Policy Research Working Paper 9529

The Technology-Employment Trade-Off

Automation, Industry, and Income Effects

Gene Kindberg-Hanlon

WORLD BANK GROUP
Prospects Group
January 2021

Policy Research Working Paper 9529

Abstract

New technologies can both substitute for and complement labor. Evidence from structural vector autoregressions using a large global sample of economies suggests that the substitution effect dominates in the short-run for over three-quarters of economies. A typical 10 percent technology-driven improvement in labor productivity reduces employment by 2 percent in advanced economies in the first year and 1 percent in emerging market and developing economies (EMDEs). Advanced economies have been more affected by employment-displacing technological change in recent decades but the disruption to the labor market in EMDEs has been more persistent. The negative employment effect is larger and more persistent in economies that have experienced a larger increase, or smaller fall, in industrial employment shares since 1990. In contrast, economies where workers have been better able to transition to other sectors have benefited more in the medium run from the positive "income effect" of new technologies. This corresponds with existing evidence that industrial jobs are most at risk of automation and reduced-form evidence that more industrially-focused economies have tended to create fewer jobs in recent decades. EMDEs are likely to face increasing challenges from automation as their share of global industry and production complexity increases.

This paper is a product of the Prospects Group. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/prwp. The author may be contacted at gkindberghanlon@worldbank.org.

The Technology-Employment Trade-Off: Automation, Industry, and Income effects*

Gene Kindberg-Hanlon†

Keywords: Employment, Productivity, Technology shocks, SVARs.
JEL classification: C11, C30, C33, E24

*I would like to thank Alistair Dieppe, John Fernald, Neville Francis, Ayhan Kose, Franziska Ohnsorge, Temel Taksin, and Lei Ye for valuable comments and Aygul Evdokimova for excellent research assistance. I gratefully acknowledge financial support from the PHRD fund. The findings, interpretations, and conclusions expressed herein are those of the author and do not necessarily reflect the view of the World Bank.

†World Bank gkindberghanlon@worldbank.org

1 Introduction

Concerns are frequently raised about how the gains from new technologies are shared, particularly through their impact on employment. Currently, concerns are highest around the automation of manufacturing jobs but historically many major innovations have been accompanied by the threat of job losses (World Bank, 2019). New technologies can be both a substitute or a complement for labor, and therefore can boost job opportunities as well as reduce them (Autor, 2015). Certain segments of the labor market can be harmed by technological change even where they result in new labor tasks. Where the skills needed to accompany new technologies are unavailable, or demand for new labor tasks does not rise sufficiently, aggregate employment can be persistently lower.

A large literature has attempted to assess the impact of technological change on employment within affected sectors but so far the effects on aggregate employment have been under-explored, particularly in EMDEs. This paper sheds new light on the question of how typical technology shocks affect aggregate employment in a broad range of 30 advanced economies and 96 emerging and developing economies (EMDEs). The large coverage relative to existing studies is enabled by new methodologies developed in Dieppe et al. (2019, 2021).[1] A further key contribution of this paper is that the large sample size of economies under consideration allows for an exploration of country-characteristics which are important in determining the size and persistence of employment impacts following a technology shock.

"Technology shocks" in the sense discussed in this paper are derived from a structural vector-auto regression (SVAR). These are identified as innovations which account for the largest share of long-term (low-frequency) changes in productivity or TFP (Dieppe et al., 2019). The resulting shocks "look-through" temporary changes in labor productivity or TFP, which are often driven by non-technological factors such as changing factor utilization (Basu et al., 2006).[2] Furthermore, this methodology captures many types of technology shock; these can be neutral or complementary to labor, resulting in increased employment, and can also be labor-saving, such as automating technology developments (Acemoglu, 2003). The shocks captured by the SVAR approach will therefore reflect the "typical" technology introduced over the estimation horizon.

Ninety percent of advanced economies and 70 percent of EMDEs experience a fall in employment in year 1 following a positive technology shock, while half of advanced economies and one-third of EMDEs experience a statistically significant fall in employment. In many economies, this fall in employment is relatively persistent, lasting 10 years on average before the effect dissipates. The finding is robust across multiple estimation methodologies for identifying technology shocks and suggests that the type of technology shocks affecting most economies have been labor-substituting.

Economies with higher average productivity levels have suffered larger initial employment losses

[1] Galí (1999) and Rujin (2019) assess the impacts of technology shocks on hours and employment in the G7.

[2] This identification does not rule out that factors other than productivity-enhancing technologies can drive long-run productivity developments but it does assume that this is the dominant driver.

following technology shocks, reflecting a higher propensity to introduce labor-saving technologies in advanced economies relative to EMDEs. EMDEs have yet to be affected by labor-substituting technological change to the same degree as advanced economies. However, EMDEs have been affected by these technologies in large number and experienced more persistent employment losses on average than advanced economies. This finding corresponds with the finding that "premature deindustrialization" has been a widespread phenomenon in EMDEs. Higher trade openness and exposure to FDI inflows are associated with *lower* initial employment losses following technology shocks, consistent with a large literature that has found FDI is associated with job gains in EMDEs; firms frequently invest in these economies to take advantage of lower labor costs, not to remove workers from the production process.

Those economies where employment shares have remained more concentrated in the industrial sector since 1990 have seen larger and more persistent employment losses on average following a technology shock. This may seem paradoxical at first: the industrial sector, which has had a higher-than-average potential for automation in recent decades, has proved more resilient as a share of employment in the face of labor-displacing technological innovation in some economies. A simple adaption of the framework of Acemoglu and Restrepo (2017, 2018) to incorporate differential income elasticities shows that the services sector will be the key beneficiary of increased demand due to the efficiency gains from automation that occur in the industrial sector. This is because the income-elasticity and price-elasticity of demand for industrial goods is low, so that the displacement effect for workers is unlikely to be offset by rising demand. In a world where the majority of economies have experienced the direct effects of the loss of employment from labor-displacing technologies, those that have been best able to transition workers to less-affected sectors such as services will have stronger aggregate employment growth and a larger fall in the share of industrial employment. Those economies that have been less successful at redeploying workers to other sectors will have lower aggregate employment but higher industrial employment shares.

This paper proposes an additional perspective on the "premature deindustrialization" hypothesis (Rodrik, 2016), where lower employment shares in industry are presented as a loss of the "escalator of development"; instead, accelerated rotation away from industrial employment may be seen as a sign of success in many cases, resulting from a smooth transition of workers and demand to new sectors following the introduction of new labor-displacing technologies in industry.

A lack of transferable skills to other sectors will reduce the growth of aggregate employment following a labor-substituting technology shock. Policies to improve the skill base of workers can aid worker transitions to sectors less affected by labor-displacing technologies but may require a prolonged period to implement and reap returns. Policies to increase demand for industrial products may provide nearer-term relief; there are large scale requirements for infrastructure investment to meet poverty reduction targets in EMDEs and emission reduction targets in both advanced economies and EMDEs. Public sector investment and increased incentives for private sector invest-

ment will boost industrial output and could help reduced prolonged falls in employment as new technologies are introduced.

2 Literature

Before turning to the new SVAR evidence on the effects of productivity-enhancing technology shocks presented in this paper, it is useful to review the evidence from the existing SVAR literature, alongside the theoretical and non-SVAR empirical evidence for the effects of automation and new technologies on employment.

2.1 Theory

Productivity-improving technologies can generate opposing forces on total employment in an economy: first, a substitution effect, where new technologies can replace the need for workers; and second, an income effect, where increases in the profitability of production increase the demand for labor, in the affected or alternative sectors (Aghion and Howitt, 1994). The ability of the income effect to offset automation will depend crucially on the type of workers required to complement new technologies and capital assets, and the supply of workers with the appropriate skills for these tasks (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018).

Search and matching models have been used to examine the employment displacement effects of introducing new technologies into the production process. Mortensen and Pissarides (1998) show that as the costs of updating skills and equipment for existing workers grow, "creative destruction", or the replacement of current labor tasks with capital embodying new technologies is likely to occur (Schumpeter, 1942). In addition, they find that employment protection and unemployment benefit levels can exacerbate the degree and persistence of unemployment following a technology shock (Mortensen and Pissarides, 1999). Restrepo (2015) finds that employment in routine jobs is likely to have declined since the 1990s due to search-and-matching frictions, as new technologies frequently require increasingly novel skills.

This paper focuses on new technologies displacing labor as a driver of falling employment following a technology shock. However, a second mechanism exists for technology shocks to reduce employment (or labor input more generally). The bulk of the literature analyzing the impact of technology shocks using SVAR identification techniques have focused on the effects of sticky prices in the canonical New Keynesian model as the reason for falling employment. Here, the mechanism is that aggregate demand remains inflexible in the short run due to sticky prices and grows by less than productivity following a technology shock. This leads firms to cut labor input in response to new technologies. As prices adjust, demand expands and employment recovers (Galí, 1999; Basu

et al., 2006).[3] In contrast, this paper finds evidence of more persistent employment effects than might be driven by sticky prices. For most economies, investment does not fall in response to a technology shock. Investment, as well as employment, is expected to fall initially in standard New Keynesian models.

In some cases, the SVAR literature has assigned the fall in employment found in some advanced economies following an SVAR-identified technology shock to the displacement of labor for new technologies, rather than the sticky-price mechanism. Canova et al. (2013) and Michelacci and Lopez-Salido (2007) attribute falling employment after a technology shock to Schumpeterian 'creative destruction' effects. This paper builds on this tradition, using a significantly larger set of economies under consideration, and using new SVAR identification techniques. However, a range of alternative and complementary approaches have also been implemented to establish the effects of new technologies on employment.

2.2 Sectoral evidence

A large body of evidence has shown that jobs have become increasingly polarized into low- and high-skill occupations in the U.S. and Europe in recent decades, as a combination of automation and offshoring has reduced demand for middle and low-skilled workers performing routine and codifiable jobs (Acemoglu, 1999; Autor, 2015; Goos et al., 2014). Many of these lost occupations were in the industrial sector, even as value-added produced by the sector remained resilient; in the United States, employment of machine operators, assemblers, and other production employees fell by over one-third every 10 years between 1980 and 2005 (Autor and Dorn, 2013). In a study of 16 European economies during 1993-2010, the share of employment accounted for by middle- and low-skilled industrial sector occupations fell by nearly 10 percentage points (Goos et al., 2014). In the United States and France, the increased use of robotics in industry is found to be inversely related to industrial employment levels since 1990 and 2010, respectively (Acemoglu et al., 2020; Acemoglu and Restrepo, 2020). Some service sector occupations are also found to have been negatively affected by this trend in both regions, notably middle-skilled jobs such as office clerks. However, codifiable middle- and low-skill jobs have been (at least partially) replaced by higher demand for both low-skill service sector jobs, which are less easy to automate, and higher-skill jobs that complement new technologies. SVAR analysis of sectoral manufacturing data for advanced economies has also found negative effects on total hours worked of developments that have driven persistent positive TFP growth (Chang and Hong, 2006; Park, 2012; Khan and Tsoukalas, 2013).

[3] A wide range of SVAR literature has tackled the response of labor inputs in response to technology to assess the presence of a sticky-price mechanism, including (to name just a few) Christiano et al. (2004), Dedola and Neri (2007), Collard and Dellas (2007), Francis and Ramey (2005), Francis et al. (2014), and Canova et al. (2010). The majority of research has found evidence of decreasing employment from technology shocks in the U.S. context.

2.3 General equilibrium impacts of technological progress on employment

Several studies in the U.S. have found that technological change has caused *aggregate* employment to fall. These studies have gone beyond an examination of the sectoral effects of technological job displacement. During the so-called 'jobless recoveries' in the US after the recessions of 1991, 2001, and 2008, when the employment rate fell or stagnated overall, middle-skilled and automatable roles declined, particularly in the manufacturing sector (Jaimovich and Siu, 2019; Charles et al., 2016). It has been further argued that even high-skilled workers have been substituted for newer technologies and pushed into lower-skilled positions, reducing overall employment (Beaudry et al., 2016). However, there remains controversy over the net impacts of technological change. In some cases, the fall in employment in the sector affected by technological change is found to be offset by employment gains in other sectors, particularly in "downstream" sectors which use inputs from the affected sector (Autor and Salomons, 2018).

2.4 EMDE evidence

There is so far little evidence of the effects of technological change on employment in EMDEs. In part this is because EMDEs have been large beneficiaries of outsourcing from advanced economies; many manufacturing and codifiable service sector jobs have moved to EMDEs to take advantage of cheap labor costs (Maloney and Molina, 2016). What technology-influenced change does appear to be occurring has increased the share of routine semi-skilled jobs in many EMDEs, in contrast to the fall in these types of jobs in advanced economies (World Bank, 2019). That said, large increases in manufacturing productivity do appear to have resulted in "premature deindustrialization" in EMDEs, where employment falls in the sector at much lower levels of income per capita than has occurred historically (Rodrik, 2016). That could suggest that productivity-enhancing technology in the manufacturing sector has reduced employment relative to a counterfactual, which would have been otherwise higher still.

