

Measuring the Bias of Technological Change*

Ulrich Doraszelski[†]

Harvard University and CEPR

Jordi Jaumandreu[‡]

Boston University, Universidad Carlos III, and CEPR

First draft: February 2009

This draft: August 2012

—Preliminary and incomplete—

Abstract

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. The literatures in industrial organization, productivity, and economic growth rest on very specific assumptions about the bias of technological change. Yet, the evidence is sparse.

In this paper we propose a general framework for estimating production functions that allows productivity to be multi-dimensional. Using firm-level panel data, we are able to directly assess the bias of technological change by measuring, at the level of the individual firm, how much of technological change is factor neutral and how much of it is labor augmenting. We further relate the speed and the direction of technological change to firms' R&D activities.

*We thank Pol Antràs, Matthias Doepke, Michaela Draganska, Chad Jones, Dale Jorgenson, Larry Katz, Pete Klenow, Zhongjun Qu, Juan Sanchis, Matthias Schündeln, and John Van Reenen for helpful discussions.

[†]Department of Economics, Harvard University, Littauer Center, 1875 Cambridge Street, Cambridge, MA 02138, USA. E-mail: doraszelski@harvard.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. As vividly illustrated by the major new technologies introduced during the Industrial Revolution and the triumph of a middle class of industrialists and businessmen over a landed class of nobility and gentry (see, e.g., Mokyr 1990), the bias of technological change determines which societal groups are the winners and which are the losers and thus their willingness to embrace technological change.

Factor-neutral (also called Hicks-neutral) technological change is assumed, either explicitly or implicitly, in 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, Akerberg, Caves & Frazer 2006). These techniques therefore do not allow us to assess whether technological change is biased towards some factors of production. Moreover, even the results of a simple decomposition can be misleading if there is a bias (Bessen 2008).

In contrast, the literature on economic growth rests on the assumption of labor-augmenting technological change. 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 increase the relative productivity of labor vis-à-vis other factors of production. 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 establishing theoretically that profit-maximizing incentives can ensure that technological change is, at least in the long run, purely labor augmenting (Acemoglu 2003, Jones 2005). Whether this is indeed the case is an empirical question that remains to be answered.

Recent research also points to biased technological change as a key driver of the diverging experiences of the continental European and U.S. and U.K. economies during the 1980s and 1990s (Blanchard 1997, Caballero & Hammour 1998, Bentolila & Saint-Paul 2004, McAdam & William 2013). The unemployment rate rose steadily in a number of continental European economies, most egregiously in Spain (Bentolila & Jimeno 2006), while the share of labor in income has fallen sharply. In contrast, the labor share has been stable in the U.S. and U.K. While these medium-run dynamics are certainly consistent with biased technological change in the continental European economies, there is little direct evidence for it.

Despite the importance of assessing the bias of technological change, the empirical evidence is relatively scarce, perhaps owing to a lack of suitable data. Following early work by Brown & de Cani (1963) and David & van de Klundert (1965), economists have esti-

mated aggregate production or cost functions that proxy for labor- and capital-augmenting technological change with time trends (see, e.g., Lucas 1969, Kalt 1978, Antràs 2004, Binswanger 1974*a*, Jin & Jorgenson 2008).¹ While this line of research has produced some evidence of labor-augmenting technological progress, 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 NIPA accounts (see, e.g., Gordon 1990, Krueger 1999).

In this paper we combine recently available firm-level panel data with advances 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 further relate the speed and the direction of technological change to firms' R&D activities.

We propose a general framework for estimating production functions that enables us to separate Hicks-neutral from labor-augmenting technological change. Our empirical framework has two key features. First, we account for firm-level heterogeneity in productivity by assuming that the evolution of productivity is subject to random shocks. We take these innovations to productivity to represent the resolution over time of all uncertainties, in particular those inherent in the R&D process such as chance in discovery and success in implementation. Because we allow the productivity innovations to accumulate over time, they can cause persistent differences in productivity 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. Measuring firm-level productivity also gives us the ability to quantify the contribution of different types of firms and turnover among firms to the average.

Second, we assess the role of R&D in determining the differences in productivity across firms and the evolution of firm-level productivity over time. We focus on R&D because it is widely seen as a major source of technological change and productivity growth (see Griliches (1998, 2000) for surveys of the empirical literature). R&D is also a natural lever for encouraging technological progress. The link between R&D and productivity is therefore of immediate interest for policy makers.

To illustrate our empirical framework we analyze an unbalanced panel of more than 1800 Spanish manufacturing firms in nine industries from 1990 to 2006. Spain is an attractive setting for examining the speed and direction of technological change and its relationship to R&D for two reasons. First, Spain is a highly developed country that became fully integrated into the European Union between the end of the 1980s and the beginning of

¹A much larger literature has estimated the elasticity of substitution in aggregate production functions whilst maintaining the assumption of Hicks-neutral technological progress, see, e.g., Arrow, Chenery, Minhas & Solow (1961), McKinnon (1962), Kendrick & Sato (1963), and Berndt (1976) for early work and Hamermesh (1993) for a survey.

the 1990s. As such, any trends in technological change that our analysis uncovers for Spain may 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. R&D is traditionally 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 importance of R&D in stimulating different types of technological change with potentially different implications for employment and income inequality.

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 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. 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 about prices and quantities, at least for a subset of industries (Roberts & Supina 1996, Foster, Haltiwanger & Syverson 2008*a*).

Our approach builds on recent advances in the structural estimation of production functions to measure Hicks-neutral and labor-augmenting technological change and their relationship with R&D at the level of the individual firm. Olley & Pakes (1996), Levinsohn & Petrin (2003), and Akerberg et al. (2006) have shown how to recover a single-dimensional productivity measure from a Cobb-Douglas production function. Doraszelski & Jaumandreu (2013) have shown how to endogenize the productivity process by incorporating R&D expenditures into the model. We extend this line of research by generalizing the Cobb-Douglas specification of the production function and allowing for a multi-dimensional productivity measure.

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 capital-biased (labor-saving) productivity from a plant's decision on a variable input. Building on Doraszelski & Jaumandreu (2013), our approach is similar in that it uses a parametric inversion to recover unobserved productivity from observed inputs, but it differs in that we recover a multi-dimensional productivity measure. This simplifies the estimation. 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 further related to the literature on skill bias and the techniques we develop may also be applied to investigate the skill bias of technological change. While we focus on

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

the differential impact of technological change on labor vs other factors of production, the literature on skill bias studies the impact on different types of labor. A question that has received considerable attention is how the mix of skilled and unskilled labor changes over time, in particular in response to computerization. Much work aims to detect correlations in the data using a reduced-form approach (see, e.g., Autor, Katz & Krueger 1998, Autor, Levy & Murnane 2003). A number of obvious issues arise. First, since the data shows the equilibrium quantity and/or price of the different types of labor, the supply side has to be held fixed in order to isolate the effect of computerization. Some recent work employs a production/cost function perspective in an attempt to use factor prices to control for supply-side changes (see, e.g. Machin & Van Reenen 1998, Black & Lynch 2001, Bloom, Sadun & Van Reenen 2007). Second, because computerization is endogenous, some recent work uses changes in the regulatory environment as instruments (see, e.g., Acemoglu & Finkelstein 2008). Third, while the data is usually fairly aggregate, some recent work uses matched employer-employee data sets to analyze changes at the level of the individual firm (see, e.g., Abowd, Haltiwanger, Lane, McKinney & Sandusky 2007).

Our approach is similar to some of the recent work on skill bias in that it starts from a production function and focuses on the individual firm. Our approach differs by explicitly modeling (and estimating) the differences in productivity across firms and the evolution of firm-level productivity over time. It is perhaps also more structural in tackling the endogeneity problems that arise in estimating production functions (see Section 3 for details).

The remainder of this paper is organized as follows: Section 2 sets out a dynamic investment model in which a firm that can invest in R&D in order to improve its productivity over time in addition to carrying out a series of investments in physical capital. Productivity comprises a Hicks-neutral component and a labor-augmenting component and the evolution of both components is governed by controlled stochastic processes that capture the uncertainties inherent in R&D. Section 3 develops an estimator for production functions that allows us to retrieve productivity and its relationship with R&D at the firm level. Sections 4 and 5 describe the data and some preliminary results. Section 6 outlines a number of robustness checks that we plan to undertake and Section 7 discusses extensions and directions for future research.

2 A dynamic investment model

A firm carries out two types of investments, one in physical capital and another in productivity through R&D expenditures. The investment decisions are made in a discrete time setting with the goal of maximizing the expected net present value of future cash flows. Capital is the only fixed (or “dynamic”) input among the conventional factors of production and accumulates according to $K_{jt} = (1 - \delta)K_{jt-1} + I_{jt-1}$, where K_{jt} is the stock of capital of firm j in period t and δ is the rate of depreciation. This law of motion implies that investment

I_{jt-1} chosen in period $t - 1$ becomes productive in period t . The productivity of firm j in period t is the tuple $(\omega_{Ljt}, \omega_{Hjt})$, where ω_{Ljt} and ω_{Hjt} is labor-augmenting productivity and Hicks-neutral productivity, respectively. The components of productivity are presumably correlated with each other and over time and possibly also correlated across firms. We assume that they are governed by controlled first-order, time-inhomogeneous Markov processes with transition probabilities $P_{Lt}(\omega_{Ljt}|\omega_{Ljt-1}, r_{jt-1})$ and $P_{Ht}(\omega_{Hjt}|\omega_{Hjt-1}, r_{jt-1})$, where r_{jt-1} is the log of R&D expenditures.³ Because these stochastic processes are time-inhomogeneous, they can accommodate secular trends in productivity. We also refer to the tuple $(\omega_{Ljt}, \omega_{Hjt})$ as unobserved productivity since it is unobserved from the point of view of the econometrician (but known to the firm).

