University of Pennsylvania
ScholarlyCommons

Marketing Papers

Wharton Faculty Research

5-2018

Measuring the Bias of Technological Change

Ulrich Doraszelski
University of Pennsylvania

Jordi Jaumandreu

Follow this and additional works at: https://repository.upenn.edu/marketing_papers

Part of the [Business Administration, Management, and Operations Commons](#), [Marketing Commons](#), and the [Technology and Innovation Commons](#)

Recommended Citation

Doraszelski, U., & Jaumandreu, J. (2018). Measuring the Bias of Technological Change. *Journal of Political Economy*, 126 (5), <http://dx.doi.org/10.1086/697204>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/marketing_papers/424
For more information, please contact repository@pobox.upenn.edu.

Measuring the Bias of Technological Change

Abstract

Technological change can increase the productivity of the various factors of production in equal terms, or it can be biased toward a specific factor. We directly address the bias of technological change by measuring, at the level of the individual firm, how much of it is labor augmenting and how much is factor neutral. To do so, we develop a framework for estimating production functions when productivity is multidimensional. Using panel data from Spain, we find that technological change is biased, with both its labor-augmenting and its factor-neutral components causing output to grow by about 1.5 percent per year.

Disciplines

Business | Business Administration, Management, and Operations | Marketing | Technology and Innovation

Measuring the Bias of Technological Change

Ulrich Doraszelski

University of Pennsylvania

Jordi Jaumandreu

Boston University

Technological change can increase the productivity of the various factors of production in equal terms, or it can be biased toward a specific factor. We directly assess the bias of technological change by measuring, at the level of the individual firm, how much of it is labor augmenting and how much is factor neutral. To do so, we develop a framework for estimating production functions when productivity is multidimensional. Using panel data from Spain, we find that technological change is biased, with both its labor-augmenting and its factor-neutral components causing output to grow by about 1.5 percent per year.

I. 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 toward a specific factor. Whether technological change fa-

We thank the editor Ali Hortaçsu, two anonymous referees, Pol Antràs, Paul David, Matthias Doepke, Michaela Draganska, José Carlos Fariñas, Ivan Fernandez-Val, Paul Grieco, Chad Jones, Dale Jorgenson, Larry Katz, Pete Klenow, Jacques Mairesse, Ariel Pakes, Amil Petrin, Zhongjun Qu, Devesh Raval, Juan Sanchis, Matthias Schündeln, Johannes Van Biesebroeck, John Van Reenen, and Hongsong Zhang for helpful comments and discussions and Tianyi Chen, Sterling Horne, Mosha Huang, Thomas O'Malley, and Dan Sacks for research assistance. We gratefully acknowledge financial support from the National Science Foundation under grants 0924380 and 0924282; Jaumandreu further acknowledges partial funding from Fundacion Rafael del Pino. Codes and instructions for obtaining the data are provided as supplementary material online.

Electronically published May 7, 2018

[*Journal of Political Economy*, 2018, vol. 126, no. 3]

© 2018 by The University of Chicago. All rights reserved. 0022-3808/2018/12603-0004\$10.00

vors some factors of production over others is central to economics. Yet, the empirical evidence is relatively sparse.

The literature on economic growth rests on the assumption that technological change increases the productivity of labor vis-à-vis the other factors of production. It is well known that for a neoclassical growth model to exhibit steady-state growth, either the production function must be Cobb-Douglas or technological change must be labor augmenting (Uzawa 1961), and many endogenous growth models point to human capital accumulation as a source of productivity increases (Lucas 1988; Romer 1990). A number of recent papers provide microfoundations for the literature on economic growth by theoretically establishing that profit-maximizing incentives can ensure that technological change is, at least in the long run, purely labor augmenting (Acemoglu 2003; Jones 2005). Whether this is indeed the case is, however, an empirical question that remains to be answered.

One reason for the scarcity of empirical assessments of the bias of technological change may be a lack of suitable data. Following early work by Brown and de Cani (1963) and David and van de Klundert (1965), economists have estimated aggregate production or cost functions that proxy for labor-augmenting technological change with a time trend (Lucas 1969; Binswanger 1974; Kalt 1978; Cain and Paterson 1986; Antràs 2004; Klump, McAdam, and Willman 2007; Jin and Jorgenson 2010).¹ While this line of research has produced some evidence of labor-augmenting technological change, the staggering amount of heterogeneity across firms in combination with simultaneously occurring entry and exit (Dunne, Roberts, and Samuelson 1988; Davis and Haltiwanger 1992) may make it difficult to interpret a time trend as a meaningful average economy- or sectorwide measure of technological change. Furthermore, this line of research pays scant attention to the fundamental endogeneity problem in production function estimation. This problem arises because a firm's decisions depend on its productivity, and productivity is not observed by the econometrician and may severely bias the estimates (Marschak and Andrews 1944).

While traditionally using more disaggregated data, the productivity and industrial organization literatures assume that technological change is factor neutral. Hicks-neutral technological change underlies, either explicitly or implicitly, most of the standard techniques for measuring productivity, ranging from the classic growth decompositions of Solow (1957) and Hall (1988) to the recent structural estimators for production functions (Olley and Pakes 1996; Levinsohn and Petrin 2003; Doraszelski and Jaumandreu 2013; Gandhi, Navarro, and Rivers 2013; Akerberg, Caves, and Frazer 2015).

¹ A much larger literature has estimated the elasticity of substitution using either aggregated or disaggregated data while maintaining the assumption of factor-neutral technological change; see Hamermesh (1993) for a survey.

In this paper, we extend the productivity and industrial organization literatures and develop a framework for estimating production functions when productivity is multidimensional and has a labor-augmenting and a Hicks-neutral component. We use firm-level panel data that are now widely available to directly assess the bias of technological change by measuring, at the level of the individual firm, how much of technological change is labor augmenting and how much of it is Hicks neutral.

To tackle the endogeneity problem in production function estimation, we build on the insight of Olley and Pakes (1996) that if the decisions that a firm makes can be used to infer its productivity, then productivity can be controlled for in the estimation. We infer the firm's multidimensional productivity from its input usage, in particular its labor and materials decisions. The key to identifying the bias of technological change is that Hicks-neutral technological change scales input usage but, in contrast to labor-augmenting technological change, does not change the mix of inputs that a firm uses. A change in the input mix therefore contains information about the bias of technological change, provided we control for the relative prices of the various inputs and other factors that may change the input mix.

We apply our framework to examine the speed and direction of technological change in the Spanish manufacturing sector in the 1990s and early 2000s. Spain is an attractive setting for two reasons. First, Spain became fully integrated into the European Union between the end of the 1980s and the beginning of the 1990s. Any trends in technological change that our analysis uncovers for Spain may thus be viewed as broadly representative for other continental European economies. Second, Spain inherited an industrial structure with few high-tech industries and mostly small and medium-sized firms. R&D is widely seen as lacking. Yet, Spain grew rapidly during the 1990s, and R&D became increasingly important (European Commission 2001). The accompanying changes in industrial structure are a useful source of variation for analyzing the role of R&D in stimulating different types of technological change.

The particular data set we use has several advantages. The broad coverage 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.²

The Spanish manufacturing sector also poses several challenges for identifying the bias of technological change from a change in the mix of

² There are other firm-level data sets such as the Colombian Annual Manufacturers Survey (Eslava et al. 2004), the Prowess data collected by the Centre for Monitoring the Indian Economy (De Loecker et al. 2016), and the Longitudinal Business Database at the US Census Bureau, which contain separate information on prices and quantities, at least for a subset of industries (Roberts and Supina 1996; Foster, Haltiwanger, and Syverson 2008, 2016).

inputs that a firm uses. First, outsourcing directly changes the input mix as the firm procures customized parts and pieces from its suppliers rather than makes them in-house from scratch. Second, the Spanish labor market manifestly distinguishes between permanent and temporary labor. We further contribute to the literature following Olley and Pakes (1996) by accounting for outsourcing and the dual nature of the labor market.

Our estimates provide clear evidence that technological change is biased. *Ceteris paribus*, labor-augmenting technological change causes output to grow, on average, in the vicinity of 1.5 percent per year. While there is a shift from unskilled to skilled workers in our data, this skill upgrading explains some but not all of the growth of labor-augmenting productivity. In many industries, labor-augmenting productivity grows because workers with a given set of skills become more productive over time.

At the same time, our estimates show that Hicks-neutral technological change plays an equally important role. In addition to labor-augmenting technological change, Hicks-neutral technological change causes output to grow, on average, in the vicinity of 1.5 percent per year.

Behind these averages lies a substantial amount of heterogeneity across industries and firms. Our estimates point to substantial and persistent differences in labor-augmenting and Hicks-neutral productivity across firms, in line with the “stylized facts” about productivity in Bartelsman and Doms (2000) and Syverson (2011). Beyond these facts, we show that, at the level of the individual firm, the levels of labor-augmenting and Hicks-neutral productivity are positively correlated, as are their rates of growth.

Our estimates further indicate that firms’ R&D activities play a key role in determining the differences in the components of productivity across firms and their evolution over time. Interestingly, labor-augmenting productivity is slightly more closely tied to firms’ R&D activities than is Hicks-neutral productivity. Through the lens of the literature on induced innovation and directed technical change (Hicks 1932; Acemoglu 2002), this may be viewed as supporting the argument that firms direct their R&D activities to conserve on labor.

Biased technological change has consequences far beyond the growth of output. To illustrate, we use our estimates to show that biased technological change is the primary driver of the decline of the aggregate share of labor in the Spanish manufacturing sector over our sample period. Similar declines have been observed in many advanced economies in past decades and have attracted considerable attention in the macroeconomics literature (Blanchard 1997; Bentolila and Saint-Paul 2004; McAdam and Willman 2013; Karabarbounis and Neiman 2014; Oberfield and Raval 2014).

The starting point of this paper is the recent structural estimators for production functions. We differ from much of the previous literature by exploiting the parameter restrictions between the production and input

demand functions, as in Doraszelski and Jaumandreu (2013). This allows us to parametrically invert from observed input usage to unobserved productivity and eases the demands on the data compared to the non-parametric inversion in Olley and Pakes (1996), Levinsohn and Petrin (2003), and Akerberg et al. (2015), especially if the input demand functions are high-dimensional.³

Our paper is related to that of Van Biesebroeck (2003). Using plant-level panel data for the US automobile industry, he estimates a plant's Hicks-neutral productivity as a fixed effect and parametrically inverts from the plant's input usage to its capital-biased (also called labor-saving) productivity. Our approach is more general in that we allow all components of productivity to evolve over time and in response to firms' R&D activities.

Our paper is further related to Grieco, Li, and Zhang (2016) and subsequent work in progress by Zhang (2015). Because their data contain the materials bill rather than its split into price and quantity, Grieco et al. (2016) build on Doraszelski and Jaumandreu (2013) and parametrically invert from a firm's input usage to its Hicks-neutral productivity and the price of materials that the firm faces.

Finally, our paper touches—although more tangentially—on the literature on skill bias that studies the differential impact of technological change, especially in the form of computerization, on the various types of labor (see Card and DiNardo [2002] and Violante [2008] and the references therein). While we focus on labor versus the other factors of production, the techniques we develop may be adapted to investigate the skill bias of technological change, although our particular data set is not ideal for this purpose.

The remainder of this paper is organized as follows: Section II explains how we identify the bias of technological change and previews our empirical strategy. Section III describes the data and some patterns in the data that inform the subsequent analysis. Section IV sets out a dynamic model of the firm. Section V develops an estimator for production functions when productivity is multidimensional. Sections VI–IX present our main results on labor-augmenting and Hicks-neutral technological change. Section X explores whether capital-augmenting technological change plays a role in our data in addition to labor-augmenting and Hicks-neutral technological change. Section XI presents conclusions. The supplementary appendix contains additional results and technical details.

Throughout the paper, we adopt the convention that uppercase letters denote levels and lowercase letters denote logs. Unless noted otherwise, we refer to output and the various factors of production in terms of

³ See Doraszelski and Jaumandreu (2013) for details on the pros and cons of the parametric inversion.

quantity and not in terms of value. In particular, we refer to the value of labor as the wage bill and to the value of materials as the materials bill.

II. Labor-Augmenting and Hicks-Neutral Productivity

We first show how to separately recover a firm's labor-augmenting and Hicks-neutral productivity from its labor and materials decisions. Then we show that the constant elasticity of substitution (CES) production function that we use in our application approximates, to a first order, the relationship between the input mix and labor-augmenting productivity that arises in a wider class of production functions. To facilitate the exposition, we proceed in a highly simplified setting. Our application extends the setting to accommodate the institutional realities of the Spanish manufacturing sector.

Consider a firm with the production function

$$Y_{jt} = F(K_{jt}, \exp(\omega_{Ljt})L_{jt}, M_{jt}) \exp(\omega_{Hjt}) \exp(e_{jt}), \quad (1)$$

where Y_{jt} is the output of firm j in period t , K_{jt} is capital, L_{jt} is labor, and M_{jt} is materials. The labor-augmenting productivity of firm j in period t is ω_{Ljt} and its Hicks-neutral productivity is ω_{Hjt} . Finally, e_{jt} is a random shock.

To relate the input ratio M_{jt}/L_{jt} to labor-augmenting productivity ω_{Ljt} , we assume that $(\exp(\omega_{Ljt})L_{jt}, M_{jt})$ is separable from K_{jt} in that the function $F(\cdot)$ in equation (1) is composed of the functions $G(\cdot)$ and $H(\cdot)$ as

$$F(K_{jt}, \exp(\omega_{Ljt})L_{jt}, M_{jt}) = G(K_{jt}, H(\exp(\omega_{Ljt})L_{jt}, M_{jt})), \quad (2)$$

where $H(\exp(\omega_{Ljt})L_{jt}, M_{jt})$ is homogeneous of arbitrary degree.⁴ Without loss of generality, we set the degree of homogeneity to one. Throughout we maintain that all functions are differentiable as needed. As in Levinsohn and Petrin (2003), we finally assume that labor and materials are static (or "variable") inputs that are chosen each period to maximize short-run profits and that the firm is a price taker in input markets, where it faces W_{jt} and P_{Mjt} as prices of labor and materials, respectively.

The input ratio M_{jt}/L_{jt} is therefore the solution to the ratio of the first-order conditions for labor and materials

$$\frac{\frac{\partial H(\exp(\omega_{Ljt})L_{jt}, M_{jt})}{\partial L_{jt}} \exp(\omega_{Ljt})}{\frac{\partial H(\exp(\omega_{Ljt})L_{jt}, M_{jt})}{\partial M_{jt}}} = \frac{\frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial L_{jt}} \exp(\omega_{Ljt})}{\frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial M_{jt}}} = \frac{W_{jt}}{P_{Mjt}}, \quad (3)$$

⁴ Equation (2) immediately implies that $F(K_{jt}, \exp(\omega_{Ljt})L_{jt}, M_{jt})$ is weakly separable in the partition $(K_{jt}, (\exp(\omega_{Ljt})L_{jt}, M_{jt}))$ (Chambers 1988, eq. 1.26). It is equivalent to $F(K_{jt}, \exp(\omega_{Ljt})L_{jt}, M_{jt})$ being weakly separable under some additional monotonicity and quasi-concavity assumptions (Goldman and Uzawa 1964).

where the first equality uses that $H(\exp(\omega_{Ljt})L_{jt}, M_{jt})$ is homogeneous of degree one and, recall, uppercase letters denote levels and lowercase letters denote logs.

Equation (3) implies that the input ratio M_{jt}/L_{jt} depends on the price ratio P_{Mjt}/W_{jt} and labor-augmenting productivity ω_{Ljt} . Importantly, the input ratio M_{jt}/L_{jt} does not depend on Hicks-neutral productivity ω_{Hjt} . This formalizes that the mix of inputs that a firm uses is related to—and therefore contains information about—its labor-augmenting productivity but is unrelated to its Hicks-neutral productivity. Intuitively, the labor and materials decisions hinge on the marginal products of labor and materials. Because the marginal products are proportional to Hicks-neutral productivity, materials per unit of labor as determined by the ratio of the first-order conditions in equation (3) are unrelated to Hicks-neutral productivity. In this sense, separating labor-augmenting from Hicks-neutral productivity does not rely on functional form beyond the separability assumption in equation (2).⁵

The following proposition further characterizes the log of the input ratio $m_{jt} - l_{jt}$.

PROPOSITION 1. The input ratio $m_{jt} - l_{jt}$ has the first-order Taylor series

$$\begin{aligned} \gamma_L^0 - \sigma(\exp(\omega_{Ljt}^0 - (m_{jt}^0 - l_{jt}^0)))(p_{Mjt} - w_{jt}) \\ + [1 - \sigma(\exp(\omega_{Ljt}^0 - (m_{jt}^0 - l_{jt}^0)))]\omega_{Ljt} \end{aligned} \quad (4)$$

around a point $(m_{jt}^0 - l_{jt}^0, p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)$ satisfying equation (3), where γ_L^0 is a constant and $\sigma(\exp(\omega_{Ljt}^0 - (m_{jt}^0 - l_{jt}^0)))$ is the elasticity of substitution between materials and labor in the production function in equation (1).

The proof can be found in appendix A.

Our application uses a CES production function

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

where ν and σ are the elasticity of scale and substitution, respectively, and β_K and β_M are the so-called distributional parameters.⁶ Depending on the

⁵ One can forgo the separability assumption by relying more on functional form. Our empirical strategy generalizes to, e.g., a translog production function that does not satisfy eq. (2).

⁶ We implicitly set the constant of proportionality β_0 to one because it cannot be separated from an additive constant in Hicks-neutral productivity ω_{Hjt} . We estimate them jointly and carefully ensure that the reported results depend only on their sum. We similarly normalize the distributional parameter β_L . Technological change can therefore equivalently be thought of as letting these parameters of the production function vary by firm and time. The nascent literature on heterogeneous production functions (Balat, Brambilla, and Sasaki 2015; Kasahara, Schrimpf, and Suzuki 2015; Fox et al. 2016) explores to what extent it is possible to let all parameters of the production function vary by firm and time.

elasticity of substitution, the CES production function encompasses the special cases of a Leontieff ($\sigma \rightarrow 0$), Cobb-Douglas ($\sigma = 1$), and linear ($\sigma \rightarrow \infty$) production function.

