DAVID AUTOR
Massachusetts Institute of Technology

ANNA SALOMONS
Utrecht University

Is Automation Labor Share–Displacing? Productivity Growth, Employment, and the Labor Share

ABSTRACT Many technological innovations replace workers with machines. But this capital–labor substitution need not reduce aggregate labor demand, because it simultaneously induces four countervailing responses: own-industry output effects; cross-industry input–output effects; between-industry shifts; and final demand effects. We quantify these channels using four decades of harmonized cross-country and industry data, whereby we measure automation as industry-level movements in total factor productivity that are common across countries. We find that automation displaces employment and reduces labor's share of value added in the industries where it originates (a direct effect). In the case of employment, these own-industry losses are reversed by indirect gains in customer industries and induced increases in aggregate demand. By contrast,

Conflict of Interest Disclosure: David Autor is a trustee and board member for the Urban Institute, and codirector of the Labor Studies Program at the National Bureau of Economic Research. He received financial support for this work from the IBM Open Collaboration Research award program, the Schmidt Family Foundation, and the Smith Richardson Foundation. Anna Salomons received financial support for this work from the Netherlands Organization for Scientific Research through the Innovational Research Incentives Scheme Veni research grant program. With the exception of the aforementioned, the authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. With the exception of the aforementioned, they are currently not officers, directors, or board members of any organization with an interest in this paper. No outside party had the right to review this paper before publication.

own-industry labor share losses are not recouped elsewhere. Our framework can account for a substantial fraction of the reallocation of employment across industries and the aggregate fall in the labor share over the last three decades. It does not, however, explain why the labor share fell more rapidly during the 2000s.

It is a widely held view that recent and incipient breakthroughs in artificial intelligence and dexterous, adaptive robotics are profoundly shifting the terms of human-versus-machine comparative advantage. In light of these advances, numerous scholars and popular writers anticipate the wholesale elimination of a vast set of currently labor-intensive and cognitively demanding tasks, leaving an ever-diminishing set of activities in which labor adds significant value (Brynjolfsson and McAfee 2014; Ford 2015; Frey and Osborne 2017). The displacement of labor from production could take (at least) two forms: employment displacement, meaning the elimination of aggregate employment; or labor share displacement, meaning the erosion of labor's share of value added in the economy.

Whether technological progress ultimately proves employment-displacing or labor share–displacing depends proximately on two factors: how technological innovations shape employment and labor's share of value added *directly* in the industries where they occur; and how these direct effects are augmented or offset by employment and labor share changes elsewhere in the economy that are *indirectly* spurred by these same technological forces. The first of these phenomena—the direct effect of technological progress on employment and labor share in the specific settings in which it occurs—is often readily observable, and we suspect that observation of these *direct* labor share–displacing effects shapes theoretical and empirical study of the aggregate impact of technological progress. The *indirect* effects of technological progress on these same outcomes, however, are likely more challenging to observe and quantify, and hence may receive short shrift in economic analysis and in the wider public debate.[1]

1. Caselli and Manning (forthcoming) observe that many recent analyses of the potential impact of new technology on workers implicitly rely on models that omit general equilibrium effects.

To see the challenge this creates, consider the two panels of figure 1, which reports bivariate scatters of the relationship between industry-level total factor productivity (TFP) growth over the 1970–2007 period and contemporaneous industry-level log employment growth (the top panel) and industry-level changes in log labor share (the bottom panel), defined as the log ratio of the wage bill to value added.[2] Both panels reveal a well-determined downward slope: Industries experiencing faster measured TFP growth on average exhibit steep relative declines in employment and labor share over this period. It would be tempting to infer from these figures that technological advances (captured by TFP growth) erode aggregate employment and labor's share of national income.

But theory makes clear that there is no direct mapping between the evolution of productivity and labor demand at the industry level and the evolution of labor demand in the aggregate (Foster and others 2017). A long-standing body of literature, starting with research by William Baumol (1967), has considered reallocation mechanisms for employment, showing that labor moves from technologically advancing to technologically lagging sectors if the outputs of these sectors are not close substitutes. Further, Rachel Ngai and Christopher Pissarides (2007) and Daron Acemoglu and Veronica Guerrieri (2008) show that such ongoing unbalanced productivity growth across sectors can nevertheless yield a balanced growth path for labor and capital shares. Indeed, one of the central stylized facts of modern macroeconomics, immortalized by Nicholas Kaldor (1961), is that during a century of unprecedented technological advancement in transportation, production, and communication, labor's share of national income remained roughly constant (Jones and Romer 2010). This empirical regularity, which John Maynard Keynes (1939) deemed "a bit of a miracle," has provided economists—though not the lay public—with grounds for optimism that, despite seemingly limitless possibilities for labor-saving technological progress, automation need not displace labor as a factor of production.

Table 1 confirms the broad relevance of these theoretical observations. Aggregate employment *grew* dramatically in all countries from 1970 to 2007, even as relative employment fell in the industries experiencing the fastest productivity growth. Yet, conversely, labor's share of value added

2. Our data sources and methods are documented in detail in section I. The figures above average across the 19 developed countries in our sample encompassing 28 market industries. Each industry is weighted by its own-country average share of employment (the top panel of figure 1) or value added (the bottom panel of figure 1) over the full time interval. Patterns are similar when instead using decadal changes in employment or labor's share and previous-decade TFP growth starting in the 1980s.

Figure 1. Industry-Level Total Factor Productivity Growth versus Employment Growth and the Change in Labor's Share of Value Added, 1990–2007[a]

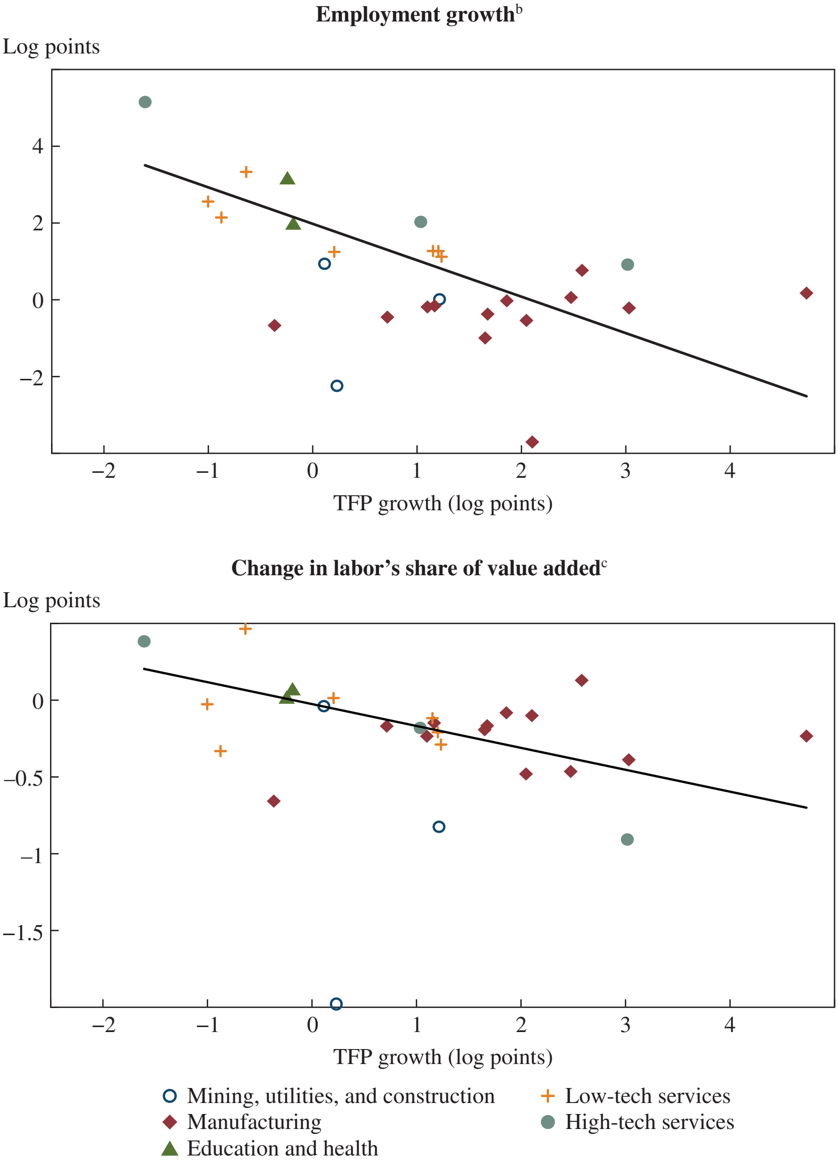

Source: EU KLEMS.

a. All values are expressed as annual, unweighted average changes across country-years in log points.

b. The line shows the linear fit weighted by industries' employment shares. Statistics: $\beta = -0.949$ (SE = 0.181), $R^2 = .515$.

c. The line shows the linear fit weighted by industries' value-added shares. Statistics: $\beta = -0.143$ (SE = 0.050), $R^2 = .238$.

Table 1. Trends in Hours Worked and Labor Share by Country and Decade, 1970–2007

	Average across years			100 × annualized change in log hours worked[a]				100 × annualized change in log labor share[b]			
Country	*Log hours*	*Labor share*	*Value-added share*	*1970s*	*1980s*	*1990s*	*2000s*	*1970s*	*1980s*	*1990s*	*2000s*
Australia	9.41	0.648	0.020	1.77	2.48	2.32	3.03	–0.22	–1.07	0.01	–0.27
Austria	8.61	0.672	0.009	0.52	0.48	1.42	1.95	–0.72	–1.12	–0.79	–1.16
Belgium	8.50	0.641	0.011	–0.96	0.23	1.82	1.79	0.92	–1.27	0.22	2.52
Canada	9.82	0.594	0.028	2.59	2.38	1.82	2.55	–0.40	–0.02	–1.13	0.42
Denmark	8.20	0.676	0.007	–0.07	0.18	1.18	2.02	0.14	–0.37	–0.96	0.59
Finland	8.10	0.683	0.006	0.26	1.34	–0.30	2.17	–0.10	0.35	–2.53	0.17
France	10.36	0.679	0.063	0.04	0.37	1.07	1.80	–0.37	–1.07	–0.81	–0.44
Germany	10.82	0.666	0.093	–0.60	0.29	1.13	0.80	0.42	–1.18	0.15	–1.52
Ireland	7.77	0.559	0.007		3.71	5.32	4.37		0.17	–2.15	0.78
Italy	10.39	0.682	0.052	1.20	1.21	0.84	1.94	0.54	–0.52	–1.82	–0.53
Japan	11.57	0.566	0.196	1.17	0.80	–0.27	0.97	2.38	–0.43	–0.76	–0.71
Luxembourg	5.82	0.554	0.001		3.52	4.86	3.99		1.21	–0.84	–0.08
Netherlands	9.06	0.683	0.017	–0.59	1.26	3.26	1.77	–1.73	–0.47	0.09	–0.85
Portugal	8.87	0.594	0.004	1.43	–1.23	0.76	0.88	3.01	2.26	–0.56	–0.55
South Korea	10.33	0.695	0.017	6.46	3.43	1.82	1.56	–0.07	0.44	–1.22	0.93
Spain	9.85	0.628	0.027	0.81	1.28	2.72	4.06	0.10	–0.11	0.31	–0.94
Sweden	8.72	0.679	0.015		1.50	0.29	1.30		–0.61	–0.91	0.40
United Kingdom	10.65	0.705	0.059	0.11	1.46	0.92	2.38	–0.34	0.36	–0.91	0.32
United States	12.08	0.637	0.366	2.39	2.70	2.50	0.70	0.12	–0.38	0.36	–1.46
Weighted average				1.424	1.699	1.553	1.350	0.513	–0.459	–0.263	–0.861

Sources: EU KLEMS; authors' calculations.

a. Changes are annualized log differences by decade weighted by time-averaged hours-worked shares.

b. Changes are annualized log differences by decade weighted by time-averaged value-added shares.

was steady or rising in the 1970s, declined modestly in the 1980s and 1990s, and then fell steeply in the 2000s in many countries. These facts thus highlight the pitfalls of extrapolating from direct, first-order technological relationships (here, observed at the industry level) to labor market outcomes in the aggregate, because the latter incorporate both direct and indirect consequences of technological progress (as well as many non-technological factors).

This paper applies harmonized cross-country and cross-industry data to explore the relationship between technological change and labor market outcomes over four decades. A first contribution of the paper is to attempt to overcome the tension, endemic to this area of work, of using microeconomic variation to afford identification while attempting to speak to macroeconomic outcomes. This tension arises here because we study the relationship between productivity growth, innovation, and labor displacement at the country-industry level. As figure 1 underscores, naively extrapolating from industry-level to aggregate-level relationships is potentially fallacious. The alternative—directly estimating effects at the macro level—often suffers from underidentification and low statistical power, and furthermore is silent on the microeconomic channels through which aggregate effects come about.

To overcome these pitfalls, we empirically model three micro–macro linkages that, in combination with the industry-level estimates, allow us to make broader inferences about aggregate labor displacement effects.[3] The first link uses harmonized data from the World Input–Output Database (Timmer and others 2015), enumerating cross-industry input–output linkages to trace the effects of productivity growth in each industry to outcomes occurring in customer industries and in supplier industries—that is, industries for which, respectively, the originating industry is upstream or downstream in the production chain.[4] The second link connects aggregate economic growth and sectoral labor demands. Recognizing that productivity growth in each industry augments aggregate income and hence indirectly raises final demand, we estimate the elasticity of sectoral demand emanating from aggregate income growth and then apply our TFP estimates to infer the indirect contribution of each industry's productivity growth to final demand. Third, our analytic framework recognizes that uneven productivity growth

3. Our approach here builds on our earlier work (Autor and Salomons 2017), in which we incorporate only one of these linkages.

4. Our analysis follows many recent works exploiting these linkages to study the propagation of trade and technology shocks (Acemoglu and others 2016; Pierce and Schott 2016; Acemoglu, Akcigit, and Kerr 2016).

across industries yields shifts in industry shares of value added, which in turn potentially alter labor's share of aggregate value added.[5]

Our net estimates of the impact of productivity growth and innovation on aggregate outcomes of interest therefore sum over (i) direct industry-level effects; (ii) indirect customer and supplier effects in linked sectors; (iii) final demand effects accruing through the effect of productivity growth on aggregate value added; and (iv) composition effects accruing through productivity-induced changes in industry shares of value added. We believe that this simple accounting framework can be usefully applied to other data sets and sources of variation.

Distinct from earlier work that focuses on specific measures of technological adoption or susceptibility (for example, robotics and routine task replacement), a second contribution of the analysis is to employ TFP—which is an *omnibus* measure of technological progress (Solow 1956). Using TFP as our baseline measure potentially overcomes the challenge for consistent measurement posed by the vast heterogeneity of innovation across sectors and periods. TFP is also applicable to our analysis for a second reason: Because all margins of technological progress ultimately induce a rise in TFP—either by increasing the efficiency of capital or labor in production or by reallocating tasks from labor to capital or vice versa—our empirical approach is not predicated on a specific mechanism through which technological progress affects outcomes of interest. But the flip side of this agnosticism is that merely observing a change in TFP in any industry or time period does not tell us *which* channel (augmentation, reallocation) is operative. Using information on output, employment, earnings, and labor's share of value added, however, we can infer these channels. Specifically, we study how changes in industry-level TFP affect output (value added) quantities and prices, employment, earnings, and labor's share of value added economy-wide, to draw inferences on both industry-level and aggregate labor-augmenting and labor share–displacing effects of technological progress.

It is well understood that estimates of TFP may also be confounded with business cycle effects, industry trends, and cross-industry differences in cyclical sensitivity (Basu and Fernald 2001). We confront these issues directly. We purge the simultaneity between an industry's estimated TFP growth and changes in other industry-level measures that serve as inputs

5. This mechanism is akin to skill-biased structural change in the framework developed by Buera, Kaboski, and Rogerson (2015), though here we focus on labor share rather than skill composition.

into the TFP calculation (for example, output, wage bill, and employment) by replacing own-country-industry TFP with the mean TFP of the corresponding industry observed in other countries in the same year.[6] We purge the potential cyclicality of TFP by including a set of distributed lags as well as country–business cycle indicators, which absorb business cycle variation in productivity measures. We address the opaqueness of TFP as a measure of technological progress by complementing it with an alternative, directly observable measure of industry-level technological advancement: patent awards by industry and country (Autor and others 2017a). Patent awards—and even more so, patent citations—prove to be strong predictors of industry-level TFP growth. Using patent awards in place of TFP growth, we obtain strongly comparable estimates of the relationships between technological progress, employment, wage bill, and value added, which we view as useful corroborative evidence.

TFP's virtue as an omnibus technology measure is also its shortcoming as a specific technology measure. Because TFP incorporates productivity growth arising from all sources, our analysis cannot directly answer the question of whether recent or specific technologies—such as industrial robotics or artificial intelligence—are more or less labor-complementing or labor share–displacing than earlier generations of technology. By the same token, our analysis cannot distinguish between the effects of automation-based versus non-automation-based sources of TFP growth, which may in turn have distinct (or even countervailing) effects on either employment or on labor's share of value added. We refer readers to recent studies focusing on specific technological advances for this evidence (Graetz and Michaels, forthcoming; Acemoglu and Restrepo 2017; Dauth and others 2017; Chiacchio, Petropoulos, and Pichler 2018).

Our work builds on an active, recent body of literature that questions the optimistic implications of the long-standing Kaldor facts by offering models where aggregate labor displacement is a potential consequence of advancing technology. Acemoglu and Pascual Restrepo (2018, forthcoming) consider models in which two countervailing economic forces determine the evolution of labor's share of income: the march of technological progress, which gradually replaces "old" labor-using tasks, reducing labor's share of output and possibly diminishing real wages; and endogenous technological progress that generates novel labor-demanding tasks, potentially reinstating

6. This strategy leverages the fact that changes in other-country, same-industry TFP are highly predictive of the evolution of own-country-industry TFP but are not intrinsically correlated with its evolution.

labor's share. The interplay of these forces need not necessarily yield a balanced growth path; that is, labor's share may decline. Daniel Susskind (2017) develops a model in which labor is ultimately immiserated by the asymptotic encroachment of automation into the full spectrum of work tasks—contrary to Acemoglu and Restrepo (forthcoming), labor immiseration is guaranteed because falling labor scarcity does not spur the endogenous creation of new labor-using tasks or labor-complementing technologies.[7]

A central empirical regularity that underscores the relevance of this recent work is that labor's share of national income has indeed fallen in many nations in recent decades, a trend that may have become more pronounced in the 2000s (Elsby, Hobijn, and Şahin 2013; Karabarbounis and Neiman 2013; Piketty 2014; Barkai 2017; Autor and others 2017b; Dao and others 2017; Gutiérrez and Philippon 2017). Reviewing an array of within- and cross-country evidence, Loukas Karabarbounis and Brent Neiman (2014) argue that labor's falling share of value added is caused by a steep drop in the quality-adjusted equipment prices of information and communication technology relative to labor. Though Karabarbounis and Neiman's work is controversial, in that it implies an aggregate capital/labor substitution in excess of 1—which is a nonstandard assumption in this literature—their work has lent empirical weight to the hypothesis that computerization may erode labor demand. Related work by Maya Eden and Paul Gaggl (2018) calibrates an aggregate production function, and similarly attributes part of the decline in the U.S. labor share to a rise in the share of income paid to information and communication technology capital.

A growing microeconometric literature presents a mixed set of findings on whether such erosion has occurred recently or in the past. Focusing on the first half of the twentieth century, Michelle Alexopoulos and Jon Cohen (2016) find that positive technology shocks raised productivity and lowered unemployment in the United States between 1909 and 1949. Using contemporary European data, Terry Gregory, Anna Salomons, and Ulrich Zierahn (2016) test whether routine-replacing technical change has

7. The conceptual frameworks of both papers build on the work of Zeira (1998), Autor, Levy, and Murnane (2003), and Acemoglu and Autor (2011), who offer models in which advancing automation reduces labor's share by substituting machines (or computers) for workers in a subset of activities (which Autor, Levy, and Murnane designate as "tasks"). Other labor-displacement mechanisms are given by Sachs and Kotlikoff (2012) and Berg, Buffie, and Zanna (2018), who develop overlapping-generation models in which rapid labor-saving technological advances generate short-run gains for skilled workers and capital owners, but in the longer run, immiserate those who are not able to invest in physical or human capital. Stansbury and Summers (2017) present time-series evidence that productivity growth and wage growth are positively correlated.

reduced employment overall across Europe, and they find that though this type of change has reduced middle-skill employment, this reduction has been more than offset by compensatory product demand and local demand spillovers. In work closely related to ours, Mai Chi Dao and others (2017) analyze sources of the trend decline in labor share in a panel of 49 emerging and industrialized countries. Using cross-country and cross-sector variation in the prevalence of occupations potentially susceptible to automation (à la Autor and Dorn 2013), Dao and others find that countries and sectors initially more specialized in routine-intensive activities have seen a larger decline in labor share, consistent with the possibility of labor displacement.[8]

Concentrating on industrial robotics, arguably the leading edge of workplace automation, Georg Graetz and Guy Michaels (forthcoming) conclude that industry-level adoption of industrial robots has raised labor productivity, increased value added, augmented workers' wages, had no measurable effect on overall labor hours, and modestly shifted employment in favor of high-skill workers within countries that belong to the European Union. Conversely, using the same underlying industry-level robotics data but applying a cross-city design within the United States, Acemoglu and Restrepo (2017) present evidence that U.S. local labor markets that were relatively exposed to industrial robotics experienced differential falls in employment and wage levels between 1990 and 2007.[9]

Our analysis proceeds as follows. Section I summarizes the data and measurement framework and presents the simple shift-share decomposition that undergirds our accounting framework. Section II presents our estimates for the direct effects of productivity growth (measured initially by TFP, in subsection II.A; and by patents in subsection II.B) on labor input, value added, and labor's share of value added, across a range of model specifications. Section III then presents our main results accounting for both direct ("own-industry") effects, and for indirect effects operating through input–output linkages and final demand. Section IV quantifies the aggregate implications of these direct and indirect effect estimates for employment, hours worked, and labor's share of value added to assess

8. Using an analogous approach, Michaels, Natraj, and Van Reenen (2014) find that information and communication technology adoption is predictive of within-sector occupational polarization in a country-industry panel sourced from EU KLEMS covering 11 countries observed over 25 years.

9. Dauth and others (2017) and Chiacchio, Petropoulos, and Pichler (2018) apply the Acemoglu–Restrepo approach to German and EU-wide data, respectively. Dauth and others find that robot adoption leads to worker reallocation but has no net impact on employment or wages. Chiacchio, Petropoulos, and Pichler affirm the Acemoglu–Restrepo results for employment though not for wages.

whether technological progress has, on net, been either augmenting or displacing of the aggregate employment or labor share. We also consider in this section whether our accounting approach can explain cross-industry patterns of employment change and aggregate, time-series changes in the evolution of the labor share between and within industries.

To briefly summarize our results, automation (as embodied in TFP growth) has been *employment-augmenting* yet *labor share–displacing* over the last four decades. As implied by the scatter plot in figure 1 (top panel), industries with persistent gains in relative productivity secularly contract as a share of aggregate employment, meaning that the *direct* effect of rising productivity has been to reduce labor input in the sectors where it originates. But this direct effect is more than fully offset by two *indirect* effects: First, rising TFP within supplier industries catalyzes strong, offsetting employment gains among their downstream customer industries; and second, TFP growth in each sector contributes to aggregate growth in real value added and hence rising final demand, which in turn spurs further employment growth across all sectors.

