on the skill mix and a prediction $\widehat{w}_{Cjt}$ based on the remaining variables. $\widehat{w}_{Qjt}$ and $\widehat{w}_{Cjt}$ are positively correlated. According to $R^2 = \frac{Var(\widehat{w}_{jt})}{Var(w_{jt})}$ in column (9), depending on the industry, the wage regression explains between 63% and 76% of the variation in the wage, with an average of 70%. The skill mix by itself explains between 2% and 20% of the variation in the wage, with an average of 10% (see $R_Q^2 = \frac{Var(\widehat{w}_{Qjt})}{Var(w_{jt})}$ in column (10)). In contrast, the remaining variables explain between 36% and 64% of the variation in the wage, with an average of 48% (see $R_C^2 = \frac{Var(\widehat{w}_{Cjt})}{Var(w_{jt})}$ in column (11)). The larger part of the variation in the wage therefore appears to be due to temporal and geographic differences in the supply of labor, the fact that firms operate in different product submarkets, and other firm characteristics.

In developing the quality index Θ_{jt}, we assume that there are Q types of permanent labor. We approximate the wage premium $\left(\frac{W_{Pqjt}}{W_{P1jt}} - 1 \right)$ of permanent labor of type q over type 1 by the estimated coefficient on S_{qjt} in the wage regression and the share $S_{Pqjt} = \frac{L_{Pqjt}}{L_{Pjt}} = \frac{L_{Pqjt}}{L_{jt}} \Big/ \frac{L_{Pjt}}{L_{jt}}$ of permanent labor of type q by $\frac{S_{qjt}}{1-S_{Tjt}}$.

References

Abowd, J., Haltiwanger, J., Lane, J., McKinney, K. & Sandusky, K. (2007), Technology and the demand for skill: An analysis of within and between firm differences, Working paper no. 13043, NBER, Cambridge.

Acemoglu, D. (2003), 'Labor- and capital-augmenting technical change', *Journal of the European Economic Association* **1**(1), 1–37.

Ackerberg, D., Benkard, L., Berry, S. & Pakes, A. (2007), Econometric tools for analyzing market outcomes, *in* J. Heckman & E. Leamer, eds, 'Handbook of Econometrics', Vol. 6A, North-Holland, Amsterdam, pp. 4171–4276.

Ackerberg, D., Caves, K. & Frazer, G. (2006), Structural identification of production functions, Working paper, UCLA, Los Angeles.

Amemiya, T. (1974), 'The nonlinear two-stage least-squares estimator', *Journal of Economic Literature* **2**, 105–110.

Antràs, P. (2004), 'Is the U.S. aggregate production function Cobb-Douglas? New estimates of the elasticity of substitution', *Contributions to Macroeconomics* **4**(1).

Bentolila, S. & Saint-Paul, G. (2004), 'Explaining movements in the labor share', *Contributions to Macroeconomics* **3**(1).

Binswanger, H. (1974), 'The measurement of technical change biases with many factors of production', *American Economic Review* **64**(6), 964–976.

Black, S. & Lynch, L. (2001), 'How to compete: The impact of workplace practices and information technology on productivity', *Review of Economics and Statistics* **83**(3), 434–445.

Blanchard, O. (1997), 'The medium run', *Brookings Papers on Economic Activity: Macroecomomics* **1997**(2), 89–158.

Bloom, N., Sadun, R. & Van Reenen, J. (2012), 'Americans do IT better: US multinationals and the productivity miracle', *American Economic Review* **102**(1), 167–201.

Brown, M. & de Cani, J. (1963), 'Technological change and the distribution of income', *International Economic Review* **4**(3), 289–309.

Buettner, T. (2005), R&D and the dynamics of productivity, Working paper, LSE, London.

Caballero, R. & Hammour, M. (1998), 'Jobless growth: appropriability, factor substitution, and unemployment', *Carnegie-Rochester Conference Series on Public Policy* **48**, 51–94.

Card, D. & DiNardo, J. (2002), 'Skill-biased technological change and rising wage inequality: Some problems and puzzles', *Journal of Labor Economics* **20**(4), 733–783.