The ratio of the first-order conditions in equation (3) implies

$$m_{jt} - l_{jt} = \sigma \ln \beta_M - \sigma(p_{Mjt} - w_{jt}) + (1 - \sigma)\omega_{Ljt}. \quad (5)$$

Comparing equations (4) and (5) shows that the CES production function approximates, to a first order, the input ratio $m_{jt} - l_{jt}$ arising from an arbitrary production function satisfying equation (2). This gives a sense of robustness to the CES production function.⁷

Our empirical strategy uses equation (5) to recover a firm's labor-augmenting productivity from its input mix. In doing so, we must control for other factors besides the relative prices of the various inputs that may change the input mix, in particular outsourcing and the dual nature of the Spanish labor market. With labor-augmenting productivity in hand, we use the first-order condition for labor to recover Hicks-neutral productivity. The remainder of our empirical strategy follows along the lines of Olley and Pakes (1996), Levinsohn and Petrin (2003), Doraszelski and Jaumandreu (2013), and Akerberg et al. (2015) by combining the inferred productivities with their laws of motion to set up estimation equations.

Equation (5) has a long tradition in the literature, although it is used in a very different way from ours. With skilled and unskilled workers in place of materials and labor, equation (5) is at the heart of the literature on skill bias (Card and DiNardo 2002; Violante 2008). With capital in place of materials, equation (5) serves to estimate the elasticity of substitution σ in an aggregate value-added production function (see Antràs 2004). More recently, Raval (2013) uses a variant of equation (5) obtained from a value-added production function with capital- and labor-augmenting productivity to estimate σ from firm-level panel data.

Equation (5) is typically estimated by ordinary least squares (OLS). The problem is that labor-augmenting productivity, which is not observed by the econometrician, is correlated over time and also with the wage. We intuitively expect the wage to be higher when labor is more productive, even if it adjusts slowly with some lag. This positive correlation induces an upward bias in the estimate of the elasticity of substitution. This is a variant of the endogeneity problem in production function estimation. Because we use equation (5) to recover labor-augmenting productivity rather than directly estimate it, we are able to tackle the endogeneity problem with a

⁷ It also suggests that our "nonparametric" estimates of labor-augmenting technological change can be fed into a growth decomposition along the lines of Solow (1957) and Hall (1988) to obtain a "nonparametric" estimate of Hicks-neutral technological change. We leave this to future research.

combination of assumptions on the timing of decisions and the evolution of the components of productivity.

III. Data

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, and spans 1990–2006. At the beginning of the survey, 5 percent of firms with up to 200 workers were sampled randomly by industry and size strata. All firms with more than 200 workers were asked to participate in the survey, and 70 percent of them complied. Some firms vanish from the sample as a result of either exit (shutdown by death or abandonment of activity) or attrition. These reasons can be distinguished in the data, and attrition remained within acceptable limits. To preserve representativeness, newly created firms were added to the sample every year. We provide details on industry and variable definitions in appendix B.

Our sample covers a total of 2,375 firms in 10 industries when restricted to firms with at least 3 years of data. Columns 1 and 2 of table 1 show the number of observations and firms by industry. Sample sizes are moderate. Newly created firms are a large fraction of the total number of firms, ranging from 26 percent to 50 percent in the different industries. There is a much smaller fraction of exiting firms, ranging from 6 percent to 15 percent and above in a few industries.

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

An attractive feature of our data is that they contain firm-specific, Paasche-type price indices for output and materials. We note that the variation in these price indices is partly due to changes over time in the bundles of goods that make up output and, respectively, materials (see Bernard, Redding, and Schott [2010] and Goldberg et al. [2010] for evidence on product turnover) and that these changes may be related to a firm's productivity. The growth of prices, averaged from the growth of prices as reported individually by each firm, is moderate. The growth of the price of output in column 7 ranges from 0.8 percent to 2.1 percent. The growth of the wage ranges from 4.3 percent to 5.4 percent and the growth of the price of materials ranges from 2.8 percent to 4.1 percent.

Biased technological change.—The evolution of the relative quantities and prices of the various factors of production already hints at an important role for labor-augmenting technological change. As columns 8 and

TABLE I
DESCRIPTIVE STATISTICS

INDUSTRY	OBSERVATIONS ^a (1)	FIRMS ^a (2)	RATES OF GROWTH ^b								
			Output (3)	Capital (4)	Labor (5)	Materials (6)	Price (7)	M/L (8)	P _M /W (9)	M/K (10)	P _M /P _K (11)
1. Metals and metal products	2,365	313	.045 (.235)	.051 (.192)	.008 (.161)	.030 (.327)	.017 (.052)	.022 (.316)	-.008 (.176)	-.021 (.373)	.049 (.099)
2. Nonmetallic minerals	1,270	163	.046 (.228)	.057 (.212)	.010 (.177)	.041 (.285)	.012 (.058)	.031 (.272)	-.012 (.147)	-.016 (.333)	.043 (.104)
3. Chemical products	2,168	299	.060 (.228)	.062 (.182)	.015 (.170)	.044 (.274)	.008 (.055)	.029 (.250)	-.015 (.153)	-.019 (.313)	.044 (.141)
4. Agricultural and industrial machinery	1,411	178	.031 (.252)	.040 (.190)	-.003 (.169)	.018 (.347)	.015 (.026)	.022 (.335)	-.015 (.155)	-.021 (.390)	.041 (.099)
5. Electrical goods	1,505	209	.059 (.268)	.041 (.173)	.010 (.205)	.048 (.359)	.008 (.046)	.038 (.344)	-.021 (.174)	.007 (.394)	.045 (.095)
6. Transport equipment	1,206	161	.060 (.287)	.043 (.164)	.004 (.201)	.051 (.375)	.008 (.031)	.047 (.343)	-.019 (.171)	.008 (.396)	.033 (.093)
7. Food, drink, and tobacco	2,455	327	.023 (.206)	.047 (.177)	.003 (.169)	.012 (.286)	.021 (.054)	.009 (.295)	-.018 (.176)	-.035 (.328)	.049 (.116)
8. Textile, leather, and shoes	2,368	335	.004 (.229)	.031 (.189)	-.015 (.180)	-.009 (.348)	.015 (.042)	.006 (.355)	-.021 (.183)	-.040 (.385)	.040 (.099)
9. Timber and furniture	1,445	207	.025 (.225)	.045 (.168)	.013 (.184)	.014 (.335)	.020 (.031)	.001 (.329)	-.019 (.171)	-.031 (.371)	.067 (.123)
10. Paper and printing products	1,414	183	.031 (.187)	.052 (.221)	-.001 (.149)	.013 (.252)	.017 (.074)	.014 (.247)	-.017 (.159)	-.039 (.326)	.046 (.122)

NOTE.—Numbers in parentheses are standard deviations.

^a Including $S_{ijt} = L_{ijt} = 0$.

^b Computed for 1991–2006.

9 of table 1 show, with the exception of industries 7, 8, and 9, the input ratio M_{jt}/L_{jt} increases much more than the price ratio P_{Mjt}/W_{jt} decreases. One possible explanation is that the elasticity of substitution between materials and labor exceeds one. To see this, recall that the elasticity of substitution (Chambers 1988, eq. 1.12) is

$$\frac{d \ln \left(\frac{M_{jt}}{L_{jt}} \right)}{d \ln (|\text{MRTS}_{MLjt}|)} = - \frac{d \ln \left(\frac{M_{jt}}{L_{jt}} \right)}{d \ln \left(\frac{P_{Mjt}}{W_{jt}} \right)},$$

where $|\text{MRTS}_{MLjt}|$ is the absolute value of the marginal rate of technological substitution between materials and labor, and the equality follows to the extent that it equals the price ratio P_{Mjt}/W_{jt} . However, because the estimates of the elasticity of substitution in the previous literature lie somewhere between zero and one (see Chirinko [2008] and the references therein for the elasticity of substitution between capital and labor and Bruno [1984], Rotemberg and Woodford [1996], and Oberfield and Raval [2014] for the elasticity of substitution between materials and an aggregate of capital and labor), this explanation is implausible. Labor-augmenting technological change offers an alternative explanation. As it makes labor more productive, equation (4) implies that it directly increases materials per unit of labor. Thus, labor-augmenting technological change may go a long way in rationalizing why the change in the input ratio M_{jt}/L_{jt} exceeds the change in the price ratio P_{Mjt}/W_{jt} .

In contrast, columns 10 and 11 of table 1 provide no evidence for capital-augmenting technological change. The investment boom in Spain in the 1990s and early 2000s was fueled by improved access to European and international capital markets. With the exception of industries 5, 6, and 8, the concomitant decrease in the input ratio M_{jt}/K_{jt} is much smaller than the increase in the price ratio P_{Mjt}/P_{Kjt} , where P_{Kjt} is the price of capital as, however, roughly measured by the user cost in our data. This pattern is consistent with an elasticity of substitution between materials and capital between zero and one. Indeed, capital-augmenting technological change can directly contribute to the decline in materials per unit of capital only in the unlikely scenario in which it makes capital less productive.

On the basis of these patterns in the data, we focus on labor-augmenting technological change in the subsequent analysis. We return to capital-augmenting technological change in Section X. In the remainder of this section we point out other features of the data that figure prominently in our analysis.

Temporary labor.—We distinguish between permanent and temporary labor and treat temporary labor as a static input that is chosen each pe-

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

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

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

Firms' R&D activities.—Columns 8–10 of table 2 show that the 10 industries differ markedly in terms of firms' R&D activities and that there is again substantial heterogeneity across the firms within an industry. Industries 3, 4, 5, and 6 exhibit high innovative activity. More than two-thirds of firms perform R&D during at least one year in the sample period, with at least 36 percent of stable performers engaging in R&D in all years (col. 8) and at least 28 percent of occasional performers engaging in R&D in some but not all years (col. 9). The R&D intensity among performers ranges, on average, from 2.2 percent to 2.9 percent (col. 10). Industries 1, 2, 7, and 8 are in an intermediate position. Less than half of firms perform R&D, and there are fewer stable than occasional performers. The R&D intensity is, on average, between 1.1 percent and 1.7 percent with a much lower value of 0.7 percent in industry 7. Finally, industries 9 and 10 exhibit low innovative

TABLE 2
DESCRIPTIVE STATISTICS (Continued)

INDUSTRY	TEMPORARY LABOR					INTRAFIRM MAX-MIN					OUTSOURCING					WITH R&D		
	Observations		Share	Share of Temporary	Hours Worked ^a	Hours per Worker ^a	Observations		Share	Stable	Occasional	R&D Intensity						
	(%)	(SD)	(2)	(3)	(4)	(5)	(%)	(SD)	(6)	(7)	(8)	(9)	(10)					
1. Metal and metal products	1,877	.260	.243	.448	.069	1,014	.200	56	109	.012								
	(79.4)	(.221)	(.197)	(.360)	(.090)	(42.9)	(.193)	(17.9)	(34.8)	(.018)								
2. Nonmetallic minerals	1,018	.231	.232	.482	.065	316	.177	20	62	.011								
	(80.2)	(.207)	(.183)	(.403)	(.063)	(24.9)	(.179)	(12.3)	(38.0)	(.022)								
3. Chemical products	1,722	.170	.203	.446	.043	924	.146	121	85	.026								
	(79.4)	(.176)	(.185)	(.427)	(.038)	(42.6)	(.183)	(40.5)	(28.4)	(.034)								
4. Agricultural and industrial machinery	1,069	.189	.227	.485	.086	808	.288	64	62	.022								
	(75.8)	(.181)	(.181)	(.419)	(.166)	(57.3)	(.263)	(36.0)	(34.8)	(.026)								
5. Electrical goods	1,221	.245	.280	.559	.063	763	.181	83	61	.029								
	(81.1)	(.206)	(.216)	(.452)	(.077)	(50.7)	(.194)	(39.7)	(29.2)	(.040)								
6. Transport equipment	962	.206	.239	.555	.131	637	.233	60	56	.028								
	(79.8)	(.198)	(.184)	(.415)	(.237)	(52.8)	(.261)	(37.3)	(34.8)	(.049)								
7. Food, drink, and tobacco	2,067	.276	.266	.468	.058	514	.142	65	86	.007								
	(84.2)	(.237)	(.215)	(.343)	(.065)	(20.9)	(.172)	(19.9)	(26.3)	(.022)								
8. Textile, leather, and shoes	1,726	.238	.291	.489	.062	1,214	.252	44	85	.017								
	(79.2)	(.260)	(.244)	(.402)	(.086)	(51.3)	(.237)	(13.1)	(25.4)	(.031)								
9. Timber and furniture	1,175	.320	.326	.523	.056	535	.183	21	44	.010								
	(81.3)	(.226)	(.234)	(.387)	(.076)	(37.0)	(.201)	(10.1)	(21.3)	(.017)								
10. Paper and printing products	1,024	.155	.221	.425	.057	679	.273	17	48	.015								
	(72.4)	(.145)	(.196)	(.346)	(.065)	(48.0)	(.253)	(9.3)	(26.2)	(.028)								

^a Computed as difference in logs.

activity. About a third of firms perform R&D, and the R&D intensity is, on average, between 1.0 percent and 1.5 percent.

IV. A Dynamic Model of the Firm

The purpose of our model is to enable us to infer a firm's productivity from its input usage and to clarify our assumptions on the timing of decisions that we rely on in estimation.

Production function.—The firm has the CES production function

$$Y_{jt} = \left\{ \beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + \left[\exp(\omega_{Ljt}) L_{jt}^* \right]^{-\frac{1-\sigma}{\sigma}} + \beta_M (M_{jt}^*)^{-\frac{1-\sigma}{\sigma}} \right\}^{-\frac{\sigma}{1-\sigma}} \times \exp(\omega_{Hjt}) \exp(e_{jt}), \quad (6)$$

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

The production function in equation (6) is the most parsimonious we can use to separate labor-augmenting from Hicks-neutral productivity. It encompasses three restrictions. First, technological change does not affect the parameters ν and σ , as we are unaware of evidence suggesting that the elasticity of scale or the elasticity of substitution varies over our sample period. Second, the elasticity of substitution between capital, labor, and materials is the same.⁸ We assess this restriction in Section VIII. For now we note that previous estimates of the elasticity of substitution between materials and an aggregate of capital and labor (Bruno 1984; Rotemberg and Woodford 1996; Oberfield and Raval 2014) fall in the same range as estimates of the elasticity of substitution between capital and labor (Chirinko 2008). Third, the productivities of capital and materials are restricted to change at the same rate and in lockstep with Hicks-neutral technological change.⁹ Treating capital and materials the same is in line with the fact that both

⁸ The elasticity of substitution between L_{Pjt} and L_{Tjt} , respectively, M_{Ijt} and M_{Ojt} , depends on the aggregators $\Lambda(L_{Pjt}, L_{Tjt})$ and $\Gamma(M_{Ijt}, M_{Ojt})$ and may differ from σ .

⁹ A production function with capital-augmenting, labor-augmenting, and materials-augmenting productivity that is homogeneous of arbitrary degree is equivalent to a production function with capital-augmenting, labor-augmenting, and Hicks-neutral productivity. Without loss of generality, we therefore subsume the common component of capital-augmenting, labor-augmenting, and materials-augmenting technological change into Hicks-neutral productivity.

are, at least to a large extent, produced goods. In contrast, labor is traditionally viewed as unique among the various factors of production, and changes in its productivity are a tenet of the literature on economic growth. The patterns in the data described in Section III further justify focusing on labor-augmenting technological change.

Laws of motion: productivity.—The components of productivity are presumably correlated with each other and over time and possibly also correlated across firms. As in Doraszelski and Jaumandreu (2013), we endogenize productivity by incorporating R&D expenditures into the model. We assume that the evolution of the components of productivity is governed by controlled first-order, time-inhomogeneous Markov processes with transition probabilities $P_{L,t+1}(\omega_{L,jt+1}|\omega_{L,jt}, R_{jt})$ and $P_{H,t+1}(\omega_{H,jt+1}|\omega_{H,jt}, R_{jt})$, where R_{jt} is R&D expenditures. Despite their parsimony, these stochastic processes accommodate correlation between the components of productivity.¹⁰ Moreover, because they are time-inhomogeneous, they accommodate secular trends in productivity.

The firm knows its current productivity when it makes its decisions for period t and anticipates the effect of R&D on its future productivity. The Markovian assumption implies

$$\omega_{L,jt+1} = E_t[\omega_{L,jt+1}|\omega_{L,jt}, R_{jt}] + \xi_{L,jt+1} = g_{L,t}(\omega_{L,jt}, R_{jt}) + \xi_{L,jt+1}, \quad (7)$$

$$\omega_{H,jt+1} = E_t[\omega_{H,jt+1}|\omega_{H,jt}, R_{jt}] + \xi_{H,jt+1} = g_{H,t}(\omega_{H,jt}, R_{jt}) + \xi_{H,jt+1}. \quad (8)$$

That is, *actual* labor-augmenting productivity $\omega_{L,jt+1}$ in period $t+1$ decomposes into *expected* labor-augmenting productivity $g_{L,t}(\omega_{L,jt}, R_{jt})$ and a random shock $\xi_{L,jt+1}$. This productivity innovation by construction is mean independent (although not necessarily fully independent) of $\omega_{L,jt}$ and R_{jt} . It captures the uncertainties that are naturally linked to productivity as well as those that are inherent in the R&D process such as chance of discovery, degree of applicability, and success in implementation. Nonlinearities in the link between R&D and productivity are captured by the conditional expectation function $g_{L,t}(\cdot)$ that we estimate nonparametrically along with the parameters of the production function. Actual Hicks-neutral productivity $\omega_{H,jt+1}$ decomposes similarly.

Laws of motion: capital.—Capital accumulates according to $K_{jt+1} = (1 - \delta) K_{jt} + I_{jt}$, where δ is the rate of depreciation. As in Olley and Pakes (1996), investment I_{jt} chosen in period t becomes effective in period $t+1$. Choosing I_{jt} is therefore equivalent to choosing K_{jt+1} .

¹⁰ Our empirical strategy generalizes to a joint Markov process $P_{t+1}(\omega_{L,jt+1}, \omega_{H,jt+1}|\omega_{L,jt}, \omega_{H,jt}, r_{jt})$. While R&D is widely seen as a major source of productivity growth (Griliches 1998), our empirical strategy extends to other sources such as technology adoption, learning by importing (Kasahara and Rodrigue 2008), and learning by exporting (De Loecker 2013). Both extensions are demanding on the data, however, as they increase the dimensionality of the functions that must be nonparametrically estimated.

Permanent labor.—Permanent labor is subject to convex adjustment costs $C_{L_r}(L_{pjt}, L_{pjt-1})$ that reflect the substantial cost of hiring and firing that the firm may incur (Hamermesh 1993). The choice of permanent labor thus may have dynamic implications. In contrast, temporary labor is a static input.