Conversely, we find that productivity growth is directly labor share–*displacing* in the industries where it originates; and it is particularly important that this direct effect is not offset by *indirect* effects spurred by input–output linkages, compositional shifts, or final demand increases. Thus, we conclude that productivity growth has contributed to an erosion of labor's share of value added. Notably, this labor share–eroding effect was not present in the first decade of our sample, the 1970s, but then became strongly evident thereafter. Our analysis therefore broadly supports the hypothesis that the decline in the labor share since the 1980s is consistent with a shift toward more labor-displacing technology commencing in the 1980s. But the acceleration in the labor share decline observed during the 2000s is left unaccounted for by this mechanism.

In section V, we briefly consider the interpretation of our findings, focusing in particular on the relationship between the industry-level and aggregate outcomes observed in our data, and the underlying unobserved firm-level dynamics that may contribute to these outcomes.

I. Data and Measurement

Our analysis draws on EU KLEMS, an industry-level panel database covering the countries that belong to the Organization for Economic Cooperation and Development since 1970 (O'Mahony and Timmer 2009). We use the 2008 release of EU KLEMS, supplemented with data from the 2007 and 2011 releases to maximize data coverage. Our primary analytic sample covers the period 1970–2007. We limit our analysis to the 19 developed countries of the European Union, excluding its Eastern European members; and we also include Australia, Canada, Japan, South Korea, and the United States. These countries and their years of data coverage are listed in online appendix table A1.[10] The EU KLEMS database contains detailed data for 32 industries in both the market and nonmarket economies, as summarized in online appendix table A2. We focus on nonfarm employment, and we omit the poorly measured private household sector, and public administration, defense, and extraterritorial organizations, which are almost entirely nonmarket sectors.[11] The end year of our analysis is dictated by major revisions to the industry definitions in EU KLEMS that were implemented from the 2013 release onward. These definitional changes inhibit us from extending our consistent 1970–2007 analysis through to the present, though we analyze 2000–15 separately using the 2017 release of EU KLEMS for a smaller subset of countries for which these data are available.[12]

Table 1 summarizes trends in aggregate hours of labor input and labor's share of value added by decade for the 19 countries in our analysis. As with all analyses in the paper, these statistics are calculated using the 28 market industries that constitute our analytic sample and are annualized to account for the fact that years of data coverage differ by country. With very few exceptions, aggregate labor hours rise in all countries and time periods. The growth rate of labor hours is most rapid in the 1980s, slower in the 1990s, and slower still in the 2000s. Distinct from aggregate labor hours, trends in labor's share of value added differ by country and time period. On average, the aggregate labor share rises in the 1970s and then falls during the subsequent three decades, with by far the sharpest annual rate of decline in the 2000s.

Table 2 reports analogous statistics for trends in hours of labor input and labor's share of value added among the 28 industries in our sample. There is a substantial diversity of experiences among industries. Employment fell steeply in mining and quarrying, textiles and related products, and refining,

10. The online appendixes for this and all other papers in this volume may be found at the *Brookings Papers* web page, www.brookings.edu/bpea, under "Past BPEA Editions."

11. Although EU KLEMS classifies health care and education as nonmarket sectors, they are a substantial and growing part of GDP across the developed world; and in many countries (for example, the United States), they also encompass a large private sector component. We therefore choose to retain these sectors in our analysis.

12. This subset includes Austria, Belgium, Denmark, Finland, France, Germany, Italy, the Netherlands, Spain, Sweden, the United Kingdom, and the United States.

Table 2. Trends in Hours Worked, Labor Share, and Total Factor Productivity, by Industry, 1970–2007

ISIC code (rev. 3)	*Description*	*Time-averaged value added share*	*100 × annual log change*		
			Hours worked[a]	*Labor share*[b]	*Total factor productivity*[b]
C	Mining and quarrying	0.015	–2.45	–1.22	0.29
15–16	Manufacture of food, beverages, and tobacco products	0.026	–0.52	–0.08	0.72
17–19	Manufacture of textiles, apparel, leather, and related products	0.012	–3.96	0.18	2.07
20	Manufacture of wood and wood products, excluding furniture	0.005	–1.34	–0.32	2.12
21–22	Manufacture of paper and paper products, printing, and publishing	0.022	–0.25	–0.19	1.10
23	Manufacture of coke, refined petroleum products, and nuclear fuel	0.006	–1.54	–1.60	–0.49
24	Manufacture of chemicals and chemical products	0.022	–0.78	–0.44	3.19
25	Manufacture of rubber and plastics products	0.010	0.67	0.21	2.56
26	Manufacture of other nonmetallic mineral products	0.009	–1.33	–0.18	1.68
27–28	Manufacture of basic and fabricated metals	0.029	–0.87	–0.22	1.72
29	Manufacture of machinery and equipment not elsewhere classified	0.023	–0.60	0.03	1.86
30–33	Manufacture of electrical and optical equipment	0.033	–0.28	–0.10	4.49
34–35	Manufacture of motor vehicles and transportation equipment	0.024	–0.12	–0.27	2.42
36–37	Manufacture of furniture and manufacturing not elsewhere classified; recycling	0.008	–0.58	–0.03	1.09
E	Electricity, gas, and water supply	0.025	–0.28	–0.65	1.29
F	Construction	0.068	0.94	0.04	0.20
50	Sale, maintenance, and repair of motor vehicles and fuel	0.014	0.95	–0.05	0.11

(*continued on next page*)

Table 2. Trends in Hours Worked, Labor Share, and Total Factor Productivity, by Industry, 1970–2007 (*Continued*)

ISIC code (rev. 3)	*Description*	*Time-averaged value added share*	*100 × annual log change*		
			Hours worked[a]	*Labor share*[b]	*Total factor productivity*[b]
51	Wholesale trade, excluding motor vehicles	0.064	0.67	–0.28	1.07
52	Retail trade, excluding motor vehicles; repair of personal and household goods	0.052	0.73	–0.16	1.18
H	Hotels and restaurants	0.028	1.80	–0.09	–0.88
60–63	Transportation activities of travel agencies	0.045	0.91	–0.17	1.24
64	Post and telecommunications	0.025	0.52	–1.18	3.04
J	Financial intermediation	0.064	1.70	–0.46	0.95
70	Real estate activities	0.113	3.08	0.70	–0.66
71–74	Renting of machinery and equipment; computer and related activities; research and development; and other business activities	0.098	4.63	0.68	–1.65
M	Education	0.057	1.67	–0.01	–0.14
N	Health and social work	0.066	2.89	0.05	–0.22
O	Other community, social, and personal service activities	0.040	2.16	0.11	–1.02

Sources: EU KLEMS; authors' calculations.
a. Changes are annualized log differences weighted by country size and hours-worked shares.
b. Changes are annualized log differences weighted by country size and value-added shares.

while growing rapidly in many business and personal services. Labor's share of value added declined in the majority of sectors, with the steepest fall in heavy industry. TFP growth, meanwhile, was most rapid in manufacturing and was negative in several service industries.

Table 3 summarizes trends in employment, hours, wages, value added, labor share, and TFP by industry over the four decades of our sample. We quantify these trends overall, by broad sector, and by decade by estimating regression models for the change in country-industry-year outcomes (multiplied by 100). In this table, and throughout the paper, regression models are weighted by time-averaged shares of the relevant weighting variable—employment, hours, or value added—within

Table 3. Within-Industry Trends in Key Variables Used in the Analysis, 1970–2007[a]

	100 × mean annual log change						
	Employment	*Hours worked*	*Nominal hourly wage*	*Real hourly wage*	*Nominal value added*	*Labor share*	*Total factor productivity*
Overall	1.337***	1.001***	6.472***	1.700***	7.058***	–0.051	0.619***
	(0.166)	(0.171)	(0.171)	(0.116)	(0.188)	(0.104)	(0.150)
1970s	2.035***	1.572***	11.556***	2.472***	11.643**	0.503**	0.440**
	(0.204)	(0.210)	(0.401)	(0.261)	(0.270)	(0.217)	(0.201)
1980s	1.661***	1.365***	6.550***	1.728***	7.689***	–0.265**	0.994***
	(0.198)	(0.214)	(0.238)	(0.170)	(0.294)	(0.123)	(0.168)
1990s	0.996***	0.656***	3.815***	1.319***	4.135***	–0.141	0.603***
	(0.190)	(0.220)	(0.178)	(0.157)	(0.253)	(0.171)	(0.154)
2000s	0.382**	0.174	3.043***	1.127***	3.876***	–0.395***	0.360***
	(0.190)	(0.210)	(0.226)	(0.176)	(0.293)	(0.152)	(0.133)
Mining, utilities, and construction	0.625	0.521	6.441***	1.641***	6.523***	–0.391*	0.405***
	(0.393)	(0.443)	(0.809)	(0.573)	(0.550)	(0.223)	(0.114)
Manufacturing	–0.810***	–0.984***	7.068***	2.246***	5.484***	–0.157**	2.185***
	(0.170)	(0.182)	(0.266)	(0.160)	(0.236)	(0.080)	(0.180)
Education and health	2.566***	2.351***	6.422***	1.664***	8.306***	–0.058	–0.190***
	(0.203)	(0.196)	(0.412)	(0.245)	(0.430)	(0.089)	(0.023)
Low-tech services	1.676***	1.227***	6.158***	1.403***	7.164***	0.162	0.150
	(0.148)	(0.168)	(0.251)	(0.168)	(0.254)	(0.237)	(0.193)
High-tech services	3.286***	3.091***	6.324***	1.608***	8.688***	–0.095	–0.022
	(0.379)	(0.379)	(0.380)	(0.274)	(0.360)	(0.222)	(0.452)
No. of observations[b]	18,062	18,062	18,062	18,062	18,062	18,062	18,062
Weights	Employment	Hours	Hours	Hours	Value added	Value added	Value added

Sources: EU KLEMS; authors' calculations.

a. The data are for all sectors of the economy, excluding agriculture, public administration, private households, and extraterritorial organizations. All models are weighted by time-averaged industry shares of the weighting variable within countries, multiplied by time-varying country shares in the total annual value of the weighting variable. Standard errors clustered by country-industry are in parentheses. Statistical significance is indicated at the *10 percent, **5 percent, and ***1 percent levels.

b. The number of observations is the number of country-industry cells multiplied by the number of years.

countries multiplied by time-varying country shares of the weighting variable. As such, we weight by country size in our main estimates, and we show in the online appendixes that our main results are not sensitive to this choice.

The top row of table 3 reports estimates for all industries and time periods. The middle rows report these relationships separately by decade, and the bottom rows report them separately for five broad sectors encompassing the 28 industries in our analysis. As detailed in online appendix table A2, these sectors are: mining, utilities, and construction; manufacturing; education and health; low-tech services (including personal services, retail, wholesale, and real estate); and high-tech services (including post and telecommunications, finance, and other business services). The reported regression coefficients, which correspond to *within-industry* changes, reflect a number of key trends in the data. Employment growth, measured in workers or hours, is positive in all decades but slows substantially across consecutive decades. Employment growth is negative in manufacturing; modestly positive in mining, utilities, and construction; and strongly positive in services—with the most rapid growth evident in high-tech services, followed by education and health, and finally low-tech services. Like employment, the growth of real hourly wages is positive in all periods but is secularly slowing.

Consistent with results reported in much recent work (Elsby, Hobijn, and Şahin 2013; Karabarbounis and Neiman 2014; Autor and others 2017b), trends in the labor share of value added vary across the decades. Labor's share of value added trends modestly upward in the 1970s, then falls in each decade of the 1980s, 1990s, and 2000s. This trend is most pronounced in manufacturing and in mining, utilities, and construction. It is modest in high-tech services, and in the education and health sector, and it is absent in the low-tech services sector.

The descriptive statistics given in table 3 focus on *within-industry* changes in the labor share of value added and its components. But of course, changes in the aggregate labor share may stem from both (i) within-industry shifts in labor's share of value added; and (ii) changes in the share of value added accounted for by industries that differ in their labor shares. Our analysis assesses the contribution of technological change to both margins. To quantify the importance of within- versus between-industry shifts, we implement a simple shift-share decomposition, as follows. Let $\bar{L}_{c,t} = \sum_i \omega_{i,c,t} l_{i,c,t}$ equal the aggregate log labor share of value added in country c in year t, defined as the weighted sum of log labor shares $l_{i,c,t}$ in each industry i, where weights $\omega_{i,c,t}$ correspond to industry i's share in value added in

Table 4. Shift-Share Analysis of the Changes in Labor Share by Decade[a]

	Weighted by country size			*Unweighted*		
Decade	*Mean*	*Between industry*	*Within industry*	*Mean*	*Between industry*	*Within industry*
1970s	0.513	−0.187	0.700	0.230	−0.146	0.376
		(−0.36)	(1.36)		(−0.63)	(1.63)
1980s	−0.459	−0.183	−0.276	−0.201	−0.121	−0.080
		(0.40)	(0.60)		(0.60)	(0.40)
1990s	−0.263	−0.075	−0.188	−0.750	−0.304	−0.446
		(0.28)	(0.72)		(0.41)	(0.59)
2000s	−0.861	−0.425	−0.436	−0.126	−0.018	−0.104
		(0.49)	(0.51)		(0.17)	(0.83)

Sources: EU KLEMS; authors' calculations.

a. The units are 100 × annualized decadal log changes in labor share by country. The values in parentheses are the shares explained by between-industry or within-industry shifts.

its respective country and year.[13] Let $\Delta \overline{L}_{c,\tau}$ equal the change in aggregate log labor share in country c over time interval τ—equal to 1970–80, 1980–90, 1990–2000, or 2000–07—where Δ is the first-difference operator. Finally, let $\overline{l}_{i,c,\tau} = (l_{i,c,t_1} - l_{i,c,t_0})/2$ and $\overline{\omega}_{i,c,\tau} = (\omega_{i,c,t_1} - \omega_{i,c,t_0})/2$. We can then decompose the observed labor share change in each decade as

$$\Delta \overline{L}_{c,\tau} = \sum_i \overline{\omega}_{i,c,\tau} \Delta l_{i,c,\tau} + \sum_i \overline{l}_{i,c,\tau} \Delta \omega_{i,c,\tau}, \tag{1}$$

where the first term to the right of the equals sign is the contribution of within-industry changes in labor share to the aggregate change, and the second term is the contribution to the aggregate change due to shifts in value-added shares across industries.

The results of this decomposition, reported in table 4, indicate that the majority, but not the entirety, of the change in aggregate labor share of value added in each decade is accounted for by within-industry shifts. Focusing first on the country size–weighted calculations (the left columns), we find that more than all of the rise in labor share in the 1970s is due to within-industry changes, whereas between 51 and 72 percent of the fall in the labor share in the subsequent three decades is accounted for by within-industry declines. If we instead weight each country equally in the shift-share decomposition, we reach similar conclusions about the importance of within-industry labor share movements (the right columns). Further, if

13. Per our convention, this calculation includes only the 28 market industries featured in our analysis.

we decompose the change in the mean *level* of labor share rather than the mean *log* level (online appendix table A3), we find a similar time pattern as for the log labor share and a similarly outsized role played by within-industry changes.

These decomposition results suggest that the within-industry determinants of changes in the aggregate labor share are of greater analytic interest compared with the between-industry drivers, though we explore both margins below. The 2000s stand out, however, for having a roughly even distribution of the aggregate labor share changes into within-industry and between-industry components. Consistent with the observations of Matthew Rognlie (2015) and Germán Gutiérrez (2017), this pattern reflects the outsized growth of the real estate industry's value added in numerous countries—particularly during the 2000s—and this industry has an extremely low share of labor in value added (see online appendix table A4). If we eliminate real estate from the analysis, however, we find that the fall in the aggregate labor share in the 2000s is reduced by less than one quarter (from –0.86 to –0.64 per year); the within-industry component of the labor share decline explains no less than 90 percent of the total in each decade; and the annual rate of decline in the labor share during the 2000s is still more than twice as rapid as in the 1990s.[14] Thus, the rising share of real estate in value added is not the primary driver of the falling labor share.

Figure 2 adds country-level detail to these calculations by plotting the evolution of the aggregate labor share of value added for all the countries in our sample. Each panel contains two series: In the first series, industry shares are permitted to vary by year; the second series holds these shares constant at their within-country, over-time averages. The fact that these series closely correspond for almost all countries reinforces the inferences from the decomposition that most of the aggregate changes in the labor share observed in the data stem from within-industry movements in this share.

II. Main Estimates

Before making estimates, we tackle two remaining issues: simultaneity and timing. The simultaneity issue arises because labor's share of value added features in the construction of TFP, inducing a mechanical correlation

14. Supplemental tables are available upon request from the authors.

Figure 2. Trends in Labor's Share of Value Added by Country, 1970–2007[a]

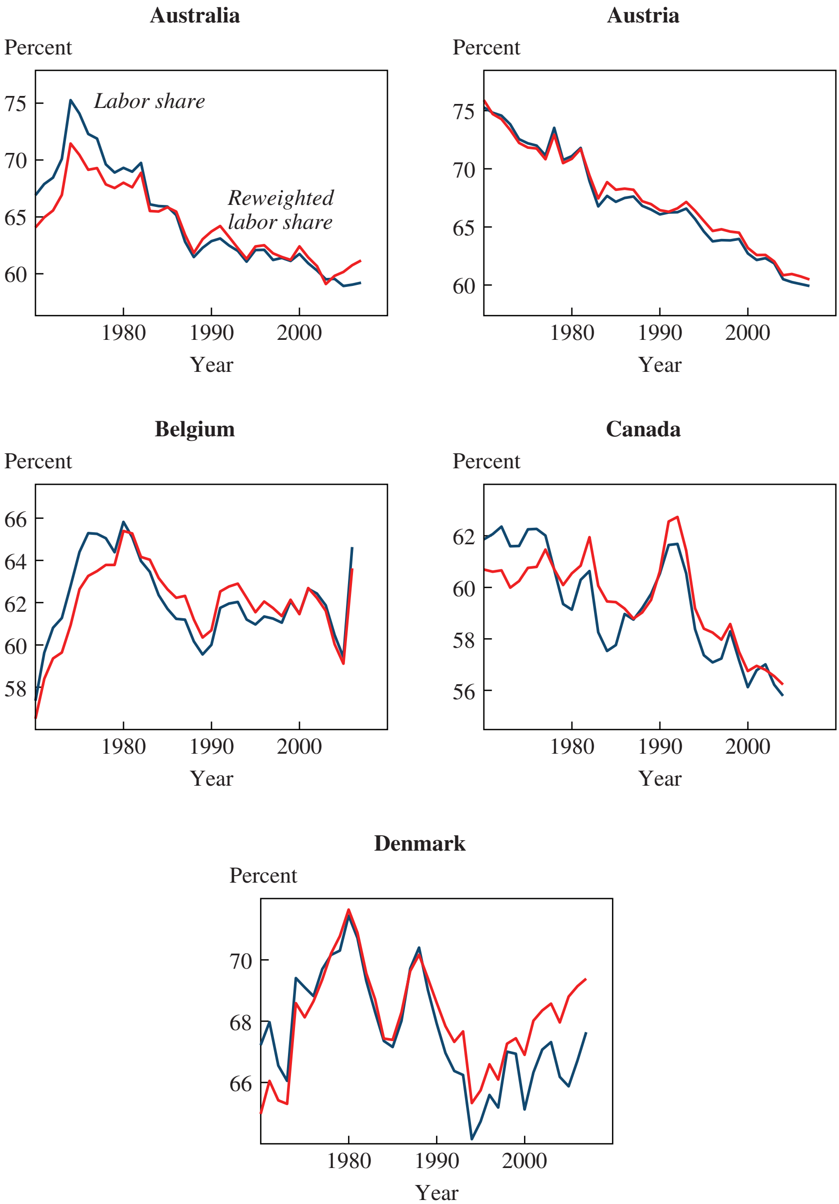

(*continued on next page*)

Figure 2. Trends in Labor's Share of Value Added by Country, 1970–2007[a] (*Continued*)

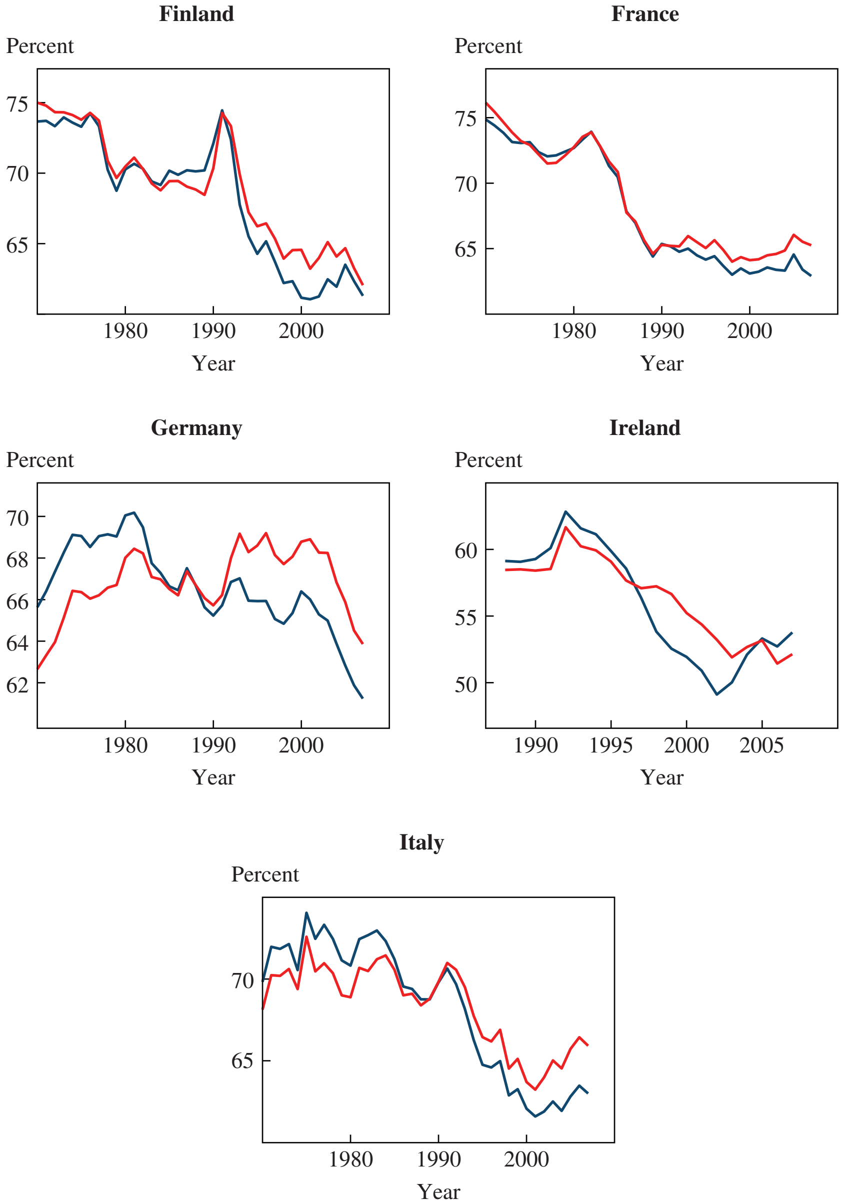

Figure 2. Trends in Labor's Share of Value Added by Country, 1970–2007[a] (*Continued*)