Caselli, F. (2005), Accounting for cross-country income differences, *in* P. Aghion & S. Durlauf, eds, 'Handbook of Economic Growth', Vol. 1A, North-Holland, Amsterdam, pp. 679–741.

Chirinko, R. (2008), 'σ: The long and short of it', *Journal of Multivariate Analysis* **30**, 671–686.

David, P. & van de Klundert, T. (1965), 'Biased efficiency growth and capital-labor substitution in the U.S., 1899-1960', *American Economic Review* **55**(3), 357–394.

Davis, S. & Haltiwanger, J. (1992), 'Gross job creation, gross job destruction, and employment reallocation', *Quarterly Journal of Economics* **107**(3), 819–863.

Dolado, J., Garcia-Serrano, C. & Jimeno, J. (2002), 'Drawing lessons from the boom of temporary jobs in Spain', *Economic Journal* **112**(480), 270–295.

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

Dunne, T., Roberts, M. & Samuelson, L. (1988), 'Patterns of firm entry and exit in U.S. manufacturing', *Rand Journal of Economics* **19**(4), 495–515.

Eslava, M., Haltiwanger, J., Kugler, A. & Kugler, M. (2004), 'The effects of structural reforms on productivity and profitability enhancing reallocation: Evidence from Colombia', *Journal of Development Economics* **75**, 333–371.

European Commission (2001), 'Statistics on innovation in Europe', Enterprise DG, Brussels.

European Commission (2004), 'European competitiveness report', Enterprise DG, Brussels.

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

Foster, L., Haltiwanger, J. & Syverson, C. (2013), The slow growth of new plants: Learning about demand?, Working paper, University of Chicago, Chicago.

Gandhi, A., Navarro, S. & Rivers, D. (2013), On the identification of production functions: How heterogeneous is productivity?, Working paper, University of Wisconsin.

Gordon, R. (1990), *The measurement of durable goods prices*, National Bureau of Economic Research Monograph Series, University of Chicago Press, Chicago.

Grieco, P., Li, S. & Zhang, H. (2014), Production function estimation with unobserved input price dispersion, Working paper, Pennsylvania State University, State College.

Griliches, Z. (1979), 'Issues in assessing the contribution of R&D to productivity growth', *Bell Journal of Economics* **10**(1), 92–116.

Griliches, Z. (1998), *R&D and productivity: The econometric evidence*, University of Chicago Press, Chicago.

Griliches, Z. (2000), *R&D, education, and productivity: A retrospective*, Harvard University Press, Cambridge.

Griliches, Z. & Mairesse, J. (1998), Production functions: The search for identification, *in* S. Strom, ed., 'Econometrics and Economic Theory in the 20th Century: The Ragnar Frisch Centennial Symposium', Cambridge University Press, Cambridge.

Grossman, G. & Helpman, E. (2002), 'Integration versus outsourcing in industry equilibrium', *Quarterly Journal of Economics* **117**(1), 85–120.

Grossman, G. & Helpman, E. (2005), 'Outsourcing in a global economy', *Review of Economic Studies* **72**(1), 135–159.

Hall, R. (1988), 'The relation between price and marginal cost in U.S. industry', *Journal of Political Economy* **96**(5), 921–947.

Hammermesh, D. (1993), *Labor demand*, Princeton University Press, Princeton.

Hammermesh, D. & Pfann, G. (1996), 'Adjustment costs in factor demand', *Journal of Economic Literature* **34**(3), 1264–1292.

Hansen, L. (1982), 'Large sample properties of generalized method of moments estimators', *Econometrica* **50**(4), 1029–1054.

Hicks, J. (1966), *The theory of wages*, 2nd edn, Macmillan, London.

Jin, H. & Jorgenson, D. (2010), 'Econometric modeling of technical change', *Journal of Econometrics* **157**(2), 205–219.

Jones, C. (2005), 'The shape of production functions and the direction of technical change', *Quarterly Journal of Economics* **120**(2), 517–549.

Jorgenson, D. (1995*a*), *Productivity. Volume 1: Postwar U.S. economic growth*, MIT Press, Cambridge.

Jorgenson, D. (1995*b*), *Productivity. Volume 2: International comparisons of economic growth*, MIT Press, Cambridge.

Judd, K. (1998), *Numerical methods in economics*, MIT Press, Cambridge.