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

Output and input markets.—The firm has market power in the output market, for example, because products are differentiated. Its inverse residual demand function $P(Y_{jt}, D_{jt})$ depends on its output Y_{jt} and the demand shifter D_{jt} .¹¹ The firm is a price taker in input markets, where it faces W_{Tjt} , W_{Ijt} , P_{Ijt} , and P_{Ojt} as prices of permanent and temporary labor and in-house and outsourced materials, respectively.

The demand shifter and the prices that the firm faces in input markets evolve according to a Markov process that we do not further specify. As a consequence, the prices that the firm faces in period $t + 1$ may depend on its productivity in period t or on an average industrywide measure of productivity. Finally, the Markov process may be time-inhomogeneous to accommodate secular trends.

Bellman equation.—The firm makes its decisions to maximize the expected net present value of cash flows. In contrast to its labor-augmenting productivity ω_{Ljt} and its Hicks-neutral productivity ω_{Hjt} , the firm does not know the random shock e_{jt} when it makes its decisions for period t . Letting $V_t(\cdot)$ denote the value function in period t , the Bellman equation for the firm's dynamic programming problem is

$$\begin{aligned}
 V_t(\Omega_{jt}) = & \max_{K_{jt+1}, L_{pjt}, L_{tjt}, Q_{Mjt+1}, M_{Ijt}, R_{jt}} P\left(X_{jt}^{-\frac{\rho\sigma}{1-\sigma}} \exp(\omega_{Hjt}), D_{jt}\right) \\
 & \times X_{jt}^{-\frac{\rho\sigma}{1-\sigma}} \exp(\omega_{Hjt}) \mu - C_t(K_{jt+1} - (1 - \delta)K_{jt}) - W_{Tjt}L_{tjt} \\
 & - C_{L_r}(L_{pjt}, L_{pjt-1}) - W_{Ijt}L_{Ijt} - (P_{Ijt} + P_{Ojt}Q_{Mjt})M_{Ijt} \\
 & - C_{Q_M}(Q_{Mjt+1}, Q_{Mjt}) - C_R(R_{jt}) + \frac{1}{1 + \rho} E_t[V_{t+1}(\Omega_{jt+1}) | \Omega_{jt}, R_{jt}],
 \end{aligned} \tag{9}$$

¹¹ In general, the residual demand that the firm faces depends on its rivals' prices. In taking the model to the data, one may replace rivals' prices by an aggregate price index or dummies, although this substantially increases the dimensionality of the functions that must be nonparametrically estimated.

where

$$X_{jt} = \beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + [\exp(\omega_{Ljt}) L_{jt}^*]^{-\frac{1-\sigma}{\sigma}} + \beta_M (M_{jt}^*)^{-\frac{1-\sigma}{\sigma}};$$

$$\mu = E_t[\exp(e_{jt})];$$

the vector of state variables is

$$\Omega_{jt} = (K_{jt}, L_{Pjt-1}, Q_{Mjt}, \omega_{Ljt}, \omega_{Hjt}, W_{Pjt}, W_{Tjt}, P_{Ijt}, P_{Ojt}, D_{jt});$$

ρ is the discount rate; and $C_I(I_{jt})$ and $C_R(R_{jt})$ are the cost of investment and R&D, respectively, and accommodate indivisibilities in investment and R&D projects. The firm's dynamic programming problem gives rise to policy functions that characterize its investment and R&D decisions (and thus the values of K_{jt+1} or, equivalently, I_{jt} and R_{jt} in period t) as well as its input usage (L_{Pjt} , L_{Tjt} , Q_{Mjt+1} , and M_{Ijt}). The labor and materials decisions are central to our empirical strategy.

Inverse functions.—From the first-order conditions for permanent labor, temporary labor, and in-house materials, we derive functions $\tilde{h}_L(\cdot)$ and $h_H(\cdot)$ that allow us to recover unobservable labor-augmenting and Hicks-neutral productivity from observables. Appendix C contains detailed derivations.

We make several assumptions. First, our data have hours worked by permanent and temporary workers $L_{jt} = L_{Pjt} + L_{Tjt}$ and the (quantity) share of temporary labor $S_{Tjt} = L_{Tjt}/L_{jt}$. To map L_{jt}^* in the production function in equation (6) to the data, we assume that the aggregator $\Lambda(L_{Pjt}, L_{Tjt})$ is linearly homogeneous. This implies $L_{jt}^* = L_{jt}\Lambda(1 - S_{Tjt}, S_{Tjt})$. Moreover, because our data combine the wages of permanent and temporary workers into $W_{jt} = W_{Pjt}(1 - S_{Tjt}) + W_{Tjt}S_{Tjt}$, we assume that $W_{Pjt}/W_{Tjt} = \lambda_0$ is an (unknown) constant.¹²

Second, our data have the materials bill $P_{Mjt}M_{jt} = P_{Ijt}M_{Ijt} + P_{Ojt}M_{Ojt}$, the (value) share of outsourced materials $S_{Ojt} = P_{Ojt}M_{Ojt}/P_{Mjt}M_{jt}$, and the price of materials P_{Mjt} . To connect the model with the data, we assume $P_{Mjt} = P_{Ijt} + P_{Ojt}Q_{Mjt}$ so that the price of materials is the effective cost of an additional unit of in-house materials. This implies $M_{jt} = M_{Ijt}$. We further assume that $P_{Ijt}/P_{Ojt} = \gamma_0$ is an (unknown) constant and that $\Gamma(M_{Ijt}, M_{Ojt})$ is linearly homogeneous. This implies

$$M_{jt}^* = M_{Ijt}\Gamma\left(1, \gamma_0 \frac{S_{Ojt}}{1 - S_{Ojt}}\right).$$

We finally normalize $\Gamma(M_{Ijt}, 0) = M_{Ijt}$.

Third, the first-order condition for permanent labor involves a gap Δ_{jt} between the wage of permanent workers W_{Pjt} and their shadow wage. We

¹² In the supplementary appendix, we use a wage regression to estimate wage premia of various types of labor. We show that the wage premia do not change much, if at all, over time in line with our assumption that the ratio W_{Pjt}/W_{Tjt} is constant.

exploit the fact that we have three first-order conditions to substitute out for Δ_{jt} in the inverse functions $\tilde{h}_L(\cdot)$ and $h_H(\cdot)$. As this presumes interior solutions for permanent and temporary labor, we exclude observations with $S_{Tjt} = 0$ and thus $L_{Tjt} = 0$ from the subsequent analysis.¹³

Taken together, our assumptions allow us to recover (conveniently re-scaled) labor-augmenting productivity $\tilde{\omega}_{Ljt} = (1 - \sigma)\omega_{Ljt}$ and Hicks-neutral productivity ω_{Hjt} as

$$\begin{aligned}\tilde{\omega}_{Ljt} &= \tilde{\gamma}_L + m_{jt} - l_{jt} + \sigma(p_{Mjt} - w_{jt}) - \sigma\lambda_2(S_{Tjt}) + (1 - \sigma)\gamma_1(S_{Ojt}) \\ &\equiv \tilde{h}_L(m_{jt} - l_{jt}, p_{Mjt} - w_{jt}, S_{Tjt}, S_{Ojt}),\end{aligned}\tag{10}$$

$$\begin{aligned}\omega_{Hjt} &= \gamma_H + \frac{1}{\sigma}m_{jt} + p_{Mjt} - p_{jt} - \ln\left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right) \\ &\quad + \left(1 + \frac{\nu\sigma}{1 - \sigma}\right)x_{jt} + \frac{1 - \sigma}{\sigma}\gamma_1(S_{Ojt}) \\ &\equiv h_H(k_{jt}, m_{jt}, S_{Mjt}, p_{jt}, p_{Mjt}, D_{jt}, S_{Tjt}, S_{Ojt}),\end{aligned}\tag{11}$$

where $\tilde{\gamma}_L = -\sigma \ln \beta_M$, $\gamma_H = -\ln(\nu\beta_M\mu)$, $\eta(p_{jt}, D_{jt})$ is the absolute value of the price elasticity of the residual demand that the firm faces,

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

and $S_{Mjt} = P_{Mjt}M_{jt}/VC_{jt}$ is the share of materials in variable cost $VC_{jt} = W_{jt}L_{jt} + P_{Mjt}M_{jt}$. Without loss of generality, we set $\beta_K + \beta_M = 1$ in what follows.

We treat $\lambda_1(S_{Tjt})$, $\lambda_2(S_{Tjt})$, and $\gamma_1(S_{Ojt})$ as (unknown) functions of the share of temporary labor S_{Tjt} , respectively, the share of outsourced materials S_{Ojt} , that must be estimated nonparametrically along with the parameters of the production function. We thus think of $\lambda_1(S_{Tjt})$, $\lambda_2(S_{Tjt})$, and $\gamma_1(S_{Ojt})$ as “correction terms” on labor and, respectively, materials that help account for the substantial heterogeneity across the firms within an industry. Because we estimate these terms nonparametrically, they can accommodate different theories about the Spanish labor market and the role of outsourcing. For example, we develop an alternative model of outsourcing in the supplementary appendix that assumes that both in-house and outsourced materials are static inputs that the firm may mix and match at will, thereby dispensing with the costly to adjust ratio of outsourced to in-house materials.

¹³ Compare cols. 1 and 2 of tables 1 and 3 with cols. 1 and 2 of table 4 for the exact number of observations and firms we exclude.

V. Empirical Strategy

We combine the inverse functions in equations (10) and (11) with the laws of motion for labor-augmenting and Hicks-neutral productivity in equations (7) and (8) into estimation equations for the parameters of the production function in equation (6).

Labor-augmenting productivity.—Substituting the inverse function in equation (10) into the law of motion in equation (7), we form our first estimation equation

$$m_{jt} - l_{jt} = -\sigma(p_{Mjt} - w_{jt}) + \sigma\lambda_2(S_{Tjt}) - (1 - \sigma)\gamma_1(S_{Ojt}) \\ + \tilde{g}_{Lt-1}(\tilde{h}_L(m_{jt-1} - l_{jt-1}, p_{Mjt-1} - w_{jt-1}, S_{Tjt-1}, S_{Ojt-1}), R_{jt-1}) + \tilde{\xi}_{Ljt}, \quad (12)$$

where the (conveniently rescaled) conditional expectation function is

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

and $\tilde{\xi}_{Ljt} = (1 - \sigma)\xi_{Ljt}$.¹⁴

We allow $\tilde{g}_{Lt-1}(\tilde{h}_L(\cdot), R_{jt-1})$ to differ between zero and positive R&D expenditures and specify

$$\tilde{g}_{Lt-1}(\tilde{h}_L(\cdot), R_{jt-1}) = \tilde{g}_{L0}(t - 1) + 1(R_{jt-1} = 0)\tilde{g}_{L1}(\tilde{h}_L(\cdot)) \\ + 1(R_{jt-1} > 0)\tilde{g}_{L2}(\tilde{h}_L(\cdot), r_{jt-1}), \quad (13)$$

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

Compared to directly estimating equation (5) by OLS, equation (12) intuitively diminishes the endogeneity problem because breaking out the part of $\tilde{\omega}_{Ljt}$ that is observable via the conditional expectation function $\tilde{g}_{Lt-1}(\cdot)$ leaves “less” in the error term. This also facilitates instrumenting for any remaining correlation between the included variables and the error term.

In our model, labor l_{jt} , materials m_{jt} , the wage w_{jt} , and the share of temporary labor S_{Tjt} are correlated with the productivity innovation $\tilde{\xi}_{Ljt}$ (since $\tilde{\xi}_{Ljt}$ is part of $\tilde{\omega}_{Ljt}$). Note that $w_{jt} = \ln(W_{Tjt}(1 - S_{Tjt}) + W_{Tjt}S_{Tjt})$ may be cor-

¹⁴ Equation (12) is a semiparametric, partially linear, model with the additional restriction that the inverse function $\tilde{h}_L(\cdot)$ is of known form. Identification in the sense of the ability to separate the parametric and nonparametric parts of the model follows from standard arguments (Robinson 1988; Newey, Powell, and Vella 1999).

related with $\tilde{\xi}_{Ljt}$ even though the firm takes the wage of permanent workers W_{Pjt} and the wage of temporary workers W_{Tjt} as given because S_{Tjt} may depend on $\tilde{\omega}_{Ljt}$ and ω_{Hjt} through equations (C1) and (C2) in appendix C. We therefore base estimation on the moment conditions

$$E[A_{Ljt}(z_{jt})\tilde{\xi}_{Ljt}] = 0, \quad (14)$$

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

In considering instruments it is important to keep in mind that equation (12) models the evolution of labor-augmenting productivity $\tilde{\omega}_{Ljt}$. As a consequence, instruments have to be uncorrelated with the productivity innovation $\tilde{\xi}_{Ljt}$ but not necessarily with productivity itself. Because $\tilde{\xi}_{Ljt}$ is the innovation to productivity $\tilde{\omega}_{Ljt}$ in period t , it is not known to the firm when it makes its decisions in period $t - 1$. All past decisions are therefore uncorrelated with $\tilde{\xi}_{Ljt}$. In particular, having been decided in period $t - 1$, l_{jt-1} and m_{jt-1} are uncorrelated with $\tilde{\xi}_{Ljt}$, although they are correlated with $\tilde{\omega}_{Ljt}$ as long as productivity is correlated over time. Similarly, because S_{Tjt-1} and thus $w_{jt-1} = \ln(W_{Pjt-1}(1 - S_{Tjt-1}) + W_{Tjt-1}S_{Tjt-1})$ are determined in period $t - 1$, they are uncorrelated with the productivity innovation $\tilde{\xi}_{Ljt}$ in period t . We therefore use lagged labor l_{jt-1} , lagged materials m_{jt-1} , and the lagged wage w_{jt-1} for instruments.

In contrast to the wage w_{jt} , in our model the price of materials $p_{Mjt} = \ln(P_{Ijt} + P_{Ojt}Q_{Mjt})$ is uncorrelated with $\tilde{\xi}_{Ljt}$ because the ratio of outsourced to in-house materials Q_{Mjt} is determined in period $t - 1$. For the same reason, the share of outsourced materials $S_{Ojt} = P_{Ojt}Q_{Mjt}/(P_{Ijt} + P_{Ojt}Q_{Mjt})$ is uncorrelated with $\tilde{\xi}_{Ljt}$. We nevertheless choose to err on the side of caution and restrict ourselves to the lagged price of materials p_{Mjt-1} and the lagged share of outsourcing S_{Ojt-1} for instruments. Finally, time t and the demand shifter D_{jt} are exogenous by construction, and we use them for instruments.

The reasoning that the timing of decisions and the Markovian assumption on the evolution of productivity taken together imply that all past decisions are uncorrelated with productivity innovations originates in Olley and Pakes (1996). The subsequent literature uses it to justify lagged input quantities as instruments (see, e.g., Akerberg et al. 2007, sec. 2.4.1). In Doraszelski and Jaumandreu (2013, 1347–48), we extend this reasoning to justify lagged output and input prices as instruments. More recently, De Loecker et al. (2016, 471) do the same to justify the lagged price of output as instrument.

A test for overidentifying restrictions in Section VI cannot reject the validity of the moment conditions in equation (14). More targeted tests and additional checks further suggest that there is limited reason to doubt that w_{jt-1} and p_{Mjt-1} are uncorrelated with $\tilde{\xi}_{Ljt}$.

Hicks-neutral productivity.—Substituting the inverse functions in equations (10) and (11) into the production function in equation (6) and the law of motion for Hicks-neutral productivity ω_{Hjt} in equation (8), we form our second estimation equation:

$$y_{jt} = -\frac{\nu\sigma}{1-\sigma}x_{jt} + g_{Ht-1}(h_H(k_{jt-1}, m_{jt-1}, S_{Mjt-1}, p_{jt-1}, p_{Mjt-1}, D_{jt-1}, S_{Tjt-1}, S_{Ojt-1}), R_{jt-1}) + \xi_{Hjt} + e_{jt}. \quad (15)$$

We specify $g_{Ht-1}(h_H(\cdot), R_{jt-1})$ analogously to $\tilde{g}_{Lt-1}(\tilde{h}_L(\cdot), R_{jt-1})$ in equation (13).¹⁵

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

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

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

Estimation.—We use the two-step generalized method of moments (GMM) estimator of Hansen (1982). Let $\nu_{Ljt}(\theta_L) = \tilde{\xi}_{Ljt}$ be the residual of estimation equation (12) as a function of the parameters θ_L to be estimated and $\nu_{Hjt}(\theta_H) = \xi_{Hjt} + e_{jt}$ the residual of estimation equation (15) as a function of θ_H . The GMM problem corresponding to equation (12) is

$$\min_{\theta_L} \left[\frac{1}{N} \sum_j A_{Lj}(z_j) \nu_{Lj}(\theta_L) \right]' \widehat{W}_L \left[\frac{1}{N} \sum_j A_{Lj}(z_j) \nu_{Lj}(\theta_L) \right], \quad (16)$$

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

The GMM problem corresponding to equation (15) is analogous but considerably more nonlinear. To facilitate estimation, we impose the estimated values of those parameters in θ_L that also appear in θ_H . We correct the standard errors as described in the supplementary appendix. Because

¹⁵ Equation (15) is again a semiparametric model with the additional restriction that the inverse function $h_H(\cdot)$ is of known form. There are other possible estimation equations. In particular, one can use the labor and materials decisions in eqq. (C3) and (C5) together with the production function in eq. (6) to recover $\tilde{\omega}_{Ljt}$, ω_{Hjt} , and e_{jt} and then set up separate moment conditions in $\tilde{\xi}_{Ljt}$, ξ_{Hjt} , and e_{jt} . This may yield efficiency gains. Our estimation eq. (15) has the advantage that it is similar to a CES production function that has been widely estimated in the literature.

they tend to be more stable, we report first-step estimates for equation (15) and use them in the subsequent analysis; however, we use second-step estimates for testing.

VI. Labor-Augmenting Technological Change

From equation (12) we obtain an estimate of the elasticity of substitution and recover labor-augmenting productivity at the firm level.

Elasticity of substitution and a test for overidentifying restrictions.—Tables 3 and 4 summarize different estimates of the elasticity of substitution. To compare with the existing literature, we begin by proxying for $\tilde{\omega}_{ljt} = (1 - \sigma)\omega_{ljt}$ in equation (5) by a time trend $\tilde{\delta}_l t$ and estimate by OLS. As can be seen from columns 3 and 4 of table 3, with the exception of industry 9, the estimates of the elasticity of substitution are in excess of one, whereas the estimates in the previous literature lie somewhere between zero and one (Bruno 1984; Rotemberg and Woodford 1996; Chirinko 2008; Oberfeld and Raval 2014). This reflects, first, that a time trend is a poor proxy for labor-augmenting technological change at the firm level and, second, that the estimates are upward biased as a result of the endogeneity problem.