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

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

where Y_{jt} is the output of firm j in period t , K_{jt} capital, L_{jt} labor, and M_{jt} materials. The parameters ν and σ are the elasticity of scale and the elasticity of substitution, respectively. The remaining parameters of the production function are γ , a constant of proportionality, and β_K , β_L , and β_M , the so-called distributional parameters. In contrast to the tuple $(\omega_{Ljt}, \omega_{Hjt})$, e_{jt} is a mean zero random shock that is uncorrelated over time and across firms. The firm does not know the value of e_{jt} at the time it makes its decisions for period t .

The CES is the simplest specification that allows for biased technological change. Depending on the elasticity of substitution, it encompasses the special cases of a Leontieff ($\sigma \rightarrow 0$), Cobb-Douglas ($\sigma = 1$), and linear ($\sigma \rightarrow \infty$) production function. As is well known, a Cobb-Douglas production function has an elasticity of substitution of one. It is therefore not possible to separate labor-augmenting from Hicks-neutral technological change. Yet, our data reject a Cobb-Douglas production function. More generally, the estimates of σ in the previous literature lie somewhere between 0 and 1 (see, e.g., Arrow et al. 1961, McKinnon 1962, Kendrick & Sato 1963, Berndt 1976), thus allowing us to identify the direction of technological change.

In modeling technological change in equation (1) we hold fixed the parameters of the production function since we are not aware of evidence that suggests that either the elasticity of scale ν or the elasticity of substitution σ vary over time. Technological change instead operates by changing the efficiencies of the various factors of production. Sato & Beckmann (1968) show that the traditional definition of Hicks neutrality and labor augmentation requires these types of technological change to be modeled as in equation (1) (see also Burmeister & Dobell 1969, Sato 1970).

The production function in equation (1) is tailored to answer the research question at

³Throughout we follow the convention that lower case letters denote logs and upper case letters levels.

hand, namely to separate Hicks-neutral from labor-augmenting technological change. It allows us to capture in a parsimonious way the evolution of the share of labor in variable cost over time in our data (see Section 4 for details). At the same time, the production function in equation (1) restricts the efficiencies of capital and materials in production to change at the same rate and in lockstep with Hicks-neutral technological change. Whether a more general formulation of technological change is warranted is an empirical question that we plan to address in future work. At this stage of the research project we just note that treating capital and materials the same is in line with the fact that they are both produced goods, at least to a large extent. In contrast, labor is traditionally viewed as unique among the various factors of production. Marshall (1920), for example, writes in great detail about the variability of workers' efforts and its relationship to productivity.

The Bellman equation for the firm's dynamic programming problem is

$$V_t(S_{jt}) = \max_{I_{jt}, R_{jt}} \Pi(K_{jt}, \omega_{Ljt}, \omega_{Hjt}, Z_{jt}, W_t, P_{Mt}) - C_I(I_{jt}, X_{jt}) - C_R(R_{jt}, X_{jt}) \\ + \frac{1}{1 + \rho} E_t [V_{t+1}(S_{jt+1}) | S_{jt}, I_{jt}, R_{jt}],$$

where $S_{jt} = (K_{jt}, \omega_{Ljt}, \omega_{Hjt}, Z_{jt}, X_{jt}, W_t, P_{Mt})$ denotes the vector of state variables (to be defined below), $\Pi(\cdot)$ per-period profits, $C_I(\cdot)$ and $C_R(\cdot)$ the cost of investment in physical capital and productivity, respectively, and ρ the discount rate.

The dynamic programming problem gives rise to two policy functions, $I_t(S_{jt})$ and $R_t(S_{jt})$ for the investments in physical capital and productivity, respectively. X_{jt} denotes anything that shifts the costs of these investments over time and across firms. For example, opportunities to invest in physical capital and the price of equipment goods are likely to vary and the marginal cost of investment in productivity depends greatly on the nature of the undertaken R&D project. The cost functions $C_I(\cdot)$ and $C_R(\cdot)$ may further capture indivisibilities in investment projects or adjustment costs, but their exact forms are irrelevant for our purposes.

When the decision about investment in productivity is made in period t , the firm is only able to anticipate the expected effect of R&D on productivity in period $t + 1$. The Markovian assumption implies

$$\omega_{Ljt+1} = E_t [\omega_{Ljt+1} | \omega_{Ljt}, r_{jt}] + \xi_{Ljt+1} = g_{Lt}(\omega_{Ljt}, r_{jt}) + \xi_{Ljt+1}, \\ \omega_{Hjt+1} = E_t [\omega_{Hjt+1} | \omega_{Hjt}, r_{jt}] + \xi_{Hjt+1} = g_{Ht}(\omega_{Hjt}, r_{jt}) + \xi_{Hjt+1}.$$

That is, *actual* labor-augmenting productivity ω_{Ljt+1} in period $t+1$ can be decomposed into *expected* labor-augmenting productivity $g_{Lt}(\omega_{Ljt}, r_{jt})$ and a random shock ξ_{Ljt+1} . Actual Hicks-neutral productivity ω_{Hjt+1} can be decomposed similarly. Our key assumption is that the impact of R&D on productivity can be expressed through the dependence of the conditional expectation functions $g_{Lt}(\cdot)$ and $g_{Ht}(\cdot)$ on R&D expenditures. In contrast,

ξ_{Ljt+1} and ξ_{Hjt+1} do not depend on R&D expenditures: by construction ξ_{Ljt+1} and ξ_{Hjt+1} are mean independent (although not necessarily fully independent) of r_{jt} . These *productivity innovations* may be thought of as the realization of the uncertainties that are naturally linked to productivity plus the uncertainties inherent in the R&D process (e.g., chance in discovery, degree of applicability, success in implementation). It is important to stress the timing of decisions in this context: When the decision about investment in productivity is made in period t , the firm is only able to anticipate the expected effect of R&D on productivity in period $t+1$ as given by $g_{Lt}(\omega_{Ljt}, r_{jt})$ and $g_{Ht}(\omega_{Hjt}, r_{jt})$ while its actual effect also depends on the realizations of the productivity innovations ξ_{Ljt+1} and ξ_{Hjt+1} that occur after the investment has been completely carried out. Of course, the conditional expectation functions $g_{Lt}(\cdot)$ and $g_{Ht}(\cdot)$ are unobserved from the point of view of the econometrician (but known to the firm) and must be estimated nonparametrically.

At this stage of the research project we focus on R&D as a source of productivity growth. The adoption of technologies such as computers and factory automation is another source of productivity growth that can be accommodated in our framework by adding the relevant expenditures into the conditional expectation functions $g_{Lt}(\cdot)$ and $g_{Ht}(\cdot)$. Our data has information on investment in computer equipment and indicators of whether a firm has adopted digitally controlled machine tools, CAD, and robots. From 1990 to 2006 the number of small firms that use each type of technology has more than doubled; the number of large firms has also increased substantially. This points to a potentially important effect of technology adoption on productivity.

In sum, our model of unobserved firm-level productivity builds on the previous literature on the structural estimation of production functions in two ways. First, Olley & Pakes (1996), Levinsohn & Petrin (2003), and Akerberg et al. (2006) specify a Cobb-Douglas production function and an exogenous Markov process that governs the evolution of productivity. In this context, productivity is single-dimensional or, equivalently, technological change is Hicks neutral by construction. Second, Doraszelski & Jaumandreu (2013) endogenize the productivity process by incorporating R&D expenditures into the dynamic investment model. Their model accommodates nonlinearities and uncertainties in the link between R&D and productivity and can be viewed as a generalization of the simplest—albeit widely used—variants of the classic knowledge capital model (see Griliches (1979, 2000)). By relaxing the Cobb-Douglas specification and allowing productivity to be multi-dimensional, our present setup is a further generalization that accounts for three major characteristics of the productivity process: nonlinearity, uncertainty, and the factor-specific nature or bias of technological change.

The firm’s dynamic programming problem—and the policy functions for the investments in physical capital and productivity it gives rise to—can easily become very complicated. We thus base our empirical strategy on the firm’s decisions on variable (or “static”) inputs. These decisions are subsumed in per-period profits $\Pi(\cdot)$. Specifically,

we assume that labor L_{jt} and materials M_{jt} are chosen to maximize per-period profits with productivity $(\omega_{Ljt}, \omega_{Hjt})$ known. This gives rise to two input demand functions $L(K_{jt}, \omega_{Ljt}, \omega_{Hjt}, Z_{jt}, W_t, P_{Mt})$ and $M(K_{jt}, \omega_{Ljt}, \omega_{Hjt}, Z_{jt}, W_t, P_{Mt})$, where Z_{jt} are demand shifters, W_t is the wage prevailing in the market, and P_{Mt} the price of materials. We allow the firm to have some market power in the output market, say because products are differentiated, but assume that it behaves as a price-taker in the input markets. Our model can be consistent with unionization as long as the individual firm has no impact on the wage prevailing in the market.⁴ Importantly, the market-wide price of labor and materials, W_t and P_{Mt} , may or may not be the same as the firm-specific price of labor and materials, W_{jt} and P_{Mjt} , that we have in our data set. We give further details on our model of inputs markets in Section 3.

3 Empirical strategy

Perhaps the major obstacle in production function estimation is that the decisions that a firm makes depend on its productivity. Because the productivity of the firm is unobserved by the econometrician, this gives rise to an endogeneity problem (Marschak & Andrews 1944).⁵ Intuitively, if a firm adjusts to a change in its productivity by expanding or contracting its production depending on whether the change is favorable or not, then unobserved productivity and input usage are correlated and biased estimates result.

Recent advances in the structural estimation of production functions, starting with the dynamic investment model of Olley & Pakes (1996), tackle this issue. The insight is that if investment is a monotone function of unobserved (single-dimensional) productivity, then this function can be inverted to back out productivity. Controlling for productivity resolves the endogeneity problem.