Kalt, J. (1978), 'Technological change and factor substitution in the United States: 1929-1967', *International Economic Review* **19**(3), 761–775.

Krueger, A. (1999), 'Measuring labor's share', *American Economic Review* **90**(2), 45–51.

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

Lucas, R. (1969), Labor-capital substitution in U.S. manufacturing, *in* A. Harberger & M. Bailey, eds, 'The taxation of income from capital', Brookings Institution, Washington, pp. 223–274.

Lucas, R. (1988), 'On the mechanics of economic development', *Journal of Monetary Economics* **22**(1), 3–42.

Machin, S. & Van Reenen, J. (1998), 'Technology and changes in skill structure: Evidence from seven OECD countries', *Quarterly Journal of Economics* **113**(4), 1215–1244.

Marschak, J. & Andrews, W. (1944), 'Random simultaneous equations and the theory of production', *Econometrica* **12**(3), 143–205.

Marshall, A. (1920), *Principles of eocnomics*, 8th edn, Macmillan, London.

McAdam, P. & Willman, A. (2013), 'Medium run redux', *Macroeconomic Dynamics* **17**(4), 695–727.

Newey, W., Powell, J. & Vella, F. (1999), 'Nonparametric estimation of triangular simultaneous equations models', *Econometrica* **67**(3), 565–603.

Oberfield, E. & Raval, D. (2014), Micro data and macro technology, Working paper, Federal Trade Commission, Washington.

OECD (2007), 'The policy mix for research, development and innovation in Spain: Key issues and policy recommendations', Directorate for Science, Technology and Industry, Paris.

Oi, W. & Idson, T. (1999), Firm size and wages, *in* O. Ashenfelter & D. Card, eds, 'Handbook of labor economics', Vol. 3, Elsevier, Amsterdam.

Olley, S. & Pakes, A. (1996), 'The dynamics of productivity in the telecommunications industry', *Econometrica* **64**(6), 1263–1297.

Raval, D. (2013), Non neutral technology differences, Working paper, Federal Trade Commission, Washington.

Roberts, M. & Supina, D. (1996), 'Output price, markups, and producer size', *European Economic Review* **40**, 909–921.

Robinson, P. (1988), 'Root-n-consistent semiparametric regression', *Econometrica* **56**(4), 931–954.

Romer, P. (1986), 'Increasing returns and long run growth', *Journal of Political Economy* **94**(5), 1002–1037.

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

Solow, R. (1957), 'Technical change and the aggregate production function', *Review of Economics and Statistics* **39**(3), 312–320.

Van Biesebroeck, J. (2003), 'Productivity dynamics with technology choice: An application to automobile assembly', *Review of Economic Studies* **70**, 167–198.

Van Reenen, J. (1997), 'Employment and technological innovation: Evidence from UK manufacturing firms', *Journal of Labor Economics* **15**(2), 255–284.

Violante, G. (2008), Skill-biased technical change, *in* S. Durlauf & L. Blume, eds, 'The New Palgrave Dictionary of Economics', 2nd edn, Palgrave Macmillan, New York.

Wooldridge, J. (2004), On estimating firm-level production functions using proxy variables to control for unobservables, Working paper, Michigan State University, East Lansing.

Wooldridge, J. (2010), *Econometric analysis of cross section and panel data*, 2nd edn, MIT Press, Cambridge.

Zhang, H. (2014*a*), Biased technology and contribution of technological change to economic growth: Firm-level evidence, Working paper, University of Hong Kong, Hong Kong.

Zhang, H. (2014*b*), Biased technology and productivity growth: Firm-level evidence from China's steel industry, Slide deck, University of Hong Kong, Hong Kong.

Table 1: Descriptive statistics.