We address the endogeneity problem by modeling the evolution of labor-augmenting productivity and estimating equation (12) by GMM. To illustrate the importance of controlling for the composition of inputs in our empirical strategy, we revert to the setting in Section II and assume that labor l_{jt} and materials m_{jt} are homogeneous inputs that are chosen each period to maximize short-run profits. This implies $\lambda_1(S_{Tjt}) = 1$, $\lambda_2(S_{Tjt}) = 0$, and $\gamma_1(S_{Ojt}) = 0$, so that the correction terms on labor and materials vanish and equation (10) reduces to equation (5). Columns 5–10 of table 3 refer to this simplified model. As expected, the estimates of the elasticity of substitution are much lower and range from 0.45 to 0.64, as can be seen from column 5. With the exception of industries 6 and 8, in which σ is either implausibly high or low, we clearly reject the special cases of both a Leontieff ($\sigma \rightarrow 0$) and a Cobb-Douglas ($\sigma = 1$) production function.

Testing for overidentifying restrictions, however, we reject the validity of the moment conditions in the simplified model at a 5 percent level in five industries and we are close to rejecting in two more industries (cols. 6 and 7). To pinpoint the source of this problem, we exclude the subset of moments involving lagged materials m_{jt-1} from the estimation. As can be seen from columns 8–10, the resulting estimates of the elasticity of substitution lie between 0.46 and 0.85 in all industries, and at a 5 percent level we can no longer reject the validity of the moment conditions in any industry.

To see why the exogeneity of lagged materials m_{jt-1} is violated contrary to the timing of decisions in our model, recall that a firm engages

TABLE 3
ELASTICITY OF SUBSTITUTION

INDUSTRY	OBSERVATIONS ^a (1)	FIRMS ^a (2)	OLS		GMM INCLUDING m_{j-1} AS INSTRUMENT			GMM EXCLUDING m_{j-1} AS INSTRUMENT		
			σ (SE)	$\bar{\delta}_L$ (SE)	σ (SE)	χ^2 (df)	p -Value	σ (SE)	χ^2 (df)	p -Value
			(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1. Metals and metal products	2,365	313	1.163 (.104)	.023 (.007)	.451 (.096)	57.846 (40)	.034	.694 (.113)	13.683 (15)	.550
2. Nonmetallic minerals	1,270	163	1.227 (.119)	.038 (.008)	.643 (.086)	46.068 (40)	.234	.603 (.126)	11.299 (15)	.731
3. Chemical products	2,168	299	1.132 (.095)	.016 (.007)	.481 (.099)	65.068 (40)	.007	.618 (.124)	7.582 (15)	.939
4. Agricultural and industrial machinery	1,411	178	1.239 (.161)	.019 (.008)	.502 (.114)	56.166 (40)	.046	.598 (.103)	8.500 (15)	.902
5. Electrical goods	1,505	209	1.402 (.162)	.017 (.009)	.469 (.108)	60.674 (40)	.019	.458 (.108)	17.457 (15)	.292
6. Transport equipment	1,206	161	1.161 (.217)	.029 (.011)	1.204 (.089)	48.449 (40)	.169	.512 (.162)	7.740 (15)	.934
7. Food, drink, and tobacco	2,455	327	1.421 (.094)	.015 (.008)	.614 (.063)	70.492 (40)	.002	.707 (.084)	15.088 (15)	.445
8. Textile, leather, and shoes	2,368	335	1.846 (.169)	.001 (.010)	.059 (.077)	55.178 (40)	.056	.724 (.162)	18.453 (15)	.240
9. Timber and furniture	1,445	207	.793 (.117)	.013 (.008)	.461 (.089)	37.357 (40)	.590	.486 (.102)	5.805 (15)	.983
10. Paper and printing products	1,414	183	1.120 (.107)	.026 (.008)	.609 (.057)	51.798 (40)	.100	.854 (.077)	7.300 (15)	.949

^a Including $S_{j\mu} = L_{j\mu} = 0$.

TABLE 4
ELASTICITY OF SUBSTITUTION (Continued)

INDUSTRY	OBSERVATIONS ^a (1)	GMM			SARGAN DIFFERENCE TESTS			GMM WITH QUALITY-CORRECTED WAGE AS INSTRUMENT		
		σ (SE) (3)	χ^2 (df) (4)	ρ -Value (5)	χ^2 (df) (6)	ρ -Value (7)	P_{M-1} (8)	σ (SE) (10)	χ^2 (df) (11)	ρ -Value (12)
1. Metals and metal products	1,759	.535 (.114)	48.882 (38)	.111 (.38)	40.773 (25)	.024 (.25)	34.044 (25)	.456 (.112)	52.058 (38)	.064
2. Nonmetallic minerals	959	.730 (.098)	46.890 (38)	.153 (.38)	36.034 (25)	.071 (.25)	35.743 (25)	.833 (.096)	45.105 (38)	.199
3. Chemical products	1,610	.696 (.102)	46.154 (38)	.171 (.38)	33.225 (25)	.126 (.25)	32.183 (25)	.695 (.072)	48.889 (38)	.111
4. Agricultural and industrial machinery	979	.606 (.196)	42.420 (38)	.286 (.38)	29.398 (25)	.248 (.25)	25.684 (25)	.762 (.206)	44.227 (38)	.225
5. Electrical goods	1,147	.592 (.123)	46.778 (38)	.155 (.38)	38.951 (25)	.037 (.25)	32.376 (25)	.624 (.125)	44.592 (38)	.214
6. Transport equipment	896	.798 (.088)	45.741 (38)	.182 (.38)	19.053 (25)	.795 (.25)	9.901 (25)	.602 (.097)	41.214 (38)	.332
7. Food, drink, and tobacco	1,963	.616 (.081)	53.931 (38)	.045 (.38)	53.454 (25)	.001 (.25)	28.523 (25)	.766 (.079)	38.379 (38)	.452
8. Textile, leather, and shoes	1,593	.440 (.186)	52.496 (38)	.059 (.38)	23.355 (25)	.557 (.25)	31.763 (25)	.462 (.149)	55.996 (38)	.030
9. Timber and furniture	1,114	.438 (.093)	39.207 (38)	.416 (.38)	28.979 (25)	.265 (.25)	22.059 (25)	.497 (.094)	36.687 (38)	.530
10. Paper and printing products	938	.530 (.088)	44.448 (38)	.219 (.38)	23.642 (25)	.540 (.25)	19.822 (25)	.449 (.085)	43.009 (38)	.265

^a Excluding $S_{ij} = L_{ij} = 0$.

in outsourcing if it can procure customized parts and pieces from its suppliers that are cheaper or better than what the firm can make in-house from scratch. Lumping in-house and outsourced materials together pushes these quality differences into the error term. As outsourcing often relies on contractual relationships between the firm and its suppliers, the error term is likely correlated over time and thus with lagged materials m_{jt-1} as well.

Our leading specification accounts for quality differences between in-house and outsourced materials, respectively, permanent and temporary labor, and differences in the use of these inputs over time and across firms. The correction term $\gamma_1(S_{Ojt})$ in equation (12) absorbs quality differences into the aggregator $\Gamma(M_{jlt}, M_{Ojt})$ and accounts for the wedge that outsourcing may drive between the relative quantities and prices of materials and labor. The correction term $\lambda_2(S_{Tjt})$ similarly absorbs quality differences into the aggregator $\Lambda(L_{jlt}, L_{Tjt})$ and accounts for adjustment costs on permanent labor. As can be seen in columns 3–5 of table 4, the correction terms duly restore the exogeneity of lagged materials m_{jt-1} as we cannot reject the validity of the moment conditions at a 5 percent level in any industry except for industry 7, in which we (barely) reject it.¹⁶ Our leading estimates of σ in column 3 of table 4 lie between 0.44 and 0.80. Compared to the estimates in column 8 of table 3, there are no systematic changes, and our leading estimates are somewhat lower in five industries and somewhat higher in five industries. In short, relaxing the assumption that labor and materials are homogeneous and static inputs is a key step in estimating the elasticity of substitution.

Sargan difference tests.—Because the lagged wage w_{jt-1} and the lagged price of materials p_{Mjt-1} play a key role in the estimation of equation (12), we supplement the omnibus test for overidentifying restrictions with two Sargan difference tests to more explicitly validate their use as instruments. In the case of w_{jt-1} , we compute the difference in the value of the GMM objective function when we exclude the subset of moments involving p_{Mjt-1} and when we exclude the subset of moments involving w_{jt-1} and p_{Mjt-1} ; in the case of p_{Mjt-1} , we proceed analogously.¹⁷ As can be seen in columns 6–9 of table 4, the exogeneity assumption on the lagged wage is rejected at a 5 percent level in three industries, while that on the lagged price of materials cannot be rejected in any industry. Viewing all these tests in conjunction, to the extent that a concern about our leading specification is warranted, it appears more related to labor than to materials.

¹⁶ As noted in Sec. IV, we exclude observations with $S_{Tjt} = 0$ and thus $L_{Tjt} = 0$. Compare cols. 1 and 2 of tables 1 and 3 with cols. 1 and 2 of table 4 for the exact number of observations and firms we exclude.

¹⁷ To use the same weighting matrix for both specifications and not unduly change variances when we exclude subsets of moments, we delete the appropriate rows and columns from the weighting matrix for our leading specification.

Additional checks.—To further probe our leading specification and assess whether quality differences at a finer level play an important role, we leverage our data on the skill mix of a firm's labor force. As we show in the supplementary appendix, in our data the larger part of the variation in the wage across firms and periods can be attributed to geographic and temporal differences in the supply of labor and the fact that firms operate in different product submarkets. This part of the variation is arguably exogenous with respect to $\tilde{\xi}_{Ljt}$. The smaller part of the variation in the wage can be attributed to differences in the skill mix and the quality of labor that may potentially be correlated with $\tilde{\xi}_{Ljt}$.¹⁸

Our estimates are robust to purging this latter variation from the lagged wage w_{jt-1} . Using $\hat{w}_{Qjt-1}$ to denote the part of the wage that depends on the skill mix of a firm's labor force, we replace w_{jt-1} as an instrument by $w_{jt-1} - \hat{w}_{Qjt-1}$. Compared to column 3 of table 4, the estimates of the elasticity of substitution in column 10 decrease somewhat in three industries, remain essentially unchanged in two industries, and increase somewhat in five industries.¹⁹ The absence of substantial and systematic changes confirms that the variation in w_{jt-1} is exogenous and therefore useful in estimating equation (12), in line with the test for overidentifying restrictions.

Below we further exploit our data on the skill mix to explicitly model quality differences at a finer level by assuming that the firm faces a menu of qualities and wages in the market for permanent labor. Taken together, these additional checks alleviate concerns about the quality and composition of labor.

Labor-augmenting technological change.—With equation (12) estimated, we recover the labor-augmenting productivity $\omega_{Ljt} = \hat{\omega}_{Ljt}/(1 - \sigma)$ of firm j in period t up to an additive constant from equation (10). In what follows, we therefore demean ω_{Ljt} by industry. Abusing notation, we continue to use ω_{Ljt} to denote the demeaned labor-augmenting productivity of firm j in period t .

To obtain aggregate measures representing an industry, we account for the survey design by replicating the subsample of small firms 70 percent/

¹⁸ A parallel discussion applies to materials. Kugler and Verhoogen (2012) point to differences in the quality of materials whereas Atalay (2014) documents substantial variation in the price of materials across plants in narrowly defined industries with negligible quality differences. This variation is partly due to geography and differences in cost and markup across suppliers that are arguably exogenous to a plant.

¹⁹ As we show in the supplementary appendix, not much changes if we isolate the part of the wage that additionally depends on firm size to try to account for the quality of labor beyond our rather coarse data on the skill mix of a firm's labor force (Oi and Idson 1999). Compared to col. 3 of table 4, the estimates of the elasticity of substitution decrease somewhat in three industries, remain essentially unchanged in three industries, and increase somewhat in four industries.

5 percent = 14 times before pooling it with the subsample of large firms. Unless noted otherwise, we report weighed averages of individual measures, where the weight $\mu_{jt} = P_{jt-2}Y_{jt-2}/\sum_j P_{jt-2}Y_{jt-2}$ is the share of sales of firm j in period $t - 2$. Using the second lag reduces the covariance between the weight and the variable of interest.

The growth of labor-augmenting productivity at firm j in period t is $\Delta\omega_{Ljt} = \omega_{Ljt} - \omega_{Ljt-1}$.²⁰ In line with the patterns in the data described in Section III, our estimates imply an important role for labor-augmenting technological change. As can be seen from column 1 of table 5, labor-augmenting productivity grows quickly, on average, with rates of growth ranging from 1.0 percent and 1.7 percent per year in industries 8 and 7 to 14.2 percent and 18.3 percent in industries 2 and 6 and above in industry 5.

Ceteris paribus,

$$\Delta\omega_{Ljt} \approx \frac{\exp(\omega_{Ljt})L_{jt-1}^* - \exp(\omega_{Ljt-1})L_{jt-1}^*}{\exp(\omega_{Ljt-1})L_{jt-1}^*}$$

approximates the rate of growth of a firm's effective labor force $\exp(\omega_{Ljt-1})L_{jt-1}^*$. To facilitate comparing labor-augmenting to Hicks-neutral productivity, we approximate the rate of growth of the firm's output Y_{jt-1} by $\epsilon_{Ljt-2}\Delta\omega_{Ljt}$, where ϵ_{Ljt-2} is the elasticity of output with respect to the firm's effective labor force in period $t - 2$ (see app. E).²¹ This output effect, while on average close to zero in industry 9, ranges from 0.7 percent per year in industry 7 to 3.1 percent, 3.2 percent, and 3.6 percent in industries 2, 4, and 6; see column 2 of table 5. Across industries, labor-augmenting technological change causes output to grow by 1.7 percent per year.

Figure 1 illustrates the magnitude of the output effect of labor-augmenting technological change and the heterogeneity in its impact across industries. The depicted index cumulates the year-to-year changes and is normalized to one in 1991. Technological change appears to have slowed in the 2000s compared to the 1990s: across industries, labor-augmenting technological change causes output to grow by 2.1 percent per year before 2000 and by 1.0 percent per year after 2000.

Dispersion and persistence.—A substantial literature documents dispersion and persistence in productivity (see Bartelsman and Doms [2000])

²⁰ Given the specification of $\tilde{g}_{L,t-1}(\tilde{h}_L(\cdot), R_{jt-1})$ in eq. (13), we exclude observations in which a firm switches from performing to not performing R&D or vice versa between periods $t - 1$ and t from the subsequent analysis. We further exclude observations in which a firm switches from zero to positive outsourcing or vice versa.

²¹ Because ϵ_{Ljt} depends on ω_{Ljt} as can be seen from eq. (E1), $\Delta\omega_{Ljt}$ is systematically negatively correlated with ϵ_{Ljt} and systematically positively correlated with ϵ_{Ljt-1} . Using ϵ_{Ljt-2} drastically reduces the correlation between the constituent parts of the output effect of labor-augmenting technological change.

TABLE 5
LABOR-AUGMENTING TECHNOLOGICAL CHANGE

INDUSTRY	FIRMS' R&D ACTIVITIES										SKILL UPGRADING		
	$\Delta\omega_L$ (1)	$\epsilon_{L,-2}\Delta\omega_L$ (2)	$\epsilon_{L,-2}\omega_L^a$		$\epsilon_{L,-2}\omega_L$		R&D-No R&D (5)	R&D (6)	No R&D (7)	σ (SE) (8)	χ^2 (df) (9)	p -Value (10)	$\Delta\omega_L$ (11)
			IQR (3)	AC (4)	R&D- ω_L (5)	$\epsilon_{L,-2}\Delta\omega_L$ (7)							
1. Metals and metal products	.091	.021	.368	.707	.189	.024	.018	.582 (.117)	44.868 (38)	.206	.104		
2. Nonmetallic minerals	.142	.031	.609	.858	.291	.022	.028	.737 (.092)	35.898 (38)	.567	.087		
3. Chemical products	.049	.014	.447	.893	.121	.018	-.002	.618 (.110)	47.832 (38)	.132	.053		
4. Agricultural and industrial machinery	.126	.032	.570	.803	.385	.028	.046	.177 (.172)	38.413 (38)	.451	.060		
5. Electrical goods	.220	.022	.523	.855	.303	.022	.012	.488 (.129)	48.365 (38)	.121	.179		
6. Transport equipment	.183	.036	.722	.840	.396	.045	.012	.781 (.101)	45.457 (38)	.189	.098		
7. Food, drink, and tobacco	.018	.007	.455	.876	.064	.009	.006	.655 (.084)	53.981 (38)	.045	-.007		
8. Textile, leather, and shoes	.010	.007	.347	.855	.165	.007	.009	.120 (.168)	41.931 (38)	.304	.000		
9. Timber and furniture	.065 ^b	.001 ^b	.243	.608	.091	.007 ^b	.000 ^b	.528 (.090)	37.674 (38)	.484	-.023		
10. Paper and printing products	.023	.014	.275	.820	.137	.007	.021	.396 (.082)	37.418 (38)	.496	-.011		
All industries	.102	.017				.020	.011				.075		

^a Without replication and weighting.

^b We trim values of $\Delta\omega_L$, respectively, $\epsilon_{L,-2}\Delta\omega_L$, below 0.25 and above 0.5.

FIG. 1.—Output effect of labor-augmenting technological change. Index normalized to one in 1991.

and Syverson [2011] and the references therein). To be able to compare labor-augmenting productivity to Hicks-neutral productivity, we focus on $\epsilon_{Ljt-2}\omega_{Ljt}$. Because ω_{Ljt} is demeaned, $\epsilon_{Ljt-2}\omega_{Ljt}$ measures the labor-augmenting productivity of firm j in period t relative to the average productivity, suitably converted into output terms. We thus refer to $\epsilon_{Ljt-2}\omega_{Ljt}$ as labor-augmenting productivity in output terms in what follows.

We measure dispersion by the interquartile range of $\epsilon_{Ljt-2}\omega_{Ljt}$. As can be seen from column 3 of table 5, the interquartile range (IQR) is between 0.24 in industry 9 and 0.72 in industry 6. This is comparable to results in the existing literature.²² Turning from dispersion to persistence, $\epsilon_{Ljt-2}\omega_{Ljt}$ is highly autocorrelated (AC; col. 4), indicating that differences in labor-augmenting productivity between firms persist over time.