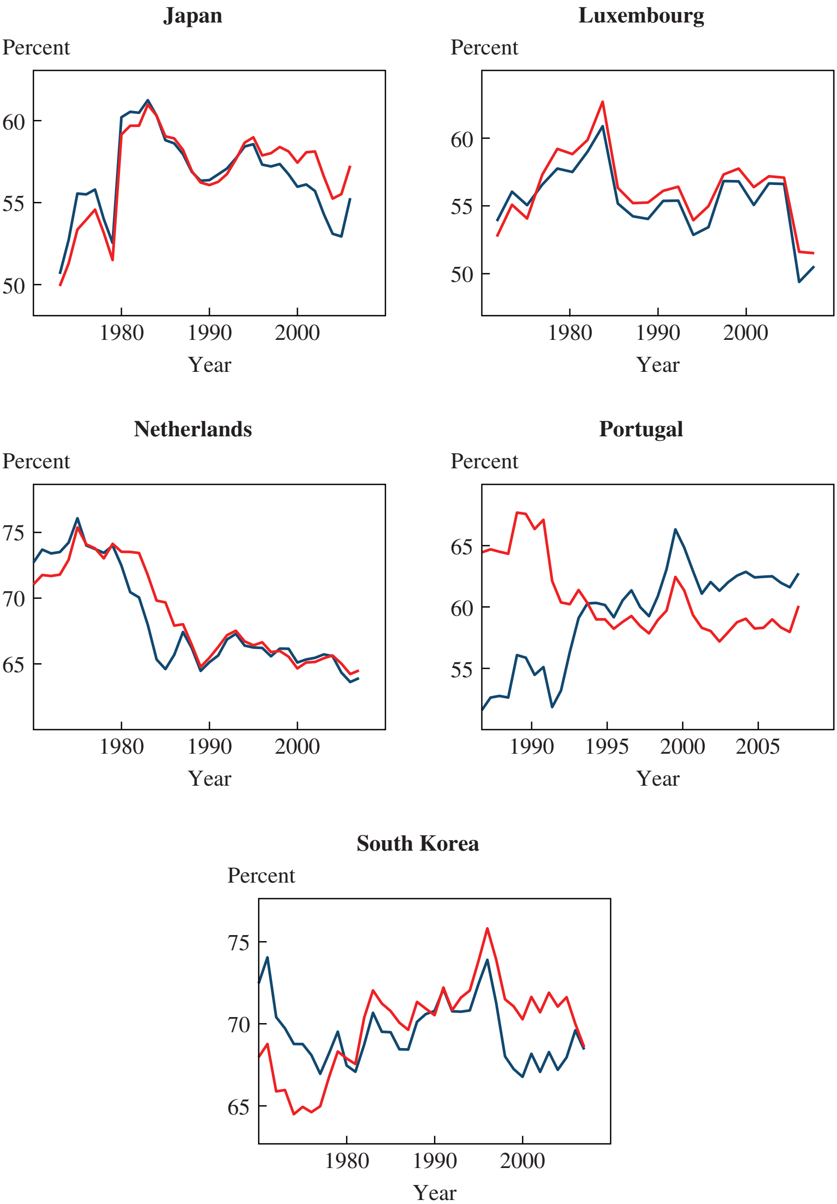

(*continued on next page*)

Figure 2. Trends in Labor's Share of Value Added by Country, 1970–2007[a] (*Continued*)

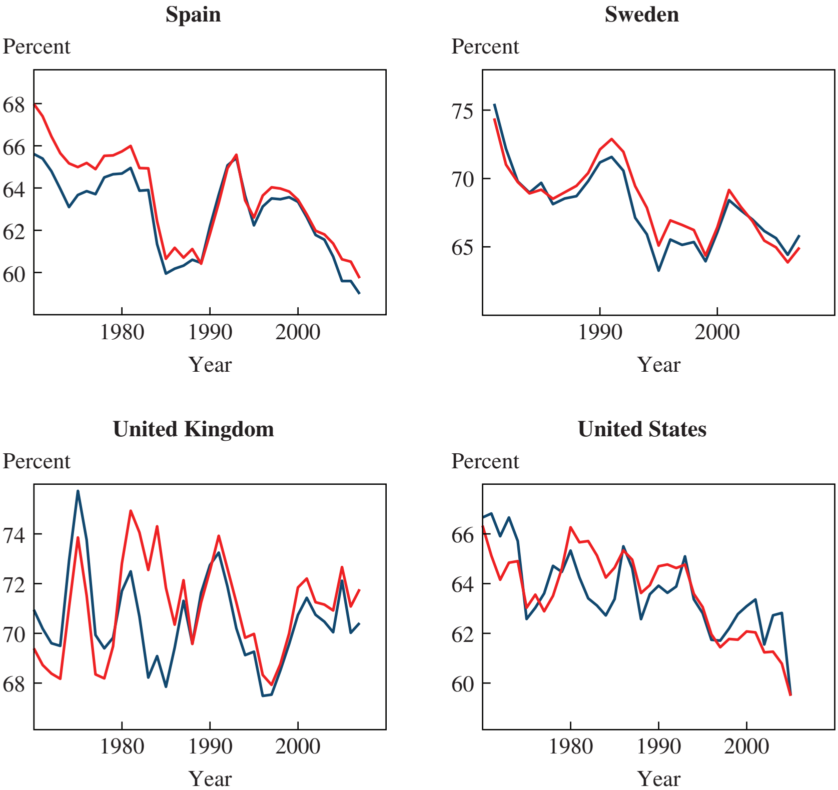

Sources: EU KLEMS; authors' calculations.

a. Labor share is labor compensation as a share of value added. Reweighted labor share is the average of industry labor shares weighted by time-averaged industry value-added shares. The data are for all sectors of the economy, excluding agriculture, public administration, private households, and extraterritorial organizations.

between TFP growth and shifts in the labor share.[15] To overcome this pitfall, we construct industry-level TFP growth for each industry–country pair as the *leave-out* mean of industry-level TFP growth in *all other* countries in the sample. This approach eliminates the mechanical correlation between TFP and labor share and arguably exploits movements in the technology

15. In EU KLEMS, TFP growth is calculated as the log change in industry value added minus the log change in labor and capital inputs, weighted by the average start and end period of their respective factor shares (Timmer and others 2007). In a regression of the change in labor share on TFP growth, the change in labor share used in the TFP calculation enters the right-hand side of the equation, leading to a mechanical relationship.

frontier that are common among industrialized economies. Confirming the utility of this strategy, we show in online appendix table A5 that other-country, same-industry TFP is a strong predictor of own-country-industry TFP: In a set of regressions of own-country-industry TFP on other-country-industry TFP that includes a large number of country, year, sector, and business cycle main effects, we obtain a prediction coefficient that ranges from 0.32 to 0.57, with a t value above 5 in all specifications. Based on this reasoning and evidence, we employ the leave-out TFP measure in place of own-industry TFP in all the analyses given below.

The second issue, timing, arises because contemporaneous productivity innovations are unlikely to induce their steady-state effects immediately, meaning that a lag structure is needed for estimating the relationship between TFP and outcomes of interest (Ramey 2016). To explore a suitable structure, we estimate simple local projection models in the spirit of Òscar Jordà (2005), which involve regressing a series of first differences of increasing length of the outcome variable of interest on the explanatory variable of interest (here, TFP growth) and a set of controls. We estimate

$$\text{(2)} \qquad \ln Y_{i,c,t+K} - \ln Y_{i,c,t-1} = \beta_0 + \beta_1 \Delta \ln TFP_{i,c\neq c(i),t-1} + \sum\nolimits_{k=0}^{K} \beta_2^k \Delta \ln TFP_{i,c\neq c(i),k}$$
$$+ \beta_3 \Delta \ln TFP_{i,c\neq c(i),t-2} + \beta_4 \Delta \ln Y_{i,c,t-2} + \alpha_{c,t} + \gamma_s + \varepsilon_{i,c,t},$$

where $\ln Y_{i,c,t+K}$ denotes the log outcome of interest in industry i, country c, and year t; and K denotes the time horizon for the local projection. The dependent variables therefore reflect the log change in outcome Y from base year $t-1$ up to year $t+K$. The impulse variable is the log change in other-country-industry TFP between years $t-2$ and $t-1$, $\Delta \ln TFP_{i,c\neq c(i),t-1}$. These effects are estimated while controlling for lagged values of both TFP growth ($\Delta \ln TFP_{i,c\neq c(i),t-2}$) and of outcome variable growth ($\Delta \ln Y_{i,c,t-2}$)—that is, conditional on the lagged history of both TFP and outcome growth at the path start time. This allows for feedback dynamics within the system and controls for them through the inclusion of the lagged variables. Each model further controls for a set of country-year fixed effects ($\alpha_{c,t}$), as well as fixed effects for five broad sectors (γ_s, as outlined in online appendix table A2). Following the approach of Coen Teulings and Nikolay Zubanov (2014), we also control for subsequent TFP innovations occurring between $t = 0$ and $t = K$, which reduce the influence of serial correlation in TFP innovations on estimates of β_1. Finally, standard errors are clustered by country-industry.

Figure 3 reports local projection estimates and confidence intervals for the relationship between a TFP innovation shock, measured as an increase

Figure 3. Local Projection Estimates of the Relationship between Total Factor Productivity Growth and Outcome Variables, 1970–2007[a]

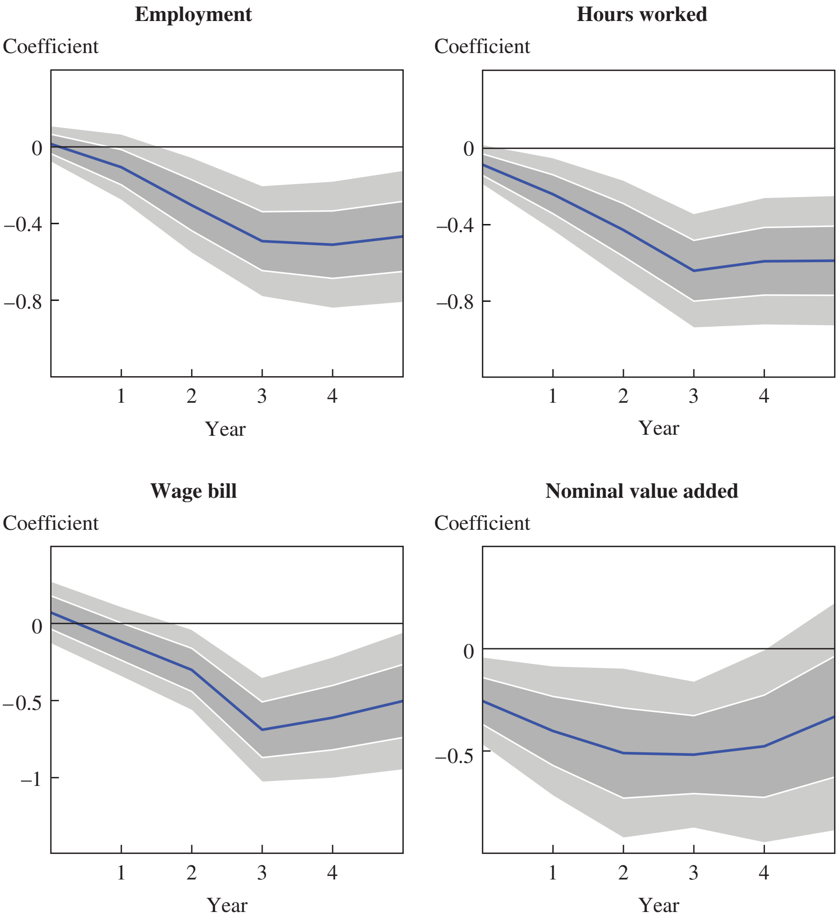

in TFP of 1 standard deviation, occurring between periods $t = -1$ and $t = 0$, and ensuing industry-level changes $\Delta_k \ln Y_{i,c} \equiv \ln Y_{i,c,t+k} - \ln Y_{i,c,t-1}$ for $K \in \{0, \ldots, 5\}$.[16] For all the outcome variables considered (employment, hours, wage bill, value added, and labor share), the local projection estimates indicate that TFP growth predicts small or negligible contemporaneous changes in the outcomes of interest that cumulate in ensuing years. In all cases, however, these effects plateau after three years, implying that

16. The standard deviation of TFP growth is 2.6 log points, as reported in online appendix table A6.

Figure 3. Local Projection Estimates of the Relationship between Total Factor Productivity Growth and Outcome Variables, 1970–2007[a] (*Continued*)

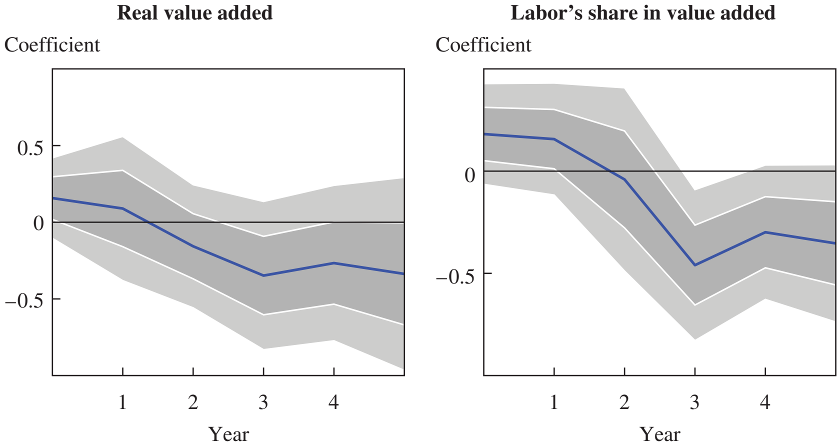

Sources: EU KLEMS; authors' calculations.

a. The coefficients are for observed, own-industry TFP shocks in year –1, and are rescaled to have a standard deviation of 1. The estimates include country-year and sector fixed effects, one lag of TFP and outcome variable growth, and controls for TFP shocks over the projection horizon. The darker shading denotes the 70 percent confidence interval and the lighter shading denotes the 95 percent confidence interval.

no more than four lags of the independent variable are needed to capture the impulse response of a contemporaneous shock. For completeness, we include five lags in our main specifications, though we shorten the lag structure when analyzing subintervals of the data.

II.A. Within-Industry Direct Effects: Own-Industry TFP and Own-Industry Outcomes

Our initial estimates, reported in table 5, consider the within-industry "direct" effects of TFP growth on own-industry outcomes. We fit ordinary least squares, first-difference models of the form

$$(3)\quad \Delta \ln Y_{i,c,t} = \beta_0 + \sum_{k=0}^{5} \beta_1^k \Delta \ln TFP_{i,c\neq c(i),t-k} + \alpha_c + \delta_t + \alpha_c \times t + \alpha_c \times (t = peak) + a_c \times (t = trough) + \varepsilon_{i,c,t},$$

where $\Delta \ln Y_{i,c,t}$ is an outcome of interest and, as above, i indexes industries, c indexes countries, and t indexes years; and the log change in TFP (contemporaneous plus five distributed lags) is the explanatory variable of interest. Because equation 3 is a first-difference specification estimated at

Table 5. Estimates of the Relationship between Total Factor Productivity Growth and Industry-Level Outcomes, 1970–2007[a]

	Annual change in log outcome variable by country-industry								
	Employment			*Hours*			*Wage bill*		
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*	*(8)*	*(9)*
$\Sigma\Delta\ln$(own-industry $TFP_{i,c,t-k}$)	−2.073***	−1.132***	−1.117***	−1.989***	−1.048***	−1.028***	−1.848***	−1.078***	−1.029***
	(0.172)	(0.144)	(0.147)	(0.187)	(0.160)	(0.162)	(0.272)	(0.220)	(0.225)
R^2	0.223	0.271	0.359	0.203	0.239	0.359	0.414	0.426	0.530
Model weights	Employment			Hours			Hours		
	Nominal value added			*Real value added*			*Labor share*		
	(10)	*(11)*	*(12)*	*(13)*	*(14)*	*(15)*	*(16)*	*(17)*	*(18)*
$\Sigma\Delta\ln$(own-industry $TFP_{i,c,t-k}$)	−1.332***	−0.629***	−0.609***	0.641	1.214***	1.238***	−0.504***	−0.571***	−0.541***
	(0.221)	(0.180)	(0.191)	(0.494)	(0.401)	(0.405)	(0.128)	(0.148)	(0.152)
R^2	0.299	0.313	0.368	0.105	0.137	0.183	0.063	0.064	0.147
Model weights	Nominal value added			Nominal value added			Nominal value added		
Fixed effects									
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sector	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Country × time trend	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No
Country × business cycle	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No
Country × year	No	No	Yes	No	No	Yes	No	No	Yes
No. of observations[b]	15,520	15,520	15,520	15,520	15,520	15,520	15,520	15,520	15,520

Sources: EU KLEMS; authors' calculations.

a. The dependent variable is the annual change in the log of the outcome variable by country-industry. TFP is other-country, within-industry TFP, and is rescaled to have a standard deviation of 1. The estimates shown are the sum of coefficients for the contemporaneous effect and five annually distributed lags. Standard errors clustered by country-industry are in parentheses. Statistical significance is indicated at the *10 percent, **5 percent, and ***1 percent levels.

b. The number of observations is equal to the number of country-industry cells multiplied by the number of years.

the industry-country-time level, it implicitly eliminates industry-country effects. We additionally include country and year indicator variables, which correspond to linear country and time trends in the first-difference model; country–time interaction terms, which allow country trends to accelerate or decelerate over the sample interval; and country-specific cyclical peak and trough indicators interacted with country indicators to account for country-specific business cycle effects. All models are weighted by industries' time-averaged shares of the relevant weighting variable—employment, hours, or value added—within countries, multiplied by time-varying country shares of the weighting variable, and standard errors are clustered at the level of country–industry pairs.

The top left panel of table 5 presents estimates for industry-level employment, measured as the log number of workers (encompassing both employees and the self-employed). We estimate that industries experiencing relative gains in productivity exhibit relative declines in employment. The point estimate of −2.07 in column 1, corresponding to the sum of the six β_1^k coefficients, implies that an increase of 1 standard deviation in own-industry TFP (2.58 log points) predicts a fall in own-industry employment of approximately 2 log points. This estimate implies that the estimated elasticity of employment to TFP growth is below 1 (0.80 = 2.07 ÷ 2.58)—that is, there is a partial industry-level demand offset (compare with Bessen 2017).

Columns 2 and 3 of table 5 stress-test this estimate by adding five major sector group fixed effects, and by replacing the country-trend and country–business cycle controls with an exhaustive set of country-year indicator variables. The inclusion of sector group trends reduces the point estimate from −2.07 to −1.13, and increases precision. This pattern suggests that TFP innovations may spill over across industries within a sector. We subsequently model these spillovers in the next section, when we add input–output linkages to the regression model; meanwhile, we add sector group dummies (reflecting sector group trends in the log-level models) to all subsequent models, so our primary identification comes from within-sector, between-industry comparisons. Conditional on the inclusion of these sector group trends, the addition of a full set of country-year dummies in column 3 has almost no impact on the magnitude or precision of the point estimates. This insensitivity is worth bearing in mind because we do *not* include exhaustive country-year dummies in our main models; these dummies would interfere with the identification of input–output linkages, which have much lower country-year variability than own-industry TFP.

The top middle panel of table 5, which reports analogous estimates for log hours of labor input, finds an almost identical slope as for employment, indicating that most of the employment adjustment to productivity changes occurs on the extensive margin. The top right panel explores the relationship between TFP and nominal industry wage bill changes. These point estimates are also similar to those for hours and employment, suggesting that industry (relative) nominal wages are not much affected by TFP changes; rather, the industry-level relationship between TFP and wage bill changes stems from employment shifts.

We turn to output measures in the bottom left and bottom middle panels of table 5. Rising industry TFP predicts significant relative declines in industry-level nominal value added (bottom left panel) and significant relative rises in real industry value added (bottom middle panel), implying (logically) that rising industry productivity lowers industry prices.

Comparing the estimates in the bottom left and bottom middle panels of table 5 reveals that a rise in industry TFP predicts a smaller (less negative) change in nominal value added than in the wage bill. This suggests that rising TFP predicts a relative fall in labor's share of industry value added.[17] The bottom right panel of the table confirms this implication: A rise in TFP of 1 standard deviation predicts a fall in an industry's labor share of value added of about 0.55 percentage point over a five-year horizon.

We have implemented a large number of tests of the robustness of these estimates, which are reported in table 6. These include weighting all countries equally rather than by their value-added shares (top rows); eliminating the contemporaneous TFP term from the distributed lag model (second group of rows); eliminating the self-employed from our employment, wage bill, and labor share models (third group of rows); imputing zeros to the TFP measures in cases where the reported values are negative (fourth group of rows);[18] estimating equation 3 using five-year-long first differences in

17. Because the wage bill regression is weighted by hours shares and the value-added regression by value-added shares, the precise impact of TFP growth on the labor share cannot be directly inferred from a comparison of these two columns.

18. Thirty-six percent of all country-industry-year TFP growth observations are negative. This is most frequently the case for renting of machinery and equipment, computer and related activities, research and development, and other business activities (codes 71–74); other community, social, and personal service activities (code O); hotels and restaurants (code H); and real estate activities (code 70). But it occurs in all industries to some extent. The likely cause is that annual frequency TFP calculations incorporate a fair amount of measurement error, leading to short-run intervals where nominal value added rises less rapidly than the share-weighted growth of labor and capital inputs.

Table 6. Robustness Tests for Estimates in Table 5[a]

	Annual change in log outcome variable by country-industry					
	Employment *(1)*	*Hours* *(2)*	*Wage bill* *(3)*	*Nominal value added* *(4)*	*Real value added* *(5)*	*Labor share* *(6)*
All countries given equal weight[b,c]						
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$	−1.038***	−0.996***	−0.888***	−0.603***	1.040***	−0.426***
	(0.123)	(0.125)	(0.146)	(0.147)	(0.182)	(0.108)
R^2	0.331	0.335	0.565	0.395	0.218	0.104
No. of observations	15,520	15,520	15,520	15,520	15,520	15,520
Excluding contemporaneous effect[c,d]						
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$	−1.038***	−0.985***	−1.039***	−0.719***	0.947**	−0.423***
	(0.142)	(0.153)	(0.198)	(0.157)	(0.367)	(0.145)
R^2	0.358	0.358	0.530	0.367	0.174	0.146
No. of observations	15,520	15,520	15,520	15,520	15,520	15,520
Excluding self-employed[b,c]						
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$	−1.156***	−1.056***	−0.996***	−0.609***	1.238***	−0.528***
	(0.156)	(0.163)	(0.218)	(0.191)	(0.405)	(0.142)
R^2	0.384	0.386	0.580	0.368	0.183	0.147
No. of observations	15,520	15,520	15,520	15,520	15,520	15,520
Setting negative TFP growth to zero[b,c]						
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$	−1.109***	−0.962***	−0.880***	−0.490**	1.880***	−0.690***
	(0.223)	(0.234)	(0.309)	(0.228)	(0.575)	(0.186)
R^2	0.350	0.352	0.528	0.367	0.186	0.145
No. of observations	15,520	15,520	15,520	15,520	15,520	15,520

(*continued on next page*)

Table 6. Robustness Tests for Estimates in Table 5[a] (*Continued*)

	Annual change in log outcome variable by country-industry					
	Employment (1)	*Hours (2)*	*Wage bill (3)*	*Nominal value added (4)*	*Real value added (5)*	*Labor share (6)*
Five-year-long first differences[c,e]						
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$	–0.683***	–0.636***	–0.713***	–0.472***	0.631***	–0.348***
	(0.090)	(0.097)	(0.119)	(0.115)	(0.231)	(0.104)
R^2	0.505	0.490	0.787	0.687	0.263	0.119
No. of observations	2,820	2,820	2,820	2,820	2,820	2,820
EU KLEMS 2000–15 data[b]						
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$	–1.194***	–0.943***	–0.904**	0.070	0.896	–0.633*
	(0.304)	(0.310)	(0.359)	(0.286)	(0.562)	(0.368)
R^2	0.365	0.492	0.331	0.272	0.304	0.094
No. of observations	3,148	3,148	3,148	3,148	3,148	3,148
Model weights	Employment	Hours	Hours	Value added	Value added	Value added

Sources: EU KLEMS; authors' calculations.

a. TFP is other-country, within-industry TFP, and is rescaled to have a standard deviation of 1. All estimates include country, year, and country-year fixed effects. Standard errors clustered by country-industry are in parentheses. Statistical significance is indicated at the *10 percent, **5 percent, and ***1 percent levels.

b. The estimates shown are the sum of coefficients for the contemporaneous effect and five annually distributed lags.

c. These panels contain sector fixed effects.

d. The estimates shown are the sum of coefficients for five annually distributed lags.

e. This panel contains country, year, and country-year fixed effects, but the years are defined as five-year intervals.

place of annual first differences (fifth group of rows);[19] and using data from the 2000–15 period from the 2017 release of the EU KLEMS data (van Ark and Jäger 2017), thus adding eight additional outcome years at the cost of dropping prior decades and several countries (bottom rows).[20] Results are remarkably stable across these many sets of estimates, though precision is much lower for models fitted using the short 2000–15 panel.