Industry	Obs.	Firms	Rates of growth[a] Output (s. d.)	Capital (s. d.)	Labor (s. d.)	Materials (s. d.)	Price (s. d.)	$\frac{M}{L}$ (s. d.)	$\frac{P_M}{W}$ (s. d.)	$\frac{M}{K}$ (s. d.)	$\frac{P_M}{P_K}$ (s. d.)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
1. Metals and metal products	2365	313	0.045	0.051	0.008	0.030	0.017	0.022	-0.008	-0.021	0.049
			(0.235)	(0.192)	(0.161)	(0.327)	(0.052)	(0.316)	(0.176)	(0.373)	(0.099)
2. Non-metallic minerals	1270	163	0.046	0.057	0.010	0.041	0.012	0.031	-0.012	-0.016	0.043
			(0.228)	(0.212)	(0.177)	(0.285)	(0.058)	(0.272)	(0.147)	(0.333)	(0.104)
3. Chemical products	2168	299	0.060	0.062	0.015	0.044	0.008	0.029	-0.015	-0.019	0.044
			(0.228)	(0.182)	(0.170)	(0.274)	(0.055)	(0.250)	(0.153)	(0.313)	(0.141)
4. Agric. and ind. machinery	1411	178	0.031	0.040	-0.003	0.018	0.015	0.022	-0.015	-0.021	0.041
			(0.252)	(0.190)	(0.169)	(0.347)	(0.026)	(0.335)	(0.155)	(0.390)	(0.099)
5. Electrical goods	1505	209	0.059	0.041	0.010	0.048	0.008	0.038	-0.021	0.007	0.045
			(0.268)	(0.173)	(0.205)	(0.359)	(0.046)	(0.344)	(0.174)	(0.394)	(0.095)
6. Transport equipment	1206	161	0.060	0.043	0.004	0.051	0.008	0.047	-0.019	0.008	0.033
			(0.287)	(0.164)	(0.201)	(0.375)	(0.031)	(0.343)	(0.171)	(0.396)	(0.093)
7. Food, drink and tobacco	2455	327	0.023	0.047	0.003	0.012	0.021	0.009	-0.018	-0.035	0.049
			(0.206)	(0.177)	(0.169)	(0.286)	(0.054)	(0.295)	(0.176)	(0.328)	(0.116)
8. Textile, leather and shoes	2368	335	0.004	0.031	-0.015	-0.009	0.015	0.006	-0.021	-0.040	0.040
			(0.229)	(0.189)	(0.180)	(0.348)	(0.042)	(0.355)	(0.183)	(0.385)	(0.099)
9. Timber and furniture	1445	207	0.025	0.045	0.013	0.014	0.020	0.001	-0.019	-0.031	0.067
			(0.225)	(0.168)	(0.184)	(0.335)	(0.031)	(0.329)	(0.171)	(0.371)	(0.123)
10. Paper and printing products	1414	183	0.031	0.052	-0.001	0.013	0.017	0.014	-0.017	-0.039	0.046
			(0.187)	(0.221)	(0.149)	(0.252)	(0.074)	(0.247)	(0.159)	(0.326)	(0.122)

[a] Computed for 1991 to 2006.

Table 2: Descriptive statistics.