Firms' R&D activities.—As can be seen from column 5 of table 5, firms that perform R&D have, on average, higher levels of labor-augmenting

²² For US manufacturing industries, Syverson (2004) reports an interquartile range of log labor productivity of 0.66.

productivity in output terms than firms that do not perform R&D in all industries. In six industries the output effect of labor-augmenting technological change for firms that perform R&D, on average, exceeds that of firms that do not perform R&D (cols. 6 and 7). Overall, our estimates indicate that firms' R&D activities are associated not only with higher levels of labor-augmenting productivity but by and large also with higher rates of growth of labor-augmenting productivity.

Firm turnover.—To assess the impact of firm turnover on the output effect of labor-augmenting technological change, we classify a firm as a survivor if it enters the industry in or before 1990 and does not exit in or before 2006, as an exitor if it enters the industry in or before 1990 and exits in or before 2006, and as an entrant otherwise. Survivors account for most of the output effect of labor-augmenting technological change. Their contribution is 80 percent in industry 6 and above, except for industry 3, where the contribution of entrants is on par with the contribution of survivors. In the remaining industries, the contribution of entrants is small. The contribution of exitors is small in all industries.

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

To answer this question—and to further alleviate concerns about the quality and composition of labor—we exploit that, in addition to the share of temporary labor S_{Tjt} , our data have the share of white-collar workers and the shares of engineers, and, respectively, technicians. We assume that there are Q types of permanent labor with qualities 1, $\theta_2, \dots, \theta_Q$ and corresponding wages $W_{P1jt}, W_{P2jt}, \dots, W_{PQjt}$. The firm, facing this menu of qualities and wages, behaves as a price taker in the labor market. In recognition of their different qualities,

$$L_{Pjt}^* = L_{P1jt} + \sum_{q=2}^Q \theta_q L_{Pqjt} = L_{Pjt} \left[1 - \sum_{q=2}^Q (\theta_q - 1) S_{Pqjt} \right]$$

is an aggregate of the Q types of permanent labor, with L_{Pqjt} being the quantity of permanent labor of type q at firm j in period t and S_{Pqjt} the corresponding share in $L_{Pjt} = \sum_{q=1}^Q L_{Pqjt}$; $L_{jt}^* = \Lambda(L_{Pjt}^*, L_{Tjt})$ is the aggregate of permanent labor L_{Pjt}^* (instead of L_{Pjt}) and temporary labor L_{Tjt} in the production function in equation (6). Permanent labor is subject to convex adjustment costs $C_{B_r}(B_{Pjt}, B_{Pjt-1})$, where $B_{Pjt} = \sum_{q=1}^Q W_{Pqjt} L_{Pqjt}$ is the wage bill for permanent labor. The state vector Ω_{jt} in the firm's dynamic pro-

gramming problem therefore includes $B_{p_{jt}-1}$, $W_{p_{1jt}}$, $W_{p_{2jt}}$, ..., $W_{p_{Qjt}}$ instead of $L_{p_{jt}-1}$ and $W_{p_{jt}}$.

In the supplementary appendix we show that our first estimation equation (12) remains unchanged except that $\lambda_2(S_{Tjt})$ is replaced by $\lambda_2(S_{Tjt}, \Theta_{jt})$, where

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

is a quality index. We use a wage regression to estimate the wage premium, $W_{p_{qjt}}/W_{p_{1jt}} - 1$, of permanent labor of type q over type 1 and construct the quality index Θ_{jt} .

The estimates of the elasticity of substitution in column 8 of table 5 continue to hover around 0.6 across industries, with the exception of industries 4 and 8, in which they are implausibly low. Compared to column 3 of table 4, they decrease somewhat in three industries, remain essentially unchanged in two industries, and increase somewhat in five industries. This further supports the notion that quality differences at a finer level than permanent and temporary labor are of secondary importance for estimating equation (12).

We develop the quality index Θ_{jt} mainly to “chip away” at the productivity residual by improving the measurement of inputs in the spirit of the productivity literature (Griliches 1964; Griliches and Jorgenson 1967). As can be seen from column 11 of table 5, skill upgrading indeed explains some, but by no means all, of the growth of labor-augmenting productivity. Compared to column 1, the rates of growth stay the same or go down in all industries. In industries 7, 8, 9, and 10, labor-augmenting productivity is stagnant or declining after accounting for skill upgrading, indicating that improvements in the skill mix over time are responsible for most of the growth of labor-augmenting productivity. In contrast, in industries 1, 2, 3, 4, 5, and 6, labor-augmenting productivity continues to grow after accounting for skill upgrading, albeit often at a much slower rate. In these industries, labor-augmenting productivity grows also because workers with a given set of skills become more productive over time.

VII. The Decline of the Aggregate Share of Labor

In many advanced economies the aggregate share of labor in income has declined in past decades. While this decline has attracted considerable attention in the academic literature (Blanchard 1997; Bentolila and Saint-Paul 2004; McAdam and Willman 2013; Karabarbounis and Neiman 2014; Oberfield and Raval 2014) and in the public discussion following Piketty (2014), its causes and consequences remain contested. We use our estimates to show that biased technological change is the primary driver of

the decline of the aggregate share of labor in the Spanish manufacturing sector over our sample period.

Let $VC_{Ljt} = W_{jt}L_{jt}$ be the wage bill, $VC_{jt} = W_{jt}L_{jt} + P_{Mjt}M_{jt}$ variable cost, and $S_{Ljt} = VC_{Ljt}/VC_{jt}$ the share of labor in variable cost of firm j in period t . Let $VC_{Lt} = \sum_j VC_{Ljt}$ and $VC_t = \sum_j VC_{jt}$ be the corresponding industrywide aggregates. We focus on the aggregate share of labor in variable cost

$$S_{Lt} = \frac{VC_{Lt}}{VC_t} = \sum_j \frac{VC_{Ljt}}{VC_{jt}} \frac{VC_{jt}}{VC_t} = \sum_j S_{Ljt} \theta_{jt},$$

where $\theta_{jt} = VC_{jt}/VC_t$ is the variable cost of firm j in period t as a fraction of aggregate variable cost. As can be seen in figure 2, the aggregate share of labor in variable cost closely tracks the aggregate share of labor in value added in the Spanish manufacturing sector in the National Accounts over our sample period.²³

The year-to-year change in the aggregate share of labor in variable cost is $S_{Lt} - S_{Lt-1}$. Cumulated over our sample period, the decline of the aggregate share of labor ranges from 0.01 and 0.05 in industries 9 and 4 to 0.15 and 0.19 in industries 2 and 5, as can be seen in column 1 of table 6.²⁴ To obtain insight into this decline, we build on Oberfield and Raval (2014) and decompose the year-to-year change as

$$S_{Lt} - S_{Lt-1} = \sum_j \theta_{jt} (S_{Ljt} - S_{Ljt-1}) + \sum_j (\theta_{jt} - \theta_{jt-1}) S_{Ljt-1}.$$

The second term captures reallocation across firms. Our model enables us to further decompose the first term. Rewriting equation (10) yields

$$\begin{aligned} S_{Ljt} = & [1 + \exp(-\tilde{\gamma}_L + (1 - \sigma)(p_{Mjt} - w_{jt} + \omega_{Ljt})) \\ & + \sigma\lambda_2 (S_{Tjt}) - (1 - \sigma)\gamma_1 (S_{Ojt})]^{-1}. \end{aligned} \quad (17)$$

The first term may thus be driven by a change in the price of materials p_{Mjt} relative to the price of labor w_{jt} , a change in labor-augmenting productivity ω_{Ljt} , a change in the share of temporary labor S_{Tjt} , and a change in the share of outsourced materials S_{Ojt} . To quantify these drivers, we use a second-order approximation to $S_{Ljt} - S_{Ljt-1}$ as described in appendix F.

We report the decomposition of the year-to-year change, cumulated over our sample period, in columns 2–7 of table 6. The small size of

²³ Contabilidad Nacional de España, Bases 1986 and 1995, Instituto Nacional de Estadística.

²⁴ We estimate $S_{Lt} - S_{Lt-1}$ as well as the various terms of the decomposition using firms that are in the sample in periods t and $t - 1$.

FIG. 2.—Aggregate share of labor in value added in National Accounts (left axis) and aggregate share of labor in variable cost in sample and counterfactual (right axis). The latter indices cumulate year-to-year changes using level in 1990 as base and average over industries using their share of total value added in column 4 of table B1 as weight.

the residual in column 7 indicates that our second-order approximation to $S_{Ljt} - S_{Ljt-1}$ readily accommodates nonlinearities. As can be seen in column 3, biased technological change emerges as the main force behind the decline of the aggregate share of labor. Changes in input prices in column 2 attenuate the decline. In contrast, the impact of temporary labor, outsourced materials, and reallocation across firms in the remaining columns is sometimes positive and sometimes negative and mostly small.

We use our model to compute the counterfactual evolution of the aggregate share of labor without biased technological change by zeroing out the change in labor-augmenting productivity ω_{Ljt} in the decomposition of the year-to-year change. As can be seen in figure 2, without biased technological change, the aggregate share of labor remains roughly constant over our sample period. We emphasize that this counterfactual holds fixed not only reallocation across firms but also the evolution of input prices, temporary labor, and outsourced materials. This may be questionable over longer stretches of time.

TABLE 6
AGGREGATE SHARE OF LABOR IN VARIABLE COST

INDUSTRY	GROWTH OF LABOR SHARE ^a (1)	DECOMPOSITION ^a				
		$p_M - w$ (2)	ω_L (3)	Temporary Labor (4)	Outsourcing (5)	Reallocation (6)
1. Metal and metal products	-.107	.011	-.086	-.034	.011	-.004
2. Nonmetallic minerals	-.153	.011	-.125	-.010	-.002	-.002
3. Chemical products	-.087	.018	-.143	.033	-.002	.008
4. Agricultural and industrial machinery	-.046	.031	-.064	-.027	-.007	.016
5. Electrical goods	-.188	.031	-.183	.025	-.017	-.040
6. Transport equipment	-.066	.023	-.139	.047	.003	.003
7. Food, drink, and tobacco	-.060	.022	-.100	-.012	.023	-.001
8. Textile, leather, and shoes	-.057	.016	-.073	-.012	.006	.016
9. Timber and furniture	-.009	.056	-.043	-.005	-.002	-.018
10. Paper and printing products	-.059	.020	-.102	.009	.001	.013
						.000

^a Computed for 1991–2006.

Our conclusion that biased technological change is the primary driver of the decline of the aggregate share of labor echoes that of Oberfield and Raval (2014). The authors develop a decomposition of the change in the aggregate share of labor in value added in the US manufacturing sector from 1970 to 2010. Perhaps the most important difference between their decomposition and ours is that we directly measure the bias of technological change at the level of the individual firm, whereas Oberfield and Raval treat it as the residual of their decomposition. Despite this difference and the different data sets used, the decompositions are complementary and both point to the overwhelming role of biased technological change in the decline of the aggregate share of labor.

VIII. Hicks-Neutral Technological Change

From equation (12) we obtain an estimate of the elasticity of substitution and recover labor-augmenting productivity at the firm level. To recover Hicks-neutral productivity and the remaining parameters of the production function, we have to estimate equation (15).

Distributional parameters and elasticity of scale.—Table 7 reports the distributional parameters β_K and $\beta_M = 1 - \beta_K$ and the elasticity of scale ν . Our estimates of β_K range from 0.07 in industry 8 to 0.31 in industry 6 (col. 1). Although the estimates of the elasticity of scale are rarely significantly different from one, taken together they suggest slightly decreasing returns to scale (col. 2). We cannot reject the validity of the moment conditions in any industry by a wide margin (cols. 3 and 4).²⁵

Price elasticity.—Column 5 of table 7 reports the average absolute value of the price elasticity $\eta(p_{jt-1}, D_{jt-1})$ implied by our estimates. It ranges from 1.79 in industry 9 to 6.04 and 9.11 in industries 5 and 2 and averages 3.20 across industries.²⁶

Elasticity of substitution: Lagrange multiplier test.—The production function in equation (6) assumes that the elasticity of substitution between capital, labor, and materials is the same. We compare our leading specification to the more general nested CES production function

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

²⁵ In light of this wide margin, we do not further probe the validity of lagged prices as instruments.

²⁶ For US manufacturing industries, Oberfield and Raval (2014) report price elasticities in a somewhat narrower range between 2.91 and 5.22 with a roughly comparable average of 3.91 across industries.

TABLE 7
DISTRIBUTIONAL PARAMETERS, ELASTICITY OF SCALE, AND PRICE ELASTICITY

INDUSTRY	GMM				LAGRANGE MULTIPLIER TEST	
	β_K (SE) (1)	ν (SE) (2)	χ^2 (df) (3)	p-Value (4)	$\chi^2(1)$ (6)	p-Value (7)
1. Metals and metal products	.232 (.073)	.941 (.029)	3.207 (8)	.921	1.023	.312
2. Nonmetallic minerals	.225 (.133)	.911 (.063)	4.528 (8)	.807	.489	.485
3. Chemical products	.137 (.059)	.933 (.041)	1.109 (8)	.997	.342	.559
4. Agricultural and industrial machinery	.138 (.125)	.806 (.088)	8.251 (9)	.509	1.126	.289
5. Electrical goods	.132 (.037)	.848 (.045)	2.960 (10)	.982	.902	.342
6. Transport equipment ^b	.307 (.182)	.923 (.061)			2.163	
7. Food, drink, and tobacco	.303 (.137)	.931 (.040)	2.415 (8)	.966	.295	.587
8. Textile, leather, and shoes	.066 (.097)	.976 (.035)	1.120 (9)	.999	.357	.550
9. Timber and furniture ^b	.103 (.107)	.932 (.066)			1.787	
10. Paper and printing products	.233 (.083)	.939 (.041)	3.846 (8)	.871	1.716	.190

^a We trim 5 percent of observations at the right tail.

^b We have been unable to compute the second-step GMM estimate.

where the additional parameter τ is the elasticity of substitution between capital and labor, respectively, materials. We show in the supplementary appendix that our first estimation equation (12) remains unchanged and generalize our second estimation equation (15). This allows us to conduct a Lagrange multiplier test for $\tau = \sigma$. As can be seen in columns 6 and 7 of table 7, we cannot reject the validity of our leading specification in any industry.

Hicks-neutral technological change.—With equation (15) estimated, we recover the Hicks-neutral productivity ω_{Hjt} of firm j in period t up to an additive constant from equation (11); in what follows, we use ω_{Hjt} to denote the demeaned Hicks-neutral productivity. We proceed as before to obtain aggregate measures representing an industry.

The growth of Hicks-neutral productivity at firm j in period t is $\Delta\omega_{Hjt} = \omega_{Hjt} - \omega_{Hjt-1}$. *Ceteris paribus*,

$$\Delta\omega_{Hjt} \approx \frac{X_{jt-1}^{-\frac{\sigma\tau}{1-\sigma}} \exp(\omega_{Hjt}) \exp(e_{jt-1}) - X_{jt-1}^{-\frac{\sigma\tau}{1-\sigma}} \exp(\omega_{Hjt-1}) \exp(e_{jt-1})}{X_{jt-1}^{-\frac{\sigma\tau}{1-\sigma}} \exp(\omega_{Hjt-1}) \exp(e_{jt-1})}$$

approximates the rate of growth of a firm's output Y_{jt-1} and is therefore directly comparable to the output effect of labor-augmenting technological change. As can be seen from column 1 of table 8, Hicks-neutral productivity grows quickly in five industries, with rates of growth ranging, on average, from 1.2 percent per year in industry 8 to 4.4 percent in industry 1. It grows much more slowly or barely at all in three industries, with rates of growth below 0.5 percent per year. While there is considerable heterogeneity in the rate of growth of Hicks-neutral productivity across industries, Hicks-neutral technological change causes output to grow by 1.4 percent per year.

Figure 3 illustrates the magnitude of Hicks-neutral technological change. The depicted index cumulates the year-to-year changes and is normalized to one in 1991.²⁷ The heterogeneity in the impact of Hicks-neutral technological change across industries clearly exceeds that of the output effect of labor-augmenting technological change (see again fig. 1). Once again, technological change appears to have slowed in the 2000s compared to the 1990s: across industries, Hicks-neutral technological change causes output to grow by 2.7 percent per year before 2000 and to shrink by 0.6 percent per year after 2000.

Dispersion and persistence.—We measure dispersion by the interquartile range of ω_{Hjt} . As can be seen from column 2 of table 8, the interquartile range is between 0.37 in industry 3 and 0.98 in industry 4.²⁸ Hicks-neutral productivity appears to be somewhat more disperse than labor-augmenting

²⁷ In industry 9, in line with col. 1 of table 8, we trim values of $\Delta\omega_H$ below -0.25 and above 0.5 .

²⁸ For Chinese manufacturing industries, Hsieh and Klenow (2009) report an interquartile range of log total factor productivity of 1.28.

TABLE 8
HICKS-NEUTRAL TECHNOLOGICAL CHANGE

INDUSTRY	FIRMS' R&D ACTIVITIES								
	$\Delta\omega_H$ (1)	ω_H^a		ω_H		$\Delta\omega_H$		TOTAL TECHNOLOGICAL CHANGE	
		IQR (2)	AC (3)	R&D-No		R&D (5)	No R&D (6)	$\epsilon_{L-2}\Delta\omega_L + \Delta\omega_H$ (7)	$\text{corr}(\epsilon_{L-2}\Delta\omega_L, \Delta\omega_H)^a$ (8)
				R&D (4)					
1. Metals and metal products	.044	.718	.736	.004	.046	.038	.065	.172	
2. Nonmetallic minerals	.005	.551	.685	.262	-.019	.041	.036	.205	
3. Chemical products	.019	.368	.872	.002	.022	.011	.032	-.029	
4. Agricultural and industrial machinery	.041	.978	.890	.445	.039	.022	.072	.076	
5. Electrical goods	.020	.696	.844	.723	.009	.055	.042	.072	
6. Transport equipment	.042	.555	.595	.135	.058	-.031	.078	.360	
7. Food, drink, and tobacco	.001	.760	.909	-.146	.007	.000	.007	.519	
8. Textile, leather, and shoes	.012	.715	.859	-.199	-.003	.032	.019	.238	
9. Timber and furniture	.021 ^b	.772	.772	-.098	.008 ^b	.035 ^b	.041 ^b	.709	
10. Paper and printing products	.002	.612	.879	-.129	.006	.006	.015	.283	
All industries	.014				.016	.011	.031		

^a Without replication and weighting.

^b We trim values of $\Delta\omega_H$ respectively, $\epsilon_{L-2}\Delta\omega_L + \Delta\omega_H$ below -0.25 and above 0.5 .

FIG. 3.—Hicks-neutral technological change. Index normalized to one in 1991.

productivity in output terms. Once again, ω_{Hjt} is highly autocorrelated (col. 3), indicating that differences in Hicks-neutral productivity between firms persist over time.