As noted by Levinsohn & Petrin (2003) and Akerberg et al. (2006), backing out unobserved productivity from the demand of a variable input such as labor or materials is a convenient alternative to backing out unobserved productivity from investment. In the tradition of Olley & Pakes (1996), however, Levinsohn & Petrin (2003) and Akerberg et al. (2006) use nonparametric methods to estimate the inverse input demand function. This forces them either to rely on a two-stage procedure or to jointly estimate a system of equations as suggested by Wooldridge (2004). The drawback of the two-stage approach is a loss of efficiency whereas the joint estimation of a system of equations is numerically more demanding.

Doraszelski & Jaumandreu (2013) recognize that a nonparametric inversion is unnecessary because, given a parametric specification of the production function, the functional

⁴Wages in Spain are determined by collective bargaining between unions and employer organizations that takes place at the level of the industry and region. Many firms, however, engage in additional bargaining with their labor force that may run counter to our price-taking assumption.

⁵See Griliches & Mairesse (1998) and Akerberg, Benkard, Berry & Pakes (2007) for reviews of this and other problems involved in the estimation of production functions.

form of the inverse input demand function is known. They propose a parametric inversion that fully exploits the structural assumptions and renders identification and estimation more tractable. Below we apply their approach to our substantially more general model. In our model a parametric inversion is particularly advantageous because unobservable productivity is multi-dimensional. Hence, we require as many input demands as there are components of productivity. Showing that these input demands are invertible is easy if the parametric specification of the production function is utilized, but much harder if it is not.

Consider a downward sloping demand function that depends on the price of the output P_{jt} and the demand shifters Z_{jt} . Profit maximization requires that the firm sets the price that equates marginal cost to marginal revenue $P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, z_{jt})}\right)$, where $\eta(\cdot)$ is the absolute value of the elasticity of demand. Hence, the conditional demands for labor and materials can be written as

$$\begin{aligned} L_{jt} &= (\gamma\nu\beta_L\mu)^\sigma \left(\frac{W_{jt}}{P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, z_{jt})}\right)} \right)^{-\sigma} X_{jt}^{-\sigma(1+\frac{\nu\sigma}{1-\sigma})} \exp(\sigma\omega_{Hjt}) \exp(-(1-\sigma)\omega_{Ljt}) \\ M_{jt} &= (\gamma\nu\beta_M\mu)^\sigma \left(\frac{P_{Mjt}}{P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, z_{jt})}\right)} \right)^{-\sigma} X_{jt}^{-\sigma(1+\frac{\nu\sigma}{1-\sigma})} \exp(\sigma\omega_{Hjt}), \end{aligned} \quad (3)$$

where W_{jt} and P_{Mjt} are the price of labor and materials, respectively, and we define the shorthands $\mu = E_t[\exp(e_{jt})]$ and

$$X_{jt} = \beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \beta_L (\exp(\omega_{Ljt}) L_{jt})^{-\frac{1-\sigma}{\sigma}} + \beta_M M_{jt}^{-\frac{1-\sigma}{\sigma}}.$$

Solving the input demands in equations (2) and (3) for ω_{Ljt} and ω_{Hjt} we obtain

$$\begin{aligned} \tilde{\omega}_{Ljt} &\equiv \tilde{h}_L(m_{jt} - l_{jt}, p_{Mjt} - w_{jt}) = -a + (m_{jt} - l_{jt}) + \sigma(p_{Mjt} - w_{jt}), \\ \omega_{Hjt} &\equiv h_H(K_{jt}, S_{Mjt}, M_{jt}, p_{Mjt}, p_{jt}, z_{jt}) \\ &= -b + \frac{1}{\sigma} m_{jt} + p_{Mjt} - p_{jt} - \ln \left(1 - \frac{1}{\eta(p_{jt}, z_{jt})} \right) \\ &\quad + \left(1 + \frac{\nu\sigma}{1-\sigma} \right) \ln \left[\beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \beta_M \frac{1}{S_{Mjt}} M_{jt}^{-\frac{1-\sigma}{\sigma}} \right], \end{aligned} \quad (4)$$

where $\tilde{\omega}_{Ljt} = (1-\sigma)\omega_{Ljt}$ is labor-augmenting productivity subject to a convenient normalization, $a = \sigma \ln \left(\frac{\beta_M}{\beta_L} \right)$, $b = \ln(\gamma\nu\beta_M\mu)$, $S_{Mjt} = \frac{P_{Mjt}M_{jt}}{W_{jt}L_{jt} + P_{Mjt}M_{jt}}$ is the share of materials in variable cost, and, recall, $m_{jt} = \ln M_{jt}$. The function $\tilde{h}_L(\cdot)$ allows us to recover unobservable labor-augmenting productivity $\tilde{\omega}_{Ljt}$ from materials per unit of labor and the price of materials relative to the price of labor. The function $h_H(\cdot)$ similarly allows us to recover unobservable Hicks-neutral productivity ω_{Hjt} from observables.⁶ From hereon we

⁶In practice the function $\eta(\cdot)$ is unknown and must be estimated nonparametrically.

refer to $\tilde{h}_L(\cdot)$ and $h_H(\cdot)$ as inverse functions and use \tilde{h}_{Ljt} and h_{Hjt} as shorthands for their value $\tilde{h}_L(m_{jt} - l_{jt}, p_{Mjt} - w_{jt})$ and $h_H(K_{jt}, S_{Mjt}, M_{jt}, p_{Mjt}, p_{jt}, z_{jt})$. As usual in the literature on structurally estimating production functions (Olley & Pakes 1996, Levinsohn & Petrin 2003, Akerberg et al. 2006, Doraszelski & Jaumandreu 2013) we rely on the ability to recover unobserved productivity from observed decisions; this, in turn, presumes that the quantities and prices of variable inputs are accurately measured.

While we use equation (4) to recover unobservable labor-augmenting productivity, much empirical work uses relationships like it to directly estimate the elasticity of substitution from observations on quantities and prices of variable inputs. Rewriting equation (4) as

$$(m_{jt} - l_{jt}) = a - \sigma(p_{Mjt} - w_{jt}) + \tilde{\omega}_{Ljt} \quad (6)$$

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, according to the price of materials relative to the price of labor. For example, if the relative wage rises, then materials per unit of labor rises. Second, labor-augmenting technological change increases materials per unit of labor. A rise in $\tilde{\omega}_{Ljt}$ directly leads *ceteris paribus* to a fall in l_{jt} . This is the displacement effect of labor-augmenting technological change.⁷ By contrast, Hicks-neutral technological change does not have a similar, relative displacement effect.

Often OLS is used on equation (6). The problem is that unobserved labor-augmenting productivity is correlated over time and also with the price of labor. The wage is likely to be higher when labor is more productive, even if it adjusts slowly with some lag. This positive correlation is likely to induce 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 estimates of the elasticity of substitution may be biased as a result (see, e.g., the discussion in Antràs 2004). Proxying for unobserved labor-augmenting productivity by a time trend, time dummies, or a measure of innovation is unlikely to completely remove the bias. Antràs (2004) shows that estimates of the elasticity of substitution in an aggregate production function can be improved by including a time trend and allowing for serial correlation, but it is doubtful that all structure has been removed from the error term. Intuitively, failure to fully account for the evolution of productivity leaves an error term that remains correlated with the ratio of prices. Using firm-level panel data, Van Reenen (1997) proxies for unobserved labor-augmenting productivity by the number of innovations commercialized in a given year. His approach to estimation assumes that

⁷There is another effect: As labor-augmenting technological change decreases the marginal cost of production, the usage of all factors of production increases. In contrast to the displacement effect, this output effect affects the various factors in equal terms. The available evidence suggests that the output effect of technological change on employment typically more than offsets the displacement effect (Harrison, Jaumandreu, Mairesse & Peters 2008). While we do not pursue this avenue, the techniques developed in the present paper can be used not only to separate labor-augmenting from Hicks-neutral technological change but also to quantify their various effects on employment.

the remaining error term is white noise and is thus unlikely to succeed if productivity is governed by a more general stochastic process. Indeed, Van Reenen (1997) obtains a positive effect of innovation on employment, contrary to what is expected from theory and the displacement effect of labor-augmenting technological change.

Our model of unobserved firm-level productivity provides a natural way to resolve the endogeneity problem and estimate equation (6) based on Doraszelski & Jaumandreu (2013). Using \tilde{h}_{Ljt-1} to replace $\tilde{\omega}_{Ljt-1}$ in the law of motion for labor-augmenting productivity, we obtain

$$(m_{jt} - l_{jt}) = a - \sigma(p_{Mjt} - w_{jt}) + \tilde{g}_{Lt-1}(\tilde{h}_L(m_{jt-1} - l_{jt-1}, p_{Mjt-1} - w_{jt-1}), r_{jt-1}) + \tilde{\xi}_{Ljt}, \quad (7)$$

where $\tilde{g}_{Lt-1}(\tilde{\omega}_{Ljt-1}, r_{jt-1}) = (1 - \sigma)g\left(\frac{\tilde{\omega}_{Ljt-1}}{1 - \sigma}, r_{jt-1}\right)$ and $\tilde{\xi}_{Ljt} = (1 - \sigma)\xi_{Ljt}$.

Our estimation equation (7) is a semiparametric, so-called partially linear, model with the additional restriction that the inverse function $\tilde{h}_L(\cdot)$ is of known form. Identification follows from standard arguments (Robinson 1988, Newey, Powell & Vella 1999). Equation (7) has the intuitive advantage over equation (6) that the included variables only have to be uncorrelated with the innovation to productivity $\tilde{\xi}_{Ljt}$ but not necessarily with productivity $\tilde{\omega}_{Ljt}$. To appreciate the difference, note that the lagged price of labor and materials is uncorrelated with $\tilde{\xi}_{Ljt}$ in equation (7) yet correlated with $\tilde{\omega}_{Ljt}$ in equation (6) (as long as productivity is correlated over time). Moreover, if the current price of labor and materials are exogenous to the firm, then none of the included variables, neither in the parametric nor in the nonparametric part of equation (7), is correlated with $\tilde{\xi}_{Ljt}$ by virtue of our timing assumptions.