Industry	Temp. labor Obs. (%)	Temp. labor Share (s. d.)	Intrafirm max-min Share of temp. (s. d.)	Intrafirm max-min Hours worked[a] (s. d.)	Intrafirm max-min Hours per worker[a] (s. d.)	Outsourcing Obs. (%)	Outsourcing Share (s. d.)	With R&D Stable (%)	With R&D Occas. (%)	With R&D R&D intens. (s. d.)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1. Metal and metal products	1877	0.260	0.243	0.448	0.069	1014	0.200	56	109	0.012
	(79.4)	(0.221)	(0.197)	(0.360)	(0.090)	(42.9)	(0.193)	(17.9)	(34.8)	(0.018)
2. Non-metallic minerals	1018	0.231	0.232	0.482	0.065	316	0.177	20	62	0.011
	(80.2)	(0.207)	(0.183)	(0.403)	(0.063)	(24.9)	(0.179)	(12.3)	(38.0)	(0.022)
3. Chemical products	1722	0.170	0.203	0.446	0.043	924	0.146	121	85	0.026
	(79.4)	(0.176)	(0.185)	(0.427)	(0.038)	(42.6)	(0.183)	(40.5)	(28.4)	(0.034)
4. Agric. and ind. machinery	1069	0.189	0.227	0.485	0.086	808	0.288	64	62	0.022
	(75.8)	(0.181)	(0.181)	(0.419)	(0.166)	(57.3)	(0.263)	(36.0)	(34.8)	(0.026)
5. Electrical goods	1221	0.245	0.280	0.559	0.063	763	0.181	83	61	0.029
	(81.1)	(0.206)	(0.216)	(0.452)	(0.077)	(50.7)	(0.194)	(39.7)	(29.2)	(0.040)
6. Transport equipment	962	0.206	0.239	0.555	0.131	637	0.233	60	56	0.028
	(79.8)	(0.198)	(0.184)	(0.415)	(0.237)	(52.8)	(0.261)	(37.3)	(34.8)	(0.049)
7. Food, drink and tobacco	2067	0.276	0.266	0.468	0.058	514	0.142	65	86	0.007
	(84.2)	(0.237)	(0.215)	(0.343)	(0.065)	(20.9)	(0.172)	(19.9)	(26.3)	(0.022)
8. Textile, leather and shoes	1726	0.238	0.291	0.489	0.062	1214	0.252	44	85	0.017
	(79.2)	(0.260)	(0.244)	(0.402)	(0.086)	(51.3)	(0.237)	(13.1)	(25.4)	(0.031)
9. Timber and furniture	1175	0.320	0.326	0.523	0.056	535	0.183	21	44	0.010
	(81.3)	(0.226)	(0.234)	(0.387)	(0.076)	(37.0)	(0.201)	(10.1)	(21.3)	(0.017)
10. Paper and printing products	1024	0.155	0.221	0.425	0.057	679	0.273	17	48	0.015
	(72.4)	(0.145)	(0.196)	(0.346)	(0.065)	(48.0)	(0.253)	(9.3)	(26.2)	(0.028)

[a] Computed as difference in logs.

Table 3: Elasticity of substitution.

Industry	Obs.	Firms	OLS σ (s. e.)	OLS $\tilde{\delta}_L$ (s. e.)	GMM incl. m_{jt-1} as instr. σ (s. e.)	GMM incl. m_{jt-1} as instr. χ^2 (df)	GMM incl. m_{jt-1} as instr. p val.	GMM excl. m_{jt-1} as instr. σ (s. e.)	GMM excl. m_{jt-1} as instr. χ^2 (df)	GMM excl. m_{jt-1} as instr. p val.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1. Metals and metal products	2365	313	1.163	0.023	0.451	57.846	0.034	0.694	13.683	0.550
			(0.104)	(0.007)	(0.096)	(40)		(0.113)	(15)	
2. Non-metallic minerals	1270	163	1.227	0.038	0.643	46.068	0.234	0.603	11.299	0.731
			(0.119)	(0.008)	(0.086)	(40)		(0.126)	(15)	
3. Chemical products	2168	299	1.132	0.016	0.481	65.068	0.007	0.618	7.582	0.939
			(0.095)	(0.007)	(0.099)	(40)		(0.124)	(15)	
4. Agric. and ind. machinery	1411	178	1.239	0.019	0.502	56.166	0.046	0.598	8.500	0.902
			(0.166)	(0.008)	(0.114)	(40)		(0.103)	(15)	
5. Electrical goods	1505	209	1.402	0.017	0.469	60.674	0.019	0.458	17.457	0.292
			(0.163)	(0.009)	(0.108)	(40)		(0.108)	(15)	
6. Transport equipment	1206	161	1.161	0.029	1.204	48.449	0.169	0.512	7.740	0.934
			(0.218)	(0.011)	(0.089)	(40)		(0.162)	(15)	
7. Food, drink and tobacco	2455	327	1.421	0.015	0.614	70.492	0.002	0.707	15.088	0.445
			(0.094)	(0.008)	(0.063)	(40)		(0.084)	(15)	
8. Textile, leather and shoes	2368	335	1.846	0.001	0.059	55.178	0.056	0.724	18.453	0.240
			(0.169)	(0.100)	(0.077)	(40)		(0.162)	(15)	
9. Timber and furniture	1445	207	0.793	0.014	0.461	37.357	0.590	0.486	5.805	0.983
			(0.117)	(0.008)	(0.089)	(40)		(0.102)	(15)	
10. Paper and printing products	1414	183	1.120	0.026	0.609	51.798	0.100	0.854	7.300	0.949
			(0.107)	(0.008)	(0.057)	(40)		(0.077)	(15)	

Table 4: Elasticity of substitution (cont'd).