Firms' R&D activities.—As can be seen from column 4 of table 8, firms that perform R&D have, on average, higher levels of Hicks-neutral productivity than firms that do not perform R&D in six industries but lower levels of Hicks-neutral productivity in four industries. While there is practically no difference in industry 10, the rate of growth of Hicks-neutral productivity for firms that perform R&D, on average, exceeds that of firms that do not perform R&D in five industries (cols. 5 and 6). Overall, our estimates indicate that firms' R&D activities are associated with higher levels and rates of growth of Hicks-neutral productivity, although firms' R&D activities seem less closely tied to Hicks-neutral than to labor-augmenting productivity. This is broadly consistent with the large literature on induced innovation that argues that firms direct their R&D activities to conserve the relatively more expensive factors of production, in particular, labor.²⁹

²⁹ More explicitly testing for induced innovation is difficult because we do not observe what a firm does with its R&D expenditures. One way to proceed may be to add interactions of R&D expenditures and input prices to the laws of motion in eqq. (7) and (8). We leave this to future research.

Firm turnover.—Similarly to the output effect of labor-augmenting technological change, survivors account for most of Hicks-neutral technological change. Their contribution is 61 percent in industry 9 and above. While the contributions of entrants and exitors are small in most industries, they are negative and more sizable in industries 2, 5, 7, 8, and 10. As a result, in these industries the rate of growth of Hicks-neutral productivity is 0.7 percent, 3.0 percent, 1.2 percent, 1.7 percent, and 1.2 percent among survivors compared to 0.5 percent, 2.0 percent, 0.1 percent, 1.2 percent, and 0.2 percent for all firms (see again col. 1 of table 8).

Total technological change and its components.—As productivity is multidimensional, we take total technological change to be $\epsilon_{Ljt-2}\Delta\omega_{Ljt} + \Delta\omega_{Hjt}$. Taken together, labor-augmenting and Hicks-neutral technological change cause output to grow by, on average, between 0.7 percent in industry 7 and 7.2 percent and 7.8 percent in industries 4 and 6, as can be seen in column 7 of table 8. Across all industries, total technological change causes output to grow by 3.1 percent per year.

The output effects of labor-augmenting technological change $\epsilon_{Ljt-2}\Delta\omega_{Ljt}$ and Hicks-neutral technological change $\Delta\omega_{Hjt}$ are positively correlated in nine industries, while the correlation is slightly negative in one industry (col. 8). The correlation between labor-augmenting productivity in output terms $\epsilon_{Ljt-2}\omega_{Ljt}$ and Hicks-neutral productivity ω_{Hjt} is positive in all industries. Overall, our estimates not only provide evidence that productivity is multi- instead of single-dimensional but also suggest that the various components of productivity are intertwined.

IX. An Aggregate Productivity Growth Decomposition

In quantifying labor-augmenting and Hicks-neutral technological change in Sections VI and VIII, we leverage our firm-level panel data to follow individual firms over time. In this section, we complement our findings by analyzing the aggregate productivity of the Spanish manufacturing sector and its growth over our sample period. To obtain insight into the drivers of growth, we decompose aggregate productivity growth along the lines of Olley and Pakes (1996).

Aggregate productivity $\phi_t = \sum_j \mu_{jt} \phi_{jt}$ in period t is a weighted average of the productivity of individual firms, where ϕ_{jt} is a measure of the productivity of firm j in period t and μ_{jt} is its weight. We separately examine labor-augmenting productivity in output terms $\epsilon_{Ljt-2}\omega_{Ljt}$, Hicks-neutral productivity ω_{Hjt} , and total productivity $\epsilon_{Ljt-2}\omega_{Ljt} + \omega_{Hjt}$. Throughout the weight $\mu_{jt} = (p_{jt} + y_{jt}) / \sum_j (p_{jt} + y_{jt})$ is the share of the log of sales of firm j in period t .

The growth in aggregate productivity from period t_1 to period t_2 is $\Delta\phi = \phi_{t_2} - \phi_{t_1}$. Following Olley and Pakes (1996) and Melitz and Polanec (2015), we decompose this growth as

$$\begin{aligned}
\Delta\phi &= \Delta\phi^S + \mu_{t_2}^E(\phi_{t_2}^E - \phi_{t_2}^S) + \mu_{t_1}^X(\phi_{t_1}^S - \phi_{t_1}^X) \\
&= \Delta\bar{\phi}^S + N^S\Delta\text{Cov}\left(\frac{\mu}{\mu^S}, \phi\right) + \mu_{t_2}^E(\phi_{t_2}^E - \phi_{t_2}^S) + \mu_{t_1}^X(\phi_{t_1}^S - \phi_{t_1}^X),
\end{aligned} \tag{18}$$

where S , E , and X index the group of survivors, entrants, and exitors, respectively; $\mu_t^G = \sum_{j \in G} \mu_{jt}$ is the total weight of group G in period t ; $\phi_t^G = \sum_{j \in G} (\mu_{jt}/\mu_t^G) \phi_{jt}$ is the weighted average restricted to group G ; $\bar{\phi}_t^G = (1/N^G) \sum_{j \in G} \phi_{jt}$ is the unweighted average restricted to group G ; N^G is the number of firms in group G ; and $\Delta\text{Cov}(\mu/\mu^S, \phi)$ is short for

$$\text{Cov}\left(\frac{\mu_{jt_2}}{\mu_{t_2}^S}, \phi_{jt_2}\right) - \text{Cov}\left(\frac{\mu_{jt_1}}{\mu_{t_1}^S}, \phi_{jt_1}\right).$$

In the first line of the decomposition, the first term captures the contribution of survivors to aggregate productivity growth, the second that of entrants, and the third that of exitors. The second line further decomposes the contribution of survivors to aggregate productivity growth into a shift in the distribution of productivity (first term) and a change in covariance that captures reallocation (second term).

As the decomposition pertains to the population of firms, applying it to the sample of firms in our firm-level panel data is subject to a caveat. As before we account for the survey design by replicating the subsample of small firms. We classify a firm as a survivor if it enters the industry in or before period t_1 and does not exit in or before period t_2 . We further classify a firm as an entrant if it enters the industry after period t_1 and as an exitor if it exits the industry in or before period t_2 . Because of attrition and the periodic addition of new firms to the sample, we observe productivity for a subset of survivors in period t_1 and for another subset of survivors in period t_2 . Because we average over potentially quite different subsets of firms, especially if periods t_1 and t_2 are far apart, our estimates of the various terms in the decomposition may be noisy.

We report the change in aggregate productivity and its decomposition in table 9 for the period 1992–2006 and the three subperiods 1992–96, 1997–2001, and 2002–6. The change in aggregate productivity in column 1 in table 9 is consistent with our findings in Sections VI and VIII. Aggregate labor-augmenting productivity in output terms grew by 21.7 percent for the period 1992–2006 or about 1.6 percent per year. Aggregate Hicks-neutral productivity grew by 19.7 percent or about 1.4 percent per year and aggregate total productivity by 41.6 percent or about 3.0 percent per year. For the later subperiods, technological change appears to have slowed down, in particular, in the case of aggregate Hicks-neutral and total productivity.

We report the various terms of the decomposition in the first line of equation (18) in columns 2, 5, and 6 of table 9 and those in the second line in columns 3, 4, 5, and 6. Turning to columns 2, 5, and 6, survivors

TABLE 9
AGGREGATE PRODUCTIVITY GROWTH DECOMPOSITION

PERIOD t_1-t_2	CHANGE IN AGGREGATE PRODUCTIVITY ^{a,b} $\Delta\phi$ (1)	DECOMPOSITION ^c				
		Survivors			Entrants $\mu_{i_t}^E(\phi_{i_t}^E - \phi_{i_t}^S)$ (5)	Exitors $\mu_{i_t}^X(\phi_{i_t}^S - \phi_{i_t}^X)$ (6)
		Total	Shift	Covariance		
		$\Delta\phi^S$ (2)	$\Delta\bar{\phi}^S$ (3)	$N^S\Delta\text{Cov}(\cdot)$ (4)		
Labor-Augmenting Productivity in Output Terms $\epsilon_{L,-2}\omega_L$						
1992-2006	.217	.153	.133	.020	.024	.040
1992-96	.110	.092	.081	.011	.023	-.005
1997-2001	.063	.057	.057	.001	.007	-.002
2002-6	.056	.051	.058	-.007	.000	.004
Hicks-Neutral Productivity ω_H						
1992-2006	.197	.150	.151	-.001	.030	.017
1992-96	.098	.062	.060	.002	.036	.001
1997-2001	.051	.055	.055	-.001	-.007	.003
2002-6	-.003	-.012	-.005	-.007	.001	.008
Total Productivity $\epsilon_{L,-2}\omega_L + \omega_H$						
1992-2006	.416	.307	.283	.023	.059	.050
1992-96	.204	.163	.141	.021	.049	-.008
1997-2001	.095	.090	.087	.003	-.001	.006
2002-6	.055	.040	.057	-.017	.001	.013

^a We trim 1 percent of observations at each tail of the productivity distribution separately for survivors, entrants, and exitors but pooled across the start and end year.

^b Changes over subperiods do not add up because of trimming and because subperiods do not overlap. Without trimming, changes over subperiods almost add up for $\epsilon_{L-2}\omega_L$; with overlapping subperiods (1992–96, 1997–2001, and 2002–6), changes over subperiods almost add up for ω_H .

^c Columns 2, 5, and 6 correspond to the first line of eq. (18) and add up col. 1; cols. 3, 4, 5, and 6 correspond to the second line of eq. (18) and add up col. 1.

account for most of the change in aggregate total productivity and its components, again in line with our findings in Sections VI and VIII.³⁰ With the possible exception of entrants for the subperiod 1992–96, the contribution of entrants and exitors appears to be limited. Homing in on survivors and further decomposing their contribution, shifts in the distribution of productivity in column 3 are substantially more important than changes in the covariance in column 4. The contribution of reallocation to the change in aggregate total productivity and its components is sometimes positive and sometimes negative and mostly small.

X. Capital-Augmenting Technological Change

As discussed in Section III, the evolution of the relative quantities and prices of the various factors of production provides no evidence for capital-

³⁰ Survivors account for less of the change for the period 1992–2006 than for the three subperiods simply because the definition of survivor is more demanding if periods t_1 and t_2 are further apart.

augmenting technological change. Our leading specification therefore restricts the productivities of capital and materials to change at the same rate and in lockstep with Hicks-neutral technological change. A more general specification allows for capital-augmenting productivity ω_{Kjt} so that the production function in equation (6) becomes

$$Y_{jt} = \left\{ \beta_K [\exp(\omega_{Kjt}) K_{jt}]^{-\frac{1-\sigma}{\sigma}} + [\exp(\omega_{Ljt}) L_{jt}^*]^{-\frac{1-\sigma}{\sigma}} + \beta_M (M_{jt}^*)^{-\frac{1-\sigma}{\sigma}} \right\}^{-\frac{\sigma}{1-\sigma}} \exp(\omega_{Hjt}) \exp(e_{jt}). \quad (19)$$

We explore the role of capital-augmenting technological change in our data in two ways.

First, we follow Raval (2013) and parts of the previous literature on estimating aggregate production functions (see Antràs [2004] and the references therein) and assume that capital is a static input that is chosen each period to maximize short-run profits. In analogy to equation (10), we recover (conveniently rescaled) capital-augmenting productivity $\tilde{\omega}_{Kjt} = (1 - \sigma)\omega_{Kjt}$ as

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

where $\tilde{\gamma}_K = -\sigma \ln(\beta_M/\beta_K)$, and we use the user cost in our data as a rough measure of the price of capital P_{Kjt} . Using our leading estimates from Section VI, we recover the capital-augmenting productivity $\omega_{Kjt} = \tilde{\omega}_{Kjt}/(1 - \sigma)$ of firm j in period t up to an additive constant; in what follows, we use ω_{Kjt} to denote the demeaned capital-augmenting productivity.³¹ *Ceteris paribus*,

$$\Delta\omega_{Kjt} \approx \frac{\exp(\omega_{Kjt})K_{jt-1} - \exp(\omega_{Kjt-1})K_{jt-1}}{\exp(\omega_{Kjt-1})K_{jt-1}}$$

in column 1 of table 10 approximates the rate of growth of a firm's effective capital stock $\exp(\omega_{Kjt-1})K_{jt-1}$ and $\epsilon_{Kjt-2}\Delta\omega_{Kjt}$ in column 2 approximates the rate of growth of the firm's output Y_{jt-1} , where ϵ_{Kjt-2} is the elasticity of output with respect to the firm's effective capital stock (see app. E). As can be seen from column 1, capital-augmenting productivity grows slowly,

³¹ As an alternative to plugging our leading estimates from Sec. VI into eq. (20), in the supplementary appendix we use eq. (20) to form the analogue to our first estimation eq. (12):

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

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

TABLE 10
CAPITAL-AUGMENTING TECHNOLOGICAL CHANGE

INDUSTRY	$\Delta\omega_K$ (1)	$\epsilon_{K-2}\Delta\omega_K$ (2)	GMM				
			β_K (SE) (3)	ν (SE) (4)	δ_K (SE) (5)	χ^2 (df) (6)	p -Value (7)
1. Metals and metal products	.056	.004	.254 (.129)	.903 (.055)	.036 (.061)	2.555 (7)	.923
2. Nonmetallic minerals	-.010	.007	.236 (.102)	.906 (.072)	.010 (.072)	3.979 (7)	.782
3. Chemical products	-.018	.001	.125 (.068)	.942 (.041)	-.031 (.092)	.598 (7)	.999
4. Agricultural and industrial machinery	-.020	.000	.182 (.177)	.801 (.081)	.031 (.122)	9.026 (8)	.340
5. Electrical goods	-.078	.000	.129 (.041)	.845 (.054)	-.004 (.056)	2.493 (9)	.981
6. Transport equipment ^a	.008	.005	.115 (.088)	.981 (.050)	-.143 (.138)		
7. Food, drink, and tobacco	-.005	.002	.282 (.286)	.918 (.058)	-.045 (.204)	2.279 (7)	.943
8. Textile, leather, and shoes	-.085	-.002	.080 (.143)	.971 (.047)	.053 (.135)	2.714 (8)	.951
9. Timber and furniture ^a	.035 ^b	.000 ^b	.088 (.119)	.924 (.067)	-.021 (.059)		
10. Paper and printing products	.022	.007	.229 (.089)	.935 (.033)	.005 (.045)	3.066 (7)	.879
All industries	-.038	.001					

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

^b We trim values of $\Delta\omega_K$, respectively, $\epsilon_{K-2}\Delta\omega_K$, below -0.5 and above 0.5 .

on average, with rates of growth of 0.8 percent per year in industry 6, 2.2 percent in industry 10, and 5.6 percent in industry 1. The rate of growth is negative in the remaining seven industries. The growth of capital-augmenting productivity is especially underwhelming in comparison to the growth of labor-augmenting productivity (see again col. 1 of table 5). The output effect of capital-augmenting technological change in column 2 is also close to zero in all industries, although this likely reflects the fact that capital is not a static input. As the user cost excludes adjustment costs, it falls short of the shadow price of capital, and using it drives down the elasticity of output with respect to the firm's effective capital stock.

Second, we return to the usual setting in the literature following Olley and Pakes (1996) and allow the choice of capital to have dynamic implications. We follow parts of the previous literature in estimating aggregate

production functions and proxy for ω_{Kjt} by a time trend $\delta_K t$. Our second estimation equation (15) remains unchanged except that

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

Columns 3–7 of table 10 summarize the resulting estimates of β_K , v , and δ_K . The estimates of β_K and v are very comparable to those in table 5. Moreover, the insignificant time trend leaves little room for capital-augmenting technological change in our data.

In sum, in line with the patterns in the data described in Section III, there is little, if any, evidence for capital-augmenting technological change in our data. Of course, our ways of exploring the role of capital-augmenting technological change are less than ideal in that they either rest on the assumption that capital is a static input or abstract from firm-level heterogeneity in capital-augmenting productivity. An important question is therefore whether our approach can be extended to treat capital-augmenting productivity on par with labor-augmenting and Hicks-neutral productivity.

Recovering a third component of productivity, at a bare minimum, requires a third decision besides labor and materials to invert. Investment is a natural candidate. In contrast to the demand for labor and materials, however, investment depends on the details of the firm's dynamic programming problem. There are two principal difficulties. First, one has to prove that the observed demands for labor and materials along with investment are jointly invertible for unobserved capital-augmenting, labor-augmenting, and Hicks-neutral productivity. Second, the inverse functions $\tilde{h}_K(\cdot)$, $\tilde{h}_L(\cdot)$, and $\tilde{h}_H(\cdot)$ are high-dimensional. Thus, estimating these functions nonparametrically is demanding on the data. In ongoing work, Zhang (2015) proposes combining a parametric inversion that exploits the parameter restrictions between production and input demand functions similarly to our paper with a nonparametric inversion of investment similarly to Olley and Pakes (1996).

XI. Conclusions

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

To this end, we develop a dynamic model of the firm in which productivity is multidimensional. At the center of the model is a CES produc-

tion function that parsimoniously yet robustly relates the relative quantities of materials and labor to their relative prices and labor-augmenting productivity. To properly isolate and measure labor-augmenting productivity, we account for other factors that affect this relationship, in particular, outsourcing and adjustment costs on permanent labor.

We apply our estimator to an unbalanced panel of 2,375 Spanish manufacturing firms in 10 industries from 1990 to 2006. Our estimates indicate limited substitutability between the various factors of production. This calls into question whether the widely used Cobb-Douglas production function with its unitary elasticity of substitution adequately represents firm-level production processes.

Our estimates provide clear evidence that technological change is biased. *Ceteris paribus*, labor-augmenting technological change causes output to grow, on average, in the vicinity of 1.5 percent per year. While skill upgrading explains some of the growth of labor-augmenting productivity, in many industries labor-augmenting productivity grows because workers with a given set of skills become more productive over time. In short, our estimates cast doubt on the assumption of Hicks-neutral technological change that underlies many of the standard techniques for measuring productivity and estimating production functions.

At the same time, however, our estimates do not validate the assumption that technological change is purely labor augmenting that plays a central role in the literature on economic growth. In addition to labor-augmenting technological change, our estimates show that Hicks-neutral technological change causes output to grow, on average, in the vicinity of 1.5 percent per year.