3 Estimating the effects of technology shocks on employment

Much of the evidence on the impact of productivity-improving technology on employment has centered around the impacts of IT and manufacturing technologies in the US and Europe in recent decades. There has yet to be an assessment of the effects of general improvements in technology on employment in a broad range of countries. To assess the impact of productivity-enhancing technology changes on a range of countries, we turn to the SVAR literature, which has already

extensively estimated the relationship between "technology shocks" and total hours worked in the United States and some European economies, finding a negative impact on total hours worked in the first year following the shock (Galí, 1999; Francis and Ramey, 2005; Francis et al., 2014).[4]

In this paper, innovations in technology ("technology shocks") are identified as shocks that drive the largest proportion of long-run developments in labor productivity, while other shocks are assumed to have only transitory effects. Therefore, technology shocks represent any structural development that results in lasting changes in labor productivity. These innovations could boost labor productivity through higher technical efficiency, in the sense of higher TFP in a standard Cobb-Douglas framework. Equally, new technologies may allow for the replacement of tasks previously completed using manual labor with capital, or "robot" inputs, altering the ratio of labor and capital inputs, without necessarily boosting the overall efficiency with which capital and labor are combined (Acemoglu and Restrepo, 2018). In both cases, output per unit of labor input is increased. The identification may capture a mix between labor-augmenting and labor-displacing technologies.

Traditionally, long-run restrictions have been used to identify technology shocks. However, long-run restrictions are susceptible to significant bias from the presence of non-technology shocks which can also drive developments in productivity (Chari et al., 2009; Dieppe et al., 2019, 2021). Instead, the Spectral identification, detailed below, is more robust, particularly when applied to volatile emerging and developing market economy data.

3.1 Spectral identification of technology shocks

The Spectral identification searches for the shock which maximizes the contribution to the variance of productivity at long-run frequencies. This abstracts from the shocks which may instead drive variation in labor productivity at business-cycle frequencies, such as demand and other short-run impact shocks. This approach effectively applies a band-pass filter to the reduced-form coefficients of a VAR containing macroeconomic variables, identifying the spectral density of the variables within a particular frequency band. The technology shock is then identified by maximizing the variance of labor productivity explained at the desired frequency.

We start by writing the Wold representation of the VAR (assuming it is invertible):

$$Y_t = \left(I - \left(B_1 L + B_2 L^2 + \ldots B_p L^p\right)\right)^{-1} u_t = D u_t \tag{1}$$

Here, B are the reduced-form auto-regressive coefficients, while u are the reduced form errors in the VAR representation and D reflects the sum of the MA-representation coefficients of the VAR. By post-multiplying Y_t by $Y_{t-\tau}$ and summing across its lags (of τ periods), the series of auto and cross covariances (γ) can be written as a function of D and the variance-covariance matrix of errors:

[4]For a detailed review of the technology SVAR literature, see Ramey (2016).

$$\sum_{\tau=-\infty}^{\infty} \gamma(\tau) = \sum_{\tau=-\infty}^{\infty} EY_t Y_{t-\tau} = D\Sigma_u D' \tag{2}$$

The spectral density of Y at frequency ω can be written as a function of D, $D\left(e^{-i\tau\omega}\right) = \left(I - \left(B_1 L e^{-i\omega} + B_2 L^2 e^{-i2\omega} + \ldots B_p L^p e^{-ip\omega}\right)\right)^{-1}$:

$$S_{YY}(\omega) = D\left(e^{-i\tau\omega}\right)\Sigma_u D\left(e^{-i\tau\omega}\right)' = \sum_{\tau=-\infty}^{\infty} \gamma(\tau)\, e^{i\tau\omega} \tag{3}$$

Due to the limited time series available for estimation in many EMDEs, the "limited spectral" identification, where D is truncated at the 10-year horizon due to the associated biases when estimating the long-run estimation on finite data (Dieppe et al., 2019, 2021).

To assess the spectral density within a frequency band, the spectral power can be integrated between $\omega = [\underline{\omega}\,, \bar{\omega}]$. The shock that maximizes the variance of labor productivity over the desired frequency is the eigenvector associated with the largest eigenvalue of $S_{YY}(\omega)$, see (Uhlig, 2003). The frequency band of interest here is from 10-∞ years so that technology shocks are identified as the shock that dominates very long-run productivity developments, while excluding short-term developments such as those at business-cycle frequencies from the maximization problem.

3.2 Assessing the impact of technology shocks in a pooled panel environment

Initially, in order to understand the typical impacts of technology shocks across countries, a simple pooled panel estimation with fixed effects is used. This is performed separately across advanced economies and EMDEs, particularly since the variance of productivity is three times larger in EMDEs than in advanced economies, given assumptions required on common parameters for panel estimations. For each economy, the VAR consists of the log-level of labor productivity, log-employment per capita (including self-employment), the log share of investment (gross-fixed capital formation), and separately consumption in GDP, and consumer price inflation.

Macroeconomic aggregates such as GDP and employment are from the World Bank's World Development Indicators (WDI) database and The Conference Board's Total Economy Database (TED) for employment. Employment and GDP data are extended where available with WDI and ILO employment estimates and Penn World Table 9.1 GDP data where possible. The majority of economies have full data coverage over the same period over which the growth accounting components are available, and the average sample length is 40 years for EMDEs and 45 years for advanced economies. Hence, here, annual data is used to estimate the VARs. This choice reduces the degrees of freedom in the VAR estimation, while at the same time significantly lengthening the period over which the VAR is estimated for many EMDEs. The span of the data is critical for identifying

technology shocks as those which drive long-term developments in productivity.

The estimation is performed using the pooled panel VAR approach:

$$Y_t^n = C^n + \sum_{\tau=1}^{k} B_\tau Y_{t-\tau}^n + u_t$$

Where C^n varies across countries, n, while B and the variance-covariance matrix of residuals Σ_u are assumed to be common across countries. The estimated parameters, B and Σ_u, can then be used to identify the effects of technology shocks using the spectral identification described above. Standard Normal-Inverse-Wishart priors are used in the Bayesian estimation. For several EMDEs, periods of high, or hyper-inflation, have occurred. Dummy variables are included for those economies during periods in which inflation has exceeded 20 percent.

In both EMDEs and advanced economies, the impulse responses of a range of macroeconomic variables behave in a way consistent with theoretical predictions. Labor productivity rises on a sustained basis over the 10-year horizon under consideration (Figure 1). The additional IRFs are scaled to the size of the effect on productivity for greater ease in comparing the impacts between advanced economies and EMDEs. The IRFs for employment, investment, consumption, and consumer price inflation can, therefore, be interpreted as the impact on those variables for 10 percent improvement in labor productivity resulting from the technology shock.

In advanced economies, employment falls in the first year following a technology shock by 2 percentage points for each 10 percentage point rise in labor productivity.[5] In EMDEs, employment falls by 1 percent, but more persistently, only becoming statistically insignificant from zero after 9 years. In both advanced economies and EMDEs, investment rises more than one-for-one with the impact on productivity for the majority of the IRF horizon, while consumption adjusts more slowly to the new steady state of higher output. Consumer prices fall in both regions as output can be produced with fewer inputs. In EMDEs, the adjustment of prices is slower, possibly reflecting less well-anchored inflation expectations.

While the initial impact on employment in EMDEs is smaller than in advanced economies, it is more persistent. One interpretation is that EMDEs have introduced fewer labor-displacing technologies with a smaller initial disruptive impact on employment than in advanced economies. However, EMDEs have coped less well at moving workers into new roles, or generating higher demand from these new technologies, even at long horizons. The smaller negative employment impact in EMDEs is consistent with the literature which has found that the replacement of low and middle-skilled workers has primarily been an advanced-economy phenomenon so far, with many new technologies implemented in EMDEs specifically designed to take advantage of lower-cost labor inputs.

[5]Rounding from 1.6 percent for advanced economies, and 1.1 percent for EMDEs.

Figure 1: Technology shock IRFs targeting labor productivity: Pooled estimation

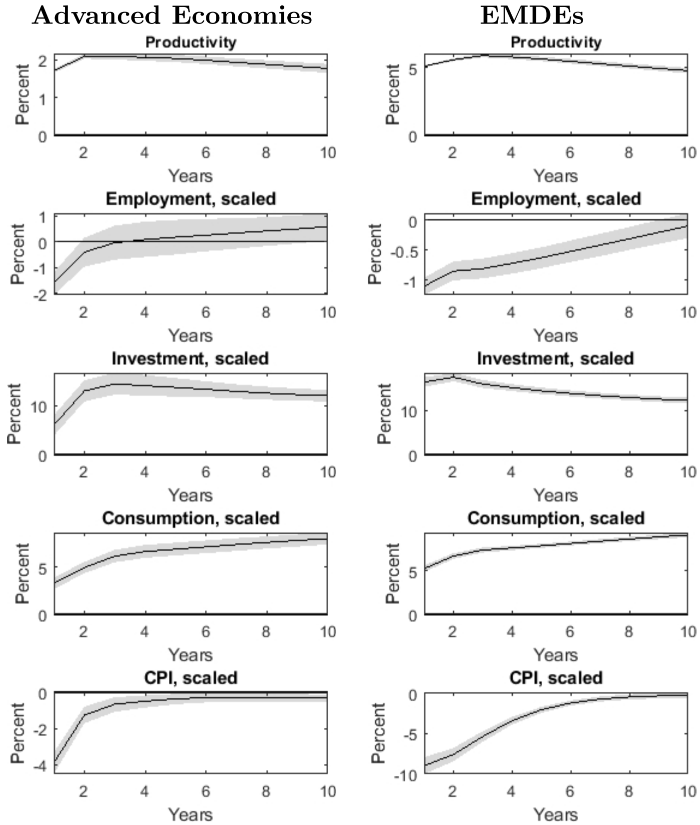

Note: With the exception of the labor productivity IRF, IRFs are scaled by the impact of the technology shock's impact on productivity. Each IRF can be interpreted as the effect of a technology shock which increases labor productivity by 10 percent. Shaded area reflects 68% confidence intervals.

As a robustness check, a new VAR specification is estimated, substituting log-TFP levels for labor productivity using data from the Penn World Table 9.1 (Appendix A, Figure 10). This second approach may more accurately capture the "technology" driver of productivity as arguably, shocks explaining the largest share of low-frequency variation in labor productivity may be subject to contamination from capital-specific shocks, such as corporate tax cuts that are expected to be permanent. There are also downsides to targeting TFP, the estimation of which is highly liable to measurement error of the capital stock, particularly in EMDEs. Under this approach, the impacts are similar in magnitude for most variables but the employment impacts become even more persistently negative in EMDEs.

3.2.1 Results of individual VAR estimations

A key assumption of the pooled panel estimation is the assumption of both a common set of autoregressive parameters (B) and a variance-covariance matrix. The finding of falling employment following a technology shock may therefore simply be an average effect that does not apply to all EMDEs and advanced economies, or be swayed by outliers.

As an alternative, individual country estimations are produced. There are several limitations to this standalone identification: annual data limits the degrees of freedom for estimation, and, measurement error, particularly in EMDEs, could limit estimation accuracy and inference. These limitations can be partly tackled with an additional panel approach using hierarchical priors that allows for a degree of heterogeneity while dealing with limited sample size issues by incorporating priors based on estimations from other economies (Jarociński, 2010). Further details of the hierarchical prior estimation are provided in Appendix A.

Individual and hierarchical-prior based estimations point to a high degree of homogeneity in the response of employment to technology shocks. Individual country results are provided in Appendix C. Results demonstrate that the fall in employment is not a result of outliers, but a widespread result of technology shocks in many economies. Around three-quarters of EMDEs and 90% of advanced economies experience a negative response of employment to technology shocks in year 1 (Figure 2). 40-50% of all economies experience a statistically significant fall in employment at the 16% confidence level when incorporating hierarchical priors into the estimation. This ratio falls to approximately 35% in EMDEs without the use of these priors.

Figure 2: Proportion of economies with negative, and statistically-significant negative impacts of technology on employment in year 1

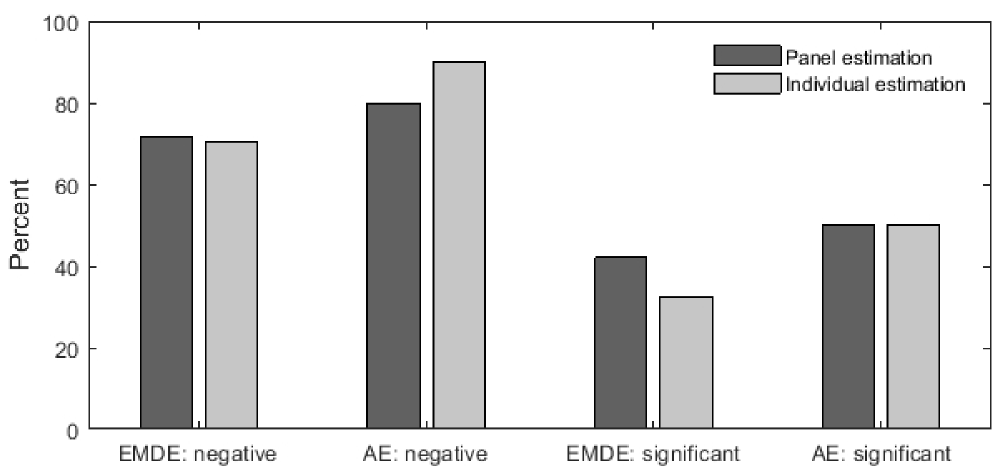

The results of the employment impacts using the Spectral identification show relatively modest

differences when incorporating hierarchical priors. This is a result of this identification methodology proving relatively robust to the issues posed by short samples and large contamination from non-technology shocks in the data. As an additional robustness exercise, the more established long-run restriction and Max-Share methodologies are also implemented (Appendix A). Using these methodologies, 60-80% of economies experience a negative impact on employment, while the use of hierarchical priors has a more broad-based effect on increasing the number of economies that are found to have statistically significant negative impacts.