Industry	Obs.	Firms	GMM σ (s. e.)	GMM χ^2 (df)	GMM p val.	GMM with quality-corrected wage as instr. σ (s. e.)	GMM with quality-corrected wage as instr. χ^2 (df)	GMM with quality-corrected wage as instr. p val.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1. Metals and metal products	1759	278	0.535	48.882	0.111	0.456	52.058	0.064
			(0.114)	(38)		(0.112)	(38)	
2. Non-metallic minerals	959	146	0.730	46.890	0.153	0.833	45.105	0.199
			(0.098)	(38)		(0.096)	(38)	
3. Chemical products	1610	269	0.696	46.154	0.171	0.695	48.889	0.111
			(0.102)	(38)		(0.072)	(38)	
4. Agric. and ind. machinery	979	164	0.607	42.420	0.286	0.762	44.227	0.225
			(0.196)	(38)		(0.206)	(38)	
5. Electrical goods	1147	191	0.592	46.782	0.155	0.624	44.592	0.214
			(0.123)	(38)		(0.125)	(38)	
6. Transport equipment	896	146	0.798	45.740	0.182	0.602	41.214	0.332
			(0.088)	(38)		(0.097)	(38)	
7. Food, drink and tobacco	1963	306	0.616	53.931	0.045	0.766	38.379	0.452
			(0.081)	(38)		(0.079)	(38)	
8. Textile, leather and shoes	1593	282	0.440	52.496	0.059	0.462	55.996	0.030
			(0.186)	(38)		(0.149)	(38)	
9. Timber and furniture	1114	188	0.438	39.204	0.416	0.497	36.687	0.530
			(0.093)	(38)		(0.094)	(38)	
10. Paper and printing products	938	162	0.525	44.508	0.217	0.449	43.009	0.265
			(0.088)	(38)		(0.085)	(38)	

Table 5: Labor-augmenting technological change.

Industry	$\Delta\omega_L$	$corr(\Delta\omega_L, \omega_L)$	$\epsilon_L \Delta\omega_L$	ω_L R&D–No R&D	$\epsilon_L \Delta\omega_L$ R&D	$\epsilon_L \Delta\omega_L$ No R&D	Skill upgrading σ (s. e.)	Skill upgrading χ^2 (df)	Skill upgrading p val.	Skill upgrading $\Delta\omega_L$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1. Metals and metal products	0.091	0.183	0.021	0.885	0.024	0.018	0.582	44.868	0.206	0.104
							(0.117)	(38)		
2. Non-metallic minerals	0.142	0.191	0.031	1.461	0.022	0.028	0.737	35.898	0.567	0.087
							(0.092)	(38)		
3. Chemical products	0.049	0.186	0.013	1.239	0.018	-0.002	0.618	47.832	0.132	0.053
							(0.110)	(38)		
4. Agric. and ind. machinery	0.126	0.209	0.032	1.537	0.028	0.046	0.177	38.413	0.451	0.060
							(0.172)	(38)		
5. Electrical goods	0.220	0.237	0.022	2.783	0.022	0.012	0.488	48.365	0.121	0.179
							(0.129)	(38)		
6. Transport equipment	0.183	0.261	0.036	0.637	0.045	0.012	0.781	45.457	0.189	0.098
							(0.101)	(38)		
7. Food, drink and tobacco	0.018	0.131	0.007	-0.044	0.009	0.006	0.655	53.981	0.045	-0.007
							(0.084)	(38)		
8. Textile, leather and shoes	0.010	0.179	0.007	0.480	0.007	0.009	0.120	41.931	0.304	0.000
							(0.168)	(38)		
9. Timber and furniture	-0.013	0.142	0.002	-0.024	0.007	0.002	0.528	37.674	0.484	-0.023
							(0.090)	(38)		
10. Paper and printing products	0.021	0.094	0.014	0.579	0.007	0.020	0.396	37.418	0.496	-0.011
							(0.082)	(38)		

Table 6: Distributional parameters and elasticity of scale.