While we are primarily interested in measuring how much of technological change is labor augmenting and how much of it is Hicks neutral, we also use our estimates to illustrate the consequences of biased technological change beyond the growth of output. In particular, we show that it is the primary driver of the decline of the aggregate share of labor in the Spanish manufacturing sector over our sample period. An interesting avenue for future research is to investigate the implications of biased technological change for employment. Recent research points to biased technological change as a key driver of the diverging experiences of the continental European, US, and UK economies during the 1980s and 1990s (Blanchard 1997; Caballero and Hammour 1998; Bentolila and Saint-Paul 2004; McAdam and Willman 2013). Our estimates lend themselves to decomposing firm-level changes in employment into displacement, substitution, and output effects and to compare these effects between labor-augmenting and Hicks-neutral technological change. This may be helpful for better understanding and predicting the evolution of employment as well as for designing labor market and innovation policies in the presence of biased technological change.

Appendix A

Proof of Proposition 1

Rewriting the ratio of first-order conditions (3) yields

$$\begin{aligned}
 0 &= \ln \frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial L_{jt}} + \omega_{Ljt} \\
 &\quad - \ln \frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial M_{jt}} + p_{Mjt} - w_{jt} \\
 &= f(m_{jt} - l_{jt}, p_{Mjt} - w_{jt}, \omega_{Ljt}).
 \end{aligned}$$

Differentiating the so-defined function $f(\cdot)$ yields

$$\begin{aligned}
 &\frac{\partial f(m_{jt} - l_{jt}, p_{Mjt} - w_{jt}, \omega_{Ljt})}{\partial (m_{jt} - l_{jt})} \\
 &= \left[-\frac{\frac{\partial^2 H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial L_{jt}^2}}{\frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial L_{jt}}} + \frac{\frac{\partial^2 H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial M_{jt} \partial L_{jt}}}{\frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial M_{jt}}} \right] \exp(\omega_{Ljt} - (m_{jt} - l_{jt})) \\
 &= \frac{H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1) \frac{\partial^2 H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial M_{jt} \partial L_{jt}}}{\frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial L_{jt}} \frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial M_{jt}}} = \frac{1}{\sigma(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})))},
 \end{aligned}$$

where the second equality uses that $H(\exp(\omega_{Ljt})L_{jt}, M_{jt})$ is homogeneous of degree one and the third equality uses that the elasticity of substitution between materials and labor (Chambers 1988, eq. 1.13) for the production function in equation (1) simplifies to

$$\sigma(\exp(\omega_{Ljt} - (m_{jt} - l_{jt}))) = \frac{\frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial L_{jt}} \frac{\partial H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial M_{jt}}}{H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1) \frac{\partial^2 H(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})), 1)}{\partial M_{jt} \partial L_{jt}}}.$$

Similarly,

$$\begin{aligned}
 &\frac{\partial f(m_{jt} - l_{jt}, p_{Mjt} - w_{jt}, \omega_{Ljt})}{\partial (p_{Mjt} - w_{jt})} = 1, \\
 &\frac{\partial f(m_{jt} - l_{jt}, p_{Mjt} - w_{jt}, \omega_{Ljt})}{\partial \omega_{Ljt}} = -\frac{1}{\sigma(\exp(\omega_{Ljt} - (m_{jt} - l_{jt})))} + 1.
 \end{aligned}$$

By the implicit function theorem, around a point $(m_{jt}^0 - l_{jt}^0, p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)$ satisfying $f(m_{jt}^0 - l_{jt}^0, p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0) = 0$, there exists a continuously differentiable

function $m_{jt} - l_{jt} = g(p_{Mjt} - w_{jt}, \omega_{Ljt})$ such that $f(g(p_{Mjt} - w_{jt}, \omega_{Ljt}), p_{Mjt} - w_{jt}, \omega_{Ljt}) = 0$ and

$$\frac{\partial g(p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)}{\partial (p_{Mjt} - w_{jt})} = - \frac{\frac{\partial f(m_{jt}^0 - l_{jt}^0, p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)}{\partial (p_{Mjt} - w_{jt})}}{\frac{\partial f(m_{jt}^0 - l_{jt}^0, p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)}{\partial (m_{jt} - l_{jt})}} = -\sigma(\exp(\omega_{Ljt}^0 - (m_{jt}^0 - l_{jt}^0))),$$

$$\frac{\partial g(p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)}{\partial \omega_{Ljt}} = - \frac{\frac{\partial f(m_{jt}^0 - l_{jt}^0, p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)}{\partial \omega_{Ljt}}}{\frac{\partial f(m_{jt}^0 - l_{jt}^0, p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)}{\partial (m_{jt} - l_{jt})}} = 1 - \sigma(\exp(\omega_{Ljt}^0 - (m_{jt}^0 - l_{jt}^0))).$$

The first-order Taylor series for $m_{jt} - l_{jt} = g(p_{Mjt} - w_{jt}, \omega_{Ljt})$ around the point $(m_{jt}^0 - l_{jt}^0, p_{Mjt}^0 - w_{jt}^0, \omega_{Ljt}^0)$ follows immediately.

Appendix B

Data

We observe firms for a maximum of 17 years between 1990 and 2006. We restrict the sample to firms with at least 3 years of data on all variables required for estimation. The number of firms with 3, 4, ..., 17 years of data is 313, 240, 218, 215, 207, 171, 116, 189, 130, 89, 104, 57, 72, 94, and 160, respectively. Table B1 gives the industry definitions along with their equivalent definitions in terms of the ESEE, National Accounts, and International Standard Industrial Classifications (ISIC; cols. 1–3). On the basis of the National Accounts in 2000, we further report the shares of the various industries in the total value added of the manufacturing sector (col. 4).

In what follows we define the variables we use for our main analysis.

- *Investment.* Value of current investments in equipment goods (excluding buildings, land, and financial assets) deflated by the price index of investment. The price of investment is the equipment goods component of the index of industry prices computed and published by the Spanish Ministry of Industry.
- *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_{jt} as

$$\tilde{K}_{jt} = (1 - \delta) \frac{P_{jt}}{P_{j,t-1}} \tilde{K}_{j,t-1} + \tilde{I}_{j,t-1},$$

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 $K_{jt} = \tilde{K}_{jt}/P_{jt}$.

- *Labor.* Total hours worked computed as the number of workers times the average hours per worker, where the latter is computed as normal hours plus average overtime minus average working time lost at the workplace.
- *Materials.* Value of intermediate goods consumption (including raw materials, components, energy, and services) deflated by a firm-specific price index of materials.

- *Output*. Value of produced goods and services computed as sales plus the variation of inventories deflated by a firm-specific price index of output.
- *Wage*. Hourly wage cost computed as total labor cost including social security payments divided by total hours worked.
- *Price of materials*. Firm-specific price index for intermediate consumption. Firms are asked about the price changes that occurred during the year for raw materials, components, energy, and services. The price index is computed as a Paasche-type index of the responses.
- *Price of output*. Firm-specific price index for output. Firms are asked about the price changes they made during the year in up to five separate markets in which they operate. The price index is computed as a Paasche-type index of the responses.
- *Demand shifter*. Firms are asked to assess the current and future situation of the main market in which they operate. The demand shifter codes the responses as 0, 0.5, and 1 for slump, stability, and expansion, respectively.
- *Share of temporary labor*. Fraction of workers with fixed-term contracts and no or small severance pay.
- *Share of outsourcing*. Fraction of customized parts and pieces that are manufactured by other firms in the value of the firm's intermediate goods purchases.
- *R&D expenditures*. R&D expenditures include the cost of intramural R&D activities, payments for outside R&D contracts with laboratories and research centers, and payments for imported technology in the form of patent licensing or technical assistance, with the various expenditures defined according to the OECD Oslo and Frascati manuals.

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

- *User cost of capital*. Computed as $P_{it}(r_{jt} + \delta - \text{CPI}_t)$, where P_{it} is the price index of investment, r_{jt} is a firm-specific interest rate, δ is an industry-specific estimate of the rate of depreciation, and CPI_t is the rate of inflation as measured by the consumer price index.
- *Skill mix*. Fraction of nonproduction employees (white-collar workers), workers with an engineering degree (engineers), and workers with an intermediate degree (technicians). Available in the year a firm enters the sample and every subsequent 4 years; assumed to be unchanging in the interim.
- *Region*. Dummy variables corresponding to the 19 Spanish autonomous communities and cities where employment is located if it is located in a unique region and another dummy variable indicating that employment is spread over several regions.
- *Product submarket*. Dummy variables corresponding to a finer breakdown of the 10 industries into subindustries (restricted to subindustries with at least five firms; see col. 5 of table B1).
- *Technological sophistication*. Dummy variable that takes the value one if the firm uses digitally controlled machines, robots, CAD/CAM, or some combination of these procedures.

- *Identification between ownership and control.* Dummy variable that takes the value one if the owner of the firm or the family of the owner hold management positions.
- *Age.* Years elapsed since the foundation of the firm with a maximum of 40 years.
- *Firm size.* Number of workers in the year the firm enters the sample.

TABLE B1
INDUSTRY DEFINITIONS AND EQUIVALENT CLASSIFICATIONS

INDUSTRY	CLASSIFICATIONS			SHARE OF VALUE ADDED (4)	NUMBER OF SUBINDUSTRIES (5)
	ESSE (1)	National Accounts (2)	ISIC (Rev. 4) (3)		
1. Ferrous and nonferrous metals and metal products	12+13	DJ	C 24+25	13.2	11
2. Nonmetallic minerals	11	DI	C 23	8.2	8
3. Chemical products	9+10	DG–DH	C 20+21+22	13.9	7
4. Agricultural and industrial machinery	14	DK	C 28	7.1	7
5. Electrical goods	15+16	DL	C 26+27	7.5	13
6. Transport equipment	17+18	DM	C 29+30	11.6	7
7. Food, drink, and tobacco	1+2+3	DA	C 10+11+12	14.5	10
8. Textile, leather, and shoes	4+5	DB–DC DD–	C 13+14+15	7.6	11
9. Timber and furniture	6+19	DN 38	C 16+31	7.0	6
10. Paper and printing products	7+8	DE	C 17+18	8.9	4
All industries				99.5	84

Appendix C

Inverse Functions

The first-order conditions for permanent and temporary labor are

$$\begin{aligned}
 \nu \mu X_{jt}^{-(1+\frac{\sigma}{1-\sigma})} \exp(\omega_{Hjt}) \exp\left(-\frac{1-\sigma}{\sigma} \omega_{Ljt}\right) (L_{jt}^*)^{-\frac{1}{\sigma}} \frac{\partial L_{jt}^*}{\partial L_{Pjt}} \\
 = \frac{W_{Pjt}(1 + \Delta_{jt})}{P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}, \tag{C1}
 \end{aligned}$$

$$\begin{aligned}
 \nu \mu X_{jt}^{-(1+\frac{\sigma}{1-\sigma})} \exp(\omega_{Hjt}) \exp\left(-\frac{1-\sigma}{\sigma} \omega_{Ljt}\right) (L_{jt}^*)^{-\frac{1}{\sigma}} \frac{\partial L_{jt}^*}{\partial L_{Tjt}} \\
 = \frac{W_{Tjt}}{P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}, \tag{C2}
 \end{aligned}$$

where, by the envelope theorem, the gap between the wage of permanent workers W_{Pjt} and the shadow wage is

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

Equations (C1) and (C2) allow the mix of permanent and temporary labor to depend on the firm's productivity and the other state variables (through Δ_{jt}).

Our assumption that $\Lambda(L_{Pjt}, L_{Tjt})$ is linearly homogeneous implies $L_{jt}^* = L_{jt}\Lambda(1 - S_{Tjt}, S_{Tjt})$, $\partial L_{jt}^* / \partial L_{Pjt} = \Lambda_P(1 - S_{Tjt}, S_{Tjt})$, and $\partial L_{jt}^* / \partial L_{Tjt} = \Lambda_T(1 - S_{Tjt}, S_{Tjt})$. Using Euler's theorem to combine equations (C1) and (C2) yields

$$\begin{aligned}\nu \mu X_{jt}^{-(1+\frac{\sigma}{1-\sigma})} \exp(\omega_{Hjt}) \exp\left(-\frac{1-\sigma}{\sigma} \omega_{Ljt}\right) L_{jt}^{-\frac{1}{\sigma}} \Lambda(1 - S_{Tjt}, S_{Tjt})^{-\frac{1-\sigma}{\sigma}} \\ = \frac{W_{jt} \left(1 + \frac{\Delta_{jt}}{1 + \frac{W_{Tjt}}{W_{Pjt}} \frac{S_{Tjt}}{1-S_{Tjt}}}\right)}{P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)} = \frac{W_{jt} \left(\frac{\Lambda_P(1-S_{Tjt}, S_{Tjt})}{\Lambda_T(1-S_{Tjt}, S_{Tjt})} + \frac{S_{Tjt}}{1-S_{Tjt}}\right)}{P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)},\end{aligned}\quad (C3)$$

where the second equality follows from dividing equations (C1) and (C2) and solving for Δ_{jt} . Our assumption that $W_{Pjt}/W_{Tjt} = \lambda_0$ is an (unknown) constant implies that

$$\frac{\frac{\Lambda_P(1-S_{Tjt}, S_{Tjt})}{\Lambda_T(1-S_{Tjt}, S_{Tjt})} + \frac{S_{Tjt}}{1-S_{Tjt}}}{\lambda_0 + \frac{S_{Tjt}}{1-S_{Tjt}}} = \lambda_1(S_{Tjt})$$

as an (unknown) function of S_{Tjt} .

Turning from the labor to the materials decision, because the firm must maintain the ratio of outsourced to in-house materials Q_{Mjt} , the first-order condition for in-house materials is

$$\nu \beta_M \mu X_{jt}^{-(1+\frac{\sigma}{1-\sigma})} \exp(\omega_{Hjt}) (M_{jt}^*)^{-\frac{1}{\sigma}} \frac{dM_{jt}^*}{dM_{Ijt}} = \frac{P_{Mjt}}{P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}, \quad (C4)$$

where $P_{Mjt} = P_{jt} + P_{Ojt} Q_{Mjt}$ is the effective cost of an additional unit of in-house materials.

Our assumption that $\Gamma(M_{Ijt}, M_{Ojt})$ is linearly homogeneous implies

$$M_{jt}^* = M_{Ijt} \Gamma\left(1, \frac{P_{Ijt}}{P_{Ojt}} \frac{S_{Ojt}}{1 - S_{Ojt}}\right)$$

and

$$\frac{dM_{jt}^*}{dM_{jt}} = \Gamma\left(1, \frac{P_{jt}}{P_{ojt}} \frac{S_{ojt}}{1 - S_{ojt}}\right).$$

Rewriting equation (C4) yields

$$\begin{aligned} v\beta_M \mu X_{jt}^{-(1+\frac{\mu}{1-\sigma})} \exp(\omega_{Hjt}) M_{jt}^{-\frac{1}{\sigma}} \Gamma\left(1, \frac{P_{jt}}{P_{ojt}} \frac{S_{ojt}}{1 - S_{ojt}}\right)^{-\frac{1-\sigma}{\sigma}} \\ = \frac{P_{Mjt}}{P_{jt} \left(1 - \frac{1}{\eta(p_{jt}, D_{jt})}\right)}. \end{aligned} \quad (C5)$$

Our assumption that $P_{jt}/P_{ojt} = \gamma_0$ is an (unknown) constant implies that

$$\ln \Gamma\left(1, \gamma_0 \frac{S_{ojt}}{1 - S_{ojt}}\right) = \gamma_1(S_{ojt})$$

as an (unknown) function of S_{ojt} .³²

Solving equations (C3) and (C5) for $\tilde{\omega}_{Ljt} = (1 - \sigma)\omega_{Ljt}$ and ω_{Hjt} yields equations (10) and (11), where

$$\lambda_2(S_{Tjt}) = \ln\left(\lambda_1(S_{Tjt})\Lambda(1 - S_{Tjt})^{\frac{1-\sigma}{\sigma}}\right).$$

Appendix D

Estimation

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

Starting with the conditional expectation functions, we specify $\tilde{g}_{L1}(\tilde{h}_L(\cdot)) = q(\tilde{h}_L(\cdot) - \tilde{\gamma}_L)$, $\tilde{g}_{L2}(\tilde{h}_L(\cdot), r_{jt}) = q_0 + q(\tilde{h}_L(\cdot) - \tilde{\gamma}_L, r_{jt})$, $g_{H1}(h_H(\cdot)) = q(h_H(\cdot) - \gamma_H)$, and $g_{H2}(h_H(\cdot), r_{jt}) = q_0 + q(h_H(\cdot) - \gamma_H, r_{jt})$, where q_0 is a constant and the function $q(\cdot)$ is modeled as described above. Without loss of generality, we absorb $\tilde{\gamma}_L$ and γ_H into the overall constants of our estimation equations. Turning to the absolute value of the price elasticity, to impose the theoretical restriction $\eta(p_{jt}, D_{jt}) > 1$, we specify $\eta(p_{jt}, D_{jt}) = 1 + \exp(q(p_{jt}, D_{jt}))$, where the function $q(\cdot)$ is modeled as described above except that we suppress terms involving D_{jt}^2 and D_{jt}^3 . Turning to the correction terms, we specify $\lambda_1(S_{Tjt}) = q(\ln S_{Tjt})$ and $\lambda_2(S_{Tjt}) =$

³² Equation (C5) presumes an interior solution for in-house materials; it is consistent with a corner solution for outsourced materials. Indeed, without outsourcing eq. (C5) reduces to the first-order condition for in-house materials.

$q(\ln S_{Tjt})$ in industries 2, 3, and 10 and $\lambda_1(S_{Tjt}) = q(\ln(1 - S_{Tjt}))$ and $\lambda_2(S_{Tjt}) = q(\ln(1 - S_{Tjt}))$ in the remaining industries.³³ Finally, we specify $\gamma_1(S_{Ojt}) = q(S_{Ojt})$; this ensures that $\gamma_1(S_{Ojt}) = 0$ if $S_{Ojt} = 0$ in line with the normalization $\Gamma(M_{jt}, 0) = M_{jt}$.

Parameters and instruments.—Our first estimation equation (12) has 36 parameters: the constant, σ , 15 parameters in $\tilde{g}_{L,0}(t-1)$ (time dummies), three parameters in $\tilde{g}_{L,1}(\tilde{h}_L(\cdot))$, 10 parameters in $\tilde{g}_{L,2}(\tilde{h}_L(\cdot), r_{jt-1})$, three parameters in $\lambda_2(S_{Tjt})$, and three parameters in $\gamma_1(S_{Ojt})$.