Our data has firm-specific price indices for output and inputs. Prices vary substantially both across firms and across periods. To the extent that this variation is due to geographic and temporal differences in the supply of labor and materials and the fact that firms operate in different product submarkets, it is exogenous to firms and can thus be exploited for estimating equation (7).

A potential concern, however, is that the variation in prices is not exogenous. In particular, suppose the labor market is segmented into different quality tiers. Firms behave as price takers but choose in which tier to operate: Some firms may elect to hire high quality workers at high wages and others to hire low quality workers at low wages. If so, then the firm-specific wage in our data is higher in the first case where labor is also more productive than in the second case. This is a standard problem in the existing literature. By modeling the evolution of productivity equation (7) mitigates this problem because prices only have to be uncorrelated with the innovation to productivity. Whether this is the case in our data is testable by means of a standard overidentification test. Our estimates so far pass this test in 8 out of 9 industries (see Section 5 for details). At this stage of the research project, we therefore maintain the assumption of exogenous variation.

In the next stage of the project we plan to instrument for prices as part of our robustness

checks in Section 6. Industry-wide price indices such as the average wage of white-collar and blue-collar workers in a given industry and region are valid instruments for firm-specific price indices and are readily available. We can also construct instruments from our data set by computing the average wage at the other firms in the same industry and region as the firm under consideration. Finally, our data has information on the characteristics of a firm's labor force that can be exploited by first regressing wages on characteristics and then using predicted wages as instruments for actual wages.⁸

Equation (7) is enough to obtain an estimate of the elasticity of substitution and to recover labor-augmenting technological change at the firm level. To also recover Hicks-neutral technological change and the remaining parameters of the production function (elasticity of scale and distributional parameters) we proceed as follows. Using equation (4) to replace ω_{Ljt} in the production function in equation (1), we obtain

$$Y_{jt} = \gamma \left[\beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \beta_M \frac{1}{S_{Mjt}} M_{jt}^{-\frac{1-\sigma}{\sigma}} \right]^{-\frac{\nu\sigma}{1-\sigma}} \exp(\omega_{Hjt}) \exp(e_{jt}). \quad (8)$$

Using h_{Hjt-1} to replace ω_{Hjt-1} in the law of motion for Hicks-neutral productivity, we further obtain

$$y_{jt} = \ln \gamma - \frac{\nu\sigma}{1-\sigma} \ln \left[\beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \beta_M \frac{1}{S_{Mjt}} M_{jt}^{-\frac{1-\sigma}{\sigma}} \right] + g_{Ht-1}(h_H(K_{jt-1}, S_{Mjt-1}, M_{jt-1}, p_{Mjt-1}, p_{jt-1}, z_{jt-1}), r_{jt-1}) + \xi_{Hjt} + e_{jt}. \quad (9)$$

Without loss of generality we assume that $\beta_K + \beta_M = 1$.

Equation (9) is our second estimation equation. It is again a semiparametric model with the additional restriction that the inverse function $h_H(\cdot)$ is of known form. Both K_{jt} , whose value is determined in period $t-1$ by I_{t-1} , and r_{jt-1} are uncorrelated with ξ_{Hjt} by virtue of our timing assumptions. In contrast, M_{jt} is correlated with ξ_{Hjt} (since ξ_{Hjt} is part of ω_{Hjt} and M_{jt} is a function of ω_{Hjt}). While the share of materials S_{Mjt} depends on the ratio $\frac{M_{jt}}{L_{jt}}$ and thus directly only on ω_{Ljt} which includes ξ_{Ljt} , it is nevertheless correlated with ξ_{Hjt} if the productivity innovations ξ_{Hjt} and ξ_{Ljt} are correlated. Nonlinear functions of the other variables can be used as instruments for M_{jt} and S_{Mjt} , as can be lagged values of these two and the other variables. Lagged prices p_{Mjt-1} and p_{jt-1} and cost shifters z_{jt-1} are uncorrelated with ξ_{Hjt} .

Our estimation equation (7) has the advantage that it is directly comparable to traditional approaches to estimating the elasticity of substitution and our estimation equation (9) that it is not unlike a CES production function, a specification that has been widely

⁸Fluctuations in the utilization of labor over the business cycle are another possible source of correlation between wages and productivity innovations. Utilization may be increased by paying workers overtime, thus giving rise to a positive correlation. On the other hand, utilization may be increased by hiring additional temporary workers, a common practice in Spain during the 1990s. In this case the lower wages of temporary workers may work in the opposite direction. Our data set contains information on overtime and temporary workers that allows us to control for both effects.

used in the literature. There are other estimation equations that can be derived. In particular, one can use the conditional demands for labor and materials in equations (2) and (3) and the production function in equation (1) to recover $\tilde{\omega}_{Ljt} = \tilde{h}_L(m_{jt} - l_{jt}, p_{Mjt} - w_{jt})$, $\omega_{Hjt} = h_H(K_{jt}, S_{Mjt}, M_{jt}, p_{Mjt}, p_{jt}, z_{jt})$, and

$$\begin{aligned} e_{jt} &\equiv h_e(y_{jt}, K_{jt}, S_{Mjt}, M_{jt}, p_{Mjt}, p_{jt}, z_{jt}) \\ &= y_{jt} + \ln(\nu\beta_M\mu) - \ln \left[\beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \beta_M \frac{1}{S_{Mjt}} M_{jt}^{-\frac{1-\sigma}{\sigma}} \right] \\ &\quad - \frac{1}{\sigma} m_{jt} - p_{Mjt} + p_{jt} + \ln \left(1 - \frac{1}{\eta(p_{jt}, z_{jt})} \right), \end{aligned}$$

and then set up separate moment conditions in $\tilde{\xi}_{Ljt} = \tilde{\omega}_{Ljt} - \tilde{g}_{Lt-1}(\tilde{\omega}_{Ljt-1}, r_{jt-1})$, $\xi_{Hjt} = \omega_{Hjt} - g_{Ht-1}(\omega_{Ljt-1}, r_{jt-1})$, and e_{jt} . This may yield efficiency gains over our current approach.

Estimation. The most efficient way to obtain estimates is to jointly estimate equations (7) and (9) subject to the relevant cross-equation parameter constraints. Since the Hicks-neutral and labor-augmenting productivity innovations are likely to be correlated, joint estimation also yields efficiency gains. The estimation problem can thus be cast in the nonlinear GMM framework

$$E \begin{bmatrix} z'_{Ljt} \tilde{\xi}_{Ljt} \\ z'_{Hjt} (\xi_{Hjt} + e_{jt}) \end{bmatrix} = E \begin{bmatrix} z'_{Ljt} v_{Ljt}(\theta) \\ z'_{Hjt} v_{Hjt}(\theta) \end{bmatrix} = 0,$$

where z_{Ljt} and z_{Hjt} are vectors of instruments and we write the error terms $v_{Ljt}(\cdot)$ and $v_{Hjt}(\cdot)$ as functions of the parameters θ to be estimated. The objective function is

$$\min_{\theta} \begin{bmatrix} \frac{1}{N} \sum_j z'_{Lj} v_{Lj}(\theta) \\ \frac{1}{N} \sum_j z'_{Hj} v_{Hj}(\theta) \end{bmatrix}' A_N \begin{bmatrix} \frac{1}{N} \sum_j z'_{Lj} v_{Lj}(\theta) \\ \frac{1}{N} \sum_j z'_{Hj} v_{Hj}(\theta) \end{bmatrix},$$

where z'_{Lj} and z'_{Hj} are $L_L \times T_j$ and $L_H \times T_j$ matrices, $v_{Lj}(\cdot)$ and $v_{Hj}(\cdot)$ are $T_j \times 1$ vectors, L_L and L_H are the number of instruments for equations (7) and (9), respectively, T_j is the number of observations of firm j , and N is the number of firms. We first use the weighting matrix

$$A_N = \begin{bmatrix} \left(\frac{1}{N} \sum_j z'_{Lj} z_{Lj} \right)^{-1} & 0 \\ 0 & \left(\frac{1}{N} \sum_j z'_{Hj} z_{Hj} \right)^{-1} \end{bmatrix}$$

to obtain a consistent estimator and then we compute the optimal estimator using the weighting matrix

$$A_N = \begin{bmatrix} \frac{1}{N} \sum_j z'_{Lj} v_{Lj}(\hat{\theta}) v_{Lj}(\hat{\theta})' z_{Lj} & \frac{1}{N} \sum_j z'_{Lj} v_{Lj}(\hat{\theta}) v_{Hj}(\hat{\theta})' z_{Hj} \\ \frac{1}{N} \sum_j z'_{Hj} v_{Hj}(\hat{\theta}) v_{Lj}(\hat{\theta})' z_{Lj} & \frac{1}{N} \sum_j z'_{Hj} v_{Hj}(\hat{\theta}) v_{Hj}(\hat{\theta})' z_{Hj} \end{bmatrix}^{-1}.$$

Series estimator. The conditional expectation functions $\tilde{g}_{Lt-1}(\cdot)$ and $g_{Ht-1}(\cdot)$ are unknown and must be estimated nonparametrically, as must be the absolute value of the elasticity of demand $\eta(\cdot)$. As suggested by Wooldridge (2004) we model an unknown function $q(\cdot)$ of one variable v by a univariate polynomial of degree Q . We model an unknown function $q(\cdot)$ of two variables v and u by a complete set of polynomials of degree Q (see Judd 1998). In the remainder of this paper we set $Q = 3$.