Industry	GMM β_K (s. e.)	ν (s. e.)	χ^2 (df)	p val.
	(1)	(2)	(3)	(4)
1. Metals and metal products	0.232	0.941	2.872	0.942
	(0.073)	(0.029)	(8)	
2. Non-metallic minerals	0.225	0.911	3.975	0.859
	(0.133)	(0.063)	(8)	
3. Chemical products	0.136	0.934	1.074	0.998
	(0.059)	(0.041)	(8)	
4. Agric. and ind. machinery	0.139	0.806	7.258	0.610
	(0.125)	(0.088)	(9)	
5. Electrical goods	0.133	0.848	3.059	0.980
	(0.038)	(0.046)	(10)	
6. Transport equipment[a]	0.308	0.923		
	(0.182)	(0.061)		
7. Food, drink and tobacco	0.303	0.931	2.006	0.981
	(0.137)	(0.040)	(8)	
8. Textile, leather and shoes	0.066	0.976	3.269	0.953
	(0.097)	(0.035)	(9)	
9. Timber and furniture	0.103	0.932	9.748	0.283
	(0.107)	(0.066)	(8)	
10. Paper and printing products	0.227	0.936	5.402	0.714
	(0.080)	(0.036)	(8)	

[a] We have been unable to compute the second-step GMM estimate.

Table 7: Hicks-neutral technological change.

Industry	$\Delta\omega_H$	$corr(\Delta\omega_H, \omega_H)$	$\epsilon_L \Delta\omega_L + \Delta\omega_H$	$corr(\Delta\omega_H, \Delta\omega_L)$	Firms' R&D activities: ω_H	Firms' R&D activities: $\Delta\omega_H$	
					R&D–No R&D	R&D	No R&D
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1. Metals and metal products	0.044	0.346	0.065	0.686	0.027	0.046	0.038
2. Non-metallic minerals	0.005	0.448	0.036	0.439	0.078	-0.019	0.041
3. Chemical products	0.019	0.220	0.032	0.717	0.182	0.022	0.011
4. Agric. and ind. machinery	0.041	0.264	0.072	0.678	0.382	0.039	0.022
5. Electrical goods	0.020	0.294	0.042	0.622	0.484	0.009	0.055
6. Transport equipment	0.042	0.714	0.078	0.549	0.121	0.058	-0.031
7. Food, drink and tobacco	0.001	0.214	0.007	0.817	-0.148	0.007	0.000
8. Textile, leather and shoes	0.012	0.295	0.019	0.612	-0.146	-0.003	0.032
9. Timber and furniture	0.021[a]	0.323	0.023[a]	0.714	-0.132	0.008	0.035
10. Paper and printing products	0.002	0.220	0.016	0.851	-0.104	0.007	0.006

[a] We trim values of $\Delta\omega_H$ below -0.25 and above 0.5. This amounts to trimming around one third of observations.

Table 8: Capital-augmenting productivity.

Industry	$\Delta\omega_K$	$\epsilon_K \Delta\omega_K$	GMM				
			β_K (s. e.)	ν (s. e.)	δ_K (s. e.)	χ^2 (df)	p val.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1. Metals and metal products	0.056	0.004	0.254	0.903	0.036	2.947	0.890
			(0.129)	(0.055)	(0.061)	(7)	
2. Non-metallic minerals	-0.010	0.007	0.236	0.906	0.010	2.921	0.822
			(0.102)	(0.072)	(0.072)	(7)	
3. Chemical products	-0.018	0.001	0.125	0.942	-0.031	0.692	0.998
			(0.068)	(0.041)	(0.092)	(7)	
4. Agric. and ind. machinery	-0.020	0.000	0.182	0.801	0.031	8.148	0.419
			(0.177)	(0.081)	(0.122)	(8)	
5. Electrical goods	-0.078	0.000	0.129	0.845	-0.004	2.969	0.966
			(0.041)	(0.054)	(0.056)	(9)	
6. Transport equipment[a]	0.008	0.005	0.115	0.981	-0.143		
			(0.088)	(0.050)	(0.138)		
7. Food, drink and tobacco	-0.005	0.002	0.282	0.918	-0.045	2.378	0.936
			(0.286)	(0.058)	(0.204)	(7)	
8. Textile, leather and shoes[a]	-0.085	-0.002	0.080	0.971	0.053		
			(0.143)	(0.047)	(0.135)		
9. Timber and furniture	-0.042	0.000	0.088	0.924	-0.021	8.861	0.263
			(0.119)	(0.067)	(0.059)	(7)	
10. Paper and printing products	0.022	0.007	0.229	0.935	0.005	5.149	0.642
			(0.089)	(0.033)	(0.045)	(7)	

[a] We have been unable to compute the second-step GMM estimate.