Figure 3: Distribution of employment impacts of technology improvement across economies

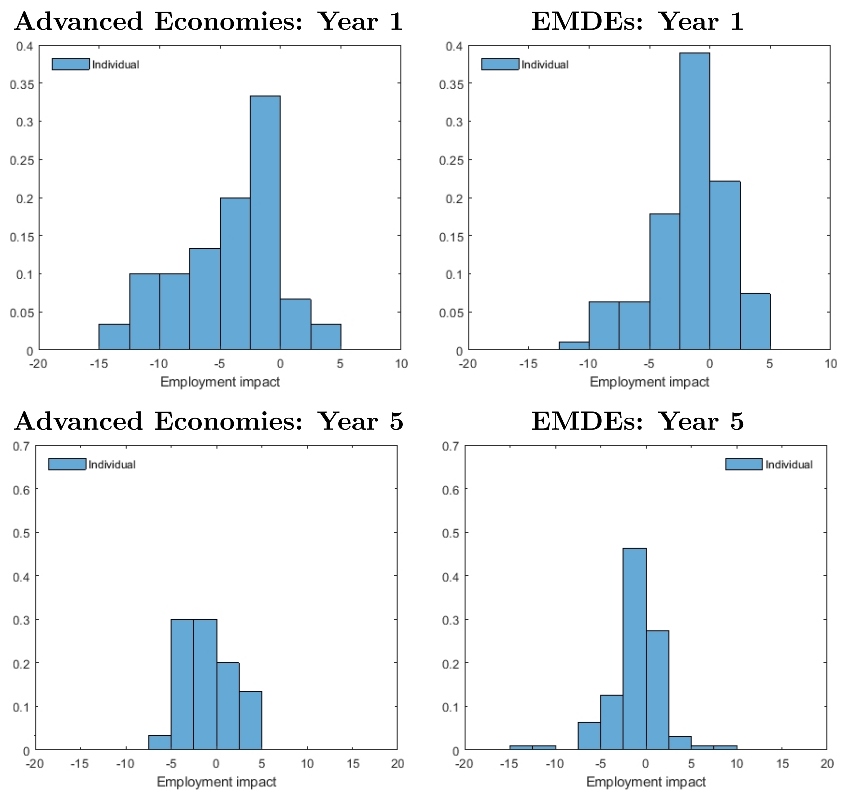

Note: Chart shows the density of median IRFs from individually-estimated VARs in advanced economies and EMDEs respectively. All IRFs are scaled by the impact of the technology shock's impact on productivity. The distribution therefore shows the range of responses of employment across countries from the effect of a technology shock that boosts labor productivity by 10 percent.

Figure 3 shows the distribution of scaled employment impacts (per one-percent labor productivity impact) at the one-year and five-year horizon. The impacts show a modal peak in both

advanced economies and EMDEs at modestly negative values. Importantly, the distribution does not display bi or multi-modal features. Therefore, the negative impact does not appear to be a categorical effect, i.e. countries that have a specific feature will see employment fall but rather the effects are a continuous distribution across economies.

A forecast error variance decomposition (FEVD) is one method of assessing the importance of technology shocks in driving the variation of employment relative to other shocks in the typical economy. For both advanced economies and EMDEs, technology shocks drive around 30-40% of the forecast error variance of employment at the 10-year horizon (Table 1).

Table 1: Forecast error variance contribution of technology shocks to employment

	Year 1	**Year 5**	**Year 10**
Advanced economy	13%	24%	31%
EMDE	19%	31%	42%

Note: Average of individual EMDE and advanced economy country-by-country estimated forecast error variance decompositions

3.3 What country-specific factors are associated with falling employment following a technology shock

The degree of labor market disruption in each economy can depend on multiple factors, including the types of technologies introduced over the sample period, and the degree to which they complement skilled or unskilled labor; the policies enacted by governments in each economy to engage workers in new roles following technology-driven labor market disruption; and measurement or data issues which may influence the results. To examine the underlying driver and correlates of employment disruption following technology shocks, the median scaled employment IRF for each country is treated as the dependent variable in cross-section regressions containing potential structural factors associated with the employment effect.

Proxies for the degree to which economies experiencing large scale technological change are rare, particularly on a consistent basis across countries. Sectoral employment shares, both in levels and how much they have changed in recent decades are one proxy for economies going under significant technological transformation. In particular, the literature has widely found that in advanced economies, manufacturing jobs have particularly been at risk of replacement, making industrial employment, which includes manufacturing and is widely available across economies, a useful covariate to test. Data on sectoral employment shares are taken from the World Bank's World Development Indicators, as are data on foreign direct investment net inflows, which may also be a proxy for the implementation of new technologies. Educational attainment is measured

in average years of schooling, taken from the Barro and Lee (2015) dataset.

Variables are taken as averages during the period 1990-2018, primarily due to widespread sectoral data availability only starting in the 1990s. This is a shorter horizon than over which the VARs are estimated. However, results are robust to using longer averages including the 1980s for non-sectoral data and shorter averages since 2000. The overall explanatory power of the covariates is modest but material for the employment impact in year 1, with an adjusted-R^2 of 0.18 in the best-fitting specification, particularly given the wide range of factors that will influence this outcome across economies. There are many statistically significant determinants of the employment effect found despite the simple nature of the proxies for structural and technological change (Table 2).

Rising industrial employment. An increasing share of employment in the industrial sector since 1990 is strongly correlated with the negative employment effects of technology shocks at all horizons.[6] A 10 percentage point rise in the industrial employment share over this period is associated with a 0.1-0.2 percentage point larger negative impact on employment on average from a technology shock. The effect is statistically significant at the 1% or 5% level in each of the horizons. The average *level* of industrial employment during this period is also separately a statistically significant determinant of the employment impact at the 5 and 10-year horizons but it loses significance when included in the same estimation as the *change* in employment share. As will be discussed in the following section, all advanced economies in the sample lost employment share in the manufacturing sector since 1990. However, a smaller loss of employment share in the industrial sector, relative to a large loss, remains a significant determinant of the size of the negative employment impact. Changes in the share of employment in agriculture and services are not statistically-significant determinants of the employment impact.

An interpretation of this finding is that economies that have made productivity gains that are more centered around the industrial sector have experienced less job intensive (or even job-reducing) productivity growth. Economies with increasing employment shares in the industrial sector since the 1990s will have been at the highest risk from automation. In addition, those economies with increasing employment shares in this industry may have had the least success in increasing employment in other sectors following job losses in the industrial sector due to automation. This finding links directly to much of the existing literature which has found that routine manufacturing jobs (although to a lesser extent routine service sector jobs) have been at the highest risk of automation.

[6] While industrial production can include types of production other than manufacturing, such as mining and extraction industries, there is no relationship between the proportion of exports made up of commodities and the employment impact, suggesting that the manufacturing component of industry is the key driver of the employment dislocation following technology shocks.

Table 2: Determinants of employment IRF at multiple time horizons

	Dependent variable:					
	Employment IRF					
	After 1 year	After 1 year	After 5 years	After 5 years	After 10 years	After 10 years
	(1)	(2)	(3)	(4)	(5)	(6)
Productivity level *log*	-0.169^{***}	-0.156^{***}	-0.036	-0.039	0.043	0.031
	(0.055)	(0.058)	(0.048)	(0.051)	(0.045)	(0.047)
Years schooling	-0.003	0.001	-0.007	-0.001	-0.018	-0.014
	(0.019)	(0.019)	(0.016)	(0.016)	(0.015)	(0.015)
Δ industry employment share	-1.613^{**}		-1.920^{***}		-1.417^{**}	
1990s to 2010-18	(0.798)		(0.706)		(0.661)	
FDI	1.008^{**}		0.503		0.288	
% GDP	(0.446)		(0.394)		(0.369)	
Productivity growth	0.129	0.103	0.257^{**}	0.250^{**}	0.301^{***}	0.294^{***}
1990s to 2010-18	(0.113)	(0.116)	(0.099)	(0.102)	(0.093)	(0.094)
Agriculture employment share	-0.300	-0.543	-0.046	-0.639	0.154	-0.368
%	(0.328)	(0.440)	(0.289)	(0.389)	(0.271)	(0.357)
Commodity exports	0.003		0.003		0.001	
	(0.002)		(0.002)		(0.002)	
% total						
Industry employment share		-0.648		-1.321^{**}		-1.069^{*}
% total		(0.720)		(0.637)		(0.584)
Trade openness		0.127^{*}		0.011		-0.022
Exports & imports % GDP		(0.070)		(0.062)		(0.056)
Change in trade openness		0.058		0.082		0.133
1990s to 2010-18		(0.107)		(0.094)		(0.087)
Constant	1.305^{**}	1.363^{*}	0.127	0.628	-0.534	-0.029
	(0.562)	(0.689)	(0.495)	(0.609)	(0.465)	(0.558)
Observations	108	108	107	107	108	108
Adjusted R^2	0.182	0.139	0.091	0.035	0.081	0.073
Residual Std. Error	0.321	0.329	0.282	0.291	0.265	0.267

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Education/Schooling levels. At each horizon, the employment impact following a technology shock is unrelated to the average years of schooling of the workforce. The insignificant effect possibly reflects opposing forces. Education levels may affect both the likelihood of domestically-generated innovation and adoption of new technologies, or cushion the impact of adopting new technologies, which more highly skilled workers able to complement new forms of production. The distribution of education may be broad, such that a highly-skilled workforce can incorporate skill-biased technologies, but have little success in reallocating labor to other tasks due to a large low-skilled segment of the workforce.

Productivity levels and productivity growth The log-level of productivity, measured in 2010 US dollars (at 2010 exchange rates) has a negative relationship with the employment impact in year 1, consistent with the larger negative impact experienced by advanced economies. In addition, the cumulative growth of productivity is positively correlated with the effect on employment: faster productivity growth since 1990 is associated with a more positive impact on employment. It is less clear how to interpret this finding. More frequent and larger positive technology shocks will result in greater cumulative productivity growth; it may be that a larger variety of introduced technologies is associated with smaller negative employment impacts, or it is plausible that the types of economies experiencing large cumulative productivity growth over this period have been better at ensuring displaced workers are successfully reintegrated into the workforce.

Foreign direct investment and trade The extent of net foreign direct investment and trade openness is positively related to the degree of employment variation following a technology shock. As FDI is one of the key vehicles for transferring technology to EMDEs, this is reassuring. A range of literature, including panel estimations and country-specific case studies, has found FDI to be associated with increased employment, particularly where it is export-focused (Hale and Xu, 2016; Waldkirch et al., 2009).

3.4 Robustness and alternative explanations for technology-driven fall in employment

This paper has attributed the fall in employment following an SVAR-identified technology shock to the introduction of labor-displacing technology. This section briefly reviews other plausible candidate drivers for this result, finding each to be an unlikely source of the result.

Short sample biases The high persistence of EMDE employment effects following a shock may fuel concerns that some of this effect may be driven by short-sample issues and limited degrees of freedom in estimation. However, short-sample issues are found to be associated with *lower* IRF persistence than the true IRF (Jarocinski and Marcet, 2010). Therefore, this effect should be in the opposite direction to the result found for emerging markets. Including the sample length for each economy under consideration as a regressor in Table 2 does not yield a statistically significant

coefficient at horizons, 1, 5, or 10 years.

Unrelated low-frequency trends in productivity and employment The second source of bias that could be present in advanced economies and EMDEs is the effects of unrelated long-run trends in productivity growth and employment Fernald (2007). In the US context, long-run trends in these variables are found to change the sign of the employment effect of technology shocks. Tests for the presence of structural breaks in unspecified time periods in the employment and productivity growth series suggests the presence of multiple breaks for all economies in the sample (Bai and Perron, 1998). The multitude of structural breaks could suggest the presence of low-frequency unrelated trends across both advanced economies and EMDEs. To ensure that these potentially unrelated trends are not driving the negative employment effect, an HP filter is applied to the log employment series, removing low-frequency trends a leaving only higher frequency movements in the employment data. The panel-estimation result continues to be a negative response of employment in both advanced economies and EMDEs (Figure 4). By construction, the removal of low-frequency variation in employment will reduce the size and persistence of the IRFs in both groups, with both IRFs now no longer significantly different from zero in the second year. The presence of low-frequency correlations between technology and employment cannot be ruled out as unrelated. However, the presence of a range of structural characteristics associated with the degree of persistence of the IRFs suggests that the relationship is unlikely to be spurious.