We allow for a different conditional expectation function when a firm adopts the corner solution of zero R&D expenditures and when it chooses positive R&D expenditures and specify

$$\begin{aligned}\tilde{g}_{Lt-1}(\tilde{h}_{Ljt-1}, r_{t-1}) &= \tilde{\beta}_{Lt-1} + 1(R_{jt-1} = 0)\tilde{g}_{L0}(\tilde{h}_{Ljt-1} + a) + 1(R_{jt-1} > 0)\tilde{g}_{L1}(\tilde{h}_{Ljt-1} + a, r_{jt-1}), \\ g_{Ht-1}(h_{Hjt-1}, r_{t-1}) &= \beta_{Ht-1} + 1(R_{jt-1} = 0)g_{H0}(h_{Hjt-1} - b) + 1(R_{jt-1} > 0)g_{H1}(h_{Hjt-1} - b, r_{jt-1}).\end{aligned}$$

The functions $\tilde{g}_{L0}(\cdot)$, $\tilde{g}_{L1}(\cdot)$, $g_{H0}(\cdot)$, and $g_{H1}(\cdot)$ are modeled as described above. Note that their constants subsume the constants of the inverse functions $\tilde{h}_L(\cdot)$ and $h_H(\cdot)$. Moreover, the constants of $\tilde{g}_{L0}(\cdot)$ and $\tilde{g}_{L1}(\cdot)$ cannot be estimated separately from a in equation (7) and, similarly, the constants of $g_{H0}(\cdot)$ and $g_{H1}(\cdot)$ cannot be estimated separately from $\ln \gamma$ in equation (9). We thus specify an overall constant and a separate dummy for firms that perform R&D. Finally, given that the Markov processes governing productivity may be time-inhomogeneous, we allow the conditional expectation functions $\tilde{g}_{Lt-1}(\cdot)$ and $g_{Ht-1}(\cdot)$ to shift over time. We represent this displacement by $\tilde{\beta}_{Lt-1}$ and β_{Ht-1} and, in practice, model it with time trends or dummies.

We specify the absolute value of the elasticity of demand as $1 + \exp(q_0 + q(p_{jt-1}, z_{jt-1}))$, where the function $q(\cdot)$ is modeled as described above, in order to impose the theoretical restriction $\eta(p_{jt-1}, z_{jt-1}) > 1$.

Estimation (cont'd). To simplify the optimization problem in the estimation, we follow a suggestion of Akerberg et al. (2006) and “concentrate out” the parameters making up the conditional expectation functions $\tilde{g}_{Lt-1}(\cdot)$ and $g_{Ht-1}(\cdot)$. Let θ_1 denote these parameters, including the time trends or dummies $\tilde{\beta}_{Lt-1}$ and β_{Ht-1} , and note that they enter linearly in our estimation equations (7) and (9). Let θ_0 denote the remaining parameters, so that $\theta = (\theta_0, \theta_1)$. Then it suffices to search over θ_0 . We proceed as follows. Given θ_0 , compute \tilde{h}_{Ljt} and h_{Hjt} from equations (4) and (5). Regress \tilde{h}_{Ljt} on a complete polynomial in \tilde{h}_{Ljt-1} and r_{jt-1} (plus a constant?) to obtain the parameters making up $\tilde{g}_{Lt-1}(\cdot)$. Regress h_{Hjt} on a complete polynomial in h_{Hjt-1} and r_{jt-1} (plus a constant?) to obtain the parameters making up $g_{Ht-1}(\cdot)$. Finally, compute $\tilde{\xi}_{Ljt}$ in equation (7) and $\xi_{Hjt} + e_{jt}$ in equation (9) and proceed as before by interacting them with z_{Ljt} and z_{Hjt} .

Instrumental variables. Our first estimation equation (7) has 16 (23) parameters: constant, σ , time trend (or eight dummies) and 13 coefficients in the series approximation of

$\tilde{g}_{Lt-1}(\cdot)$. In addition to the constant we use $p_{Mjt} - w_{jt}$ as instrument, the time trend (or dummies) and the dummy for performers. We further use a complete set of polynomials in $m_{jt-1} - l_{jt-1}$, and $p_{Mjt-1} - w_{jt-1}$ (9 instruments) to instrument the part of the series approximation corresponding to the nonperformers and a complete set of polynomials in $m_{jt-1} - l_{jt-1}$, $p_{Mjt-1} - w_{jt-1}$, and r_{jt-1} (19 instruments) for the part corresponding to the performers.

Our second estimation equation (9) has 28 (35) parameters: constant, σ , ν , β_K , β_M , time trend (or eight dummies), 13 coefficients in the series approximation of $g_{Ht-1}(\cdot)$, and nine coefficients in the series approximation of $\eta(\cdot)$. In the parametric part, in addition to the constant we use K_{jt} , M_{jt-1} , and S_{jt-1} as instruments. In the nonparametric part, we use the time trend (or dummies) and a dummy for performers as instruments. We further use a complete set of polynomials in the time trend, m_{jt-1} , $p_{Mjt-1} - p_{jt-1}$, K_{jt-1} , M_{jt-1} for nonperformers plus r_{jt-1} for performers. Finally, we use the complete set of polynomials in the variables p_{jt-1} and z_{jt-1} (9 instruments).

At this stage of the research project we use a broad set of instruments. Further experimentation with subsets of these instruments is likely to yield efficiency gains.

Testing. The value of the GMM objective function for the optimal estimator, multiplied by N , has a limiting χ^2 distribution with $L - P$ degrees of freedom, where $L = L_L + L_H$ 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 based on the instruments.

We test whether the model satisfies certain restrictions by computing the restricted estimator using the weighting matrix for the optimal estimator and then comparing the values of the properly scaled objective functions. The difference has a limiting χ^2 distribution with degrees of freedom equal to the number of restrictions.

4 Data

We use an unbalanced panel of Spanish manufacturing firms during the 1990s and 2000s. We are currently using data from 1990 to 1999 and are in the process of incorporating additional data from 2000 to 2006.

Our data come 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 surveyed is the firm, not the plant or the establishment. At the beginning of this 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, and 70% of all firms of this size chose to respond. Some firms vanish from the sample, due to both exit and attrition. The two reasons can be distinguished, and attrition remained within acceptable

limits. To preserve representativeness, samples of newly created firms were added to the initial sample every year. Details on industry and variable definitions can be found in Appendix A.

Given that our estimation procedure requires a lag of one year, we restrict the sample to firms with at least two years of data. The resulting sample covers a total of 1879 firms in nine industries. Columns (1) and (2) of Table 1 show the number of observations and firms by industry. The samples are of moderate size. Firms tend to remain in the sample for short periods, ranging from a minimum of two years to a maximum of 10 years between 1990 and 1999.⁹

An attractive feature of our data is that it contains firm-specific price indices for output and inputs. Prices vary both across firms and across periods. The coefficient of variation for the price of output and the price of materials ranges from 0.12 to 0.25 across industries. The larger part of this variation is across firms: The coefficient of variation in the cross section ranges from 0.11 to 0.21 across industries compared to from 0.04 to 0.13 in the time series. The coefficient of variation for wages ranges from 0.31 to 0.45 across industries, and again the larger part of this variation is across firms.

Biased technological change. Columns (3)–(7) of Table 1 show that the 1990s were a period of rapid output growth, coupled with stagnant or at best slightly increasing employment and intense investment in physical capital. The growth of output prices, averaged from the growth of output prices as reported individually by each firm, is moderate.

The evolution of the prices and quantities of variables inputs already hint at the important role of labor-augmenting technological change. As columns (8) and (9) of Table 1 show, with the exceptions of industries 7, 8, and 9, the increase of materials per unit of labor is much larger than the decrease in the price of materials relative to the price of labor. In the absence of labor-augmenting technological change, explaining the evolution of variable inputs therefore requires an elasticity of substitution in excess of one (see equation (6)). But this is implausible because the estimates of σ in the previous literature lie somewhere between 0 and 1, despite possibly being upward biased. On the other hand, labor-augmenting technological change directly increases materials per unit of labor and may thus go a long way in explaining the pattern in the data. Indeed, it is easy to show that, in the presence of labor-augmenting technological change, if all variable inputs grow, then labor is the input that grows slowest, and if all variable inputs shrink, then labor is the input that shrinks fastest. Importantly, it can even be the case that labor shrinks while the other inputs grow, precisely as we observe in Table 1.¹⁰

⁹Newly created firms are a large share of the total number of firms, ranging from 15% to one third in the different industries. In each industry there is a significant proportion of exiting firms, ranging from 5% to above 10% in a few cases.

¹⁰An alternative explanation for the shift from labor to materials that comes to mind is outsourcing. However, the value of the outsourced task as a fraction of the value of materials is modest (ranging from 5% to 15%), decreases in 4 out of 9 industries between 1991 and 1999, and in the 5 industries where it increases,




Figure 1: Share of labor in variable cost by industry.

To model labor-augmenting technological change we cannot proceed with a Cobb-Douglas production function. Once again the data by themselves indicate that this customary specification is inappropriate. Figure 1 illustrates the evolution of the share of labor in variable cost over time and columns (10) and (11) of Table 1 provide summary measures. The average labor share is relatively uniform across industries, ranging from 0.26 to 0.37 (column (10)). There is a clear downward trend in the labor share from 1991 to 1999, as can be seen in Figure 1. The labor share decreases from its maximum to its minimum by between 0.06 and 0.12 depending on the industry (column (11)). These are substantial decreases given that the range across industries of the average labor share is 0.11. Overlaying the trend are short-term fluctuations. In most industries the labor share peaks in 1993, the year of a sharp (but short) slowdown in Europe, and shows an upward swing in the final years of the sample period. Overall, the pattern in the data is inconsistent with a Cobb-Douglas production function that implies that the labor share is necessarily constant.

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).¹¹ The manufacturing sector consists partly of transnational

does so at a slow pace.