Table A1: Industry definitions and equivalent classifications.

Industry	Classifications: ESSE	National Accounts	ISIC (Rev. 4)	Share of value added	Number of subindustries
	(1)	(2)	(3)	(4)	(5)
1. Ferrous and non-ferrous metals and metal products	12+13	DJ	C 24+25	13.2	11
2. Non-metallic minerals	11	DI	C 23	8.2	8
3. Chemical products	9+10	DG-DH	C 20+21+22	13.9	7
4. Agricultural and industrial machinery	14	DK	C 28	7.1	7
5. Electrical goods	15+16	DL	C 26+27	7.5	13
6. Transport equipment	17+18	DM	C 29+30	11.6	7
7. Food, drink and tobacco	1+2+3	DA	C 10+11+12	14.5	10
8. Textile, leather and shoes	4+5	DB-DC	C 13+14+15	7.6	11
9. Timber and furniture	6+19	DD-DN 38	C 16+31	7.0	6
10. Paper and printing products	7+8	DE	C 17+18	8.9	4
Total				99.5	84

Table A2: Variation in the wage and its determinants.

Industry	Wage CV	Var	Within (%)	Betw. (%)	Wage regression Temp. (s. e.)	White (s. e.)	Engin. (s. e.)	Tech. (s. e.)	R^2	R^2_Q	R^2_C
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
1. Metals and metal products	0.425	39.025	9.779	29.246	-0.425	0.127	1.106	0.316	0.651	0.094	0.480
			(25.1)	(74.9)	(0.057)	(0.097)	(0.298)	(0.094)			
2. Non-metallic minerals	0.441	36.252	10.072	26.180	-0.098	0.124	0.896	0.246	0.742	0.020	0.643
			(27.8)	(72.2)	(0.065)	(0.159)	(0.280)	(0.181)			
3. Chemical products	0.440	54.332	9.673	44.659	-0.465	0.461	0.592	0.203	0.755	0.197	0.376
			(17.8)	(82.2)	(0.066)	(0.074)	(0.137)	(0.099)			
4. Agric. and ind. machinery	0.354	30.980	11.472	19.508	-0.273	0.285	0.803	-0.028	0.631	0.082	0.484
			(37.0)	(63.0)	(0.067)	(0.105)	(0.226)	(0.125)			
5. Electrical goods	0.383	31.047	8.461	22.586	-0.374	0.219	1.092	0.312	0.661	0.200	0.356
			(27.3)	(72.7)	(0.058)	(0.073)	(0.264)	(0.087)			
6. Transport equipment	0.393	40.666	12.876	27.790	-0.377	0.220	0.402	0.274	0.709	0.066	0.552
			(31.7)	(68.3)	(0.079)	(0.108)	(0.300)	(0.166)			
7. Food, drink and tobacco	0.502	36.590	5.952	30.638	-0.451	0.115	1.292	0.357	0.753	0.097	0.481
			(16.3)	(83.7)	(0.053)	(0.053)	(0.265)	(0.154)			
8. Textile, leather and shoes	0.449	16.565	3.654	12.911	-0.260	0.646	1.584	0.346	0.683	0.140	0.389
			(22.1)	(77.9)	(0.048)	(0.084)	(0.402)	(0.241)			
9. Timber and furniture	0.392	14.646	3.643	11.003	-0.356	0.173	0.288	0.002	0.697	0.061	0.525
			(24.9)	(75.1)	(0.051)	(0.089)	(0.377)	(0.164)			
10. Paper and printing products	0.464	51.667	10.003	41.664	-0.477	0.188	0.444	0.277	0.702	0.070	0.505
			(19.4)	(80.6)	(0.099)	(0.084)	(0.210)	(0.127)			