The robust negative industry-level relationships between TFP and both employment and labor's share of value added seen in tables 5 and 6 are central inputs into our subsequent analysis. We stress that these findings do *not* by themselves imply that productivity growth depresses either employment or the labor share in the aggregate. Indeed, these direct within-industry relationships do not at present incorporate any of the potentially countervailing effects operating through other channels, including input–output linkages, compositional shifts, and final demand effects. Before incorporating these links in the next section, we perform a validity test on our main technology measure.

II.B. Applying Direct Measures of Technological Progress

Our omnibus measure of productivity-augmenting technological change, TFP, has the advantage of not being bound to a specific set of technologies or their associated measurement challenges. But TFP's strength is also its weakness. Because it is an accounting residual, one can only speculate on the underlying sources of technological progress that contribute to rising TFP. To partially address this concern, we test whether our key results above hold when we focus on a specific margin of technological advancement: industry-level patenting flows (Acemoglu, Akcigit, and Kerr 2016).

Using data from Autor and others (2017a), who match patent grants to their respective corporate owners, and then to industry codes based on corporate owners' industry affiliations, we construct counts of patent grants and patent citations by year for patents granted to both U.S. and non-U.S. inventors using data from the U.S. Patent and Trademark Office that use U.S. Standard Industrial Classification codes, cross-walked to the EU KLEMS industry level. Aggregate summary statistics for standardized

19. These estimates are obtained from full-length five-year intervals (1970–75, 1975–80, . . . , 2000–05) only; and the reported coefficients reflect the effect of TFP growth occurring over the previous five-year interval.

20. More recent EU KLEMS releases cover a smaller set of countries and rely on backcasting data preceding 1995. We use a balanced panel of 12 countries—Austria, Belgium, Denmark, Finland, France, Germany, Italy, the Netherlands, Spain, Sweden, the United Kingdom, and the United States—over the period 2000–15.

patent counts and patent citations are reported in online appendix table A6, while online appendix table A7 reports the mean log number of patent grants and patent citations by industry and by inventor nationality (U.S. versus non-U.S.), and online appendix table A8 summarizes industry-level trends by decade and sector. These tables highlight the substantial heterogeneity in patent flows across industries and over time, with the highest levels of patenting occurring in chemicals and electrical equipment, and the lowest occurring in education. Patent grants rise across the decades while citations fall in the most recent decade, reflecting the substantial lag between patent grants and patent citations. Although citations are likely a better measure of innovation than the raw count of patent grants (Trajtenberg 1990), citations may understate innovation in the final years of the sample because they arrive with a lag. In what follows, we report results using both measures of patenting activity.

Given that patenting activity is an input into the industry-level innovation and automation process, it should predict TFP growth. To verify this supposition, we estimate industry-level descriptive regressions of the form

$$(4) \qquad \Delta \ln TFP_{i,c,t} = \beta_0 + \sum_{k=0}^{3} \beta_1^k \ln PAT_{i,c\neq c(i),t-k} + \alpha_c + \delta_t + \alpha_c \times (t = peak) + a_c \times (t = trough) + \varepsilon_{i,c,t},$$

where $\Delta \ln TFP_{i,c,t}$ is the measured change in industry-level TFP, and $\ln PAT_{i,c\neq c(i),t}$ is the log count of industry-level patents, which are normalized to have a standard deviation of 1. Paralleling the specifications given above, we include both contemporaneous patenting activity and a set of annually distributed lags. Analogous to our strategy of using other-country ("leave out") TFP growth by industry, we use patenting activity by *non-U.S.* inventors as predictors of U.S. TFP growth and, similarly, use patenting activity by U.S. inventors as predictors of *non-U.S.* TFP growth.

The estimates of equation 4, reported for patent counts in the upper rows of table 7 and for patent citations in the lower rows, confirm that patent flows are a strong predictor of industry TFP growth. A rate of industry patents or patent citations that is 1 standard deviation higher predicts about 0.6 log point faster industry TFP growth ($t = 2.9$). This relationship is robust; adding year effects (column 2), country–business cycle effects (column 3), and country-year effects (column 4) to these first-difference models has almost no impact on the magnitude or precision of the predictive relationship.

Table 8 explores the relationship between patenting activity and the evolution of industry-level labor input, value added, and factor payments.

Table 7. Predictive Relationships between Industry Patenting Activity and Total Factor Productivity Growth, 1970–2007[a]

	100 × annual change in log TFP by country-industry			
	(1)	*(2)*	*(3)*	*(4)*
$\Sigma\ln(\text{patents}_{i,c,t-k})$	0.574***	0.602***	0.602***	0.603***
	(0.197)	(0.202)	(0.202)	(0.204)
R^2	0.061	0.137	0.138	0.142
No. of observations[b]	16,518	16,518	16,518	16,518
$\Sigma\ln(\text{patent citations}_{i,c,t-k})$	0.608***	0.647***	0.648***	0.649***
	(0.208)	(0.229)	(0.230)	(0.233)
R^2	0.054	0.139	0.140	0.143
No. of observations[b]	16,479	16,479	16,479	16,479
Fixed effects				
Country	Yes	Yes	Yes	Yes
Year	No	Yes	Yes	Yes
Country × time trend	No	No	Yes	No
Country × business cycle	No	No	Yes	No
Country × year	No	No	No	Yes

Sources: EU KLEMS; U.S. Patent and Trade Office; authors' calculations.

a. Log patents and log patent citations are rescaled to have a standard deviation of 1. The estimates shown are the sum of coefficients for the contemporaneous effect and three annually distributed lags. Standard errors clustered by country-industry are in parentheses. Statistical significance is indicated at the *10 percent, **5 percent, and ***1 percent levels.

b. The number of observations is equal to the number of country-industry cells multiplied by the number of years.

Following the template of the tables presented above, we report regressions of industry-level first differences in outcome variables on log industry patent counts or patent citations—contemporaneous and five annually distributed lags—and the full set of controls used in table 7.[21] Comparable to the pattern of results for TFP, we find that industry-level patent citation flows predict a fall in own-industry employment and hours, a decline in nominal value added, a rise in real value added, and, most important, a fall in own-industry labor share.[22] These findings hold for both measures of patenting activity—patent counts and patent citations. Though precision is far lower for the patent-based estimates than TFP-based estimates—likely because we effectively have patenting data for only two countries, U.S. and non-U.S.—we view these findings as supportive of our main results.

21. Since the majority of variation in patenting reflects stable, cross-industry differences rather than over-time, within-industry fluctuations, we exclude sector-specific indicators from these models (which would otherwise absorb most identifying variation). Due to this limited variation, we confine our patent analysis to direct (own-industry) effects.

22. Due to the differences in underlying units, the magnitude of coefficients cannot be directly compared between the TFP and patents models.

Table 8. The Relationship between Patenting and Industry-Level Outcomes, 1970–2007[a]

	100 × annual change in log outcome variable by country-industry								
	Employment			*Hours*			*Wage bill*		
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*	*(8)*	*(9)*
$\Sigma\ln(\text{patents}_{i,c,t-k})$	–0.328*	–0.261	–0.267	–0.303	–0.243	–0.243	–0.420*	0.039	0.029
	(0.187)	(0.197)	(0.201)	(0.192)	(0.205)	(0.209)	(0.219)	(0.222)	(0.226)
R^2	0.036	0.130	0.223	0.032	0.137	0.261	0.117	0.387	0.492
No. of observations	15,456	15,456	15,456	15,456	15,456	15,456	15,456	15,456	15,456
$\Sigma\ln(\text{patent citations}_{i,c,t-k})$	–0.327	–0.239	–0.246	–0.338	–0.206	–0.211	–0.769***	0.097	0.087
	(0.208)	(0.230)	(0.235)	(0.213)	(0.239)	(0.244)	(0.272)	(0.263)	(0.269)
R^2	0.089	0.280	0.336	0.023	0.104	0.150	0.006	0.062	0.147
No. of observations	15,417	15,417	15,417	15,417	15,417	15,417	15,417	15,417	15,417
Model weights	Employment			Hours			Hours		
	Nominal value added			*Real value added*			*Labor share*		
	(10)	*(11)*	*(12)*	*(13)*	*(14)*	*(15)*	*(16)*	*(17)*	*(18)*
$\Sigma\ln(\text{patents}_{i,c,t-k})$	–0.437*	–0.133	–0.135	0.607**	0.678***	0.672**	–0.263**	–0.213*	–0.222
	(0.230)	(0.209)	(0.210)	(0.250)	(0.256)	(0.261)	(0.133)	(0.129)	(0.135)
R^2	0.089	0.280	0.336	0.023	0.104	0.150	0.006	0.062	0.147
No. of observations	15,456	15,456	15,456	15,456	15,456	15,456	15,456	15,456	15,456

$\Sigma\ln(\text{patent citations}_{i,c,t-k})$	−0.729***	−0.099	−0.121	0.553**	0.738***	0.731**	−0.329**	−0.242*	−0.235*
	(0.259)	(0.215)	(0.217)	(0.272)	(0.285)	(0.291)	(0.145)	(0.141)	(0.142)
R^2	0.116	0.284	0.341	0.021	0.106	0.152	0.008	0.066	0.153
No. of observations	15,417	15,417	15,417	15,417	15,417	15,417	15,417	15,417	15,417
Model weights	Nominal value added			Nominal value added			Nominal value added		
Fixed effects									
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Country × time trend	No	Yes	No	No	Yes	No	No	Yes	No
Country × business cycle	No	Yes	No	No	Yes	No	No	Yes	No
Country × year	No	No	Yes	No	No	Yes	No	No	Yes

Sources: EU KLEMS; U.S. Patent and Trade Office; authors' calculations.

a. Log patents and log patent citations are rescaled to have a standard deviation of 1. The estimates shown are the sum of coefficients for the contemporaneous effect and five annually distributed lags. Standard errors clustered by country-industry are in parentheses. Statistical significance is indicated at the *10 percent, **5 percent, and ***1 percent levels.

III. Linking Micro to Macro

As underscored by the top panel of figure 1, it would be erroneous to conclude that because *relative* employment declines in industries experiencing rising productivity, *aggregate* employment falls as productivity rises. To move from this cautionary observation to a rigorous quantification of how industry-level productivity growth affects the aggregate employment and labor share, we next add three micro–macro linkages to our estimation and accounting framework: customer–supplier linkages; final demand effects; and composition effects.

III.A. Accounting for Customer–Supplier Linkages

The effect of productivity growth occurring in an industry is unlikely to be confined to the sector in which it originates. Industries facing lower input prices or higher-quality inputs from their suppliers may increase purchases; similarly, industries whose customers are experiencing rising productivity may face rising or falling output demands. We account for these input–output linkages by adding two terms to equation 3:

$$\text{(5)} \quad \Delta \ln Y_{i,c,t} = \beta_0 + \sum_{k=0}^{5} \beta_1^k \Delta \ln TFP_{i,c \neq c(i),t-k} + \sum_{k=0}^{5} \beta_2^k \Delta \ln \widetilde{TFP}^{SUP}_{j \neq i,c,t-k} + \sum_{k=0}^{5} \beta_3^k \Delta \ln \widetilde{TFP}^{CUST}_{j \neq i,c,t-k} + \alpha_c + \delta_t + \gamma_s + \alpha_c \times t + a_c \times (t = peak) + a_c \times (t = trough) + \varepsilon_{i,c,t}.$$

These additional terms, $\widetilde{TFP}^{SUP}_{j \neq i,c,t}$ and $\widetilde{TFP}^{CUST}_{j \neq i,c,t}$, measure the weighted sum of TFP growth in all other domestic industries $j \neq i$, which are, respectively, the suppliers and customers of industry i:[23]

$$\text{(6)} \quad \Delta \ln \widetilde{TFP}^{L}_{j \neq i,c,t} = \sum_{j=1}^{J} weight^{L}_{j \neq i,c} \times \Delta \ln TFP^{L}_{j \neq i,c,t}, \ \forall \ L \in \{SUP, CUST\}.$$

The supplier and customer weights used for this calculation are obtained from input–output coefficients from the World Input–Output Database and are averaged over the period 1995–2007. The supplier weights are equal to each domestic supplier industry j's value added as a share of the value added of industry i, capturing the importance of supplier industries j in the

23. We eliminate the on-diagonal (own-industry) term from the input–output measures because these are captured by the direct TFP terms (β_1^k).

production of industry i's output. Analogously, the customer weights are the shares of value added of each industry i that are used in domestic industry j's final products, capturing the importance of industries j as end consumers of industry i's output.[24] These weights account not only for shocks to an industry's immediate domestic suppliers or buyers but also for the full set of input–output relationships among all connected domestic industries (that is, the Leontief inverse). We renormalize both the customer and supplier TFP terms to have a standard deviation of 1, with summary statistics reported in online appendix table A6. As with our main (direct) measure of TFP, these supplier and customer TFP linkage terms are calculated using industry-level, leave-out means of TFP growth in all other countries in the sample.

The estimates of equation 5, reported in the top half of table 9, indicate that productivity growth emanating from *supplier* industries predicts steep increases in the employment and hours of labor input of *customer* industries (though not in their nominal wage bill, value added, or labor share). Specifically, the point estimate of 0.97 on the supplier-industry TFP term in column 1 indicates that a rise of 1 standard deviation in an industry's supplier productivity predicts an employment gain of 97 log points. This effect is almost identical in magnitude but opposite in sign to the estimated direct effect of TFP growth of –0.95 on own-industry employment. Thus, this input–output linkage reveals a first channel by which direct effects of productivity growth on own-industry outcomes may be offset by effects accruing outside the originating sector.

Conversely, productivity growth emanating from customer industries (the third row of the top half of table 9) generally has negligible and always insignificant estimated effects on employment, hours, wage bill, value added, and labor share in supplier industries. This result is consistent with the simple Cobb–Douglas input–output framework developed by Acemoglu, Ufuk Akcigit, and William Kerr (2016), where productivity innovations in a given industry lead to output gain in its customer industries—those benefiting from its price declines—but have no net effect on its supplier sectors, where price and quantity effects are offsetting.

A third important pattern revealed by table 9 is that our earlier estimates of the relationship between TFP growth and own-industry outcomes are essentially unaffected by the inclusion of the customer and supplier terms (compare the point estimates in tables 5 and 9). Thus, our initial findings

24. Although every industry is potentially both a customer and supplier to every other industry, the terms "customer" and "supplier" refer to the direction of flows of inputs and outputs: Suppliers produce outputs that are purchased by (downstream) customers; and customers purchase inputs produced by (upstream) suppliers.

Table 9. Estimates of the Relationship between Total Factor Productivity Growth and Industry-Level Outcomes, 1970–2007[a]

	Annual change in log outcome variable by country-industry					
	Employment (1)	*Hours (2)*	*Wage bill (3)*	*Nominal value added (4)*	*Real value added (5)*	*Labor share (6)*
Industry effects						
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$	–0.951***	–0.869***	–1.052***	–0.579***	1.243***	–0.584***
	(0.144)	(0.160)	(0.233)	(0.201)	(0.398)	(0.171)
$\Sigma\Delta\ln(\text{supplier-industry TFP}_{j\neq i,c,t-k})$	0.971***	1.028***	0.196	0.376	0.269	–0.029
	(0.223)	(0.237)	(0.313)	(0.291)	(0.426)	(0.269)
$\Sigma\Delta\ln(\text{customer-industry TFP}_{j\neq i,c,t-k})$	0.097	0.159	–0.121	–0.410*	0.253	–0.110
	(0.128)	(0.152)	(0.202)	(0.243)	(0.221)	(0.178)
Fixed effects						
Country	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Sector	Yes	Yes	Yes	Yes	Yes	Yes
Country × time trend	Yes	Yes	Yes	Yes	Yes	Yes
Country × business cycle	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.280	0.252	0.428	0.317	0.142	0.069
No. of observations[b]	15,520	15,520	15,520	15,520	15,520	15,520
Model weights	Employment	Hours	Hours	Value added	Value added	Value added

Aggregate elasticities						
$\Sigma\Delta\ln$(aggregate real $VA_{j\neq i,c,t-k}$)	0.633***	0.558***			0.907***	
	(0.073)	(0.083)			(0.084)	
$\Sigma\Delta\ln$(aggregate nominal $VA_{j\neq i,c,t-k}$)			1.083***	1.030***		0.071***
			(0.026)	(0.024)		(0.025)
Sector fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.227	0.194	0.414	0.300	0.110	0.006
No. of observations[b]	15,520	15,520	15,520	15,520	15,520	15,520
Model weights	Employment	Hours	Hours	Value added	Value added	Value added

Sources: EU KLEMS; World Input–Output Database; authors' calculations.

a. TFP is other-country TFP, and is rescaled to have a standard deviation of 1. The estimates shown are the sum of coefficients for the contemporaneous effect and five annually distributed lags. Standard errors clustered by country-industry are in parentheses. Statistical significance is indicated at the *10 percent, **5 percent, and ***1 percent levels.

b. The number of observations is equal to the number of country-industry cells multiplied by the number of years.

for the relationship between TFP growth and own-industry employment and labor share are unaltered.

III.B. Accounting for Final Demand Effects

The lower half of table 9 adds a third channel of response: final demand effects accruing through the contribution of productivity growth to aggregate value added. To capture these final demand effects, we estimate the relationship between country-specific aggregate economic growth (contemporaneous and five distributed lags) and industry-specific inputs using the following specification:

$$(7) \qquad \Delta \ln Y_{i,c,t} = \lambda_0 + \sum_{k=0}^{5} \lambda_1^k \Delta \ln VA_{j \neq i,c,t-k} + \alpha_s + \varepsilon_{i,c,t}.$$

The explanatory variable of interest in this equation, $\Delta \ln VA_{j \neq i,c,t}$, is the growth of own-country real or nominal value added, where the subscript $j \neq i$ highlights that we exclude own-industry output from the explanatory measure for each industry to eliminate any mechanical correlation between aggregate growth and industry outcomes. These stacked first-difference regression models drop the country, year, trend, and business cycle indicators used in equation 5, so that identification largely arises from country and year time series. Because these are first-difference models, however, they implicitly eliminate industry-country effects.

The estimates of equation 7, reported in the lower half of table 9, document a second countervailing effect of industry-specific productivity innovations on aggregate outcomes: Each log point gain in country-level real value added predicts an approximately 0.6 log point rise in same-country, other-industry employment and hours. Similarly, each log point gain in country-level nominal value added predicts essentially a one-for-one rise in same-country, other-industry wage bill and nominal value added, as well as a very modest but statistically significant rise in same-country, other-industry labor share (the estimated elasticity is 0.071). Because TFP growth emanating from any one sector raises the real aggregate value added in the country where it occurs, these estimates imply that each industry's productivity growth contributes to aggregate labor demand across all other sectors.[25]

25. We report a pure stacked country-level time series version of these estimates in online appendix table A9, in which we eliminate industry-level variation entirely and instead use only country-year observations. These point estimates are similar to those used in the bottom half of table 9, which we prefer because they eliminate the mechanical relationship between own-industry and country-level aggregate outcomes.

Figure 4. Industry-Level Total Factor Productivity Growth versus Industries' Shares of Country-Level Nominal Value Added, 1970–2007[a]

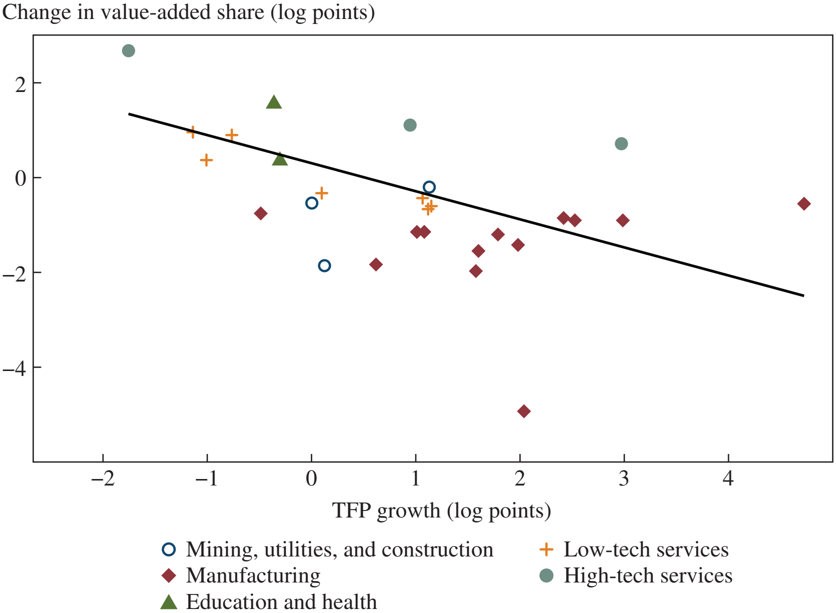

Source: EU KLEMS.

a. All values are expressed as annual, unweighted average changes across country-years in log points. The line shows the linear fit weighted by industries' value-added shares. Statistics: $\beta = -0.606$ (SE = 0.158), $R^2 = .361$.

III.C. Accounting for Compositional (Between-Sector) Effects

The estimates given in table 9 reveal one further mechanism by which sectoral productivity gains affect the aggregate labor share: by shifting relative sector sizes. Column 4, in the top half of table 9, shows that a rise in own-industry TFP growth predicts a *fall* in industry-level nominal value added with an elasticity of –0.58. This finding implies that sectors with rising productivity will tend to shrink as a share of nominal value added. Figure 4 confirms this intuition by depicting a scatter plot of the bivariate relationship between industry-level TFP growth and the change in industries' log shares of own-country nominal value added (averaged over years and across countries). On average, industries that experience 1 log point faster TFP growth than the economy-wide average lose about 0.6 log point as a share of nominal economy-wide value added.

Applying this observation to the Oaxaca decomposition equation above (equation 1), it is immediately clear that uneven productivity growth across industries will shift the aggregate labor share through changes in relative sector sizes. If rapid productivity growth occurs in industries with relatively low labor shares (for example, manufacturing industries), this will indirectly *raise* the aggregate labor share; conversely, relatively rapid productivity growth in labor-intensive sectors (for example, education and health) will have the opposite effect.[26]

IV. Quantitative Implications

With these estimates in hand, we now quantify the implied contribution of TFP growth to the evolution of the aggregate employment and labor shares accruing through the four channels outlined above: own-industry, supplier and customer, final demand, and composition. We start with employment and hours, then turn to the labor share.

IV.A. Aggregate Employment and Hours Effects

The effect of TFP growth on employment and hours combines the first three of these effects: the own-industry (or "direct") effect, the supplier and customer effects, and the final demand effect. The first (own-industry) effect is equal to the sum of the β_1^k coefficients in equation 5 multiplied by their corresponding $\Delta \ln TFP_{i,c\neq c(i),t}$ terms, and aggregated by weighting these industry-level predictions by the time-averaged share of each industry in total employment or hours:

$$(8) \qquad \Delta \ln Y_{c,t}^{OWN} \equiv \frac{\partial \ln Y_{c,t}}{\partial \ln TFP_{i,c\neq c(i),t}} = \sum_{k=0}^{5} \beta_1^k \times \sum_{i=1}^{I} \omega_{ic} \times \Delta \ln TFP_{i,c\neq c(i),t}.$$

Here, $\ln Y_{c,t}$ is log employment or hours in country c in year t; $\sum_{k=0}^{5}\beta_1^k$ is the sum of coefficients in equation 5; $\omega_{i,c}$ is the time-averaged employment or hours share of industry i in its respective country; and $\Delta \ln TFP_{i,c\neq c(i),t}$ is own-industry TFP growth.