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

After imposing the estimated values from equation (12), our second estimation equation (15) has 40 parameters: the constant, β_K , v , 15 parameters in $g_{H0}(t-1)$ (time dummies), three parameters in $g_{H1}(h_H(\cdot))$, 10 parameters in $g_{H2}(h_H(\cdot), r_{jt-1})$, three parameters in $\lambda_1(S_{Tjt})$, and six parameters in $\eta(P_{jt}, D_{jt})$.

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

Estimation.—From the GMM problem in equation (16) with weighting matrix

$$\hat{W}_L = \left[\frac{1}{N} \sum_j A_{Lj}(z_j) A_{Lj}(z_j)' \right]^{-1},$$

³³ To incorporate skill upgrading, we instead specify $\lambda_1(S_{Tjt}, \Theta_{jt}) = q(\ln S_{Tjt}, \ln \Theta_{jt})$ and $\lambda_2(S_{Tjt}, \Theta_{jt}) = q(\ln S_{Tjt}, \ln \Theta_{jt})$ in industries 2, 3, and 10 and $\lambda_1(S_{Tjt}, \Theta_{jt}) = q(\ln(1 - S_{Tjt}), \ln \Theta_{jt})$ and $\lambda_2(S_{Tjt}, \Theta_{jt}) = q(\ln(1 - S_{Tjt}), \ln \Theta_{jt})$ in the remaining industries, where the function $q(\cdot)$ is modeled as described above except that we suppress terms involving $(\ln \Theta_{jt})^2$ and $(\ln \Theta_{jt})^3$.

we first obtain a consistent estimate $\hat{\theta}_L$ of θ_L . This first step is the nonlinear two-stage least-squares estimator of Amemiya (1974). In the second step, we compute the optimal estimate with weighting matrix

$$\hat{W}_L = \left[\frac{1}{N} \sum_j A_{Lj}(z_j) \nu_{Lj}(\hat{\theta}_L) \nu_{Lj}(\hat{\theta}_L)' A_{Lj}(z_j)' \right]^{-1}.$$

Throughout the paper, we report standard errors that are robust to heteroskedasticity and autocorrelation. We further correct standard errors as described in the supplementary appendix to reflect the fact that our estimates of equation (15) are conditional on those of equation (12).

Implementation.—We use Gauss 14.0.9 and Optmum 3.1.7. To reduce the number of parameters to search over in the GMM problem in equation (16), we “concentrate out” the parameters that enter it linearly as described in the supplementary appendix. To guard against local minima, we have extensively searched over the remaining parameters, often using preliminary estimates to narrow down the range of these parameters.

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

Appendix E

Output Effect

Direct calculation starting from equation (6) yields the elasticity of output with respect to a firm’s effective labor force:

$$\begin{aligned} \epsilon_{Ljt} &= \frac{\partial Y_{jt}}{\partial \exp(\omega_{Ljt}) L_{jt}^*} \frac{\exp(\omega_{Ljt}) L_{jt}^*}{Y_{jt}} \\ &= \frac{\nu [\exp(\omega_{Ljt}) L_{jt}^*]^{-\frac{1-\sigma}{\sigma}}}{\beta_K K_{jt}^{-\frac{1-\sigma}{\sigma}} + [\exp(\omega_{Ljt}) L_{jt}^*]^{-\frac{1-\sigma}{\sigma}} + \beta_M (M_{jt}^*)^{-\frac{1-\sigma}{\sigma}}}. \end{aligned} \quad (\text{E1})$$

Using equation (10) to substitute for ω_{Ljt} and simplifying, we obtain

$$\epsilon_{Ljt} = \frac{\nu^{\frac{1-S_{Mjt}}{S_{Mjt}}} \lambda_1(S_{Tjt})}{\frac{\beta_K}{\beta_M} \left[\frac{K_{jt}}{M_{jt} \exp(\gamma_1(S_{Ojt}))} \right]^{-\frac{1-\sigma}{\sigma}} + \frac{1-S_{Mjt}}{S_{Mjt}} \lambda_1(S_{Tjt}) + 1}. \quad (\text{E2})$$

Recall from equation (C3) that

$$\lambda_1(S_{Tjt}) = 1 + \frac{\Delta_{jt}}{1 + \frac{W_{Tjt}}{W_{Ojt}} \frac{S_{Tjt}}{1-S_{Tjt}}},$$

where Δ_{jt} is the gap between the wage of permanent workers W_{pjt} and the shadow wage. To facilitate evaluating equation (E2), we abstract from adjustment costs and set $\lambda_1(S_{Tjt}) = 1$.

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

$$\begin{aligned} \epsilon_{Kjt} &= \frac{\partial Y_{jt}}{\partial \exp(\omega_{Kjt}) K_{jt}} \frac{\exp(\omega_{Kjt}) K_{jt}}{Y_{jt}} \\ &= \frac{\nu [\exp(\omega_{Kjt}) K_{jt}]^{-\frac{1-\sigma}{\sigma}}}{[\exp(\omega_{Kjt}) K_{jt}]^{-\frac{1-\sigma}{\sigma}} + [\exp(\omega_{Ljt}) L_{jt}^*]^{-\frac{1-\sigma}{\sigma}} + \beta_M (M_{jt}^*)^{-\frac{1-\sigma}{\sigma}}} \\ &= \frac{\nu}{1 + \frac{P_{Mjt} M_{jt}}{P_{Kjt} K_{jt}} \left[\frac{1-S_{Mjt}}{S_{Mjt}} \lambda_1(S_{Tjt}) + 1 \right]}, \end{aligned} \quad (E3)$$

where we use equations (10) and (20) to substitute for ω_{Ljt} and ω_{Kjt} , respectively. As with equation (E2), we set $\lambda_1(S_{Tjt}) = 1$ to evaluate equation (E3).

Appendix F

Second-Order Approximation

Let

$$\Upsilon_{jt} = -\tilde{\gamma}_L + (1 - \sigma)(p_{Mjt} - w_{jt} + \omega_{Ljt}) + \sigma\lambda_2(S_{Tjt}) - (1 - \sigma)\gamma_1(S_{Ojt})$$

and $\Delta\Upsilon_{jt} = \Upsilon_{jt} - \Upsilon_{jt-1}$. Using equation (17) we write

$$\begin{aligned} S_{Ljt} - S_{Ljt-1} &= -S_{Ljt}(1 - S_{Ljt-1})[\exp(\Delta\Upsilon_{jt}) - 1] \\ &\approx -S_{Ljt}(1 - S_{Ljt-1}) \left[\Delta\Upsilon_{jt} + \frac{1}{2} (\Delta\Upsilon_{jt})^2 \right], \end{aligned}$$

where we replace $\exp(\Delta\Upsilon_{jt}) - 1$ by its second-order Taylor series approximation around $\Delta\Upsilon_{jt} = 0$. We allocate the interactions in $(\Delta\Upsilon_{jt})^2$ in equal parts to the variables involved.

References

- Acemoglu, Daron. 2002. "Directed Technical Change." *Rev. Econ. Studies* 69:781–809.
- . 2003. "Labor- and Capital-Augmenting Technical Change." *J. European Econ. Assoc.* 1 (1): 1–37.
- Akerberg, Daniel, C. Lanier Benkard, Steve Berry, and Ariel Pakes. 2007. "Econometric Tools for Analyzing Market Outcomes." In *Handbook of Econometrics*, vol. 6A, edited by James J. Heckman and Edward E. Leamer, 4171–4276. Amsterdam: North-Holland.
- Akerberg, Daniel, Kevin W. Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica* 83 (6): 2411–51.
- Amemiya, Takeshi. 1974. "The Nonlinear Two-Stage Least-Squares Estimator." *J. Econ. Literature* 2:105–10.

- Antràs, Pol. 2004. "Is the U.S. Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution." *Contributions to Macroeconomics* 4 (1).
- Atalay, Enghin. 2014. "Materials Prices and Productivity." *J. European Econ. Assoc.* 12 (3): 575–611.
- Balat, Jorge F., Irene Brambilla, and Yuya Sasaki. 2015. "Heterogeneous Firms: Skilled-Labor Productivity and Export Destinations." Working paper, Johns Hopkins Univ.
- Bartelsman, Eric J., and Mark E. Doms. 2000. "Understanding Productivity: Lessons from Longitudinal Microdata." *J. Econ. Literature* 38 (3): 569–94.
- Bentolila, Samuel, and Gilles Saint-Paul. 2004. "Explaining Movements in the Labor Share." *Contributions to Macroeconomics* 3 (1).
- Bernard, Andrew B., Steven J. Redding, and Peter K. Schott. 2010. "Multiple-Product Firms and Product Switching." *A.E.R.* 100 (1): 70–97.
- Binswanger, Hans P. 1974. "The Measurement of Technical Change Biases with Many Factors of Production." *A.E.R.* 64 (6): 964–76.
- Blanchard, Olivier. 1997. "The Medium Run." *Brookings Papers Econ. Activity: Macroeconomics* (2): 89–158.
- Brown, Murray, and John S. de Cani. 1963. "Technological Change and the Distribution of Income." *Internat. Econ. Rev.* 4 (3): 289–309.
- Bruno, Michael. 1984. "Raw Materials, Profits, and the Productivity Slowdown." *Q.J.E.* 99 (1): 1–30.
- Caballero, Ricardo J., and Mohamad L. Hammour. 1998. "Jobless Growth: Appropriability, Factor Substitution, and Unemployment." *Carnegie-Rochester Conf. Ser. Public Policy* 48:51–94.
- Cain, Louis P., and Donald G. Paterson. 1986. "Biased Technical Change, Scale, and Factor Substitution in American Industry, 1850–1919." *J. Econ. Hist.* 46 (1): 153–64.
- Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *J. Labor Econ.* 20 (4): 733–83.
- Chambers, Robert G. 1988. *Applied Production Analysis: A Dual Approach*. Cambridge: Cambridge Univ. Press.
- Chirinko, Robert S. 2008. " σ : The Long and Short of It." *J. Multivariate Analysis* 30:671–86.
- David, Paul A., and Theo van de Klundert. 1965. "Biased Efficiency Growth and Capital-Labor Substitution in the U.S., 1899–1960." *A.E.R.* 55 (3): 357–94.
- Davis, Steven J., and John C. Haltiwanger. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Q.J.E.* 107 (3): 819–63.
- De Loecker, Jan. 2013. "Detecting Learning by Exporting." *American Econ. J.: Microeconomics* 5 (3): 1–21.
- De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. 2016. "Prices, Markups and Trade Reform." *Econometrica* 84 (2): 445–510.
- Dolado, Juan J., Carlos Garcia-Serrano, and Juan F. Jimeno. 2002. "Drawing Lessons from the Boom of Temporary Jobs in Spain." *Econ. J.* 112 (480): 270–95.
- Doraszelski, Ulrich, and Jordi Jaumandreu. 2013. "R&D and Productivity: Estimating Endogenous Productivity." *Rev. Econ. Studies* 80 (4): 1338–83.
- Dunne, Timothy, Mark Roberts, and Larry J. Samuelson. 1988. "Patterns of Firm Entry and Exit in U.S. Manufacturing." *Rand J. Econ.* 19 (4): 495–515.
- Eslava, Marcela, John C. Haltiwanger, Adriana Kugler, and Maurice Kugler. 2004. "The Effects of Structural Reforms on Productivity and Profitability Enhancing Reallocation: Evidence from Colombia." *J. Development Econ.* 75:333–71.

- European Commission. 2001. "Statistics on Innovation in Europe." Directorate General Enterprise, Brussels.
- Foster, Lucia, John C. Haltiwanger, and Chad Syverson. 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *A.E.R.* 98 (1): 394–425.
- . 2016. "The Slow Growth of New Plants: Learning about Demand?" *Economica* 83 (329): 91–129.
- Fox, Jeremy T., Valentin Haddad, Stefan Hoderlein, Amil K. Petrin, and Robert P. Sherman. 2016. Heterogeneous Production Functions, Panel Data, and Productivity Dispersion. Slide deck, Rice Univ.
- Gandhi, Amit, Salvador Navarro, and David A. Rivers. 2013. "On the Identification of Production Functions: How Heterogeneous Is Productivity?" Working paper, Univ. Wisconsin.
- Goldberg, Pinelopi K., Amit K. Khandelwal, Nina Pavcnik, and Petia B. Topalova. 2010. "Multiproduct Firms and Product Turnover in the Developing World: Evidence from India." *Rev. Econ. and Statis.* 92 (4): 1042–49.
- Goldman, Steven M., and Hirofumi Uzawa. 1964. "A Note on Separability in Demand Analysis." *Econometrica* 32 (3): 387–98.
- Grieco, Paul L., Shengyu Li, and Hongsong Zhang. 2016. "Production Function Estimation with Unobserved Input Price Dispersion." *Internat. Econ. Rev.* 57 (2): 665–90.
- Griliches, Zvi. 1964. "Research Expenditures, Education and the Aggregate Agricultural Production Function." *A.E.R.* 54 (6): 961–74.
- . 1998. *R&D and Productivity: The Econometric Evidence*. Chicago: Univ. Chicago Press.
- Griliches, Zvi, and Dale W. Jorgenson. 1967. "The Explanation of Productivity Change." *Rev. Econ. Studies* 34 (3): 249–83.
- Grossman, Gene M., and Elhanan Helpman. 2002. "Integration versus Outsourcing in Industry Equilibrium." *Q.J.E.* 117 (1): 85–120.
- . 2005. "Outsourcing in a Global Economy." *Rev. Econ. Studies* 72 (1): 135–59.
- Hall, Robert E. 1988. "The Relation between Price and Marginal Cost in U.S. Industry." *J.P.E.* 96 (5): 921–47.
- Hamermesh, Daniel S. 1993. *Labor Demand*. Princeton, NJ: Princeton Univ. Press.
- Hansen, Lars P. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50 (4): 1029–54.
- Hicks, John R. 1932. *The Theory of Wages*. London: Macmillan.
- Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Q.J.E.* 124 (4): 1403–48.
- Jin, Hui, and Dale W. Jorgenson. 2010. "Econometric Modeling of Technical Change." *J. Econometrics* 157 (2): 205–19.
- Jones, Charles I. 2005. "The Shape of Production Functions and the Direction of Technical Change." *Q.J.E.* 120 (2): 517–49.
- Kalt, Joseph P. 1978. "Technological Change and Factor Substitution in the United States: 1929–1967." *Internat. Econ. Rev.* 19 (3): 761–75.
- Karabarbounis, Loukas, and Brent Neiman. 2014. "The Global Decline of the Labor Share." *Q.J.E.* 129 (1): 61–103.
- Kasahara, Hiroyuki, and Joel Rodrigue. 2008. "Does the Use of Imported Intermediates Increase Productivity? Plant-Level Evidence." *J. Development Econ.* 87: 106–18.
- Kasahara, Hiroyuki, Paul Schrimpf, and Michio Suzuki. 2015. "Identification and Estimation of Production Functions with Unobserved Heterogeneity." Working paper, Univ. British Columbia.

- Klump, Rainer, Peter McAdam, and Alpo Willman. 2007. "Factor Substitution and Factor-Augmenting Technical Progress in the United States: A Normalized Supply-Side System Approach." *Rev. Econ. and Statis.* 89 (1): 183–92.
- Kugler, Maurice, and Eric Verhoogen. 2012. "Prices, Plant Size, and Product Quality." *Rev. Econ. Studies* 79:307–39.
- Levinsohn, James, and Amil K. Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Rev. Econ. Studies* 70 (2): 317–41.
- Lucas, Robert E., Jr. 1969. "Labor-Capital Substitution in U.S. Manufacturing." In *The Taxation of Income from Capital*, edited by A. Harberger and M. Bailey, 223–74. Washington, DC: Brookings Inst.
- . 1988. "On the Mechanics of Economic Development." *J. Monetary Econ.* 22 (1): 3–42.
- Marschak, Jacob, and William H. Andrews. 1944. "Random Simultaneous Equations and the Theory of Production." *Econometrica* 12 (3): 143–205.
- McAdam, Peter, and Alpo Willman. 2013. "Medium Run Redux." *Macroeconomic Dynamics* 17 (4): 695–727.
- Melitz, Marc J., and Saso Polanec. 2015. "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit." *Rand J. Econ.* 46 (2): 362–75.
- Newey, Whitney K., James L. Powell, and Francis Vella. 1999. "Nonparametric Estimation of Triangular Simultaneous Equations Models." *Econometrica* 67 (3): 565–603.
- Oberfield, Ezra, and Devesh R. Raval. 2014. "Micro Data and Macro Technology." Working paper, Fed. Trade Comm., Washington, DC.
- Oi, Walter Y., and Todd L. Idson. 1999. "Firm Size and Wages." In *Handbook of Labor Economics*, vol. 3, edited by O. Ashenfelter and D. Card. Amsterdam: Elsevier.
- Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Industry." *Econometrica* 64 (6): 1263–97.
- Piketty, Thomas. 2014. *Capital in the Twenty-First Century*. Cambridge, MA: Belknap.
- Raval, Devesh R. 2013. "Non Neutral Technology Differences." Working paper, Fed. Trade Comm., Washington, DC.
- Roberts, Mark J., and Dylan Supina. 1996. "Output Price, Markups, and Producer Size." *European Econ. Rev.* 40:909–21.
- Robinson, Peter M. 1988. "Root- n -Consistent Semiparametric Regression." *Econometrica* 56 (4): 931–54.
- Romer, Paul M. 1990. "Endogenous Technological Change." *J.P.E.* 98, no. 5, pt. 2 (October): S71–S102.
- Rotemberg, Julio J., and Michael Woodford. 1996. "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity." *J. Money, Credit, and Banking* 28 (4): 549–77.
- Solow, Robert M. 1957. "Technical Change and the Aggregate Production Function." *Rev. Econ. and Statis.* 39 (3): 312–20.
- Syverson, Chad. 2004. "Product Substitutability and Productivity Dispersion." *Rev. Econ. and Statis.* 86 (2): 534–50.
- . 2011. "What Determines Productivity?" *J. Econ. Literature* 49 (2): 326–65.
- Uzawa, Hirofumi. 1961. "Neutral Inventions and the Stability of Growth Equilibrium." *Rev. Econ. Studies* 28 (2): 117–24.
- Van Biesebroeck, Johannes. 2003. "Productivity Dynamics with Technology Choice: An Application to Automobile Assembly." *Rev. Econ. Studies* 70:167–98.
- Violante, Giovanni L. 2008. "Skill-Biased Technical Change." In *The New Palgrave Dictionary of Economics*, 2nd ed., edited by Steven N. Durlauf and Lawrence E. Blume. New York: Palgrave Macmillan.
- Zhang, Hongsong. 2015. "Non-neutral Technology, Firm Heterogeneity, and Labor Demand." Working paper, Univ. Hong Kong.