¹¹R&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).

companies with production facilities in Spain and huge R&D expenditures and partly of small and medium-sized companies that invested heavily in R&D in a struggle to increase their competitiveness in a growing and already very open economy.¹²

Table 2 reveals that the nine industries differ markedly in terms of firms' R&D activities. Chemical products (3), agricultural and industrial machinery (4), and transport equipment (6) exhibit high innovative activity. The share of firms that perform R&D during at least one year in the sample period is about two thirds, with slightly more than 40% of stable performers that engage in R&D in all years (column (2)) and slightly more than 20% of occasional performers that engage in R&D in some (but not all) years (column (3)). The average R&D intensity among performers ranges from 2.2% to 2.7% (column (4)). The standard deviation of R&D intensity is substantial and shows that firms engage in R&D to various degrees and quite possibly with many different specific innovative activities. Metals and metal products (1), non-metallic minerals (2), food, drink and tobacco (7), and textile, leather and shoes (8) are in an intermediate position. The share of firms that perform R&D is below one half and there are fewer stable than occasional performers. The average R&D intensity is between 1.1% and 1.5% with a much lower value of 0.7% in industry 7. Finally, timber and furniture (9) and paper and printing products (10) exhibit low innovative activity. The share of firms that perform R&D is around one quarter and the average R&D intensity is 1.4%.

The fact that the nine industries differ markedly in terms of firms' R&D activities suggests that it is interesting to explore how the speed and direction of technological change differ across industries and how they are related to firms' R&D activities.

5 Preliminary results

At this stage of the research project we have successfully estimated equation (7). This already allows us to assess the elasticity of substitution. Moreover, we are able to test for the presence of labor-augmenting technological change, quantify it, and relate it to firms' R&D activities. The next step of the project is to incorporate equation (9) in the estimation in order to quantify Hicks-neutral technological change and relate it to firms' R&D activities.

Elasticity of substitution. Table 3 summarizes different estimates of the elasticity of substitution. To facilitate the comparison with the existing literature columns (1) and (2) report the elasticity of substitution and the time trend as estimated by running OLS on equation (6). With the exception of industry 9, the estimates of the elasticity of substitution

¹²At most a small fraction of the firms that engaged in R&D receive subsidies that typically cover between 20% and 50% of R&D expenditures. The impact of subsidies is mostly limited to the amount that they add to the project, without crowding out private funds (see Gonzalez, Jaumandreu & Pazo 2005). This suggests that R&D expenditures irrespective of their origin are the relevant variable for explaining productivity.

are in excess of one. 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 described in Section 3. Nevertheless, the positive time trend hints at the importance of labor-augmenting technological change.

We therefore account for firm-level heterogeneity in productivity and correct for endogeneity by estimating equation (7) by GMM. As expected the so-obtained estimates of the elasticity of substitution are much lower and range from 0.37 to 0.63 as can be seen from column (3) of Table 3. We clearly reject the special cases of both a Leontieff ($\sigma \rightarrow 0$) and a Cobb-Douglas ($\sigma = 1$) production function.

We test for overidentifying restrictions or validity of the moment conditions.¹³ With the exception of industry 8, the validity of the moment conditions cannot be rejected, see columns (4) and (5). Since the test is close in a number of industries, however, our assumption that the variation in prices is exogenous may be questionable. In the next stage of the project we thus plan to instrument for prices as part of our robustness checks in Section 6.

Labor-augmenting technological change. Once equation (7) is estimated we can recover unobserved labor-augmenting productivity $\omega_{Ljt} = \frac{\tilde{\omega}_{Ljt}}{1-\sigma}$ of firm j in period t up to a constant. We can therefore estimate $\omega_{Ljt} - \omega_{Ljt-1}$, the growth of labor-augmenting productivity for an individual firm or, by aggregating appropriately, for an entire industry. Note that a change in ω_{Ljt} shifts the firm’s “effective labor.” If we consider a change from ω_{Ljt} to ω'_{Ljt} , then *ceteris paribus* $\omega'_{Ljt} - \omega_{Ljt}$ approximates the effect of this change in productivity on effective labor in percentage terms, i.e., $(\exp(\omega'_{Ljt})L_{jt} - \exp(\omega_{Ljt})L_{jt}) / (\exp(\omega_{Ljt})L_{jt}) = \exp(\omega'_{Ljt} - \omega_{Ljt}) - 1 \simeq \omega'_{Ljt} - \omega_{Ljt}$.

Columns (6)–(8) of Table 3 report a weighted average of the growth of labor-augmenting productivity for the entire sample and for the subsamples of firms that perform R&D and those that do not. The weights $\mu_{jt} = Y_{jt-2} / \sum_j Y_{jt-2}$ are given by the share of output of a firm two periods ago.¹⁴ The rate of growth of labor-augmenting productivity is high in most industries, ranging from 4% to 11% per year. It is zero or slightly negative in industries 8 and 9. As may be expected the rate of growth of labor-augmenting productivity tends to be higher in the more capital-intensive industries 6 and, perhaps to a lesser extent, 3. Overall, our estimates clearly indicate the presence of biased technological change.¹⁵

¹³The value of the GMM objective function for the optimal estimator, multiplied by N , has a limiting χ^2 distribution with $L - P$ degrees of freedom, where L is the number of instruments and P the number of parameters to be estimated.

¹⁴In what follows we account for the survey design as follows. First, to compare the productivities of firms that perform R&D to those of firms that do not perform R&D we conduct separate tests on the subsamples of small and large firms. Second, to be able to interpret some of our descriptive statistics as aggregates that are representative for an industry as a whole, we replicate the subsample of small firms $\frac{70}{5} = 14$ times before merging it with the subsample of large firms.

¹⁵The growth of labor-augmenting productivity is in line with the literature on skill bias. Indeed, in our data there is a shift from unskilled to skilled workers. For example, the share of university graduates in

In the next step of the research project we will further assess the importance of labor-augmenting technological change relative to Hicks-neutral technological change. To provide a preview, consider again a change in labor-augmenting productivity from ω_{Ljt} to ω'_{Ljt} . The effect of this change on output is roughly a third of its effect on effective labor. Hence, labor-augmenting technological change causes output to grow in the vicinity of 2% per year.

Firms' R&D activities. The rate of growth in labor-augmenting productivity is, on average, higher for firms that engage in R&D than for firms that do not in 7 industries, sometimes considerably so. Taken together these industries account for over two thirds of manufacturing output. Hidden behind these averages, however, is a substantial amount of heterogeneity across firms that we will explore in more detail in the next step of the research project.

The percentage contributions to labor-augmenting productivity growth in columns (7) and (8) are telling. The contribution of firms that perform R&D accounts for between 70% and 120% of productivity growth in the industries with high innovative activity and between 55% and 80% in the industries with intermediate innovative activity (with the exception of industry 8). This is all the more remarkable since in these industries between 35% and 45% and between 10% and 20%, respectively, of firms engage in R&D. While these firms manufacture between 70% and 75% of output in the industries with high innovative activity and between 30% and 55% in the industries with intermediate innovative activity, their contribution to productivity growth is often much larger than their share of output. Firms' innovative activities are thus a primary source of labor-augmenting technological change.

We can further decompose unobserved labor-augmenting productivity ω_{Ljt} of firm j in period t into a part that can be anticipated when the decision about investment in productivity is made and a part that cannot. Both $g_{Lt-1}(\omega_{Ljt-1}, r_{jt-1}) = \frac{\tilde{g}_{Lt-1}((1-\sigma)\omega_{Ljt-1}, r_{jt-1})}{1-\sigma}$ and $\xi_{Ljt} = \frac{\tilde{\xi}_{Ljt}}{1-\sigma}$ can be interpreted in percentage terms and decompose the change in labor-augmenting productivity.

Hicks-neutral technological change. Once equation (9) is estimated we can recover Hicks-neutral productivity ω_{Hjt} up to a constant. We can therefore estimate $\omega_{Hjt} - \omega_{Hjt-1}$, the growth of Hicks-neutral productivity. Note that a change in ω_{Hjt} shifts the firm's production function. If we consider a change from ω_{Hjt} to ω'_{Hjt} , then *ceteris paribus* $\omega'_{Hjt} - \omega_{Hjt}$ approximates the effect of this change in productivity on output in percentage terms, i.e., $(Y'_{jt} - Y_{jt})/Y_{jt} = \exp(\omega'_{Hjt} - \omega_{Hjt}) - 1 \simeq \omega'_{Hjt} - \omega_{Hjt}$.

the labor force of small firms increases from 6.1% in 1991 to 10.7% in 2006 and from 9.4% to 18.3% for large firms. While it has to be seen against the backdrop of a general increase of university graduates in Spain during the 1990s and 2000s, the skill upgrading in our data corroborates the consensus view that technological change is complementary with skill (see, e.g. Hammermesh 1993).

Some very preliminary estimates of equation (9) suggest that Hicks-neutral technological change causes output to grow in the vicinity of 1% per year.

Firms' R&D activities. We can further decompose unobserved Hicks-neutral productivity ω_{Hjt} of firm j in period t into a part that can be anticipated when the decision about investment in productivity is made and a part that cannot. Both $g_{Ht-1}(\omega_{Hjt-1}, r_{Hjt-1})$ and ξ_{Hjt} can be interpreted in percentage terms and decompose the change in Hicks-neutral productivity.

6 Robustness checks

In this section we outline a number of robustness checks that we plan to undertake.

Estimation equations. As we already noted in Section 3, there are other estimation equations that can be derived and that may yield efficiency gains over our current approach.

Exogenous variation in prices. A potential concern is that the variation in prices that we exploit heavily in the estimation is not exogenous. In Section 3 we suggested various ways around this problem that remain to be explored.