The supplier and customer effects are, analogously, equal to the sum of the β_2^k and β_3^k coefficients multiplied by their corresponding $\widetilde{TFP}_{j\neq i,c,t}^{SUP}$ and $\widetilde{TFP}_{j\neq i,c,t}^{CUST}$ terms, and then aggregated to the national level

26. The upstream and downstream linkages estimated in equation 5 can also contribute to the between-industry component of the labor share change through their effects on industry nominal output shares, though we estimate these effects to be comparatively small and statistically insignificant.

by weighting each by its time-averaged industry employment or hours shares ($\omega_{i,c}$):

$$(9) \qquad \Delta \ln Y_{c,t}^{SUP} \equiv \frac{\partial \ln Y_{c,t}}{\partial \ln \widetilde{TFP}_{j \neq i,c,t}^{SUP}} = \sum_{k=0}^{5} \beta_2^k \times \sum_{i=1}^{I} \omega_{i,c} \times \Delta \ln \widetilde{TFP}_{j \neq i,c,t}^{SUP},$$

and

$$\Delta \ln Y_{c,t}^{CUST} \equiv \frac{\partial \ln Y_{c,t}}{\partial \ln \widetilde{TFP}_{j \neq i,c,t}^{CUST}} = \sum_{k=0}^{5} \beta_3^k \times \sum_{i=1}^{I} \omega_{i,c} \times \Delta \ln \widetilde{TFP}_{j \neq i,c,t}^{CUST}.$$

The third component that we calculate is the final demand effect of TFP growth in each industry on employment or hours economy-wide, $\Delta Y_{c,t}^{FD}$. For any one industry, this contribution is equal to the product of four terms: (i) the effect of TFP growth in i on i's real value added ($\sum_{k=0}^{5} \beta_{1,VA}^k$); (ii) the effect of growth in i's real value added on total value added ($\phi_{i,c}$); (iii) the effect of growth in real value added on employment or hours in each industry $j \neq i$ ($\sum_{k=0}^{5} \lambda_1^k$); and (iv) the size of industry j relative to overall employment or hours in the economy ($\omega_{i,c}$).[27] To obtain the aggregate effect (summing across industries), we calculate:

$$(10) \qquad \Delta \ln Y_{c,t}^{FD} \equiv \frac{\partial \ln Y_{c,t}}{\partial \ln VA_{c,t}} \times \frac{\partial \ln VA_{c,t}}{\partial \ln TFP_{i,c \neq c(i),t}}$$

$$= \sum_{k=0}^{5} \lambda_1^k \times \sum_{i=1}^{I} \omega_{i,c} \left(\frac{\partial \ln VA_{c,t}}{\partial \ln VA_{i,c,t}} \times \frac{\partial \ln VA_{i,c,t}}{\partial \ln TFP_{i,c \neq c(i),t}} \right)$$

$$= \left(\sum_{k=0}^{5} \lambda_1^k \times \sum_{k=0}^{5} \beta_{1,VA}^k \right) \sum_{i=1}^{I} \omega_{i,c} \times \phi_{i,c}.$$

In this expression, $\ln Y_{c,t}$ is log employment or hours in country c in year t as before; $\sum_{k=0}^{5} \lambda_1^k$ is the estimated effect of the aggregate real value added on outcome Y from equation 7 reported in column 5 of the lower half of table 9; $\sum_{k=0}^{5} \beta_{1,VA}^k$ is the estimated direct effect of $\Delta \ln TFP$ in equation 5 on own-industry real value added (reported in column 5 in the lower half of table 9); $\omega_{i,c}$ is the time-averaged employment or hours share of industry i

27. In calculating $\Delta Y_{c,t}^{FD}$, we also include the customer and supplier TFP effects estimated in equation 5, though we suppress those terms above to conserve notation.

in its respective country; and $\phi_{i,c}$ is the time-averaged value-added share of industry i in country c.[28]

Figure 5 displays the results of these calculations for overall employment and hours of labor input, respectively. The first bar in the top panel of figure 5 corresponds to the direct effect of TFP growth on own-industry employment. Its height of –0.22 implies that, on average, productivity growth reduced own-industry employment by approximately 8.2 percent over the full 37-year period (0.22×37). The second bar (supplier effect), with a height of 0.35, indicates that the countervailing effect of rising supplier productivity on employment in customer industries more than offset this direct effect. The third bar (customer effect), with a height of 0.05, indicates that rising productivity in customer industries exerted a very modest positive employment effect in supplier industries. The fourth bar, with a height of 0.30, reflects the substantial contribution of rising productivity to overall employment operating through final demand. The fifth bar (net effect) sums these four components to estimate a net *positive* effect of productivity gains on aggregate employment, totaling about 18 log points ($0.48 \times 37 = 17.8$) over the outcome period.

The bottom panel of figure 5 reports the analogous exercise for hours of labor input rather than employment. We find comparable effects on this outcome: Although rising productivity reduces relative employment in the sectors in which it occurs, it augments employment in (downstream) customer sectors (as captured by the supplier effect) and boosts aggregate demand through its contribution to overall value added. As with employment, the net effect on hours is strongly positive.

To provide insight into how rising TFP spurs relative employment declines in directly affected industries while simultaneously generating rising employment in the aggregate, online appendix tables A11 and A12 report the contributions to employment growth by industry operating through each channel estimated above: direct effects, input–output linkages, and final demand effects. These contributions, underlying the aggregate employment growth predictions in figure 5, can be analyzed from two complementary perspectives. The first, reported in online appendix table A11, calculates the contribution of TFP growth *originating* in each industry to the predicted aggregate change in employment. The second,

28. This last term, $\phi_{i,c}$, is derived by differentiating the sum of industry log value added at the country level with respect to the log value added of industry i in country c, which is simply equal to i's share in country c's value added. Note that the sum of industry shares is less than 1, because we exclude nonmarket industries from the analysis, though they are logically included in aggregate national value added.

Figure 5. Predicted Effects of Total Factor Productivity Growth on Aggregate Employment and Hours of Labor Input, 1970–2007[a]

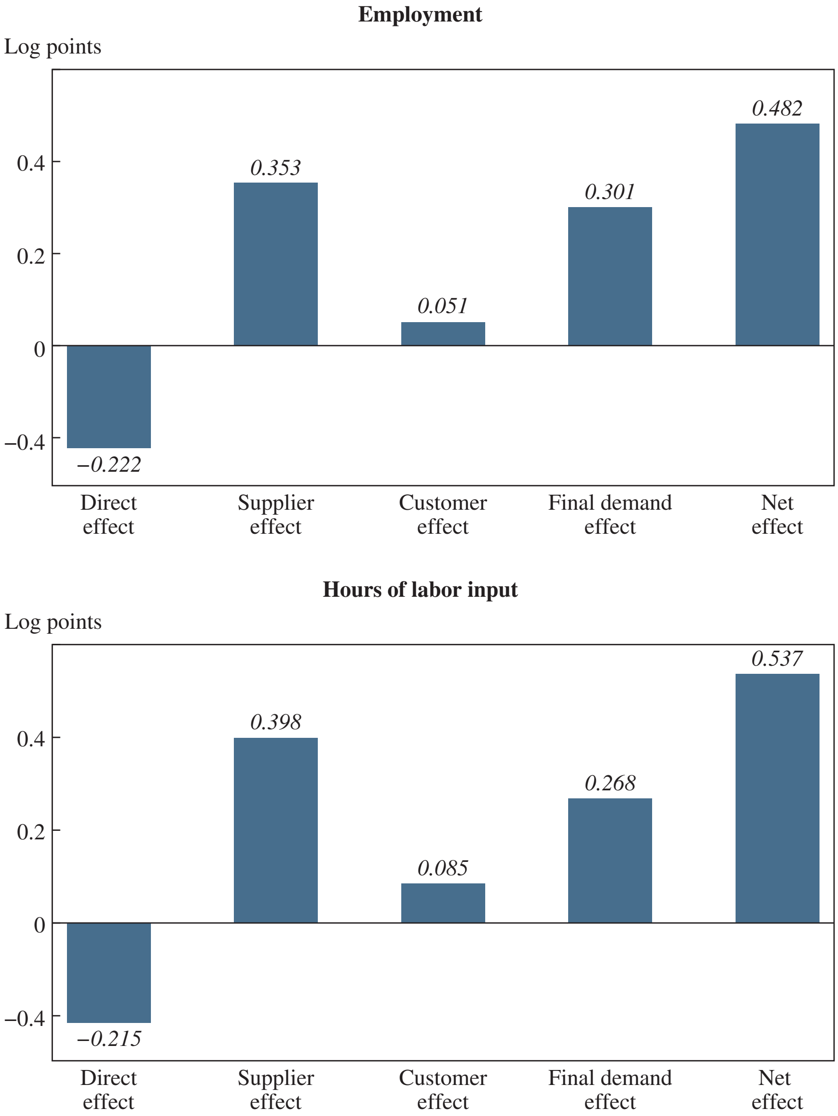

Source: Authors' calculations, based on table 9.

a. The units are the predicted annual change in the outcome variable expressed in log points. See the notes to table 9.

reported in online appendix table A12, enumerates the predicted effect of TFP growth originating throughout the economy on predicted employment growth in each *destination* industry, scaled by that industry's weight in aggregate employment.[29]

For the direct effect, the contributions to employment in the originating and destination industry are the same by definition because these direct effects operate only within industries. As shown in online appendix tables A11 and A12, the negative direct effects that we estimate for employment originate in industries that have experienced strong TFP growth (such as electrical and optical equipment, and transportation and storage), or industries that make up a large share of total value added (such as retail), or both.

Conversely, TFP growth originating in supplier and customer industries leads to employment and hours growth elsewhere in the economy through input–output linkages. The supplier/customer contribution of any given industry to aggregate employment depends on three terms: the industry's rate of TFP growth; the weight that industry has as a supplier or customer of other industries; and, in turn, the weight that those customer and supplier industries have in aggregate employment. Industries such as post and telecommunications, wholesale trade, financial intermediation, and transportation and storage produce important positive employment spillovers to other industries, in part because they are suppliers to a variety of service industries, which are themselves a large share of total employment. These industries highlight the potential of productivity growth in service industries to induce sizable positive employment spillovers. Conversely, other business activities—an important supplier industry—exhibits declining productivity, and thus contributes a meaningful negative employment spillover. Finally, manufacturing industries—such as chemicals, basic and fabricated metals, and electrical and optical equipment—make a large indirect contribution to employment in customer industries, due to their rapid productivity growth.[30]

Finally, each industry's TFP growth potentially contributes to employment via its effect on final demand. This effect depends on two terms:

29. We do not separately report contributions for hours worked because they are nearly identical to those for employment.

30. The indirect employment contribution made by productivity gains in customer industries is much smaller than the corresponding effect of productivity gains in supplier industries, and it is primarily driven by TFP growth in electrical and optical equipment, transportation equipment, and machinery.

the originating industry's rate of TFP growth, and its share in national value added. Hence, productivity growth in industries that make up a large share of value added has a larger effect on overall income. Manufacture of electrical and optical equipment (codes 30–33); post and telecommunications (code 64); financial intermediation (code J); manufacture of motor vehicles and transportation equipment (codes 34–35); manufacture of chemicals and chemical products (code 24); and wholesale trade, excluding motor vehicles (code 51) are the largest contributors by TFP source to final demand, reflecting their rapid productivity growth and substantial weight in aggregate value added.[31]

How successful is our approach in capturing the evolution of employment observed in the data? Figure 6 answers this question by comparing the industry-level employment predictions of our statistical model to observed employment changes, averaged across country-years. In each panel, employment growth predictions, obtained by summing across all channels in the model, are reported on the horizontal axis, while observed employment growth is reported on the vertical axis. The top panel of figure 6 plots the predicted versus observed log employment change by industry, while the bottom panel plots the predicted versus observed contribution that each industry makes to aggregate employment growth.[32] This figure makes evident that our model can account for a substantial part of the variation in employment growth by industry (the top panel), and the extent to which these industry effects contribute to aggregate job growth (the bottom panel). Each of the three channels featured in the model contributes to its predictive power. A regression of the observed contribution of each industry to aggregate employment growth on its predicted value based *only* on the direct (own-industry) effect yields an R^2 of .34. Adding customer and supplier effects to this prediction raises this R^2 to .45. Incorporating the final demand effect raises it further to .61. Given that the model exclusively uses variation in TFP across industries to form predictions, we consider this as strong confirmation of the utility of our accounting framework.

31. Observe that the contribution of final demand growth to employment and hours worked in *destination* industries reported in online appendix table A12 is directly proportional to the size of the industry in total employment.

32. The predicted versus observed employment contribution (the bottom panel of figure 6) depends on the proportional growth in each industry multiplied by its weight in overall employment, whereas the predicted versus observed employment change (the top panel of figure 6) depends on only the first of these terms.

Figure 6. Predicted versus Observed Employment Growth for Industry-Level Changes and Industry-Level Contributions to Aggregate Changes, 1970–2007[a]

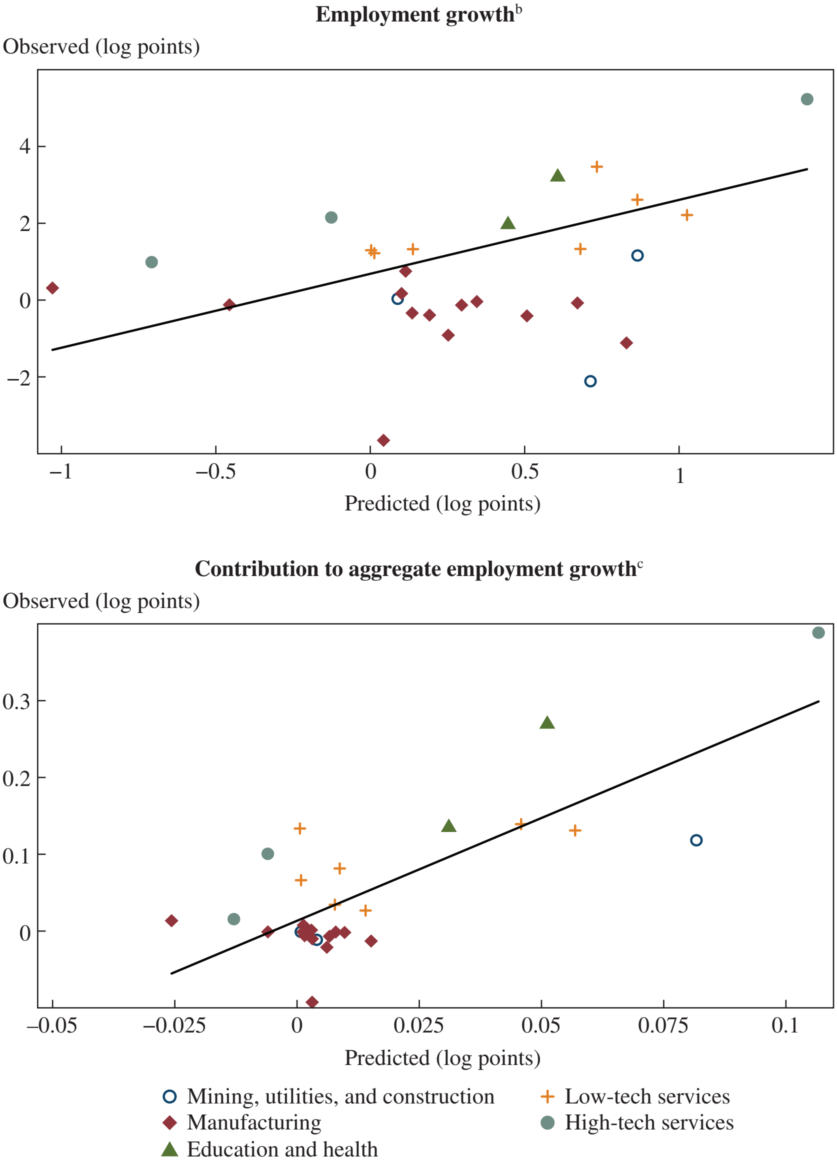

Sources: EU KLEMS; authors' calculations.

a. All values are expressed as annual, unweighted average changes across country-years in log points.

b. The line shows the linear fit weighted by industries' employment shares. Statistics: β = 1.925 (SE = 0.539), R^2 = .329.

c. The line shows the unweighted linear fit. Statistics: β = 2.676 (SE = 0.418), R^2 = .612.

IV.B. Aggregate Labor Share Effects

We now perform the analogous exercise for the implied effect of rising TFP on labor's share of value added. In this calculation, the own-industry, interindustry, and final demand effects are obtained analogously to those for employment and hours.[33] However, the labor share calculation includes a fourth channel: TFP-induced compositional shifts in value-added shares across industries. This between-industry composition effect is calculated as

$$(11)\quad \Delta \ln Y_{c,t}^{COMP} \equiv \sum_{i}^{I}\left(\Delta\hat{\omega}_{i,c} \times \overline{l}_{i,c}\right)$$

$$= \sum_{i}^{I}\left(\left\{\frac{\omega_{i,c}\exp\left(\sum_{k=0}^{5}\beta_{1,VA}^{k} \times \Delta \ln TFP_{i,c\neq c(i),t}\right)}{\sum_{i}^{I}\left[\omega_{i,c}\exp\left(\sum_{k=0}^{5}\beta_{1,VA}^{k} \times \Delta \ln TFP_{i,c\neq c(i),t}\right)\right]} - \omega_{i,c}\right\} \times \overline{l}_{i,c}\right).$$

Here, $\Delta\hat{\omega}_{i,c}$ is the predicted change in the value-added share of industry i in country c, and $\overline{l}_{i,c}$ is the time-averaged log labor share in industry i in country c. The terms $\omega_{i,c}$ and $\beta_{1,VA}^{k}$ are defined as in equation 10, again adjusted for the labor share model: The time-averaged weights $\omega_{i,c}$ are shares of nominal value added rather than shares of employment or hours worked, and the coefficients $\beta_{1,VA}^{k}$ reflect nominal rather than real value-added coefficients (shown in column 4 of table 9). Concretely, this prediction reflects the sum of induced shifts in each industry's share of own-country nominal value added ($\Delta\omega_{i,c}$, the expression in braces) multiplied by that industry's labor share.[34]

We report quantitative implications for labor's share of value added in figure 7. The first bar reflects the labor share effect associated with

33. The weights ($\omega_{i,c}$) used in equations 8, 9, and 10 are now time-averaged, industry value-added shares rather than employment or hours shares; and the final demand effect is calculated using aggregate increases in nominal rather than real value added. Hence, the coefficients $\sum_{k=0}^{5}\lambda_{1}^{k}$ and $\sum_{k=0}^{5}\beta_{1,VA}^{k}$ in equation 10 are taken from column 4 (rather than column 5) of, respectively, the lower and upper rows of table 9.

34. As with prior calculations, we incorporate customer and supplier TFP effects into this calculation but suppress them from the equation for simplicity. We have also estimated models that allow the aggregate income elasticities, estimated in the lower half of table 9, to vary by broad sector (thereby potentially admitting nonhomotheticities). This has negligible effects on the predicted composition changes, and we therefore do not report these specifications.

Figure 7. Predicted Effects of Total Factor Productivity Growth on the Aggregate Labor Share, 1970–2007[a]

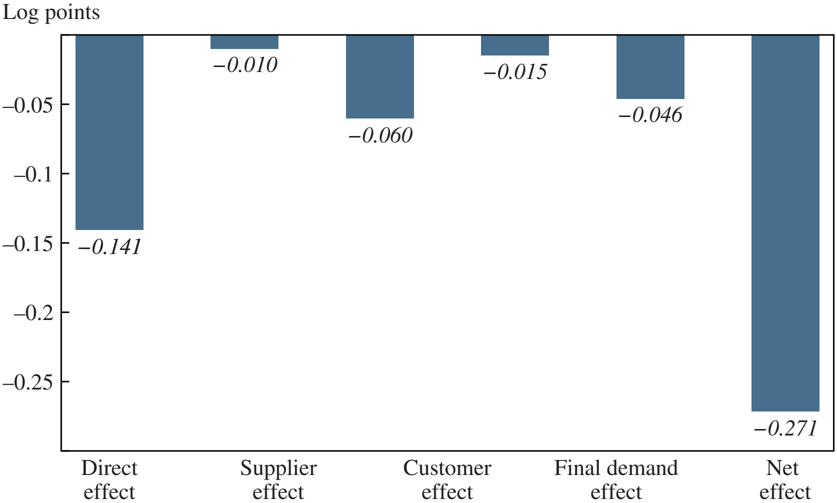

Source: Authors' calculations, based on table 9.

a. The units are the predicted annual change in the outcome variable expressed in log points. See the notes to table 9.

own-industry productivity growth. Its height of –0.14 suggests that, on average, own-industry productivity growth reduced the labor share by some 5.2 percent over the 37-year period (0.14×37). Unlike employment and hours worked, however, there are no positive countervailing effects from interindustry linkages or final demand; rather, these additional channels also serve to decrease the aggregate labor share, albeit by small amounts compared with the direct effect (–0.01, –0.06, and –0.02 log point annually for, respectively, the supplier, customer, and final demand effects). Finally, industry composition shifts resulting from a reallocation of value added across industries also predict a small net labor share decline: This effect amounts to about 1.7 percent over the entire period (0.046×37).

Taken together, all four channels operating on the labor share—direct, supplier/customer, final demand, and composition—predict a decline of –0.27 log point annually, or about 10 percent over the entire period (0.27×37). Most of this effect stems from the direct labor share–displacing effect operating within industries, combined with an absence of countervailing effects operating within industries. Compositional shifts modestly reinforce this trend. Given an initial average labor share of about 67 percent in our 19 countries (table 1), this corresponds to a nonnegligible predicted decline of 6 percentage points over the period 1970–2007, of which the large majority (0.225 log point annually—that is, 8.3 percent, or about 5.5 percentage points, over the entire period) is predicted to occur within industries.

Table 10 reports the separate industry-level contributions made to these overall predictions.[35] The first column shows each industry's contribution to the total predicted within-industry effects (that is, the predicted effects for own-industry TFP growth, interindustry linkages, and final demand taken together, which are largely driven by the own-industry effect). The second column analogously shows the contribution of each industry to the predicted between-industry effect shown in figure 7. Table 10 highlights that most industries experience a negative within-industry labor share effect. Predictably, some of the largest contributions are made by industries that have witnessed strong productivity growth, such as electrical and optical equipment, chemicals, basic and fabricated metals, and post and telecommunications. However, industries with more modest productivity growth but comprising relatively large shares of value added—such as wholesale trade, and transportation and storage—also contribute substantially to the aggregate within-industry effect. Real estate and other business activities are the only industries that contribute a small countervailing effect; here, positive within-industry labor share changes are predicted because these sectors have experienced negative TFP growth on average. Finally, several (public) services—such as education, health and social work, and other personal services—contribute almost nothing to the predicted aggregate labor share decline, because they have experienced virtually no measured productivity growth.

Table 10 also shows that the industry-specific contributions to the *composition* (that is, between-industry) effect are quite heterogeneous. In general, the predicted shift away from capital-intensive mining, utilities, and manufacturing industries tends to increase labor's share: In isolation, these industries contribute a predicted increase in the labor share of about 1.6 percent cumulated over the period (0.036×37). This is reinforced by contributions from (mostly high-tech) services, such as post and telecommunications, financial intermediation, and transportation and storage. However, real estate single-handedly contributes a large negative compositional effect of, on average, 0.086 log point annually, or over 3 percent

35. Unlike for employment and hours, most effects for the labor share are driven by the direct effect. As a result, there is no need to separately consider the industry contributions by source of TFP growth.