New Keynesian explanations of negative employment response to technology shocks The initial technology shock SVAR literature focused on determining whether the economy reacted to technology shocks consistent with an RBC or New Keynesian description of the economy. In the New Keynesian explanation, all factor inputs to production should contract initially following a technology shock, as prices are inflexible in the short run, restricting increases in production. And the same quantity of production could be achieved with fewer inputs. This framework predicts that both investment and employment or hours should fall in the short term. As documented above, it is clear that employment falls in the short-term for around three-quarters of advanced economies and EMDEs. However, the same exercise finds that investment falls in just one-third of advanced economies and less than 10% of EMDEs in the sample. Just 20% of the negative impacts are statistically significant in advanced economies, falling to 2% in EMDEs.

Figure 4: Employment impacts when estimation uses HP-filtered log employment in VAR estimation

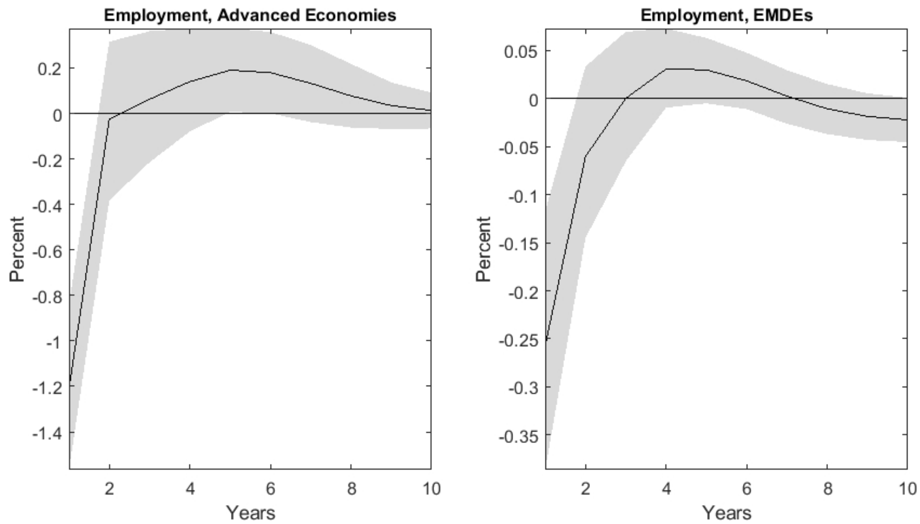

News shocks A growing literature has explored the macroeconomic impacts of "news" about new technologies (Barsky and Sims, 2011; Beaudry and Portier, 2006; Miranda-Agrippino et al., 2019; Alexopoulos, 2011). Frequently, this literature uses the Max-Share identification approach combined with an orthogonalized "surprise" technology shock which is not anticipated. The identification here does not separate technology shocks into those which are anticipated and those which are "surprise" shocks. For this identification to be enacted, the VAR is usually applied to utilization-adjusted measures of TFP. With standard TFP measures, "surprise" shocks would largely capture non-technology features of TFP. At this time, there are no publicly available utilization-adjusted TFP series outside of the United States. Regardless, it is not clear whether the absence of a separately identified news shock would bias the results of this exercise. A news shock is generally found to be associated with long-run improvements in productivity, as is the technology shock identified here. The primary difference is the imposed lack of contemporaneous impact. The news shock literature has also generally found similar evidence of an initial decline in hours-worked, such that conflating the news and surprise shocks is unlikely to bias the identified technology shock in a particular direction (Barsky and Sims, 2011; Kurmann and Sims, 2017; Miranda-Agrippino et al., 2019).

Table 3: Determinants of employment effect: employment protection

	Dependent variable:					
	Employment IRF					
	After 1 year	After 1 year	After 5 years	After 5 years	After 10 years	After 10 years
	(1)	(2)	(3)	(4)	(5)	(6)
Productivity level *log*	-0.229^{***}	-0.115^{***}	-0.076	-0.038	0.023	-0.003
	(0.070)	(0.038)	(0.059)	(0.032)	(0.049)	(0.030)
Years Schooling	0.039	0.019	0.017	0.021	-0.011	0.008
	(0.031)	(0.018)	(0.026)	(0.015)	(0.022)	(0.014)
WEF: Hiring and firing	-0.013	0.021	-0.035	-0.026	-0.008	-0.009
	(0.073)	(0.049)	(0.062)	(0.041)	(0.052)	(0.038)
OECD: Collective dismissal	-0.007		-0.045		-0.006	
	(0.109)		(0.092)		(0.077)	
OECD: Individual dismissal	-0.015		-0.009		-0.022	
	(0.044)		(0.037)		(0.031)	
Constant	1.800^{**}	0.628^{*}	0.796	0.218	-0.091	-0.056
	(0.741)	(0.366)	(0.626)	(0.307)	(0.523)	(0.286)
Observations	56	103	55	102	56	103
Adjusted R^2	0.188	0.105	-0.045	-0.006	-0.073	-0.022
Residual Std. Error	0.357	0.333	0.301	0.279	0.252	0.260
Note:						*p<0.1; **p<0.05; ***p<0.01

Note: Regression with the scaled employment response to technology shocks as the dependent variable, and a range of variables reflecting labor market frictions as the regressors. WEF variable refers to the World Economic Forum's survey of firm's ability to easily hire and fire workers, ranked from 1 (difficult) to 7 (easy). OECD indicators are the OECD's ranking of the strictness of employment legislation to protect against individual and collective dismissal, with a higher index value indicating increased strictness

Labor market frictions and ease of hiring and firing Cross-country variations in the degree of labor market flexibility and employment protections provide an alternative explanation for the variation in the persistence of employment changes following a technology shock (Rujin, 2019). The OECD's Indicators of Employment Protection cover 73 economies, 55 of which are included in the dataset of SVAR estimations. This database includes relative measures across economies of the strictness of protections against individual and collective dismissal. A separate database produced by the World Economic Forum produces a measure of labor market flexibility on a comparative scale across all economies where SVAR estimates of the employment impact are available. While labor market structures may play an important role in governing the persistence of the employment response in theory, these variables are found to be uncorrelated with the employment impact at all horizons (Table 3).

4 Industrial employment and employment losses following technology shocks

Several characteristics have been identified as being associated with the size and persistence of the impact of technology shocks on employment. Only the change in the share of industrial employment since 1990 is found to be a key determinant of the employment effect at all horizons. This section examines the SVAR and reduced-form evidence on the importance of the industrial employment share and its association with technology-driven employment losses. It then outlines a theoretical model to explain the empirical findings.

4.1 Pooled estimation of high and low industrial employment share economies

This section examines the differences between economies that have experienced higher changes in industrial employment shares in recent decades, and how their employment impact has differed from economies experiencing falling industrial employment using separate panel VAR estimations.[7]

All advanced economies in our sample have experienced a fall in the industrial share of employment since the 1990s (Figure 5). Those with a "high" employment share change can be regarded as having experienced the smallest fall in industrial employment shares in the advanced economy cohort. Those with a "high" employment share change include the United States, Australia, Finland, France, and Italy. Those that have experienced a "low" change in the industrial share of employment include Hong Kong, SAR, Singapore, Spain, Ireland, and Slovenia. These economies have more rapidly de-industrialized. In the EMDE sample, more than half of the 96 economies have increased in their share of industrial employment between 1990-99 and 2010-18. These include economies such as Vietnam, China, Thailand, Indonesia, and Oman. Those EMDEs with "low" or falling manufacturing shares include Bulgaria, the Russian Federation, South Africa, and Zimbabwe.

It is clear that the services sector has been the sector where the most new jobs have been created in both advanced economies and EMDEs. However, not all of this change has been generated by technological-displacement. There have been large flows in EMDEs from agriculture to services, and also a component of job flows due to changing income elasticities of demand for these products (Comin et al., 2015; Rodrik, 2016). This additional factor may be the reason that changes and the level of employment in agriculture and services are not found to be related to employment disruption changes following a technology shock; structural change has dominated employment flows in these sectors. This structural change has occurred regardless of whether increasing income has been driven by rising participation rates, employment-augmenting technologies, or employment displacing technologies.

[7] This approach has been used in different contexts to identify characteristic dependent fiscal multipliers (Ilzetzki et al., 2013).

Figure 5: Change in sectoral employment shares, 1990-99 and 2010-18

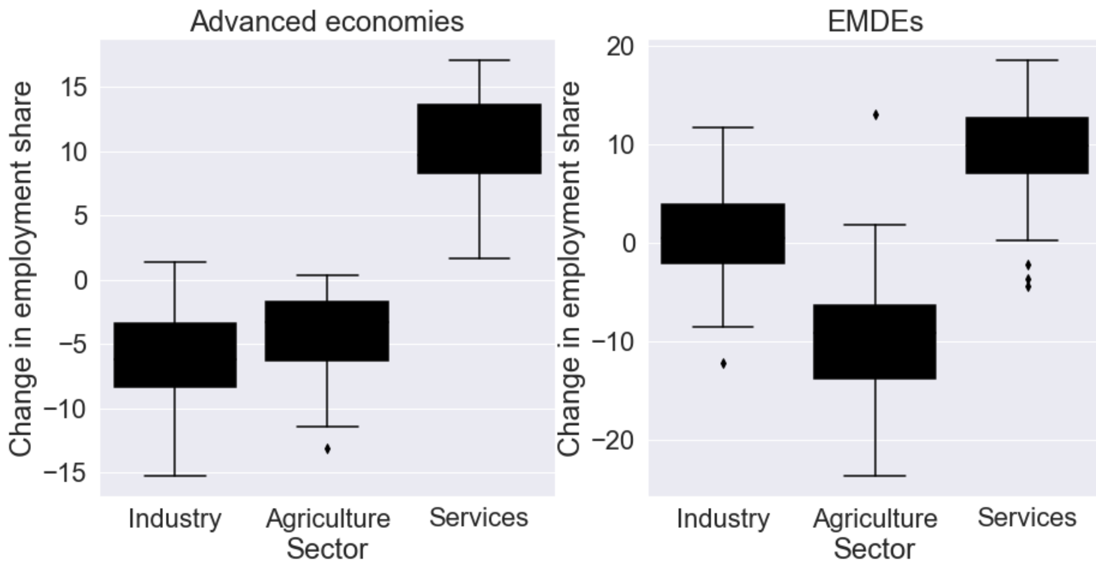

Note: Change in the employment share between 1990-99 and 2010-18. Shaded region shows the interquartile range, while the 'whiskers' shows the range excluding outliers (diamonds).

Panel-VAR estimations are performed on groups of economies that have experienced large falls in industrial employment, relative to those that have experienced increases, or smaller falls in industrial employment. First, economies are ordered by their change in employment share in the industrial sector since the 1990s. In separate groups, advanced economies and EMDEs in the highest quartile are listed as "high" change in employment share, and those in the lowest listed as "low" change. In both advanced economies and EMDEs, there is a marked difference between the response of employment to a technology shock.

The results show that the negative employment impact in advanced economies is 4 times as large for those economies where industrial employment has fallen the least (Figure 6). The median impulse response does not return to zero until year 6 in these economies, while it becomes neutral in year 2 in economies that have seen a bigger decline in the share of industrial employment. In EMDEs, the two groups of economies have only a 100% differential in the initial employment impact, but the persistence for high-manufacturing economies is much greater. Even at the 10-year mark, the negative impact remains close to -0.1 percent. For those EMDEs and advanced economies with small increases on decreases in industrial employment shares, IRFs become positive, with impacts +0.1-0.2 percent at the 10-year horizon.

Figure 6: Technology shock employment IRFs: High and low change in industrial employment share since the 1990s

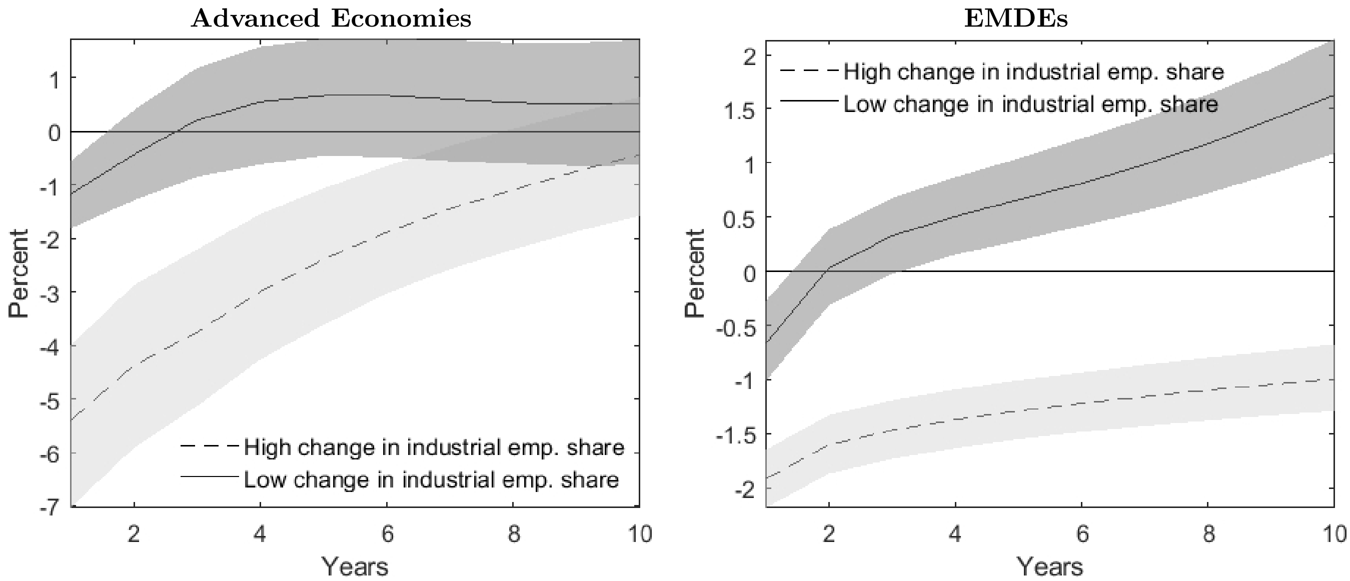

Note: For advanced economies and EMDEs separately, economies are ordered by the change in industrial employment share from 1990-99 to 2010-18. The IRF for those economies in the upper quartile is given as "High change in employment share", while those in the bottom quartile are given as "Low change in employment share". IRFs are scaled to show the effects of a technology shock which increases labor productivity by 10 percent.