Other sources of technological change. As we already noted in Section 2, other sources of technological change besides R&D may be important and can be accommodated in our framework by incorporating the relevant variables into the laws of motion for labor-augmenting and Hicks-neutral productivity.

Production function. Krusell, Ohanian, Rios-Rull & Violante (2000) argue that in a production function with capital, skilled labor, and unskilled labor it makes a difference whether one assumes a common elasticity of substitution between all factors of production or allows for capital-skill complementarities. The more general point here is that our results may be sensitive to the assumed CES specification of the production function. We can use a nested CES specification to relax the assumption of a common elasticity of substitution. Less restrictive specifications such as a translog or generalized Leontief production function may also be worth exploring.

7 Other extensions

In this section we discuss a number of extensions and directions for future research.

Nonparametric inversion. Our approach exploits the known form of the inverse input demand functions. Whether this parametric inversion is consistent with the data is testable by allowing for a more flexible functional form.

An interesting question is whether it is possible to recover a multi-dimensional productivity measure using a nonparametric inversion as in Olley & Pakes (1996), Levinsohn & Petrin (2003), and Akerberg et al. (2006). This may allow input decisions to have dynamic consequences. In particular, our current approach rules out that it is costly for a firm to adjust its labor force. There are two principal difficulties. First, the inverse input demand functions $\tilde{h}_L(\cdot)$ and $h_H(\cdot)$ have 2 and 6 arguments, respectively. Estimating these functions nonparametrically is thus quite demanding on the data. Second, one has to prove that input demands are invertible functions of unobserved productivity. At a bare minimum, this requires analyzing a dynamic programming problem in which the firm controls adjustments to labor in addition to investments in physical capital and productivity. Given the difficulties Buettner (2005) encountered in a much simpler dynamic programming problem, this is not an easy task.

General formulation of technological change. As already noted in Section 2, the production function in equation (1) is tailored to answer the research question at hand and restricts the efficiencies of capital and materials in production to change at the same rate and in lockstep with Hicks-neutral technological change. Answering other questions may require different specifications. The most general formulation is a production function that allows for separate efficiencies for the various factors of production. In an extension of our framework one may be able to use the demand for labor to back out the efficiency of labor, the demand for materials to back out the efficiency of materials, and investment to back out the efficiency of capital. Since the policy function for investment in physical capital is the result of a dynamic programming problem, it may have to be inverted nonparametrically, with the difficulties mentioned above.

Employment implications. As mentioned in Sections 1 and 3, different types of technological change have different implications for employment. In particular, labor-augmenting technological displaces labor relative to the other factors of production. Both labor-augmenting and Hicks-neutral technological change also have an output effect in that the usage of all factors of production increases as the marginal cost of production decreases. Our estimates of the elasticity of demand allow us to quantify, at least roughly, the output effect. Hence, our approach can be extended to analyze and decompose the impact of technological change on employment. This may be of interest both to accurately predict the evolution of employment and to design optimal policies in the presence of different types of technological change. Previously Nickell & Kong (1989) have shown that labor-augmenting technological change can cause employment to fall. In contrast, the majority of the existing literature has focused on the differential impact of product vs process innovations and tried to estimate

the total effect of technological change on employment. The emerging consensus is that product innovations always cause employment to rise (see, e.g., Van Reenen 1997, Harrison et al. 2008).

Induced innovation. There is a large literature on induced innovation, dating back at least to Hicks (1966). One way to get at this issue is to add factor prices to the laws of motion for labor-augmenting and Hicks-neutral productivity and ask whether the pace of labor augmentation responds to high wages.

While factor prices play a key role in the debate about induced innovation, they are just one piece of the puzzle: Technological change responds not only to factor prices but also to the cost of the various types of research as Binswanger (1974*b*) argues quite forcefully. There is also a market size or scale effect (Binswanger 1974*b*, Acemoglu 2002). Finally, there are the general equilibrium effects that Acemoglu (2003) and Acemoglu (2007) identifies. Whether our approach can be extended to account for these effects is an important question that we plan to pursue in future research.

Product market vs production function. A firm has two broad motives to invest in R&D, namely to decrease the cost of production and to develop new products. The existing literature largely treats these two motives separately.¹⁶ It may be interesting to model a firm's position in the product market as another unobservable besides the firm's productivity. Our approach may be extended to recover both unobservables and hence to speak to the long-standing distinction between process and product innovations.

Appendix A

Our data come from the ESEE survey. We observe firms for a maximum of ten years between 1990 and 1999. We restrict the sample to firms with at least two years of data on all variables required for estimation. Because of data problems we exclude industry 5 (office and data-processing machines and electrical goods). Our final sample covers 1879 firms in 9 industries. The number of firms with 2, 3, . . . , 10 years of data is 260, 377, 246, 192, 171, 135, 128, 155, and 215, 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)). We finally report the shares of the various industries in the total value added of the manufacturing sector in 1995 (column (4)).

The ESEE survey provides information on the total R&D expenditures of firms. Total 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 consider a firm to be performing R&D if it reports positive expenditures. While total R&D expenditures vary widely across

¹⁶Notable recent exceptions are Aw, Roberts & Xu (2011) and Foster, Haltiwanger & Syverson (2008*b*).

firms, it is quite likely even for small firms that they exceed nonnegligible values relative to firm size. In addition, firms are asked to provide many details about the combination of R&D activities, R&D employment, R&D subsidies, and the number of process and product innovations as well as the patents that result from these activities. Taken together, this supports the notion that the reported expenditures are truly R&D related.

In what follows we define the remaining variables.

- *Investment.* Value of current investments in equipment goods (excluding buildings, land, and financial assets) deflated by the price index of investment. 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 \tilde{K}_{jt} is computed recursively from an initial estimate and the data on current investments in equipment goods \tilde{I}_{jt} . We update the value of the past stock of capital by means of the price index of investment P_{It} as $\tilde{K}_{jt} = (1 - \delta) \frac{P_{It}}{P_{It-1}} \tilde{K}_{jt-1} + \tilde{I}_{jt-1}$, 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 $K_{jt} = \frac{\tilde{K}_{jt}}{P_{It}}$.
- *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 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.
- *Price of investment.* Equipment goods component of the index of industry prices computed and published by the Spanish Ministry of Industry.
- *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.
- *Market dynamism.* Firms are asked to assess the current and future situation (slump, stability, or expansion) of up to 5 separate markets in which they operate. The market dynamism index is computed as a weighted average of the responses.

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. (2002), 'Directed technical change', *Review of Economic Studies* **69**, 781–809.
- Acemoglu, D. (2003), 'Labor- and capital-augmenting technical change', *Journal of the European Economic Association* **1**(1), 1–37.
- Acemoglu, D. (2007), 'Equilibrium bias of technology', *Econometrica* **75**(5), 1371–1409.
- Acemoglu, D. & Finkelstein, A. (2008), 'Input and technology choices in regulated industries: Evidence from the health care sector', *Journal of Political Economy* **forthcoming**.
- Akerberg, 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.
- Akerberg, D., Caves, K. & Frazer, G. (2006), Structural identification of production functions, Working paper, UCLA, Los Angeles.
- Antràs, P. (2004), 'Is the U.S. aggregate production function Cobb-Douglas? New estimates of the elasticity of substitution', *Contributions to Macroeconomics* **4**(1).
- Arrow, K., Chenery, H., Minhas, B. & Solow, R. (1961), 'Capital-labor substitution and economic efficiency', *Review of Economics and Statistics* **43**(3), 225–250.
- Autor, D., Katz, L. & Krueger, A. (1998), 'Computing inequality: Have computers changed the labor market?', *Quarterly Journal of Economics* **113**(4), 1169–1213.
- Autor, D., Levy, F. & Murnane, R. (2003), 'The skill content of recent technological change: An empirical exploration', *Quarterly Journal of Economics* **118**(4), 1279–1333.
- Aw, B., Roberts, M. & Xu, D. (2011), 'R&D investment, exporting, and productivity dynamics', *American Economic Review* **101**(4), 1312–1344.
- Bentolila, S. & Jimeno, J. (2006), Spanish unemployment: The end of the wild ride?, in M. Werding, ed., 'Structural Unemployment in Western Europe: Reasons and Remedies', MIT Press, Cambridge, pp. 317–344.
- Bentolila, S. & Saint-Paul, G. (2004), 'Explaining movements in the labor share', *Contributions to Macroeconomics* **3**(1).
- Berndt, E. (1976), 'Reconciling alternative estimates of the elasticity of substitution', *Review of Economics and Statistics* **58**(1), 59–68.
- Bessen, J. (2008), Accounting for productivity growth when technical change is biased, Working paper, Boston University, Boston.
- Binswanger, H. (1974a), 'The measurement of technical change biases with many factors of production', *American Economic Review* **64**(6), 964–976.

- Binswanger, H. (1974*b*), ‘The microeconomic approach to induced innovation’, *Economic Journal* **84**(336), 940–958.
- 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: Macroeconomics* **1997**(2), 89–158.
- Bloom, N., Sadun, R. & Van Reenen, J. (2007), Americans do I.T. better: US multinationals and the productivity miracle, Working paper no. 13085, NBER, Cambridge.
- 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.
- Burmeister, E. & Dobell, A. (1969), ‘Disembodied technological change with several factors’, *Journal of Economic Theory* **1**, 1–8.
- Caballero, R. & Hammour, M. (1998), ‘Jobless growth: appropriability, factor substitution, and unemployment’, *Carnegie-Rochester Conference Series on Public Policy* **48**, 51–94.
- 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.
- Doraszelski, U. & Jaumandreu, J. (2013), ‘R&D and productivity: Estimating endogenous productivity’, *Review of Economic Studies* **forthcoming**.
- 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*a*), ‘Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?’, *American Economic Review* **98**(1), 394–425.
- Foster, L., Haltiwanger, J. & Syverson, C. (2008*b*), The slow growth of new plants: Learning about demand, Working paper, University of Chicago, Chicago.
- Gonzalez, X., Jaumandreu, J. & Pazo, C. (2005), ‘Barriers to innovation and subsidy effectiveness’, *Rand Journal of Economics* **36**(4), 930–950.