Table 10. Industry-Level Contributions to Predicted Within- and Between-Industry Components of the Change in Aggregate Labor Share, 1970–2007

ISIC code (rev. 3)	*Description*	*Within industry*	*Between industry*
C	Mining and quarrying	–0.003	0.001
15–16	Manufacture of food, beverages, and tobacco products	–0.006	–0.005
17–19	Manufacture of textiles, apparel, leather, and related products	–0.009	0.001
20	Manufacture of wood and wood products, excluding furniture	–0.004	0.001
21–22	Manufacture of paper and paper products, printing, and publishing	–0.009	0.001
23	Manufacture of coke, refined petroleum products, and nuclear fuel	0.000	0.000
24	Manufacture of chemicals and chemical products	–0.019	0.010
25	Manufacture of rubber and plastics products	–0.008	0.002
26	Manufacture of other nonmetallic mineral products	–0.005	0.001
27–28	Manufacture of basic and fabricated metals	–0.021	0.008
29	Manufacture of machinery and equipment not elsewhere classified	–0.013	0.000
30–33	Manufacture of electrical and optical equipment	–0.038	0.009
34–35	Manufacture of motor vehicles and transportation equipment	–0.016	0.000
36–37	Manufacture of furniture and manufacturing not elsewhere classified; recycling	–0.003	0.000
E	Electricity, gas, and water supply	–0.010	0.009
F	Construction	–0.006	–0.008
50	Sale, maintenance, and repair of motor vehicles and fuel	–0.002	0.000
51	Wholesale trade, excluding motor vehicles	–0.023	0.008
52	Retail trade, excluding motor vehicles; repair of personal and household goods	–0.018	0.002
H	Hotels and restaurants	0.003	–0.003
60–63	Transportation activities of travel agencies	–0.018	0.005
64	Post and telecommunications	–0.018	0.012
J	Financial intermediation	–0.017	0.009
70	Real estate activities	0.013	–0.086
71–74	Renting of machinery and equipment; computer and related activities; research and development; and other business activities	0.017	–0.008
M	Education	0.001	–0.002
N	Health and social work	0.001	–0.005
O	Other community, social, and personal service activities	0.006	–0.004
Total		–0.225	–0.046

Source: Authors' calculations, based on table 9.

Table 11. The Contribution of Total Factor Productivity Growth to the Within- and Between-Industry Components of the Change in Aggregate Labor Share, by Decade, 1970–2007

	Actual annual change in labor share in log points			*Predicted annual change in labor share in log points*		
Decade	*Total*	*Between industry*	*Within industry*	*Total*	*Between industry*	*Within industry*
1970s	0.513	–0.187	0.700	–0.294	–0.124	–0.169
1980s	–0.459	–0.183	–0.276	–0.365	–0.005	–0.360
1990s	–0.263	–0.075	–0.188	–0.202	0.005	–0.207
2000s	–0.861	–0.425	–0.436	–0.231	–0.091	–0.140

Source: Authors' calculations, based on table 9.

across the entire period. This prediction is consistent with the aggregate labor decomposition reported in table 4 and stems from three distinctive features of the real estate industry: a very low labor share relative to the economy-wide average, a rising share of value added, and zero or negative TFP growth.

IV.C. Why Has the Fall in Labor Share Accelerated?

Our results imply that technological progress, broadly construed, has been *employment-augmenting* but *labor share–displacing*—that is, generating net employment gains while serving to reallocate value added away from labor and toward other factors. But this observation raises a puzzle: If automation has been consistently labor share–displacing, why has the evolution of labor's share differed so sharply across the decades—rising during the 1970s, declining in the 1980s and 1990s, and then falling more steeply in the 2000s? We briefly take up this question here.

Table 11 reports our baseline model's predictions separately by decade. The first three columns report the observed annual log labor share change in each decade, both within and between industries, while the last three columns report the changes predicted by our baseline model. This table highlights the fact that, although our baseline approach explains a substantial part of the aggregate labor share fall observed since the 1980s, it fails to match two key features of the decade-specific patterns: the positive sign of the within-industry effect operating in the 1970s, and the observed acceleration of the within-industry log labor share decline in the 2000s. The proximate reason for both mismatches is clear: The bulk of the model's explanatory power for the labor share derives from the so-called direct effect—the differential decline of the labor share in industries with faster TFP growth;

thus, for the baseline approach to explain the time pattern of rising and then falling labor share across the decades, it would need to be the case that TFP growth was negative in the 1970s, became positive in the 1980s and 1990s, and then accelerated in the 2000s. This does not match the time pattern of TFP growth, however (see table 3). The model does slightly better at capturing the time pattern of between-industry effects—predicting larger compositional shifts in the 1970s and 2000s, which is approximately consistent with the data—but our explanatory power is limited here as well.[36]

Our empirical framework admits several mechanisms through which the effect of technological progress on the labor share may differ over time. One mechanism is that an acceleration of TFP growth will lead to a more rapid fall in the labor share. But as noted above, this explanation is a nonstarter because TFP growth decelerated in the 2000s, even as the fall in the labor share accelerated. Second, the locus of productivity growth may be differently distributed among industries in different eras. To the extent that industries experiencing rapid TFP gains are more (or less) labor-intensive or make up a larger (or smaller) share of the total economy, the aggregate labor share will decline more (or less) strongly through, respectively, compositional effects and within-industry effects. But table 11 suggests that these explanations have a limited bite. Allowing the sources of TFP growth to differ across decades, as we do in the table, does not explain the sharp decadal differences in the between- and within-industry contributions to the fall in the labor share.

A third possibility is that, all else being equal, a given amount of overall productivity growth might have different effects in different eras if the source of that productivity growth is changing—for example, if productivity growth increasingly stems from technologies that are relatively less labor-augmenting and relatively more labor share–displacing. Figure 8 suggests that this explanation has some promise. Akin to figure 1 above, figure 8 presents bivariate scatters of the relationship between industry-level TFP growth and changes in, respectively, industry-level log employment (the top panel) and industry-level log labor share (the bottom panel). Distinct from earlier figures, figure 8 depicts separate slopes by decade. The top panel shows a consistently stable, downward-sloping relationship between industry-level TFP growth and relative declines in employment, with a somewhat steepening slope after the 1970s. By contrast, the bottom panel

36. The substantial between-industry component of the falling labor share in the 2000s is, as above, due to the rapid growth of the real estate industry in value added, a phenomenon that is unlikely to be attributable to technological progress.

Figure 8. Industry-Level Total Factor Productivity Growth versus Industry-Level Employment and Labor Share by Decade, 1970–2007[a]

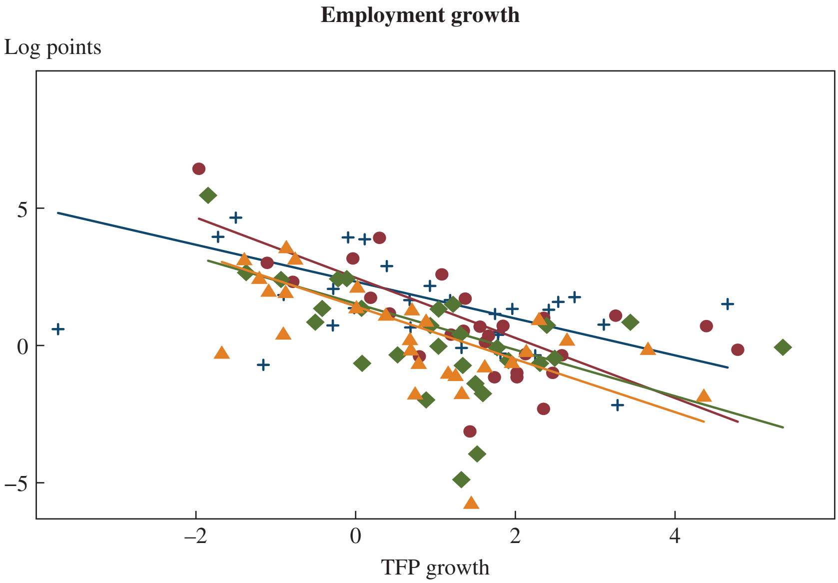

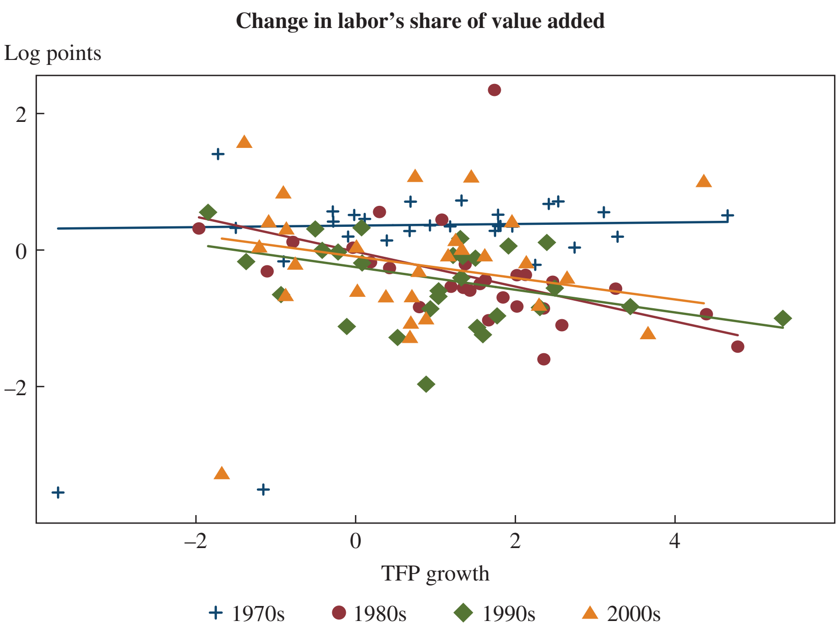

Sources: EU KLEMS; authors' calculations.

a. All values are expressed as annual, unweighted average changes across country-years in log points. The lines are the weighted linear fits by decade.

shows a much more noticeable shift in the relationship between productivity and labor's share over time. During the 1970s, there is no appreciable link between industries' productivity growth and their labor share changes. A clear negative relationship emerges in the 1980s, however, and remains in place during the 1990s and 2000s. This pattern suggests that a shift toward more labor share–displacing productivity growth is a possible explanation for the fall in the labor share commencing in the 1980s.

To explore this possibility more rigorously, we estimate a set of distributed lag models where the own-industry impact of TFP growth is allowed to vary by decade. Across a range of specifications, we find that the 1970s stand out as a period when own-industry TFP growth had a less negative effect on labor's share. We do not find much evidence of statistically significant heterogeneity in coefficients for the decades thereafter, consistent with the broad patterns shown in the bottom panel of figure 8. Table 12 provides estimates of the direct effect of TFP growth on our range of outcomes, estimated separately for the 1970s and the three subsequent decades. As shown in columns 11 and 12, there is a statistically insignificant positive relationship between own-industry TFP growth and own-industry labor share changes during the 1970s, which turns statistically significant and negative for the three more recent decades. Online appendix table A13 provides additional detail by estimating these models separately by decade, applying a five-year lagged long-difference specification.[37]

To assess the quantitative importance of these decadal differences, table 13 reports a set of decade-specific predictions based on table 12. These predictions are constructed by allowing the β_1^k coefficients in equation 8 and the $\beta_{1,VA}^k$ coefficients in equation 11 to be different in the 1970s compared with the other three decades, thereby allowing both the effect of TFP growth on the within-industry and between-industry components of the aggregate labor share to change over time.[38] A drawback of performing predictions with these estimates is that, relative to our main estimates, the estimated TFP slopes are shallower across all periods, likely because identification of the distributed lag terms is weak in short panels. Nevertheless, the predicted within-industry pattern now qualitatively matches the turnaround after the 1970s: Productivity growth is predicted to modestly *increase* labor's share during the 1970s and to decrease it thereafter.

37. We are severely limited in our ability to estimate distributed lag models for the decade of the 1970s because no country enters the EU KLEMS data before 1970, and several enter later (see online appendix table A1).

38. We restrict our attention here to the direct effect because we find this to be the main driver of aggregate labor share changes, irrespective of the time period under consideration.

Table 12. The Relationship between Productivity Growth and Industry-Level Outcomes Allowing for Decade-Specific Direct Effects[a]

	Annual change in log outcome variable by country-industry					
	Employment		*Hours*		*Wage bill*	
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$						
1970s[b]	–0.834***	0.042	–0.774***	0.071	–0.464	0.244
	(0.194)	(0.148)	(0.205)	(0.154)	(0.283)	(0.259)
1980s–2000s[c]	–2.135***	–1.182***	–2.062***	–1.109***	–1.855***	–1.014***
	(0.170)	(0.148)	(0.183)	(0.166)	(0.261)	(0.208)
Model weights	Employment		Hours		Hours	
	Nominal value added		*Real value added*		*Labor share*	
	(7)	*(8)*	*(9)*	*(10)*	*(11)*	*(12)*
$\Sigma\Delta\ln(\text{own-industry TFP}_{i,c,t-k})$						
1970s[b]	–0.469*	–0.089	0.627*	1.126***	0.146	0.289
	(0.244)	(0.329)	(0.364)	(0.363)	(0.234)	(0.335)
1980s–2000s[c]	–1.452***	–0.685***	0.640	1.224***	–0.386***	–0.423***
	(0.221)	(0.191)	(0.487)	(0.415)	(0.102)	(0.144)
Model weights	Nominal value added		Nominal value added		Nominal value added	
Fixed effects						
Country	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Sector	No	Yes	No	Yes	No	Yes
Country × time trend	Yes	Yes	Yes	Yes	Yes	Yes
Country × business cycle	Yes	Yes	Yes	Yes	Yes	Yes
Country × year	No	No	No	No	No	No

Sources: EU KLEMS; authors' calculations.

a. The models are estimated separately by subperiod. TFP is other-country, within-industry TFP, and is rescaled to have a standard deviation of 1. Standard errors clustered by country-industry are in parentheses. Statistical significance is indicated at the *10 percent, **5 percent, and ***1 percent levels.

b. The estimates shown are the sum of coefficients for the contemporaneous effect and two annually distributed lags. The number of observations is 3,520.

c. The estimates shown are the sum of coefficients for the contemporaneous effect and five annually distributed lags. The number of observations is 13,341.

Table 13. The Contribution of Total Factor Productivity Growth to the Within- and Between-Industry Components of the Change in Aggregate Labor Share, by Decade, 1970–2007

	Actual annual change in labor share in log points			*Predicted annual change in labor share in log points*		
Decade	*Total*	*Between industry*	*Within industry*	*Total*	*Between industry*	*Within industry*
1970s	0.513	–0.187	0.700	0.030	–0.020	0.050
1980s	–0.459	–0.183	–0.276	–0.201	–0.022	–0.179
1990s	–0.263	–0.075	–0.188	–0.125	–0.016	–0.109
2000s	–0.861	–0.425	–0.436	–0.150	–0.085	–0.065

Source: Authors' calculations, based on table 12.

The model is also somewhat successful at predicting the increase in the *between-industry* component of the falling labor share in the 2000s. The model is not successful, however, in explaining the acceleration of the within-industry fall in the labor share in the 2000s.

Summarizing, our analysis broadly supports the hypothesis that the decline in the labor share since the 1980s is consistent with a shift toward more labor share–displacing technology commencing in the 1980s. But the acceleration in the labor share decline observed during the 2000s is left unaccounted for by this mechanism. We hypothesize that a closer study of specific technologies may yield additional insights into these periods. At the same time, we do not assume that technological factors are the sole contributor to the changing secular pattern of the labor share decline or its recent deceleration. Instead, what our findings make clear is that technological progress has been broadly employment-augmenting and labor share–displacing for at least three decades. The consistency of the evidence, rather than its over-time acceleration or deceleration, is what gives us confidence in the utility of our approach for tracing through the within-industry, between-industry, and aggregate consequences of productivity growth originating in all industries.

V. Concluding Remarks

Theory makes clear that there is no direct mapping between the evolution of productivity and labor demand at the industry level and the evolution of labor demand in the aggregate. Theory gives less guidance about how to draw this indirect mapping. We present an empirical approach for mapping the industry-level effects of technological progress on aggregate employment and labor share outcomes, taking into account both the direct effects of productivity growth in advancing industries and the indirect effects of interindustry demand linkages, between-industry compositional change, and increases in final demand. Our findings indicate that these indirect effects are sizable and are countervailing for employment. We find that technological progress is broadly employment-augmenting in the aggregate. But this is not so for labor's share of value added, where direct labor share–displacing effects dominate. Our simple framework can account for a substantial fraction of both the reallocation of employment across industries and the aggregate fall in the labor share over the last three decades. It does not, however, explain why the share of labor in value added fell more rapidly during the 2000s than in prior decades. Nor can it distinguish between the contributions of automation-based versus non-automation-based sources of productivity growth, which may plausibly exert distinct effects on either employment or on labor's share of value added.

Although our empirical exploration of labor displacement has linked effects at the industry level to aggregate outcomes, this high-level representation is consistent with a variety of within- and between-firm adjustments. At one extreme, every firm in an industry undergoing technological progress might substitute capital for labor in a subset of tasks. Alternatively, absent any within-firm change in task allocation, a technological advance might spur an increase in industry market share among relatively capital-intensive firms, and a concomitant decline among relatively labor-intensive firms.[39] Under either scenario, labor's share in industry value added would fall. Our analysis cannot speak to these within-firm versus between-firm dynamics. Nevertheless, we believe that the scope of the evidence presented here complements more granular, but narrower, firm-level and establishment-level studies.

ACKNOWLEDGMENTS We are deeply grateful to discussants John Haltiwanger and Richard Rogerson. We thank Daron Acemoglu, Jim Bessen, Uwe Blien, Janice Eberly, Maarten Goos, Pascual Restrepo, James Stock, Coen Teulings, John Van Reenen, and Xianjia Ye, whose valuable input improved the paper. We are grateful to Pian Shu for sharing data on approved patent applications and citations by industry, year, and country.

39. For further explorations of the linkage between firm-level dynamics and aggregate productivity, see Decker and others (2017), Autor and others (2017b), and Foster and others (2017, 2018).

References

Acemoglu, Daron, Ufuk Akcigit, and William Kerr. 2016. "Networks and the Macroeconomy: An Empirical Exploration." *NBER Macroeconomics Annual* 31: 276–335.

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics, Volume 4B*, edited by David Card and Orley Ashenfelter. Amsterdam: North-Holland.

Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price. 2016. "Import Competition and the Great US Employment Sag of the 2000s." *Journal of Labor Economics* 34, suppl. 2: S141–S198.

Acemoglu, Daron, and Veronica Guerrieri. 2008. "Capital Deepening and Nonbalanced Economic Growth." *Journal of Political Economy* 116, no. 3: 467–98.

Acemoglu, Daron, and Pascual Restrepo. 2017. "Robots and Jobs: Evidence from US Labor Markets." Working Paper no. 23285. Cambridge, Mass.: National Bureau of Economic Research.

———. 2018. "Artificial Intelligence, Automation and Work." Working Paper no. 24196. Cambridge, Mass.: National Bureau of Economic Research.

———. Forthcoming. "The Race between Machine and Man: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review*.

Alexopoulos, Michelle, and Jon Cohen. 2016. "The Medium Is the Measure: Technical Change and Employment, 1909–1949." *Review of Economics and Statistics* 98, no. 4: 792–810.

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103, no. 5: 1553–97.

Autor, David, David Dorn, Gordon H. Hanson, Gary Pisano, and Pian Shu. 2017a. "Foreign Competition and Domestic Innovation: Evidence from U.S. Patents." Working Paper no. 22879. Cambridge, Mass.: National Bureau of Economic Research.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2017b. "The Fall of the Labor Share and the Rise of Superstar Firms." Working Paper no. 23396. Cambridge, Mass.: National Bureau of Economic Research.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118, no. 4: 1279–333.

Autor, David, and Anna Salomons. 2017. "Robocalypse Now: Does Productivity Growth Threaten Employment?" In *Proceedings of the ECB Forum on Central Banking: Investment and Growth in Advanced Economies*. Sintra: European Central Bank.

Barkai, Simcha. 2017. “Declining Labor and Capital Shares.” Job market paper, University of Chicago. http://home.uchicago.edu/~barkai/doc/BarkaiDecliningLaborCapital.pdf

Basu, Susanto, and John Fernald. 2001. “Why Is Productivity Procyclical? Why Do We Care?” In *New Developments in Productivity Analysis*, edited by Charles R. Hulten, Edwin R. Dean, and Michael J. Harper. University of Chicago Press.

Baumol, William J. 1967. “Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis.” *American Economic Review* 57, no. 3: 415–26.

Berg, Andrew, Edward F. Buffie, and Luis-Felipe Zanna. 2018. “Should We Fear the Robot Revolution? (The Correct Answer Is Yes).” Working Paper no. 18/116. Washington: International Monetary Fund.

Bessen, James. 2017. “Automation and Jobs: When Technology Boosts Employment.” Law & Economics Paper no. 17-09. Boston: Boston University School of Law.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton.

Buera, Francisco J., Joseph P. Kaboski, and Richard Rogerson. 2015. “Skill-Biased Structural Change.” Working Paper no. 21165. Cambridge, Mass.: National Bureau of Economic Research.

Caselli, Francesco, and Alan Manning. Forthcoming. “Robot Arithmetic: New Technology and Wages.” *American Economic Review: Insights*.

Chiacchio, Francesco, Georgios Petropoulos, and David Pichler. 2018. “The Impact of Industrial Robots on EU Employment and Wages: A Local Labour Market Approach.” Working Paper no. 2. Brussels: Bruegel.

Dao, Mai Chi, Mitali Das, Zsoka Koczan, and Weicheng Lian. 2017. “Why Is Labor Receiving a Smaller Share of Global Income? Theory and Empirical Evidence.” Working Paper no. 17/169. Washington: International Monetary Fund.

Dauth, Wolfgang, Sebastian Findeisen, Jens Südekum, and Nicole Wöflner. 2017. “German Robots: The Impact of Industrial Robots on Workers.” Working paper. https://sites.google.com/site/dauthecon/research

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2017. “Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown.” *American Economic Review* 107, no. 5: 322–26.

Eden, Maya, and Paul Gaggl. 2018. “On the Welfare Implications of Automation.” *Review of Economic Dynamics* 29: 15–43.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2013. “The Decline of the U.S. Labor Share.” *Brookings Papers on Economic Activity*, Fall: 1–52.

Ford, Martin. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future.* New York: Basic Books.

Foster, Lucia S., Cheryl A. Grim, John Haltiwanger, and Zoltan Wolf. 2017. "Macro and Micro Dynamics of Productivity: From Devilish Details to Insights." Working Paper no. 23666. Cambridge, Mass.: National Bureau of Economic Research.

———. 2018. "Innovation, Productivity Dispersion, and Productivity Growth." Working Paper no. 24420. Cambridge, Mass.: National Bureau of Economic Research.

Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114: 254–80.

Graetz, Georg, and Guy Michaels. Forthcoming. "Robots at Work." *Review of Economics and Statistics*.

Gregory, Terry, Anna Salomons, and Ulrich Zierahn. 2016. "Racing with or against the Machine? Evidence from Europe." Discussion Paper no. 16-05. Utrecht: Utrecht School of Economics, Tjalling C. Koopmans Research Institute.

Gutiérrez, Germán. 2017. "Investigating Global Labor and Profit Shares." Working paper. https://ssrn.com/abstract=3040853

Gutiérrez, Germán, and Thomas Philippon. 2017. "Declining Competition and Investment in the U.S." Working Paper no. 23583. Cambridge, Mass.: National Bureau of Economic Research.

Jones, Charles I., and Paul M. Romer. 2010. "The New Kaldor Facts: Ideas, Institutions, Population, and Human Capital." *American Economic Journal: Macroeconomics* 2, no. 1: 224–45.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95, no. 1: 161–82.