4.2 Reduced-form evidence of the relationship between aggregate employment growth and industrial employment shares

A correlation exists between the growth of employment since the 1990s as a proportion of the working-age population and the change in the workforce employed in industry, a key covariate of the estimated employment IRFs. This relationship is weak when including oil and metals exporting economies, where industrial employment developments are strongly influenced by commodity price changes. However, the relationship strengthens when these economies are excluded, both for employment and the size of the labor force relative to the population (Figure 7).

Figure 7: Employment and labor force growth and change in the industrial employment share

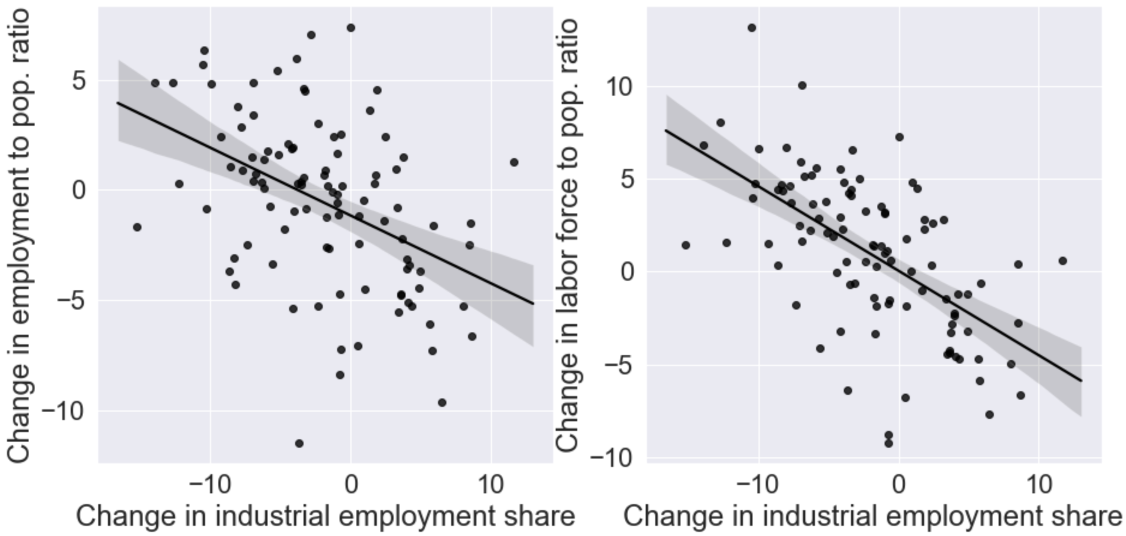

Note: Change in the industrial employment share between 1990-99 and 2010-18 compared to the change in employment as a share of the working-age population over the same time period, and the change in the labor force as a share of the working-age population. Economies where over 20% of exports consist of oil or metals are excluded.

Those economies that have been less successful in increasing the aggregate workforce over this time period have had a more robust share of industrial employment. In the next section, one plausible reason provided for this finding is that in economies where automation has triggered large positive income effects, aggregate employment will increase but demand will be reallocated to other sectors such as services, further reducing the industrial employment share. In economies with smaller income effects from labor-substituting technologies, or frictions in reallocating workers to new sectors, aggregate employment will grow by less and industrial employment will remain a larger share of total employment. The reduced-form relationship is stronger for the labor force as a whole (employment plus the unemployed) than for employment. Cortes et al. (2020) find evidence of smaller flows of job-seekers into the labor force, as well as flows out of the labor force in response to automation.

The relationship is robust even after accounting for factors relating to workforce aging, the gender composition of the workforce, and the level and growth rate of GDP per capita. The latter variable controls for whether the fall in industrial employment is simply a by-product of structural change in fast-growing economies, and not necessarily associated with technological displacement.

The change in the industrial employment share remains a statistically significant determinant of aggregate employment growth controlling for all of these additional factors (Table 4).

Table 4: Covariates of the change in employment and labor force to population ratio, 1990-99 to 2010-18

	Dependent variable:	
	Employment	Labor force
	(1)	(2)
Change in industry employment share	-0.148^{*}	-0.227^{***}
Change in share over 50 in working-age population	-0.03	-0.159
Change in female employment share	0.458^{***}	0.517^{***}
Change in log GDP per capita	0.835	0.311
Log GDP per capita	0.271	0.944^{***}
Constant	-3.70^{***}	-8.489^{**}
Observations	100	100
Adjusted R^2	0.30	0.51
Note:	*p<0.1; **p<0.05; ***p<0.01	

Note: Changes are calculated as the difference between the period averages for each economy during 1990-99 and 2010-18

4.3 Theoretical framework: industrial employment share and employment growth

This section outlines a static theoretical model that explains the empirical findings concerning technology shocks, employment growth, and the industrial employment share. It uses a simplified version of the framework of Acemoglu and Restrepo (2018, 2020), and closely follows the notation of Acemoglu and Restrepo (2017), which assesses the overall impact of automation on aggregate job creation. In their framework, the job-displacing effect from automation can in part be offset by two factors: higher relative demand for the products in the affected industry as goods are produced

more cheaply due to automation; and secondly, higher overall income driving higher demand for all goods. Automation can increase aggregate employment if the latter effects are large enough.

This section shows that aggregate job gains can result from labor-substituting technologies, but likely alongside an even larger fall in employment shares in the sector most affected historically, industry. The finding applies where low price elasticities of substitution and low income elasticities of demand exist for those sectors subject to automation. Where there is little or no offsetting income or substitution effect from automation affecting the industrial sector, aggregate employment will rise by less or fall by more, but employment shares in industry will remain more robust. This framework purely considers demand-side mechanisms, while there are a range of supply-side factors that could contribute to this finding, complementing this approach.

This framework augments that of Acemoglu and Restrepo (2018, 2020) by also considering the effects of different income-elasticities of demand between sectors; it is well documented that the demand for industrial goods falls relative to that of services as income rises (Rodrik, 2016; Comin et al., 2015). This helps to explain why the industrial sector shrinks more in economies with large positive income and employment effects from automation, but less so in economies that have not experienced as large a positive effect.

In a simplification, a model is created of an industrial sector that is subject to labor-substituting technological change, and a services sector that is not. While it is true that certain jobs in the services sector have been automated, analysis of U.S. and European job markets has found that industrial middle and low-skilled jobs have been most affected in recent decades by automation (Goos et al., 2014; Autor, 2015).

Firms are split into sectors k that produce industrial goods (I), and those that produce non-automatable services (NAS). Industrial goods and services producers produce output $X(k)$ by combining tasks x, such that production is performed as

$$X(k) = A \min_{s \in (0,S)} (x_k(s))$$

Tasks, S, are divided into those that can only be performed by workers (w), and those that can be performed by labor or capital (r). Specifically, in the industrial sector tasks $S \in [0, M]$ are technologically automated and can be performed by capital. The output of capital in every task is set to 1, and to further simplify the model it is assumed that the productivity of labor in each task is constant as well and equal to γ. Consequently, the production function for task s in industry k can be written as:

$$x_I(s) = \begin{cases} r(s) + \gamma w(s) & \text{if } s < M \\ \gamma w(s) & s > M \end{cases}$$

There is therefore perfect substitution of labor for capital in the automatable tasks $s < M$. It

is assumed that capital is always cheaper than employing labor, such that automation increases the cost-effectiveness of production. The savings from switching a unit of production from labor to capital is reflected as: $\pi = (1 - \frac{Q\gamma}{W})$, where Q is the price of capital, and W reflects the wage rate. Therefore, the demand for labor from industry is:

$$L^d(I) = \frac{(1-M)}{\gamma A} X(I)$$

In the services sector, no tasks are automated (M=0). All tasks require labor for production such that the process can be described as a single type of task.

$$x_{NAS}(S) = \gamma w(s) \text{ for } 0 \leq s \leq 1$$

In the services sector, total labor demand will be equal to:

$$L^d(NAS) = \frac{1}{\gamma A} X(NAS)$$

Given that prices are equivalent to marginal costs, it is assumed that:

$$P(NAS) = \frac{W}{\gamma A}, \quad P(I) = \frac{1}{A}(MQ + (1-M)\frac{W}{\gamma})$$

It is clear that an increase in M, the proportion of tasks that can be produced by capital, will lower the price of goods in industry, and reduce the quantity of labor required to produce a given level of output. An increase in M is therefore an example of labor-productivity enhancing technological change considered in the empirical SVAR exercise.

Households have the following consumption and labor preferences:

$$\mathcal{U}(Y, L) = logY - \frac{1}{1+\phi} L^{1+\phi} \tag{4}$$

Where Y, income, is equivalent to total consumption, and L is the total hours worked. The consumption bundle is defined over the variety of goods produced by the industrial and services sectors.

A constraint allows for a non-homothetic aggregate with different income elasticities of demand for the goods and services produced, $Y(I)$ and $Y(NAS)$, as in Comin et al. (2015):

$$1 = \sum_{k \in [I, NAS]} (\alpha_k Y)^{\epsilon_k} Y(k)^{\frac{\sigma-1}{\sigma}}$$

where σ measures the degree of substitutability between industrial and non-automatable services and $\alpha_I + \alpha_S = 1$. The parameter ϵ_k governs the degree of income elasticity for each sector, such

that consumption of a sector's produce rises by more with respect to income for higher ϵ_k.

Maximisation is subject to the following sequence of period budget constraints:

$$\sum_{k\in[I,NAS]} P(k)Y(k) = WN + QR$$

where D is the nominal payoff of a portfolio held from the previous period, Q is the stochastic discount factor for one-period ahead nominal payoffs for the domestic household, W is the nominal wage for each sector.

Optimality results in the following demand functions for industrial and services sector products:

$$Y(k) = \alpha_k \left(\frac{P(k)}{E}\right)^{-\sigma} Y^{\epsilon_k} \tag{5}$$

Where E is total expenditure, defined as:

$$E = \left[\sum_{k\in(I,NAS)} \alpha_k Y^{\epsilon_k} P(k)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}$$

Defining $P = \frac{E}{Y}$ and setting this aggregate price index to be a numeraire allows equation 5 to be rewritten:

$$Y(k) = \alpha_k \left(P(k)\right)^{-\sigma} Y^{\epsilon_k+\sigma}$$

The remaining optimality conditions concern labor supply and intertemporal allocations. Labour supply satisfies:

$$YL^{\phi} = \frac{W_{NAS}}{P_t},$$

Finally, automating capital, R, is supplied at price Q, and is assumed to have an upward sloping supply curve:

$$Q = \left(\frac{R}{Y}\right)^{\eta}$$

Income and substitution effects of automation. In response to a change in the proportion of production tasks that can be automated, $\left(\frac{dM}{1-M}\right)$, total labor demand in this economy depends on several effects (Annex A.2).

$$dlnL^d = \left(\frac{1+\eta}{1+\epsilon}\right)\left(\underbrace{l(I)\frac{dM}{1-M}}_{\text{Direct displacement}} + \underbrace{s_{L,I}\pi l(I)\sigma\frac{dM}{1-M} + (1-s_L)\pi l(I)\frac{dM}{1-M}}_{\text{Substitution and income effects}}\right)$$

The first is direct displacement, where a share of the workers in industry $(l(I))$ are directly replaced by capital. The additional effects are a substitution and income effect, where the cost savings from automation (π) reduces the price of goods in the industrial sector, increasing demand for them. This effect restores some workers in industry but crucially depends on the price elasticity of demand for these goods (σ), and the proportion of costs that were accounted for by labor in industry $(s_{L,I})$. Finally, there is an income effect, where the gains from cheaper production are spent on goods in industry and the non-automatable services sector. This also crucially depends on the savings from the automation of industrial jobs (π). Those economies where automation does not produce large cost savings, perhaps because the supply of automating capital is inelastic or expensive to install, or where the price elasticity of demand is low (σ), will not see aggregate employment increase in response to automation.