- Gordon, R. (1990), *The measurement of durable goods prices*, National Bureau of Economic Research Monograph Series, University of Chicago Press, Chicago.
- 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.
- 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.
- Harrison, R., Jaumandreu, J., Mairesse, J. & Peters, B. (2008), Does innovation stimulate employment? A firm-level analysis using comparable micro data on four European countries, Working paper, University College, London.
- Hicks, J. (1966), *The theory of wages*, 2nd edn, Macmillan, London.
- Jin, H. & Jorgenson, D. (2008), Econometric modeling of technical change, Working paper, Harvard University, Cambridge.
- Jones, C. (2005), 'The shape of production functions and the direction of technical change', *Quarterly Journal of Economics* **120**(2), 517–549.
- 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.
- Kendrick, J. & Sato, R. (1963), 'Factor prices, productivity, and economic growth', *American Economic Review* **53**(3), 974–1003.
- Krueger, A. (1999), 'Measuring labor's share', *American Economic Review* **90**(2), 45–51.
- Krusell, P., Ohanian, L., Rios-Rull, V. & Violante, G. (2000), 'Capital-skill complementarity and inequality: A macroeconomic analysis', *Econometrica* **68**(5), 1029–1053.
- 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 economics*, 8th edn, Macmillan, London.
- McAdam, P. & William, A. (2013), ‘Medium run redux: Technical change, factor shares and frictions in the euro area’, *Macroeconomic Dynamics* **forthcoming**.
- McKinnon, R. (1962), ‘Wages, capital costs, and employment in manufacturing: A model applied to 1947–58 U.S. data’, *Econometrica* **30**(2), 501–521.
- Mokyr, J. (1990), *The levers of riches: Technological creativity and economic progress*, Oxford University Press, Oxford.
- Newey, W., Powell, J. & Vella, F. (1999), ‘Nonparametric estimation of triangular simultaneous equations models’, *Econometrica* **67**(3), 565–603.
- Nickell, S. & Kong, P. (1989), Technical progress and jobs, Working paper no. 366, Centre for Labour Economics, London.
- OECD (2007), ‘The policy mix for research, development and innovation in Spain: Key issues and policy recommendations’, Directorate for Science, Technology and Industry, Paris.
- Olley, S. & Pakes, A. (1996), ‘The dynamics of productivity in the telecommunications industry’, *Econometrica* **64**(6), 1263–1297.
- 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.
- Sato, R. (1970), ‘The estimation of biased technical progress and the production function’, *International Economic Review* **11**(2), 179–208.
- Sato, R. & Beckmann, M. (1968), ‘Neutral inventions and production functions’, *Review of Economic Studies* **35**(1), 57–66.
- 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.

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

Table 1: Descriptive statistics.

Industry	Obs. ^a	Firms ^a	Rates of growth ^a						Labor share ^a		
			Output (s. d.)	Labor (s. d.)	Capital (s. d.)	Materials (s. d.)	Price (s. d.)	$\frac{M}{L}$ (s. d.)	$\frac{PM}{W}$ (s. d.)	Mean (s. d.)	Max-min (s. d.)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
1. Metals and metal products	1235	289	0.050 (0.238)	0.010 (0.183)	0.086 (0.278)	0.038 (0.346)	0.012 (0.055)	0.028 (0.328)	-0.011 (0.175)	0.336 (0.154)	0.100
2. Non-metallic minerals	670	140	0.039 (0.209)	0.002 (0.152)	0.065 (0.259)	0.040 (0.304)	0.011 (0.057)	0.039 (0.287)	-0.015 (0.154)	0.350 (0.136)	0.067
3. Chemical products	1218	275	0.068 (0.196)	0.007 (0.146)	0.093 (0.238)	0.054 (0.254)	0.007 (0.061)	0.047 (0.246)	-0.024 (0.164)	0.268 (0.120)	0.070
4. Agric. and ind. machinery	576	132	0.059 (0.275)	0.010 (0.170)	0.078 (0.247)	0.046 (0.371)	0.013 (0.032)	0.035 (0.357)	-0.014 (0.163)	0.369 (0.162)	0.098
6. Transport equipment	637	148	0.087 (0.0354)	0.011 (0.207)	0.114 (0.255)	0.087 (0.431)	0.007 (0.037)	0.075 (0.366)	-0.029 (0.191)	0.315 (0.145)	0.121
7. Food, drink and tobacco	1408	304	0.025 (0.224)	-0.003 (0.186)	0.094 (0.271)	0.019 (0.305)	0.022 (0.065)	0.021 (0.323)	-0.027 (0.197)	0.256 (0.155)	0.056
8. Textile, leather and shoes	1278	293	0.020 (0.233)	-0.007 (0.192)	0.059 (0.235)	0.012 (0.356)	0.016 (0.040)	0.019 (0.360)	-0.023 (0.194)	0.338 (0.200)	0.061
9. Timber and furniture	569	138	0.038 (0.278)	0.014 (0.210)	0.077 (0.257)	0.029 (0.379)	0.020 (0.035)	0.015 (0.376)	-0.021 (0.202)	0.314 (0.142)	0.089
10. Paper and printing products	665	160	0.035 (0.183)	-0.000 (0.140)	0.099 (0.303)	0.026 (0.265)	0.019 (0.089)	0.026 (0.262)	-0.016 (0.161)	0.345 (0.152)	0.079

^aComputed for 1991 to 1999 excluding the first observation for each firm.

Table 2: Descriptive statistics.

Industry	With R&D ^a			
	Obs. (%)	Stable (%)	Occas. (%)	R&D intensity (s. d.)
	(1)	(2)	(3)	(4)
1. Metals and metal products	420 (34.0)	63 (21.8)	72 (24.9)	0.0126 (0.0144)
2. Non-metallic minerals	226 (33.7)	22 (15.7)	44 (31.4)	0.0112 (0.0206)
3. Chemical products	672 (55.2)	124 (45.1)	55 (20.0)	0.0268 (0.0353)
4. Agric. and ind. machinery	322 (55.9)	52 (39.4)	35 (26.5)	0.0219 (0.0275)
6. Transport equipment	361 (56.7)	62 (41.9)	35 (23.6)	0.0224 (0.0345)
7. Food, drink and tobacco	386 (27.4)	56 (18.4)	64 (21.1)	0.0071 (0.0281)
8. Textile, leather and shoes	378 (29.6)	39 (13.3)	66 (22.5)	0.0152 (0.0219)
9. Timber and furniture	66 (12.6)	7 (5.1)	18 (13.8)	0.0138 (0.0326)
10. Paper and printing products	113 (17.0)	21 (13.1)	25 (13.8)	0.0143 (0.0250)

Table 3: Estimates and labor-augmenting productivity growth.

Industry	OLS		GMM		Overidentif.		Productivity growth ^a	
	σ (std. err.) (1)	trend (std. err.) (2)	σ (std. err.) (3)	χ^2 (df) (4)	p val. (5)	Total (6)	R&D obs. (% contrib.) (7)	No R&D obs. (% contrib.) (8)
1. Metals and metal products	1.013 (0.030)	0.064 (0.012)	0.479 (0.065)	19.31 (16)	0.253	0.088	0.116 (78.9)	0.040 (21.1)
2. Non-metallic minerals	1.263 (0.069)	0.038 (0.031)	0.524 (0.070)	18.59 (16)	0.291	0.096	0.119 (54.1)	0.067 (45.9)
3. Chemical products	1.017 (0.035)	0.050 (0.018)	0.500 (0.057)	25.27 (16)	0.065	0.078	0.078 (71.4)	0.077 (28.6)
4. Agric. and ind. machinery	1.236 (0.053)	0.055 (0.022)	0.376 (0.118)	21.34 (16)	0.166	0.037	0.058 (119.5)	-0.003 (-19.5)
6. Transport equipment	0.876 (0.073)	0.082 (0.023)	0.570 (0.080)	24.10 (16)	0.087	0.114	0.145 (97.7)	0.047 (2.3)
7. Food, drink and tobacco	1.198 (0.020)	0.037 (0.009)	0.631 (0.063)	24.28 (16)	0.083	0.059	0.115 (64.2)	0.03 (35.8)
8. Textile, leather and shoes	1.347 (0.011)	0.042 (0.007)	0.517 (0.090)	31.40 (16)	0.012	-0.012	-0.024 (86.6)	0.000 (13.4)
9. Timber and furniture	0.644 (0.051)	0.060 (0.016)	0.633 (0.080)	15.91 (16)	0.459	-0.007	0.049 (-25.0)	-0.017 (125.0)
10. Paper and printing products	1.055 (0.032)	0.046 (0.023)	0.366 (0.064)	25.28 (16)	0.065	0.068	0.055 (21.1)	0.069 (78.9)

^aWe trim 2.5% of observations at each tail of the distribution.

Table A1: Industry definitions, equivalent classifications, and shares.

Industry	Classifications		Share of value added	
	ESSE (1)	National Accounts (2)		ISIC (3)
1. Ferrous and non-ferrous metals and metal products	1+4	DJ	D 27+28	12.6
2. Non-metallic minerals	2	DI	D 26	7.8
3. Chemical products	3+17	DG-DH	D 24+25	13.7
4. Agricultural and industrial machinery	5	DK	D 29	5.9
6. Transport equipment	8+9	DM	D 34+35	11.0
7. Food, drink and tobacco	10+11+12	DA	D 15+16	16.5
8. Textile, leather and shoes	13+14	DB-DC	D 17+18+19	7.9
9. Timber and furniture	15	DD-DN	D 20+30	6.3
10. Paper and printing products	16	DE	D 21+22	8.2
Total				90.0