Kaldor, Nicholas. 1961. "Capital Accumulation and Economic Growth." In *The Theory of Capital*, edited by Friedrich Lutz and Douglas C. Hague. New York: St. Martin's Press.

Karabarbounis, Loukas, and Brent Neiman. 2014. "The Global Decline of the Labor Share." *Quarterly Journal of Economics* 129, no. 1: 61–103.

Keynes, John Maynard. 1939. "Relative Movements of Real Wages and Output." *Economic Journal* 49, no. 193: 34–51.

Michaels, Guy, Ashwini Natraj, and John Van Reenen. 2014. "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years." *Review of Economics and Statistics* 96, no. 1: 60–77.

Ngai, L. Rachel, and Christopher A. Pissarides. 2007. "Structural Change in a Multisector Model of Growth." *American Economic Review* 97, no. 1: 429–43.

O'Mahony, Mary, and Marcel P. Timmer. 2009. "Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database." *Economic Journal* 119, no. 538: F374–F403.

Pierce, Justin R., and Peter K. Schott. 2016. "The Surprisingly Swift Decline of US Manufacturing Employment." *American Economic Review* 106, no. 7: 1632–62.

Piketty, Thomas. 2014. *Capital in the Twenty-First Century*. Belknap Press.

Ramey, Valerie A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics, Volume 2*, edited by John B. Taylor and Harald Uhlig. Amsterdam: North-Holland.

Rognlie, Matthew. 2015. "Deciphering the Fall and Rise in the Net Capital Share: Accumulation or Scarcity?" *Brookings Papers on Economic Activity*, Spring: 1–54.

Sachs, Jeffrey D., and Laurence J. Kotlikoff. 2012. "Smart Machines and Long-Term Misery." Working Paper no. 18629. Cambridge, Mass.: National Bureau of Economic Research.

Solow, Robert M. 1956. "A Contribution to the Theory of Economic Growth." *Quarterly Journal of Economics* 70, no. 1: 65–94.

Stansbury, Anna M., and Lawrence H. Summers. 2017. "Productivity and Pay: Is the Link Broken?" Working Paper no. 24165. Cambridge, Mass.: National Bureau of Economic Research.

Susskind, Daniel. 2017. "A Model of Technological Unemployment." Discussion Paper no. 819. Oxford: University of Oxford, Department of Economics.

Teulings, Coen N., and Nikolay Zubanov. 2014. "Is Economic Recovery a Myth? Robust Estimation of Impulse Responses." *Journal of Applied Econometrics* 29, no. 3: 497–514.

Timmer, Marcel P., Erik Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J. de Vries. 2015. "An Illustrated User Guide to the World Input–Output Database: The Case of Global Automotive Production." *Review of International Economics* 23, no. 3: 575–605.

Timmer, Marcel, Ton van Moergastel, Edwin Stuivenwold, Gerard Ypma, Mary O'Mahony, and Mari Kangasniemi. 2007. "EU KLEMS Growth and Productivity Accounts Version 1.0: Part I Methodology." http://www.euklems.net/data/euklems_growth_and_productivity_accounts_part_i_methodology.pdf

Trajtenberg, Manuel. 1990. "A Penny for Your Quotes: Patent Citations and the Value of Innovations." *RAND Journal of Economics* 21, no. 1: 172–87.

van Ark, Bart, and Kirsten Jäger. 2017. "Recent Trends in Europe's Output and Productivity Growth Performance at the Sector Level, 2002–2015." *International Productivity Monitor* 33: 8–23.

Zeira, Joseph. 1998. "Workers, Machines, and Economic Growth." *Quarterly Journal of Economics* 113, no. 4: 1091–117.

Comments and Discussion

COMMENT BY

JOHN HALTIWANGER One of the oldest questions and concerns among economists is the impact of innovation and productivity growth on employment. Over the centuries, technological progress has raised productivity dramatically, enabling far greater output per unit of labor input. Moreover, accompanying product innovations have enabled associated rapid increases in the quality and range of products. Although this raises GDP per capita, concerns have frequently been raised about the workers left behind by technological advances. Recently, these concerns have arisen in the context of the impact on the workforce of perceived rapid changes in automation. Artificial intelligence, robotics, customized software, and specialized machinery have already become embedded in many production technologies, enabling the replacement of tasks once performed by workers.

Our understanding of these issues, even after much attention in the literature, remains relatively incomplete. Part of this reflects the fact that the type of data that is ideally needed to understand the labor-displacing nature of technology is not readily available.[1] Partly this reflects the view that each new wave of innovation is potentially different in both the nature and speed of any disruption and displacement that occurs. In

1. For example, firm-level evidence on technology adoption accompanied by detailed information about the size and mix of the workforce in terms of skills is needed. Occasionally, modules have been added to firm-level surveys that provide very helpful information, such as the Survey of Manufacturing Technology in the 1980s and 1990s. We need a new wave of such modules for more recent advances in technology. In addition, we need to integrate such data with longitudinal matched employer–employee data to investigate the impact on the workforce.

addition, regardless of data limitations, it is challenging to sort through the complex mechanisms at the firm, industry, country, and global levels. Economic theory reminds us that technological improvements in one sector may yield a reallocation of labor to sectors with less rapid technological change (Ngai and Pissarides 2007), depending on the elasticity of substitution across sectors. The impact on aggregate employment also inherently depends on the elasticity of the labor supply. However, this latter perspective is in the long run, and there might be much disruption, displacement, and reallocation along the way.

This paper by David Autor and Anna Salomons weighs in on this ongoing debate by using pooled data, industry by country by time (mostly annual), on the relationship between outcomes such as employment growth, value-added growth, and the labor share with indicators of technological innovation. The primary focus in this paper is to use measures of total factor productivity (TFP) to capture the latter. The main findings in terms of employment growth are that there is an own-industry negative impact of rising TFP but offsetting indirect effects arising from the input–output linkages as well as the overall positive impact of rising TFP on aggregate value added and final demand. Taken at face value, the answer to the question posed in the paper's title is that though there may be sectoral reallocation induced by technological innovation, the aggregate effect on employment growth is positive. This answer is in principle reassuring to those who have continued to express concerns about the impact of innovation on employment outcomes.

Although I am sympathetic to the overall message of the paper and the careful analysis of rich industry-by-country data, I think there are several challenges in interpreting the paper's results. First, the empirical approach is entirely reduced form, which can be very useful for helping guide future analysis; but the nature of the reduced-form approach taken here does not provide much guidance about the mechanisms underlying the estimated results.[2] Second, there are many details about the measurement, specifications, and estimation that raise a host of questions about what we learn from using these industry-by-country-by-time data. Most of my remaining comment focuses on these latter issues.

A core concern is whether TFP growth at the industry level is a good proxy for innovation. In my figure 1, average annual TFP growth rates

2. Richard Rogerson's comment provides a detailed perspective on these issues.

Figure 1. Industry-Level Total Factor Productivity Growth, by Industry, 1970–2007

Source: U.S. KLEMS.

for the 28 industries used by Autor and Salomons from the U.S. KLEMS data are presented. Several remarks are warranted. First, the outsized role of information and communication technology (ICT) in productivity growth over this period is reflected in the electrical and optical equipment industry. Second, even during this period of rapid, ICT-induced productivity growth, 45 percent of industries have negative long-run productivity growth in the U.S. data. Autor and Salomons acknowledge the relatively high fraction of observations with negative productivity growth (and do some limited robustness analysis) but attribute much of this to high-frequency measurement error.[3] My figure 1 shows that this is a pervasive issue, even over long time intervals. Moreover, my figure 1 shows that it is not just the difficult-to-measure sectors, such as finance and services or the nonprofit sectors, that exhibit these patterns.

This pattern of pervasive industry-level negative productivity growth is also present at medium-run frequencies. My figure 2 presents a scatter plot of the employment growth rates and TFP growth rates at the industry level for the United States for these 28 industries using peak-to-peak variation (using the two-year average at each peak) based on reference cycles developed by the National Bureau of Economic Research. About 45 percent of the industry-by-time observations have negative productivity growth at this frequency. It is also evident that the negative own-industry effect estimated by Autor and Salomons is present in my figure 2. The elasticity of employment growth with respect to TFP growth using this medium-run variation is –0.207 (with a standard error of 0.090), which includes industry and time period controls, and the standard error is clustered at the industry level.

This pervasive finding of negative productivity growth at the industry level is not new. Dale Jorgenson, Mun Ho, and Jon Samuels (2018) summarize their views and the literature by highlighting four competing explanations: measurement error; resource depletion, relevant for sectors such as oil and gas extraction and mining; misallocation and regulation; and finally, sectors that especially deviate from private sector profit maximization (for example, health and education). All four of these explanations raise questions about the relevance of using TFP at the industry level as an indicator of innovation and technological progress. In addition, a few other factors also raise questions.

3. I did not find the robustness analysis in the paper's table 6 compelling in terms of treating the negative TFP growth observations as zero. The discussion here highlights many different reasons that might underlie the observed negative TFP growth. Those same factors (such as reallocation dynamics) may be influencing the positive observations as well, and accordingly raise questions about the interpretation of the findings.

Figure 2. Peak-to-Peak U.S. Employment and Total Factor Productivity Growth Rates for the 28 Industries Considered by Autor and Salomons, 1972–2007[a]

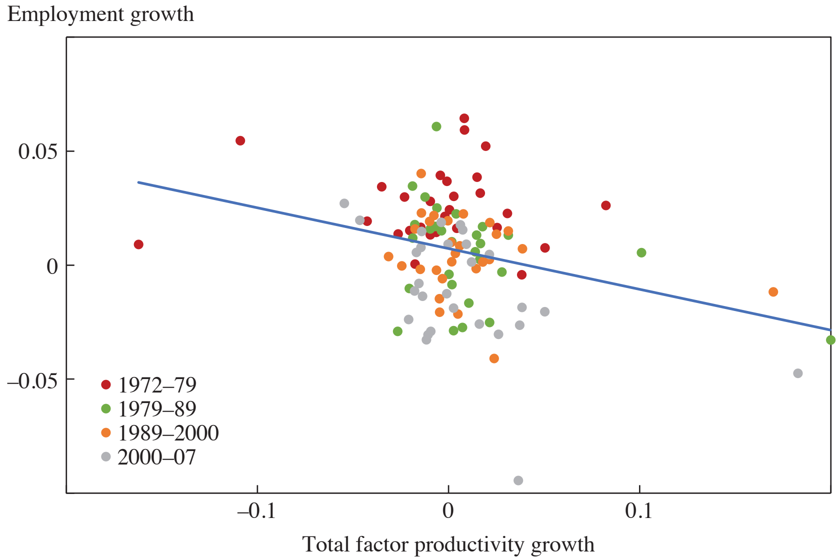

Source: U.S. KLEMS.
a. The line shows the linear fit.

In particular, Michael Gort and Steven Klepper (1982) hypothesized long and variable lags between innovation and productivity growth. They argue that a surge in innovation in an industry is accompanied by a surge in the entry of new firms that engage in experimentation with new products and processes. During this period of experimentation, productivity growth might actually fall rather than rise. It is only later that productivity growth is observed, after successful entrants grow, while less successful entrants contract and exit. Gort and Klepper used relatively crude data on business formation and exits, but they did show patterns consistent with their hypothesis. Recent evidence from Lucia Foster and others (2018) provides more direct confirming evidence. We find that a surge in entry within an industry in one three-year period yields a decline in within-industry productivity growth and an accompanying rise in productivity dispersion across firms in the industry in the next three-year period. It is only in the subsequent years that productivity growth is observed, along with an accompanying decline in productivity dispersion. Gort and Klepper's firm dynamics are in some respects related to the misallocation hypothesis mentioned above, but in this case it is a more benign form of misallocation. Namely, Gort and Klepper argue that the firm dynamics and shakeout process inherent in innovation may lead to a decrease in productivity growth during the experimentation period, but this is part of the investment needed to eventually achieve successful innovations and productivity growth.

The firm dynamics hypothesized by Gort and Klepper (1982) also highlight another limitation of using industry-level as opposed to firm-level data to investigate the main questions of interest. Industry-level fluctuations in productivity reflect not only the within-firm innovations but also between-firm reallocation dynamics that may take some time to work through. A related issue is that many firm-level studies find a strong positive relationship between TFP and employment growth at the firm level (Decker and others 2018; Ilut, Kehrig, and Schneider, forthcoming). Reconciling the firm-level evidence with the industry-level evidence considered here likely requires distinguishing within-firm from between-firm innovations. That is, the successful innovators within an industry may be increasing employment but require less employment than unsuccessful firms that contract and exit. The overall impact at the industry level may be negative, but this may be entirely attributable to reallocation. Interestingly, Autor and others (2017) find that the decline in the within-industry labor share is primarily accounted for by reallocation, and not by within-firm declines in the labor share. In addition, even at the industry level, Daron Acemoglu, Ufuk Akcigit, and William Kerr (2016), using a similar specification to that used by Autor and Salomons, find a positive own-industry effect of TFP growth on employment growth (using a one-period lag of own-industry TFP growth).

All these issues raise questions about whether Autor and Salomons are providing much guidance about the impact of innovation on the potential displacement of labor. If nothing else, the timing and dynamics are complex, and a five-year lag specification is likely inadequate—especially from the perspective and findings of Gort and Klepper (1982) and of Foster and others (2018). In addition, there are interesting and complex issues in these dynamics. If productivity growth lags innovation substantially, at what point does any displacement of workers occur? Does it occur during the innovation or experimentation phase, or does it occur in the shakeout phase? The answer is likely all of the above.

These complex dynamics are important for more reasons than getting the frequency and lag structure of the empirical specification correct. At the core of concerns of the impact of automation on displacement is the speed of the transition dynamics. The Gort and Klepper firm dynamics suggest that implementation lags are long and variable. If implementation lags have shortened, then this has potentially important implications for the reallocation and displacement dynamics that might arise, even if there is no long-run adverse impact on employment.

There are other related concerns about the details of the implementation. The authors use a leave-out-mean approach for measuring within-industry-by-country TFP growth. This approach is intended to avoid the potential mechanical relationship between TFP and employment growth. The latter concern is potentially nontrivial, but does depend on the presence of measurement error in the labor input. The latter is likely among the best measured inputs in production. Moreover, others have overcome this concern using the relationship between contemporaneous employment growth and lagged TFP growth (Acemoglu, Akcigit, and Kerr 2016). I think using the leave-out-mean approach in this context also has other problems. For one, there is much evidence that innovation and productivity growth at the industry level is quite different across countries. The 1990s were a period when ICT innovation and productivity took off in the United States. relative to the rest of the world, including Europe. The approach taken by Autor and Salomons would distort this variation. The U.S. surge in productivity in the ICT sector in the 1990s would be left out of the U.S. measure, but it is so large that it would contribute substantially to the leave-out-mean of all other countries. Relatedly, I am skeptical that the leave-out-mean approach captures the technological frontier at the industry level.

Finally, I found the analysis of upstream and downstream industries interesting but difficult to interpret. Partly, it was difficult to interpret because of technical issues. The appropriate model and measurement methodology with input–output linkages is to use gross output production functions and explicit modeling and measurement of intermediate input usage. In addition, the results as presented are a bit of a black box. It would be interesting to explore and understand what types of supply chain links are especially important in this context. Acemoglu, Akcigit, and Kerr (2016) suggest some ways of exploring these issues.

REFERENCES FOR THE HALTIWANGER COMMENT

Acemoglu, Daron, Ufuk Akcigit, and William Kerr. 2016. "Networks and the Macroeconomy: An Empirical Exploration." *NBER Macroeconomics Annual* 31: 273–335.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2017. “The Fall of the Labor Share and the Rise of Superstar Firms.” Working Paper no. 23396. Cambridge, Mass.: National Bureau of Economic Research.

Decker, Ryan A., John C. Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2018. “Changing Business Dynamism: Shocks vs. Responsiveness.” Working Paper no. 24236. Cambridge, Mass.: National Bureau of Economic Research.

Foster, Lucia, Cheryl Grim, John C. Haltiwanger, and Zoltan Wolf. 2018. “Innovation, Productivity Dispersion, and Productivity Growth.” Working Paper no. 24420. Cambridge, Mass.: National Bureau of Economic Research.

Gort, Michael, and Steven Klepper. 1982. “Time Paths in the Diffusion of Product Innovations.” *Economic Journal* 92, no. 367: 630–53.

Ilut, Cosmin, Matthias Kehrig, and Martin Schneider. Forthcoming. “Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News.” *Journal of Political Economy*.

Jorgenson, Dale W., Mun S. Ho, and Jon D. Samuels. 2018. “Educational Attainment and the Revival of U.S. Economic Growth.” In *Education, Skills, and Technical Change: Implications for Future US GDP Growth*, edited by Charles R. Hulten and Valerie A. Ramey. University of Chicago Press.

Ngai, L. Rachel, and Christopher A. Pissarides. 2007. “Structural Change in a Multisector Model of Growth.” *American Economic Review* 97, no. 1: 429–43.

COMMENT BY

RICHARD ROGERSON Although the long-run effect of technological change on aggregate labor market outcomes has long been of interest to economists, concern about this issue has recently intensified, perhaps motivated in part by the decline in labor’s share that has been observed in the United States and elsewhere in recent decades and by the sense that it might be due, at least in part, to increases in automation that reflect recent trends in technological change. This paper by David Autor and Anna Salomons seeks to assess the aggregate effects of automation on employment and labor’s share since 1970 using sectoral data from a large set of developed economies. The paper provides much information that is useful in the effort to better explain the dynamics of employment and the labor share, complementing the earlier contribution to the *Brookings Papers* by Michael Elsby, Bart Hobijn, and Ayşegül Şahin (2013), which focused entirely on the United States. However, though I think Autor and Salomons present a lot of interesting evidence, I nonetheless feel they are largely unsuccessful in their effort to offer compelling and credible evidence on the causal effects of automation on employment and the labor share at the aggregate level. The reduced-form methods employed by the authors essentially document conditional correlations. These correlations can serve as valuable diagnostics and provide suggestive evidence to help us distinguish between competing explanations. But these reduced-form methods are not well suited to delivering quantitative estimates of causal effects.

In the brief space that I have available here, I first describe why I think the empirical approach employed by the authors is unable to deliver reliable estimates of the causal effects of growth in total factor productivity (TFP) on employment and the labor share. I follow this with several shorter comments about details of the specification adopted by the authors.

USING SECTORAL DATA TO ESTIMATE AGGREGATE RESPONSES A key component of the paper's analysis is to recover the aggregate effects of TFP on employment and the labor share using reduced-form estimates of sectoral relationships. My main comment relates to the basis for interpreting these aggregate effects as the causal effects of TFP. As is well known, there is a long history of debate arguing the pros and cons of structural versus reduced-form approaches to uncovering causal effects, especially at the aggregate level. I do not want to get into this debate here, so I take as given that the goal is to learn what we can using reduced-form methods.

Before getting into specifics, I think it is important to first back up a bit to consider how the authors arrived at an analysis of sectoral data in their attempt to uncover aggregate effects. In particular, suppose we start with the premise that the goal is to use reduced-form methods to understand the effect of TFP growth on either employment or the labor share at the aggregate level in a particular country. Given this goal, it seems natural that one might first consider attempts to uncover these effects using reduced-form methods on aggregate data.

The simplest exercise that one might start with is to regress either of these aggregate outcome variables on aggregate TFP (perhaps including several lags, as the authors do). But if one simply ran the regression of either the employment–population ratio or the labor share on TFP for a single country, no one would view the coefficients as a reliable estimate of the causal effect of aggregate TFP on either outcome of interest. The reason is that potentially many other factors are at play that are also affecting these outcomes, these other factors may be correlated with TFP, and the regression is projecting all these effects onto changes in TFP.

One possible response is to try to include measures of the other potentially important factors on the right-hand side. Although there might be a few channels that we could capture this way, many of them are likely not easily measured, so there will always remain some concern that one is not isolating the effect of TFP. In such a situation, it is standard practice to include time effects as a way to control for unobserved factors; but using data for a single country, these time effects would explain all the variation in the left-hand-side variable.

If one thought that the key time effects were constant across countries, then expanding the analysis to include data from several countries would solve the problem. Note that one could of course also allow for country-level fixed effects in the analysis. But if one thought that the unobserved driving forces were specific to the country and year, expanding the analysis to many countries would not solve the basic problem of needing to isolate the effects of TFP from those of other factors, given that this would require a full set of interacted country and time effects.

A nice strategy adopted by Autor and Salomons is to use TFP from other countries as a proxy for TFP in the country being studied. This turns out to be a good proxy and, at least at first pass, would seem to eliminate the need for a fully interacted set of country and time fixed effects in the previous analysis. To the extent that global factors influence either labor share or employment across countries and are correlated with average movements in TFP, it would still be necessary to use time fixed effects to control for these other factors, but the use of other countries' TFP eliminates the effect of country-specific, non-TFP factors that might be correlated with country-specific TFP. But upon further reflection, it should be apparent that this strategy only solves the problem if it is assumed that the global non-TFP factors that are correlated with average TFP have identical effects on all countries. If not, we would still need a fully interacted set of country and time effects to control for these effects.

Why might one think that employment and labor share responses to a given shock might differ across countries? An old idea in the literature on cross-country differences in labor market evolutions—initially put forth by Michael Bruno and Jeffrey Sachs (1985), and later taken up by Paul Krugman (1994) and Olivier Blanchard and Justin Wolfers (2000)—is that country-specific factors (for example, labor market institutions) lead to differential responses across countries to a given shock. There seems every reason to think that this idea is relevant in the current context, when one seeks to estimate how output, employment, and wages respond to various driving forces. For example, the factors leading to increased global trade are plausibly correlated with TFP and plausibly have differential effects across economies, not only because of different labor market institutions, but also because different economies might have varying exposures to a given trade shock.[1]

To summarize, a key impediment to obtaining estimates of the aggregate effects of TFP on employment and the labor share from aggregate data using reduced-form methods is the need to include a set of fully interacted country and time fixed effects as a way to control for non-TFP factors.

Why might one turn to sector-level data in an attempt to uncover the aggregate effects? If one accepts that fully interacted country and time fixed effects are needed to properly control for non-TFP factors, then a sector-level analysis seems to offer a way around the issue, because one could now allow for a set of fully interacted country and time fixed effects and still have variation to consider. Two issues arise, however. First, why would we think that there are not important time and country effects at the sector level? The strategy of adding another layer of data to get around the need to have a full set of fixed effects presumes that we can rule out variation in driving forces or their impact at this new layer. But what is the rationale for this belief? My own view is that, in general, the more we disaggregate, the larger and more varied are the sources of idiosyncratic variation. Put differently, if one believes that interacted country and time fixed effects are important, why would it seem reasonable to assume that these effects do not vary at the sectoral level? Of course, these country-time-sector effects would not be a problem if they were uncorrelated with TFP, but I see no basis for assuming this. The previous example of increased trade would certainly lead one to expect country-time-sector effects.

This issue aside, the second issue with moving to sector-level data is that each sector does not represent an economy. That is, even if we properly identify the effect of own-sector TFP changes on own-sector outcomes, we still need to determine how to aggregate the effects. This requires that we need to isolate not only the causal effect of TFP growth in sector i on outcomes in sector i, but also the causal effects of TFP growth in sector i on outcomes in all other sectors.

When one moves to sectoral data as a way to estimate aggregate effects, the implicit claim is that it is straightforward or easy to identify all these cross-sector, general equilibrium effects using reduced-form analyses. In fact, the authors characterize their analysis as building up the aggregate effects by quantifying each of several underlying effects, suggesting that it is easier to compute these individual underlying pieces than it is to directly

1. This same logic, of course, suggests that it is not appropriate to impose a common response to TFP across countries.

compute the aggregate response. I argue below via a simple example that I believe this is not the case. That is, moving the analysis to the sectoral level does not get around any of the issues that led one to move from aggregate to sectoral data in the first place.