Labor substituting technological change and employment shares. In this economy, because automation is specified to only occur in the industrial sector, the share of labor employment in industry falls in response to labor-substituting technological change for plausible parameter values. As a counter-example, in order for the share of industrial employment to remain constant, the price effect must offset both the direct displacement effect and the fact that industry has a lower income-elasticity than services. The change in demand for labor in industry with response to increased automation (M) can be written:

$$\begin{aligned} dlnL_I^d &= dln\left(\frac{(1-M)}{\gamma A}\alpha_k\left(P(I)\right)^{-\sigma}Y^{\epsilon_I+\sigma}\right) \\ &= -\frac{dM}{1-M} + \sigma s_{L,I}\pi\frac{dM}{1-M} + (\epsilon_I+\sigma)dlnY \end{aligned} \tag{6}$$

For non-automatable services, the effect on labor demand is purely due to the income effect.

$$\begin{aligned} dlnL_{NAS}^d &= dln\left(\frac{1}{\gamma A}\alpha_k\left(P(NAS)\right)^{-\sigma}Y^{\epsilon_{NAS}+\sigma}\right) \\ &= (\epsilon_{NAS}+\sigma)dlnY \end{aligned} \tag{7}$$

In order for the *share* of employment in industry not to fall, the price effect must offset both the displacement effect and the difference in income elasticities between the two sectors.

$$dlnL_I^d - dlnL_{NAS}^d = \underbrace{-\frac{dM}{1-M}}_{\text{Direct displacement}} + \underbrace{\sigma s_{L,I}\pi\frac{dM}{1-M}}_{\text{Price effect offset}} + \underbrace{(\epsilon_I - \epsilon_{NAS})l(I)s_L\pi\frac{dM}{1-M}}_{\text{Differential income elasticities}} \quad (8)$$

First, in the simplified case where the income elasticities are equivalent, this requires:

$$\frac{1}{\sigma} \leq \pi s_{L,I} \quad (9)$$

Even in cases where the income elasticity of demand in both sectors was equal, where the labor share was equivalent to 0.5 in industry (s_L), and the cost savings from switching to capital is 20 % (π), the elasticity of substitution parameter, σ, would have to reach 10 to ensure that the share of labor in industry did not fall. Consumers would have to be highly price elastic in response to changes in the prices of manufactured goods and services, which is at odds with the data. σ is estimated to be well below unity in almost all economies (Comin et al., 2015). Therefore, even in cases where automation had positive aggregate employment effects, we should expect the most affected industries to lose employment shares. This is consistent with the SVAR results, where industrial employment falls in economies experiencing higher employment impacts from technology shocks.

Why are economies with lower job creation following the introduction of new technologies associated with higher industrial shares? The above demonstrates a mechanism through which economies will lose employment shares in industry if that sector experiences labor-displacing technology improvements, even where aggregate employment increases. However, it does not explain why economies with smaller employment losses from new technologies will experience relatively larger falls in their industrial employment share. This section outlines the conditions through which a larger positive income effect from a labor-substituting technology shock results in a larger reallocation of demand away from the industrial sector while boosting aggregate employment.

Differential income elasticities play a role in boosting consumption and employment by more in the service sector than industry following an improvement in income, further reducing the recovery of employment in the sector most affected by automation, industry. Where the income-elasticity differential between the two sectors outweighs the effect of lower prices in industry from the productivity improvement (low σ), aggregate employment will increase due to automation while shrinking the industrial employment share.

The income elasticity of services is found to be higher than manufacturing and industry across OECD and emerging market economies (Comin et al., 2015). Aggregate and household estimations find that the difference between the manufacturing and services income elasticity has ranged between -0.18 and -0.57, broadly equivalent to the parameters $\epsilon_I - \epsilon_{NAS}$. σ is found to be inelastic, ranging

from $0.2 - 0.63$.

The larger the rise in incomes due to the efficiency savings from the technology shock (driven by the savings, π), the larger the fall in the employment share of industry in cases where:

$$\sigma s_{L,I} \leq s_L l(I)(\epsilon_I - \epsilon_{NAS})$$

Therefore, for a sufficiently low price elasticity of substitution, or a sufficiently high income elasticity differential between the sectors, the more industrial employment shares will fall for a given saving from automation, π. The function also depends on the difference in the labor share between the industrial sector ($s_{L,I}$) and the aggregate economy (s_L).

In economies where aggregate employment has fallen by more following a technology shock, industrial employment shares have remained robust since 1990. This plausibly reflects a lower efficiency improvement from automation in these economies, and thus a smaller reallocation of demand to the services sector. It could also point to a failure to accommodate rising services demand with domestic employment and production.

Additional supply-side channels. This simple model only considers demand-side drivers of this result. Several supply-side factors could have contributed to weaker industrial employment shares in those economies with larger aggregate employment growth due to new technologies and smaller aggregate employment growth in those with more robust industrial employment shares. Several studies have documented that rising market power has increased markups and reduced the labor share in those sectors most affected by technological change, including within manufacturing (Autor et al., 2020; De Loecker et al., 2020). In this case, savings produced by automation would not flow through to lower prices, producing a smaller increase in demand for goods in that sector.

Secondly, this simple framework does not account for skills mismatch. Where labor-substituting technology shocks result in income gains that increase demand for products and services requiring high- or niche-skills, supply constraints may prevent an increase in employment in these other sectors, and aggregate employment may not rise (Restrepo, 2015). This effect would be compounded by the demand channels described above. A low income elasticity for industrial goods would reduce demand for re-employment in roles requiring similar skill-sets. Economies that have the skill-base to meet the demand for workers in other sectors would see rising aggregate employment and a more rapid decline in industrial employment.

5 Summary, future risks, and policy options

This paper has found that labor-substituting technological change has been widespread across advanced economies and EMDEs in recent decades. That is to say, while not all productivity-enhancing technologies are labor displacing, they have accounted for a sizeable proportion of

productivity-enhancing technologies, such that employment falls in the majority of economies following an "average" technology shock. Other explanations for falling employment, such as a New Keynesian sticky-price mechanism or labor market regulations and frictions, are ruled out as drivers of this result.

Technology-driven employment losses are larger in economies with higher productivity levels. This plausibly reflects a higher propensity to invest in labor-substituting technologies in economies with more sophisticated production capabilities and those with a labor force more able to design and implement these production methods. However, job losses from new technologies have often been more persistent in lower-income EMDEs. Trade openness and FDI inflows appear to be associated with reduced employment losses from new technologies, consistent with a large literature that has found that export-focused FDI is often designed to take advantage of lower labor costs in EMDEs and is associated with higher employment levels.

Higher industrial employment share growth since the 1990s is a key correlate of larger and more persistent falls in aggregate employment from productivity-enhancing technologies. A model adapted from Acemoglu and Restrepo (2018) shows that where the income effects from automation are large, low-income elasticities of demand could reduce relative demand for employment in industry even where aggregate employment improves. Therefore, those economies with a successful reallocation of labor and large income effects from job-displacing technologies have seen a larger fall in industrial employment. The reallocation of demand to sectors other than industry, where job-displacing technologies have been more prevalent, could exacerbate difficulties in matching workers with new roles suitable for their skill-sets.

This analysis has examined historical relationships. However, several studies find a pronounced risk from new technologies to employment in the future. In advanced economies, a wide range of estimates have been provided for the proportion of jobs which are at risk of future automation: Arntz et al. (2016) find that 9 percent of jobs across 21 OECD economies are at high risk of automation in the future. A broader study of 32 economies, including several EMDEs, has found that on average 14 percent of jobs are at high risk of automation, with a further 32 percent at risk of significant change due to automation, primarily in the manufacturing sector, consistent with the finding that industrial concentration has been associated with lower job growth in this paper (Nedelkoska and Quintini, 2018; Organisation for Economic Co-operation and Development, 2019). Increasingly, many service sector roles will be at risk according to these studies. These include roles such as food preparation and some sales roles. This literature has not considered the effect of new roles that could be created by the introduction of technologies, or the macroeconomic implications stemming from higher demand, so are gross rather than net impacts on total jobs. So far, no studies have estimated the impact of anticipated technological change on a large sample of EMDE labor markets. As EMDEs acquire an increasing share of global industrial and manufacturing employment, it is likely that they will increasingly face challenges from automation.

Reforms aimed at increasing workforce skills to complement new technologies and service sector roles appear to be key in reducing lost employment due to automation but could require significant lengths of time to result in improved workforce capabilities. In the nearer term, policies aimed at restoring demand in industrial sectors most affected by job displacing technologies may be appropriate to sustain employment levels.

Policies that help enhance workforce skills and educational attainment could potentially reduce the fall in employment following the incorporation of production processes which remove the need for lower- or middle-skilled employment. Many EMDEs will need to make improvements at earlier levels of education in order to build a foundation for more advanced levels of education which will adequately complement new technologies (World Bank, 2018, 2019). EMDE universities and on-the-job training are also often under-provided in many EMDEs, but also display high returns in terms of wage premia where they exist, in addition to enabling better adaptation to changing production technologies.

The persistence of declining employment in EMDEs and large scale of displacement in advanced economies highlights the need for adequate social protection to ensure that those who are displaced from their employment can increase their opportunities to transition into new industries. Encouraging both private savings and social insurance for unemployment is needed in the formal and informal sectors will act as a safety net for displaced workers.

While supply-side policies to improve skill matching may take decades, policymakers may respond to job displacement by directly increasing demand for industrial goods or incentivizing private sector demand or investment in the sector. Substantial investment requirements have been identified in both advanced economies and EMDEs to meet climate and poverty reduction targets, particularly in infrastructure investment which would require a significant increase in industrial production.[8]

[8] For example, the European Commission has estimated that investment amounting to 2% of European Union GDP would be required each year to meet current climate and energy targets in 2030 (European Commission, 2020). In addition, Rozenberg and Fay (2019) find that EMDEs may need to invest between 2 and 8 percent of GDP each year, primarily in infrastructure for energy, sanitation, agriculture, and transportation to meet the United Nation's Sustainable Development Goals by 2030.

A Appendix: Results using long-run and Max-Share technology identifications, and targeting TFP instead of labor productivity

This appendix provides robustness checks on the estimations of the employment effect of technology in the main text. In the main text, a spectral identification is implemented, which is found to be more robust to the short-sample and volatile data used across countries. Here, the Max-Share and long-run identifications of technology shocks (Francis et al., 2014; Galí, 1999) are also implemented, finding similar results. These identifications, and those in the main text, are also applied to VARs estimated using the hierarchical-prior methodology of Jarociński (2010) to aid robustness given the short sample nature of the data for many economies.

A.0.1 Max-Share SVAR identification methodology

The Max-Share identification is a similar approach to the Spectral identification which does not utilize the frequency domain. It is a more established methodology of which variants have also been used to identify technology "news" shocks (Francis et al., 2014; Barsky and Sims, 2011). It is therefore used as a robustness check on the results of the Spectral SVAR. Instead, it assumes that technology shocks are the *predominant* driver of productivity around the 10-year horizon of the forecast error variance. In this identification, the technology shock is that which drives the largest proportion of the forecast error variance of labor productivity at this horizon, as in (Francis et al., 2014).

10 years is longer than the period over which the business cycle occurs (typically assumed to be 2-8 years) but short enough to reduce challenges related to estimation on a finite sample. Francis et al. (2014) imposes this restriction in a VAR containing productivity, hours, consumption and investment as a share of GDP. The forecast error at horizon k can be written:

$$y_{t+k} - \hat{y}_{t+k} = \sum_{\tau=0}^{k-1} B^{\tau} u_{t+k-\tau} \tag{10}$$

By defining an orthonormal matrix A_0 with columns α, and e as a selection vector (size 1×n), we find the shock j which maximizes the contribution to the total forecast error variance of variable i at horizon k

$$\max \omega\left(\alpha\right) = \frac{e_i' \left(\sum_{\tau=0}^{k-1} B^{\tau} \alpha \alpha\prime B^{\tau\prime}\right) e_i}{e_i' \left(\sum_{\tau=0}^{k-1} B^{\tau} \Sigma_u B^{\tau\prime}\right) e_i} \; s.t. \;\; \alpha' \alpha = 1 \tag{11}$$

The technology shock at this maximized value is then: $\epsilon_t^{tech} = \alpha' chol\left(\Sigma_u\right)^{-1} u_t$. Following Uhlig (2003), identifying the structural shock that maximizes the contribution to the forecast error variance of productivity is solved by identifying the eigenvector associated with the maximum eigenvalue of V_τ, where V_τ is the FEVD of the target variable based on reduced-form shocks and the denominator of $\omega(\alpha)$.