But before describing the simple example to make this point, I do want to emphasize that I nonetheless think sectoral analyses can be very useful. The reason is that they potentially provide additional information about driving forces and mechanisms. To be concrete, let us focus on the issue of the aggregate decline in the labor share. One may have various candidate driving forces or mechanisms in mind. The reason that sectoral data may be very useful is that there may be considerable variation across sectors, and this variation may prove to be a useful diagnostic to help us evaluate the promise of various driving forces. For example, in the current context, Autor and Salomons are interested in assessing the role of automation as a driving force. Variation in both investment in equipment (especially, perhaps, computing equipment), and the change in the labor share at the sectoral level might reveal something about the promise of a story that stresses automation. I say more about this below.

However, though I think sectoral-level data are therefore very valuable for *qualitatively* assessing different explanations, I do not think that moving to sectoral data provides any advantage in helping us to tease *quantitative* effects out of the data using reduced-form methods. And to think otherwise is basically wishful thinking. To see why, I consider the reduced-form methods employed by the authors in the context of a simple structural model.

In particular, consider an economy that captures the basic economics of William Baumol (1967). There is a representative household with preferences each period, given by

$$C_t^{1-1/\sigma} - BH_t^{1+1/\gamma}.$$

There are N sectors, each with a constant-returns-to-scale production function:

$$y_{i,t} = A_{i,t} k_{i,t}^{\theta} h_{i,t}^{1-\theta}.$$

Aggregate consumption and investment are produced by combining the outputs of the N sectors:

$$C_t + I_t = \left(\sum_{i=1}^{N} a_i y_{i,t}^{1-1/\rho} \right)^{\rho/\rho-1}.$$

Consider the competitive equilibrium for this economy. Even without doing any analysis, one might already sense something curious vis-à-vis the authors' analysis. The specification given above suggests that the two preference parameters σ and γ are surely going to be important for determining the response of aggregate employment to changes in the profile of sector TFPs. In particular, if we assume the limiting case of $\sigma = 1$, then we have offsetting income and substitution effects, and it is easy to show that in the competitive equilibrium, aggregate employment is independent of the TFP profile across sectors. One may well ask how the empirical specification adopted by the authors is incorporating this key parameter and the associated labor supply effects, because the equations estimated by the authors all have the feel of being motivated by labor demand considerations, with no role for labor supply. Although it is common (even if not warranted) to abstract from labor supply considerations in the context of short-run fluctuations, there seems to be no basis for thinking that labor supply considerations do not factor into long-run labor market outcomes.

In what follows, I simply posit that aggregate employment is some unspecified function of the sector TFPs and capital stock, without imposing that equilibrium employment is consistent with desired labor supply of the household, taking all prices as given. Note, first, that if we normalize the wage to 1, then the sector i price is just the inverse of sector i TFP. Second, it is easy to show that maximization yields the following expressions:

$$h_{i,t} = \frac{a_i^{\rho} A_{i,t}^{\rho}}{\sum_j a_j A_{j,t}^{\rho-1}} H_t \equiv a_i^{\rho} A_{i,t}^{\rho} D_t H_t,$$

where D_t is defined by

$$D_t = \frac{1}{\sum_j a_j A_{j,t}^{\rho-1}}.$$

Taking the logs and first differences, we end up with

$$\Delta \log h_{i,t} = (\rho - 1)\Delta \log A_{i,t} - \Delta \log D_t + \Delta \log H_t.$$

Recall that H_t is implicitly a function of the profile of period t TFPs and the capital stock in the country being studied. The fact that it enters with a coefficient of 1 reflects the fact that preferences are homothetic, so that an increase in aggregate labor increases the output of each good proportionately. If one wanted to consider preferences that were nonhomothetic, then the coefficient on aggregate labor would vary across sectors, but the appropriate average of these effects would still be 1. A common coefficient of less than 1 would imply an inconsistency, given that aggregate labor must be the sum of the sectoral labors. The key point, however, is that the need to control for aggregate hours on the right-hand side surely suggests that one would need to include a fully interacted set of country and time effects to control both for the effects of TFP on total hours and for potential non-TFP factors.[2]

The approach taken by the authors is to replace H_t with a measure of either nominal or real value added. Although this implicitly allows for a particular form of country-time fixed effect, this is appropriate only if aggregate hours and value added move one-for-one across all countries. But we know from growth accounting exercises that this is simply not the case. It also explains why they obtain the troubling result that a given percentage increase in aggregate value added at a single point in time, holding all else constant, leads to a smaller percentage increase in value added in all sectors.

The significance of the previous derivation is that even in a model with no driving forces beyond TFP, and no heterogeneity across countries other than TFP, one would need to include a fully interacted set of country and time fixed effects in order to properly estimate the sectoral relationship between TFP and employment. In reality, the time and country fixed effects will also pick up non-TFP effects. It follows that when one wants to use the estimates of this equation to compute the effect of TFP changes on employment, one needs to include the component of the country-time fixed effects that reflects TFP effects as opposed to non-TFP effects. That is, in order to trace out the causal effect of changes in TFP on employment using these estimates, one would need to be able to decompose the change in estimated country-time fixed effects into the parts that come from changes in TFP as opposed to non-TFP factors. But the whole reason for moving to sectoral data was because we did not know how to isolate the effect of TFP from non-TFP effects that were country and time varying. This sectoral approach ultimately requires that we have a solution to the problem that led us to the sectoral data in the first place!

Although my example is admittedly very simple, the key point derives from a very basic and robust property: In a multisector model, the allocation

2. The second term on the right-hand side also depends on the country-specific profiles of TFP, though readers familiar with these types of models will see that D is simply the model-implied aggregate price index. For present purposes, we can ignore this term.

of labor to any sector will depend on the total amount of labor supplied in equilibrium, which will be a function of the profile of sector TFPs as well as any non-TFP factors that influence labor markets. The only way to credibly capture the evolution of this term over time and across countries in a linear regression model is to allow for a fully interacted set of country-time fixed effects.[3]

Let me summarize. I have raised two basic reasons for why the use of sectoral data should not be viewed as a path for obtaining credible estimates for causal aggregate effects using reduced-form methods. First, it presupposes that there are not important biases associated with country-time-sector effects that are correlated with the driving force of interest, in this case TFP. But second, this path requires that one be able to isolate TFP from non-TFP effects captured by country-time fixed effects. If one could do this credibly, then one would not need to go to sectoral data in the first place.

Although the discussion above explains what I view as the main limitation of the paper in terms of its ability to deliver credible estimates of the causal aggregate effects of TFP, in the remainder of my space I point out a few additional issues with the specification that the authors adopt.

TFP AS THE DRIVING FORCE The paper's title and some of its exposition suggest that the purpose of the analysis is to uncover the effects of automation on employment and labor's share. Automation reflects a particular type of technological progress. Because not all technological progress reflects automation, it would seem necessary for any study that seeks to isolate the effects of automation to first attempt to isolate the component of technological progress that might best be associated with automation. The authors instead choose to focus purely on the effects of an "omnibus" measure of technological progress—namely, TFP.

At one level, this is a simple issue of semantics—perhaps the authors' goals are really to assess the effects of TFP growth on labor market outcomes, rather than the effects of automation per se on labor market outcomes. But if we accept this alternative framing of the analysis, it seems that the analysis is implicitly abstracting from what surely must be the

3. My simple example does not include the linkages that Autor and Salomons include. Doing so would destroy the tractability of my simple example, but the basic point remains valid: A key determinant of hours in a given sector is the total amount of work being carried out.

most important question. Two simple observations explain why I say this. First, until recently, there was a consensus that the so-called Kaldor facts (Kaldor 1961) provided a good description of aggregate economic outcomes in developed economies. Namely, both the employment–population ratio and the labor share were roughly constant. Second, we know from standard growth accounting exercises that changes in TFP are the dominant source of growth. Together, these observations tell us that for a long period in many economies, steady growth in TFP has been accompanied by stable values for both the employment–population ratio and the labor share.

It follows that if TFP growth is found to have significant effects on either the employment–population ratio or the labor share in the post-1970 period, this must surely reflect a change in the effects of TFP on these outcomes. The simplest explanation for why the effects of TFP might have changed surely lies in the possibility that the nature of technological progress has changed, and the authors clearly note this. But to my mind, this suggests that any study seeking to link changes in technology to recent changes in labor market outcomes must also make some effort to isolate the potentially different components of technological progress. Moreover, sectoral data would potentially be of particular importance in this regard, because we might think that there is heterogeneity both over time and across sectors in the composition of technological change and that this variation would prove to be important.

In fact, the EU KLEMS database that the authors use for their analysis does provide information on different categories of investment, and, to the extent that the authors wish to assess the effects of automation per se on labor market outcomes, a key limitation of the analysis is that they have not integrated this additional information into the analysis. On this point, I would again note the important earlier contribution to the *Brookings Papers* by Elsby, Hobijn, and Şahin (2013). Like the present paper, it sought to shed light on the causes behind the declining labor share by examining data at the sectoral level, though it focused solely on the United States. In their analysis, Elsby, Hobijn, and Şahin did consider the role of investment in equipment, and they found that it had little explanatory power for understanding the dynamics of the labor share over time and across sectors.

LEVELS VERSUS FIRST DIFFERENCES The basic premise in using sectoral data from several countries to estimate effects is that there is something common about how a given *change* in TFP affects outcomes across countries. Although this is a standard approach, the basis for it in the present context is not entirely clear. Two features of the data are notable. First, at any point in time, there are large differences in TFP across the countries in the authors' sample. Second, there are also large differences in labor share in a given sector across countries. If the labor share we observe is related to the technology being used, which is a basic premise of the analysis in this paper, then we might think that it is the level of TFP (that is, the technology being used) that is related to the labor share, and that one cannot assume that a given change in TFP has the same effect on employment and the labor share independent of the initial level of TFP.

To pursue this further, suppose that one country has TFP that is only 80 percent that of the leader in some sector. Suppose both this country and the leader experience an improvement in TFP of 5 percentage points. Assuming that this will have the same response in both countries is to assume that the effect of TFP on these variables is linear. But if we think that changes in the nature of technological progress are influencing these effects, it seems unclear that this is a reasonable assumption. Perhaps the trailing country should have effects that resemble those of the leading economy when it moved from 80 to 85 percent of its current level.

BENCHMARK SPECIFICATION Although not stated explicitly in the paper, I think it is understood that the goal of this analysis is to uncover the "long-run" effects of automation on employment and the labor share. This is a key point to note, given that there is good reason to believe that short-run effects might be very different. In particular, we know that individuals displaced from certain industries often experience long spells of nonemployment, including early retirement. The labor share is implicitly affected by the response of both prices and wages, and, to the extent that these variables respond very differently in the short and long runs, it is clearly important to distinguish between short- and long-run responses. Also, the dynamics of TFP changes may have varying serial correlation over time. Although I appreciate that the authors have taken some care to isolate the long-run effects of changes in TFP, I remain somewhat skeptical about the extent to which they have purged their results of short-run effects.

My own preference would have been for them to focus on consecutive long-period differences to generate their benchmark results. In the robustness section, they do present results for differences over consecutive five-year periods. The results for this case were only about two-thirds as large as for their benchmark specification, suggesting that this difference is potentially significant. Moreover, even in this setting they do not exclusively rely on five-year differences, because they retain observations for the period 2005–07 and they use observations for periods in which a country's data start in between the two endpoints. I suspect that about 20 percent of their observations in this exercise are not from five-year differences. This detail aside, my own preference would be to focus on ten-year differences. The authors' own calculations lead them to conclude that effects require about five years, and using five-year differences implies that any changes after the initial year will not have realized their full effect at the end of the interval. To retain data from the post-2000 period, they could define the three periods as 1975–85, 1985–95, and 1995–2005, or, alternatively, as 1977–87, 1987–97, and 1997–2007. I would find results from this specification to be both more transparent and more compelling.

VALUE-ADDED TFP VERSUS GROSS OUTPUT TFP In their initial specification, the authors run regressions of the outcome of interest on value-added TFP at the sector level. They later suggest that they want to incorporate sectoral input–output linkages into the analysis and use this to motivate the inclusion of value-added TFP terms from other sectors on the right-hand side, distinguishing them in terms of being upstream or downstream. When doing this, the authors continue to use value-added TFP measures as their TFP measure. I think it is problematic to continue to use value-added TFP measures in a context where one aims to measure effects propagated through input–output connections. Unfortunately, I think this is an issue that many researchers seem not to appreciate, so this is one of the reasons I raise it here.

The first point to realize is that value-added TFP and gross output TFP are two truly distinct objects. It is particularly important that value-added TFP already incorporates the effects of technological progress in supplier sectors.[4] Relatedly, when one chooses to represent the production side of the economy via value-added production functions, one is not assuming that there are no input–output relationships; rather, they are embedded in the value-added TFPs.

One reason for preferring a specification in which one starts with gross output production functions is that one might reasonably think that this provides a better description of how primitive technology shocks appear—that is, they affect the ability of a given sector to produce gross output. Propagation through the input–output network will imply that value-added

4. See Moro (2012) for derivations that explicitly link value-added and gross output TFP in a simple setting.

TFP will change in other sectors, so that a representation using value-added production functions cannot easily be used to trace out the effect of a primitive technology shock to the gross output production function in one sector. However, for a given set of gross output TFP shocks, the valued-added representation will capture the same set of overall equilibrium responses, so there is no benefit to adopting one approach versus the other if one wants to study the overall effect of observed shocks. But using value-added TFP when explicitly studying input–output linkages makes it impossible to disentangle direct effects from effects operating through input–output linkages.

CONCLUDING COMMENTS One of the goals of Autor and Salomons's paper was to provide estimates of the effects of TFP on aggregate outcomes. To be sure, this is a challenging goal, and I do not think the profession is yet able to produce reliable estimates of these effects. In particular, for the reasons I have described, I do not feel that the approach taken in this paper is particularly promising in this regard. Nonetheless, the authors are to be commended for compiling a large amount of evidence about the relationships between key labor market outcomes at the sector level for a large set of countries. I think this information is valuable in the effort to learn more about the driving forces and mechanisms at work, and it will surely be useful to future researchers working on this important issue. But given the limitations of the methods used for uncovering quantitative causal relationships, I would have preferred if the authors had focused more on how the cross-country evidence shapes our priors about the plausibility of technological factors compared with other factors.

REFERENCES FOR THE ROGERSON COMMENT

Baumol, William J. 1967. "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." *American Economic Review* 57, no. 3: 415–26.

Blanchard, Olivier, and Justin Wolfers. 2000. "The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence." *Economic Journal* 110, no. 462: 1–33.

Bruno, Michael, and Jeffrey D. Sachs. 1985. *The Economics of Worldwide Stagflation*. Harvard University Press.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2013. "The Decline of the U.S. Labor Share." *Brookings Papers on Economic Activity*, Fall: 1–52.

Kaldor, Nicholas. 1961. "Capital Accumulation and Economic Growth." In *The Theory of Capital*, edited by Friedrich Lutz and Douglas C. Hague. New York: St. Martin's Press.

Krugman, Paul. 1994. "Past and Prospective Causes of High Unemployment." *Economic Review* (Federal Reserve Bank of Kansas City) 79, no. 4: 23–43.

Moro, Alessio. 2012. "The Structural Transformation between Manufacturing and Services and the Decline in the US GDP Volatility." *Review of Economic Dynamics* 15, no. 3: 402–15.

GENERAL DISCUSSION Robert Gordon noted that a large number of industries in the Bureau of Labor Statistics data have negative total factor productivity (TFP) growth, as noted by commenter John Haltiwanger. One might be tempted to believe that negative TFP growth is a result of measurement error. But consider the example of the higher education industry, one in which negative TFP growth could actually be real, and a result of various processes unrelated to innovation. At a university, undergraduate students are the "output," which the university produces at a relatively fixed level and with a stable quality over time. "Inputs" include professors, administrative resources, and information technology. But other inputs include investments in expensive new buildings and sports facilities, which require maintenance over time. With the output of students fixed and with the value of inputs rising over time, TFP growth in higher education could indeed be negative. This framework could also apply to other industries. In retail, for example, e-commerce may not contribute enough to productivity growth to offset the decline in traditional bricks-and-mortar stores. In health care, hospitals have upgraded their facilities and hired additional staff, but without an increase in patient care. Taken together, such examples may explain why negative TFP growth could be a real phenomenon and not a product of measurement error.

Martin Baily made two comments. First, as commenter Richard Rogerson had suggested, labor supply is most likely the main determinant of employment at an aggregate level, not TFP growth. TFP growth might have effects on employment at a micro or industry level, but not at the aggregate level. Second, though the instrumental variables used by the authors as proxies for innovation—for example, patents and robot adoption—were interesting, Baily suggested they were probably poor instruments. Robot adoption, he reasoned, is still a recent phenomenon, and patent flows should probably only have a small effect on productivity. In his own research, Baily has found that many of the changes in productivity over the past few decades were not actually due to automation, but rather to scaling and business organization. In retail, for example, the transition from small mom-and-pop stores to big box–style department stores was a major driver of productivity growth. In the automobile industry, companies using similar technology and equipment were able to increase productivity mostly by organizing production more effectively. Research by the McKinsey Global Institute suggests that much of the decline in the labor share of income in the United States was due to changes in the manufacturing industry.[1]

Katharine Abraham disagreed with Gordon, arguing that there are significant parts of the economy for which mismeasurement of TFP and productivity growth are a real concern. The health care industry is an example of a major industry for which the difficulty of measuring output causes serious problems for measurement of productivity. She suggested that such measurement issues could pose major problems for the broad conclusions reached by the authors. Baily had suggested that mismeasurement might not be an issue as long as relative productivity growth—how the rate of productivity growth in any one industry compares to that in other industries—was not systematically affected by measurement bias. Instead, she suggested that measurement bias might shift the levels of TFP, productivity, and output by different amounts in different industries.

Valerie Ramey recommended that the authors revisit past research that showed a negative effect of productivity growth on employment and hours. Similar results were found in work by Olivier Blanchard and Danny Quah; Jordi Galí; Neville Francis and Ramey; and Susanto Basu, John Fernald, and Miles Kimball.[2] She recommended using the models developed in these papers to study the effect of technology on the labor share of income to see if they gave similar answers.

Robert Hall mentioned a few identification issues with the authors' empirical model. First, he noted that the authors assume there is no unobserved covariate simultaneously affecting both the labor share of income (the dependent variable) and TFP growth (the independent variable). If a covariate existed, it would bias the model's results. Second, the authors did not, in Hall's view, adequately demonstrate that the

1. Sree Ramaswamy, James Manyka, Gary Pinkus, Katy George, Jonathan Law, Tony Gambell, and Andrea Serafino, "Making It in America: Revitalizing US Manufacturing" (McKinsey Global Institute, 2017).

2. Olivier Jean Blanchard and Danny Quah, "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review* 79, no. 4 (1989): 655–73; Jordi Galí, "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89, no. 1 (1999): 249–71; Neville Francis and Valerie A. Ramey, "Measures of Per Capita Hours and Their Implications for the Technology-Hours Debate," *Journal of Money, Credit and Banking* 41, no. 6 (2009): 1071–97; Susanto Basu, John G. Fernald, and Miles S. Kimball, "Are Technology Improvements Contractionary?" *American Economic Review* 96, no. 5 (2006): 1418–48.

number of patent claims is a viable instrumental variable; specifically, it was not shown to be statistically independent of an unobserved covariate. Hall suggested one possible unobserved covariate might be the existence of market power in a given industry, which he suspected was of first-order importance.

Salomons first addressed data questions. She noted that although the authors present results for higher-frequency data, they also estimated models for long time intervals over decades, which yielded quantitatively and qualitatively similar results to those presented at higher frequencies. Second, she clarified the identification strategy used to estimate the aggregate effects of TFP growth across countries and industries. The authors did not look only at interactions between time and countries or take a set of country-year fixed effects in a regression to scale up the broad effects of automation on employment. Rather, they measured the effect of TFP growth on value added at the country, industry, and year levels, and then measured the total value-added effects across countries by using own-industry coefficient estimates on various macroeconomic outcomes, such as labor's share of income and employment. The empirical model could not include country-year fixed effects because doing so would absorb any variation in TFP growth across countries.

Salomons acknowledged that measurement issues associated with TFP growth could be a real concern with regard to the authors' conclusions. The reason the authors used TFP as a proxy for automation was because it is a broad measure, and does not erroneously focus on some very specific, idiosyncratic trend. Robotics, for example, might be an interesting measure of automation, but it has only a limited effect on select sectors of the economy. She acknowledged that a major drawback of using TFP growth is that it might be *too* broad, which is why they tried using robotics and patents as instrumental variables. Finally, Salomons commented on the possible drivers of the reallocation of resources between industries or firms, which lead to differing levels of productivity growth. The paper is silent on the causes of this reallocation, but Salomons noted that it could certainly be due to market power, as suggested by Hall and by Autor in previous work.[3]

Autor acknowledged that he and Salomons were sensitive to the issues of measurement and the omnibus definition of TFP as a measure

3. David Autor, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," Working Paper no. 23396 (Cambridge, Mass.: National Bureau of Economic Research, 2017).

of innovation and automation. He suggested that the instrumental variables used in their paper—robotics and patents—might be thought of as a "rescaling" of TFP. Patent flows, in particular, should be a good measure of innovation, and the data are highly correlated across countries. Although Autor acknowledged that these measures should probably not be treated as excludable instrumental variables, he suggested that they are still useful proxies for automation. In the final version of the paper, the authors use patents as a proxy for TFP—for which it has strong predictive power—but not as an instrument for TFP.

Blanchard and Rogerson had asked about identification, particularly related to the omission of country-year fixed effects. Echoing Salomons's comments, Autor noted that including country-year fixed effects would absorb the variation in the data required to provide a measure of the aggregate effects of TFP growth on macroeconomic outcomes across countries and over time. The authors do, however, measure the country-year effects more broadly by using the chain rule to get the effect of a particular industry on an entire country, and then they estimate the effect of the country on broad outcomes. Autor conceded that this still might not be the perfect identification method.

Autor agreed with Baily's comments about the problem of excluding the labor supply from the model. However, he argued that if the approach used in the paper was wrong, and the labor supply was in fact the main driver of employment, then one would expect productivity growth to be unrelated to employment, which is largely what they find. More surprising, however, is that they find that productivity growth is also negatively related to the change in labor's share of income, and that this effect changes over time.

Autor emphasized that the ultimate goal of the paper is not to exposit the driving forces behind TFP growth, a rather broad, omnibus measure of productivity, but rather to explore how productivity growth affects industry-level and aggregate employment, sectoral reallocation, and the evolution of labor's share of national income. Because TFP is difficult to measure, the nature of productivity growth is often unclear. One model of productivity growth, proposed by Daron Acemoglu and Pascual Restrepo, posits that productivity growth could be labor-intensive and capital-augmenting, thereby complementing the use of labor and expanding the number of available tasks, rather than reducing it.[4] Alternatively, productivity growth

4. Daron Acemoglu and Pascual Restrepo, "Artificial Intelligence, Automation and Work," Working Paper no. 24196 (Cambridge, Mass.: National Bureau of Economic Research, 2018).

could be labor-displacing, meaning it reduces the share of output paid to labor. Autor emphasized that the exercise in the present paper is designed to tease out the nature of productivity growth by examining its effect on employment outcomes, which provides information about the nature of the productivity growth occurring. He also acknowledged the limitation of their cross-country, industry-level panel data set. Their main source of data, the EU KLEMS database, does allow for better analysis across countries than would be possible working with more detailed micro-level data specific to individual countries. He noted that the paper would likely be insufficient to satisfy macro or labor economists, but that the authors hoped to connect the two disciplines in an informative way—or at least to untie them in their shared disapproval of the methods and conclusions of the present paper.