A.0.2 Long-run restrictions

Isolating the long-run components of labor productivity ($prod_t$) and employment ($hours_t$), labeled LP_{LR} and $Employ_{LR}$, respectively, this methodology imposes the restriction that only the technology shock can impact labor productivity in the long-run.

$$\begin{Bmatrix} LP_{LR} \\ Employ_{LR} \end{Bmatrix} = \begin{Bmatrix} * & 0 \\ * & * \end{Bmatrix} \begin{Bmatrix} \epsilon_{tech.} \\ \epsilon_{non-tech.} \end{Bmatrix} \tag{12}$$

Assuming the structural AR matrix polynomial,

$$A\left(L\right) = I_2 - A_1 L - A_2 L^2 \ldots - A_p L^p \tag{13}$$

The long-run counterpart is therefore,

$$A\left(1\right) = \; I_2 - A_1 - A_2 \ldots - A_p \tag{14}$$

In a stationary VAR containing the log-difference series of productivity and hours, the long-run effect of the technology shock on growth will dissipate. The long run impact of each shock on the level of the target variable can be written as:

$$\begin{bmatrix} LP_{LR} \\ Hours_{LR} \end{bmatrix} = A(1)^{-1} \begin{bmatrix} \epsilon_{tech.} \\ \epsilon_{non-tech.} \end{bmatrix} = B(1)^{-1} A_0^{-1} \begin{bmatrix} \epsilon_{tech.} \\ \epsilon_{non-tec.} \end{bmatrix} = \begin{bmatrix} \Theta_{11} & 0 \\ \Theta_{21} & \Theta_{22} \end{bmatrix} \begin{bmatrix} \epsilon_{tech.} \\ \epsilon_{non-tech.} \end{bmatrix} \tag{15}$$

where $B(L)$ is the reduced-form VAR polynomial. Restricting the loading of the non-technology shock onto productivity to be zero can be accomplished by ensuring the long-run impact matrix is lower triangular. This is accomplished by solving for A_0^{-1} as follows:

$$A_0^{-1} = B(1) chol[B\left(1\right)^{-1} \Sigma_u B(1)^{-1'}] \tag{16}$$

Where Σ_u is the reduced-form variance-covariance matrix.

A.0.3 Heterogenous Panel VAR identification of technology shock and employment impact

There are several limitations to a standalone identification procedure for individual economies: annual data use limits the degrees of freedom for estimation; and, measurement error, particularly in EMDEs, could limit estimation accuracy and inference. These limitations can be partially tackled using a panel structural VAR approach which allows for some heterogeneity across economies. Here, estimation priors for each economy are informed by the mean group estimator of Pesaran and Smith (1995), and implemented as in Jarociński (2010).

Each economy has a unique set of reduced-form coefficients β_c and variance-covariance matrix Σ_c. The VAR estimation has the standard likelihood function

$$p(y_c \mid \beta_c, \Sigma_c) = (\Sigma_c)^{-\frac{1}{2}} \, exp\left(-\frac{1}{2}\,(y_c - X_x\beta_c)\,(\Sigma_c)^{-1}\,(y_c - X_x\beta_c)\right)$$

In addition, the country-specific parameters have a common mean across countries:

$$\beta_c \sim \mathcal{N}\,(b, \Sigma_b)$$

The prior density is defined as:

$$p(\beta_c \mid b, \Sigma_b) = (\Sigma_b)^{-\frac{1}{2}} \, exp\left(-\frac{1}{2}\,(\beta_c - b)\,(\Sigma_b)^{-1}\,(\beta_c - b)\right)$$

The $q \times q$ matrix Σ_b further embed prior variances that are similar to the Minnesota prior through the matrix Ω_b:

$$\Sigma_b = (\lambda_1 \otimes I_q)\Omega_b$$

Here, the parameter λ_1 governs the tightness of the application of the pooled-group priors to individual country parameters (lower λ_1 lowers the variance of the priors, and therefore lowers the variation of country parameters more rigorously).

Additional priors within Ω_b replicate the Minnesota approach:

λ_3 guides the variance on coefficients on own lags:

$$\sigma^2_{aii} = \left(\frac{1}{l^{\lambda_3}}\right)^2$$

λ_2 changes the variance applied to lags of other endogenous variables:

$$\sigma^2_{a_{ij}} = \left(\frac{\sigma_i^2}{\sigma_j^2}\right)\left(\frac{\lambda_2}{l^{\lambda_3}}\right)^2$$

λ_4 governs the tightness on exogenous variables including the constant.

$$\sigma_{c_i}^2 = \sigma_i^2 \lambda_4^2$$

Additionally, because some EMDE economies have experienced periods of high, or even hyper, inflation, additional dummies are included in countries experiencing above 20% annual inflation.

The priors are applied separately to the group of advanced economies and EMDEs given their characteristic differences.

$$V_\tau = \; e_i' \left(\sum_{\tau=0}^{k-1} B^\tau \Sigma_u B^{\tau\prime} \right) e_i \tag{17}$$

λ_2 is set to 0.5, λ_3 is set to 1, and λ_4 is set to 100. λ_1 is set to an inverse-gamma distribution on each draw, with the variance dependent on the variation of parameters around the mean-group estimator.

A.1 Results of Long-run and Max-Share identifications

Results for both the long-run (Figure 8) and Max-Share identifications (Figure 9) show a broad-based finding of negative employment impacts in the initial year following a technology shock. When estimated using hierarchical priors using the long-run identification, over 70% of economies experience a negative impact. Using the Max-Share identification, over 70% of advanced economies experience a negative initial impact, while nearly 60% of EMDEs do. However, for advanced economies, this rises to 80% when the VARs are estimated without hierarchical priors. The use of hierarchical priors substantially increases the proportion of economies where the negative impact is statistically significant in both identifications. Both identifications are shown to be more biased by non-technology shocks than the Spectral approach used in the main text (Dieppe et al., 2019).

Figure 8: Long-run identification: Proportion of economies with negative, and statistically significant negative impacts of technology on employment

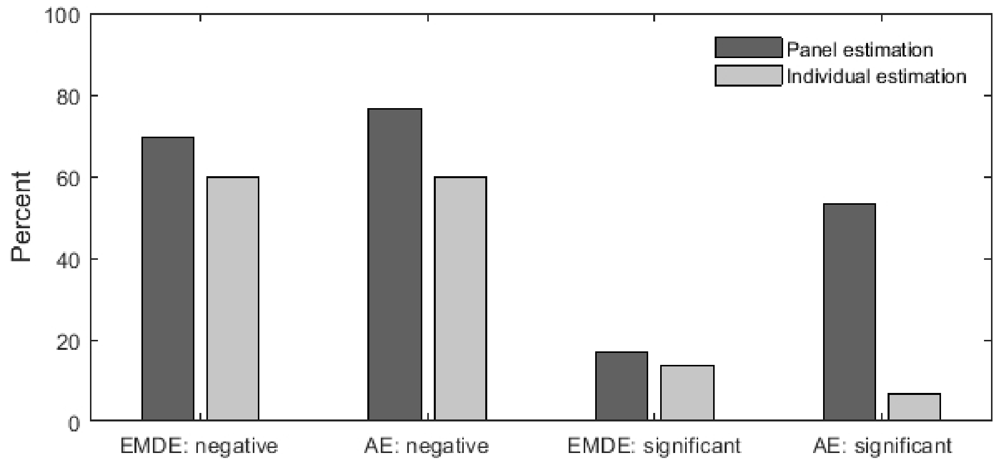

Figure 9: Max-Share estimation: Proportion of economies with negative, and statistically significant negative impacts of technology on employment

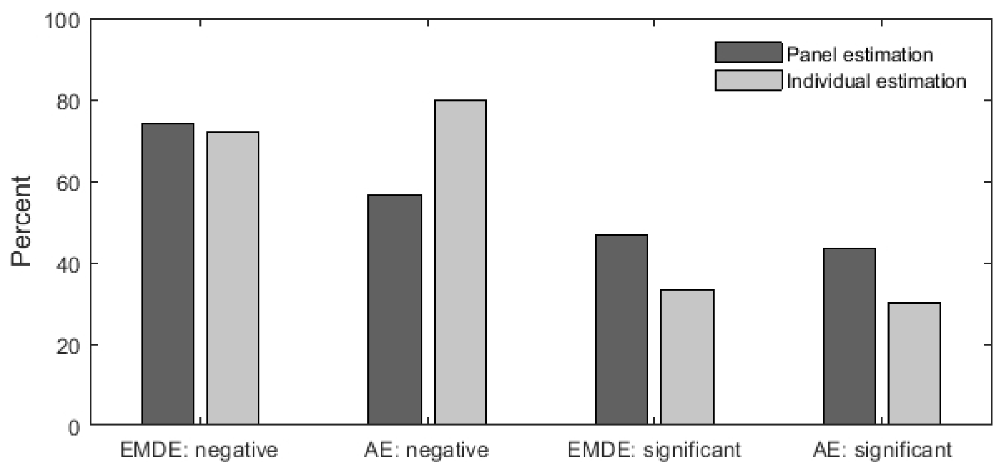

A.2 Targeting TFP to identify technology shocks

Identifying technology shocks as those which drive the largest proportion of low-frequency variation of labor productivity may in part capture lasting capital-specific shocks. For example, institutional changes that encourage a higher degree of capital deepening or permanent tax cuts. As an alternative robustness check, in this section, the level of TFP (from the Penn World Tables, rtfpna) is substituted for labor productivity in the VAR. The Spectral identification then identifies technology as the largest driver of long-run TFP variation. This does not result in a sample-reduction

on average. However, this approach also sufferers from shortcomings. The capital stock is difficult to measure, particularly in EMDEs, and often relies on initializing the capital stock at an average value (Feenstra et al., 2015). In contrast, labor productivity reflected simply as output per worker, requires fewer assumptions.

Figure 10: Technology shock IRFs targeting TFP levels: Pooled estimation

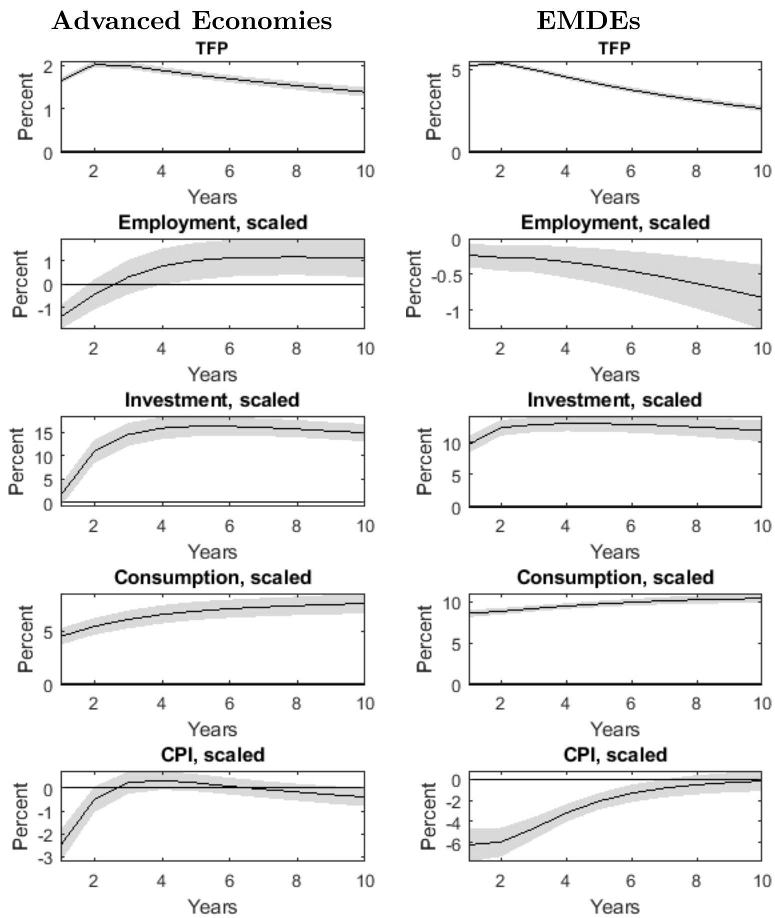

Note: All IRFs are scaled by the impact of the technology shock's impact on productivity. Each IRF can be interpreted as the effect of a technology shock which boosts labor productivity by 10 percent

The TFP-targeting IRFs show a very similar pattern to the main estimations using labor productivity, both in magnitudes and direction of impact (Figure 10). One of the primary differences is that the employment impact is even more persistent in EMDEs, and increasingly negative over time. As in the labor productivity specification, the initial impact of technology on employment

is smaller in EMDEs than in advanced economies. In addition, a similar pattern emerges for the proportion of economies with negative employment impacts initially. Around 60% of economies face negative impacts on employment in year 1, while in the case of using hierarchical priors, 90% of advanced economies show a negative impact. 20% or fewer advanced economies and EMDEs show a statistically-significant impact, rising to 80% when including the priors (Figure 11). In EMDEs, far fewer show statistically negative employment impacts, even in the hierarchical priors estimation. This could reflect greater difficulties in estimating the capital stock and TFP in these economies.

Figure 11: Spectral estimation targeting TFP, not labor productivity: Proportion of economies with negative, and statistically significant negative impacts of technology on employment

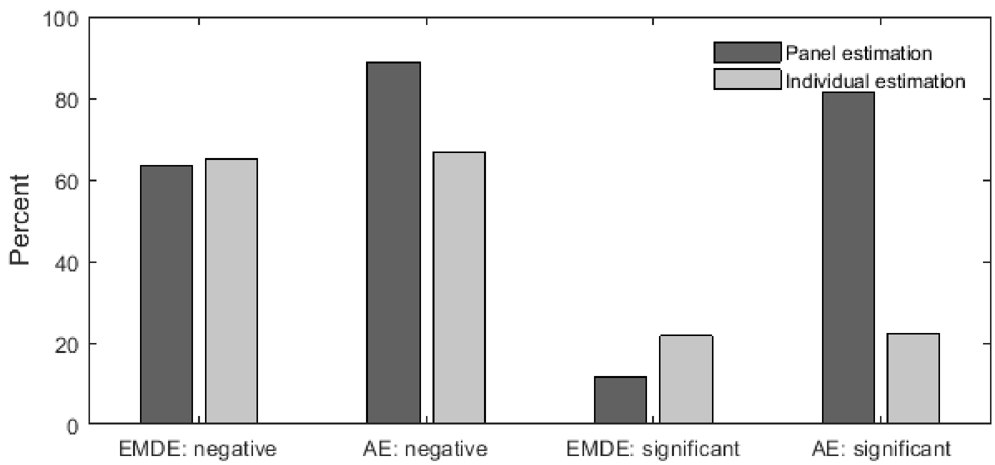

B Theoretical framework: Sectors affected and unaffected by automation.

Labor demand and supply

To disentangle how changing automation (M) in the industrial sector will affect the aggregate demand for labor, we uncover how automation will affect aggregate output and relative prices. The demand for each product $Y(k)$ is assumed to be equal to output, such that $X(k) = Y(k)$. Therefore, the demand for the products produced by the services and industrial sector's respectively can be written:

$$X(k) = \alpha \left(P(k)\right)^{-\sigma} Y^{\epsilon_k+\sigma}$$

Following the concepts in Acemoglu and Restrepo (2018), and closely following the notation of Acemoglu and Restrepo (2017), the envelope theorem is first used to disentangle the influence of automation (M) on aggregate income. Holding quantities $L(NAS), L(I)$, and R fixed

$$\begin{aligned} dlnY|_{L(NAS),L(I),R} &= \sum_{k\in[I,NAS]} s_{Y_k} dlnX(k)|_{L(NAS),L(I),R} \\ &= \sum_{k\in[I,NAS]} s_{Y_k} \left(-\alpha_k \sigma P(k)^{-1-\sigma} \frac{dP(k)}{dM} Y^{_k+\sigma}|_{L(NAS),L(I),R}\right) \\ &= s_{Y_I} \frac{-\alpha_I \left(P(I)\right)^{-1-\sigma} Y^{_k+\sigma}}{\alpha_I \left(P(I)\right)^{-\sigma} Y^{_k+\sigma}} \left(\frac{1}{A}(Q - \frac{W}{\gamma})\right) \\ &= s_{Y_I} \frac{1}{AP(I)} \left(\frac{W}{\gamma} - Q\right) dM \end{aligned}$$

Since the price index and demand for services products does not directly feature M, only the demand for industrial sector goods is affected to the first order.[9] The change in output is therefore a function of the cost savings from switching from paying wages to paying for capital, the elasticity of substitution between industrially-produced and services products, and the share of the product in aggregate production $s_{Y_I} = \frac{P(I)X(I)}{Y}$.

Substituting for the share of industry in total output and the equation for labor demand for a given $X(I)$ yields:

[9] In Acemoglu and Restrepo (2020), σ is eliminated from this equation using the idealized price index: $1 = \sum_K \alpha_k Y^{\epsilon_k} P_{X_k}^{1-\sigma}$. Differentiating yields: $0 = \sum \frac{\alpha_k(1-\sigma)P(k)^{-\sigma}}{1}$, which in turn leads to $\sum \alpha P^{-\sigma} = \sum \alpha\sigma P^{-\sigma}$

$$
\begin{aligned}
dlnY|_{L,R} &= \frac{P(I)X(I)}{Y}\frac{1}{AP(I)}\left(\frac{W}{\gamma} - Q\right)dM \\
&= \frac{\gamma L_I}{Y}\left(\frac{W}{\gamma} - Q\right)\frac{dM}{1-M} \\
&= \frac{\gamma s_L L_I}{WL}\left(\frac{W}{\gamma} - Q\right)\frac{dM}{1-M} \\
&= s_L \pi l(I)\frac{dM}{1-M}
\end{aligned}
$$

Where $s_L = \frac{WL}{Y}$, the income share of all labor, and $l_I = \frac{L(I)}{L}$ is the share of workers in the industrial sector. $\pi = (1 - \frac{Q\gamma}{W})$, or the cost savings from changing from labor to capital. Intuitively, this equation can be thought of reflecting the product of the cost savings from switching to capital/robot-driven production (π), multiplied by the elasticity of substitution for consuming industrial products (σ), multiplied by the income share of industrial workers (s_I), multiplied by the share of industrial workers actually employed in industry (and not the services sector), $l(I)$.

This partial differentiation holds inputs of labor (L) and capital (R) constant. Using factor income shares, $Y = WL + QR$, total differentiation yields:

$$dlnY = dlnY|_{L,R} + \frac{WdL}{Y} + \frac{QdR}{Y}$$

Multiplying by numerator and denominator by labor or capital inputs in each component of the equation (and using the fact that $dlnX = \frac{dX}{X}$) allows it to be rewritten as:

$$dlnY = s_L \pi l(I)\frac{dM}{1-M} + s_L dlnL + s_R dlnR$$

Where s reflects the share of each factor of production in total output (i.e. $s_L = \frac{WL}{Y}$). The share of output allocated to each factor is clearly dependent on the demand for that factor and wages in that sector.

How do prices respond to automation? Here we assess the changes in the price of industrial goods in response to an increase in automation. Starting with the fact that the price of industrial goods is equal to their marginal cost:

$$P(I) = \frac{1}{A}(MQ + (1-M)\frac{W}{\gamma})$$

$$dlnP(I) = \frac{1}{AP(I)}\left(Q - \frac{W}{\gamma}\right)dM + \frac{(1-M)W}{A\gamma P(I)}dlnW + \frac{MQ}{AP(I)}dlnQ$$

Using the definition of the cost saving from automation: $\pi = (1 - \frac{Q\gamma}{W})$

$$dlnP(I) = \frac{-W(1-M)}{A\gamma P(I)}\pi\frac{dlnM}{1-M} + \frac{(1-M)W}{A\gamma P(I)}dlnW + \frac{MQ}{AP(I)}dlnQ$$

Output in the industrial sector can be written as the sum of wages and payments to capital, $P(I)Y(I) = WL_I + QR$. Given that $L_I = \frac{(1-M)}{\gamma^I A}Y(I)$ and $R = \frac{(M)}{A}P(I)Y(I)$, $Y(I) = Y(I)\left(W\frac{(1-M)}{\gamma^I A} + \frac{(M)Q}{A}\right)$. Using the definition $s_{L,I} = W\frac{(1-M)}{\gamma AP(I)}$ and $s_{R,I} = \frac{(M)Q}{AP(I)}$:

$$dlnP(I) = -s_{L,I}\pi\frac{dlnM}{1-M} + s_{L,I}dlnW + s_{R,I}dlnQ$$

Prices are therefore a function of the cost savings from switching from labor to capital for production, and the changes in the cost of wages and capital used in production. Prices in the non-automatable service sector do not depend on the substitution away from labor to machines, but simply wages:

$$P(NAS) = \frac{W}{\gamma A}$$

$$dlnP(NAS) = \frac{W\,dlnW}{\gamma AP(NAS)} = s_{L,NAS}dlnW$$

Now that the determinants of changing income and prices have been determined, demand for labor can therefore be written as a combination of the impacts on jobs immediately affected by automation and jobs that are not:

$$\begin{aligned} L_I^d &= \frac{(1-M)}{\gamma A}X(I) \\ &= \frac{(1-M)}{\gamma A}\alpha_k^\sigma\left(P(k)\right)^{-\sigma}Y^{\epsilon_k+\sigma} \end{aligned} \tag{18}$$

$$\begin{aligned} dlnL^d = -l(I)\frac{dM}{1-M} + \sigma s_{L,I}\pi l(I)\frac{dlnM}{1-M} - \underbrace{(s_L dlnW + s_R dlnQ)}_{\sum_{k\in(I,NAS)} l(k)\sigma(s_{L,k}dlnW + s_{R,k}dlnQ)} \\ + s_L\pi l_I\frac{dM}{1-M} + s_L dlnL + s_R dlnR \end{aligned} \tag{19}$$

The fact that the income elasticities do not enter this equation is due to the assumption that it is assumed that $\sum_k l(k)(\epsilon_k + \sigma) = 1$, or that the income effect across all sectors average to unity, such that aggregate consumption rises in line with income.

Finally, this can be combined with the additional conditions:

$$dlnY = s_L \pi l(I) \frac{dM}{1-M} = s_L dW + s_R dR$$

Which comes from the partial differentiation of the identities $Y_I = L_I W + R_I Q$, $Y = LW + RQ$, reflecting changes in factor prices but while keeping factor quantities unchanged.

Secondly, the labor and capital supply schedules can be differentiated as

$$dW = s_{cL}(dlnL + dlnW) + s_{cR}(dlnR + dlnQ) + \phi dlnL$$

$$dlnQ = \eta dlnR - \eta s_{cL}(dlnL + dlnW) - \eta s_{cR}(dlnR + dlnQ)$$

The equation for $dlnL^d$, after substituting in the income identity, becomes:

$$dlnL^d(1-s_L) = -l(I)\frac{dM}{1-M} + \sigma s_{L,I} \pi l(I) \frac{dM}{1-M} - s_L \pi l(I) \frac{dM}{1-M} + s_L \pi l(I) \frac{dM}{1-M} + s_R dlnR$$

$dlnR$ now needs to be replaced using an expression purely in terms of M and $dlnL$ in order to solve for the effects of automation on labor demand. Substituting M in for the wage and capital supply equations

$$dlnW = s_L \pi l(I) \frac{dM}{1-M} + (s_L dlnL + s_R dlnR) + \phi dlnL$$

$$dlnQ = \eta dlnR(1-s_R) - \eta s_L dlnL - \eta s_L \pi l(I) \frac{dM}{1-M}$$

From the differentiated income identify, $dlnQ = \frac{s_L}{s_R}\pi l(I)\frac{dM}{1-M} - \frac{s_L}{s_R} dlnW$

$$\frac{s_L}{s_R}\pi l(I)\frac{dM}{1-M} - \frac{s_L}{s_R} dlnW = \eta dlnR(1-s_R) - \eta s_L dlnL - \eta s_L \pi l(I) \frac{dM}{1-M}$$

$$\frac{s_L}{s_R}\pi l(I)\frac{dM}{1-M} - \frac{s_L}{s_R}\left(s_L \pi l(I) \frac{dM}{1-M} + (s_L dlnL + s_R dlnR) + \phi dlnL\right) = \eta dlnR(1-s_R) - \eta s_L dlnL - \eta s_L \pi \frac{dM}{1-M}$$

$$\pi l(I) \frac{dM}{1-M}\left(1 - s_L + \eta s_R\right) + \left(-\phi - s_L + \eta s_R\right) dlnL = (1+\eta) s_R dlnR$$

$$\pi l(I) \frac{dM}{1-M} \frac{(1 - s_L + \eta s_R)}{(1+\eta)s_R} + \frac{(-\epsilon - s_L + \eta s_R)}{(1+\eta)s_R} dlnL = dlnR$$

$$\pi l(I)\frac{dM}{1-M}\underbrace{\frac{(1-s_L+\eta s_R)}{(1+\eta)s_R}}_{=1}+\frac{(-\phi-s_L+\eta s_R)}{(1+\eta)s_R}dlnL = dlnR$$

This can now be substituted into the $dlnR$ term in the labor demand equation.

$$dlnL^d\left(\frac{(1-s_L)(1+\eta)s_R}{(1+\eta)s_{s_R}}-\frac{s_R(-\phi-s_L+\eta s_R)}{(1+\eta)s_R}\right) = -l(I)\frac{dM}{1-M} + \sigma s_{L,I}\pi l(I)\frac{dM}{1-M}+s_{cR}\pi l(I)\frac{dM}{1-M} \quad (20)$$

$$dlnL^d\left(\frac{1+\phi}{1+\eta}\right) = \underbrace{l(I)\frac{dM}{1-M}}_{\text{Direct displacement}} + \underbrace{s_{L,I}\pi l(I)\sigma\frac{dM}{1-M}+(1-s_L)\pi l(I)\frac{dM}{1-M}}_{\text{Substitution and income effects}}$$

An inelastic supply of capital ($\frac{1}{\eta}$) magnifies the employment effect, while an inelastic supply of labor reduces it ($\frac{1}{\phi}$). For example, a large fall in labor demand due to automation could significantly reduce wages and the supply of labor in other sectors if labor supply is highly elastic, but this effect is dulled if labor supply is inelastic.

C Appendix: Country employment impact results

In this appendix, individual scaled employment responses to technology shocks are shown for each economy in the sample. The responses are scaled to a technology shock which boosts labor productivity growth by 10 percent. Dashed lines show the median impulse response for the effects of technology shocks on employment when the VARs are estimated individually, with lines with circles reflecting the 16th and 84th percentile error bands. Solid lines and shaded regions show the results of the VARs estimated using hierarchical priors.

Figure 12: Advanced economy employment impacts

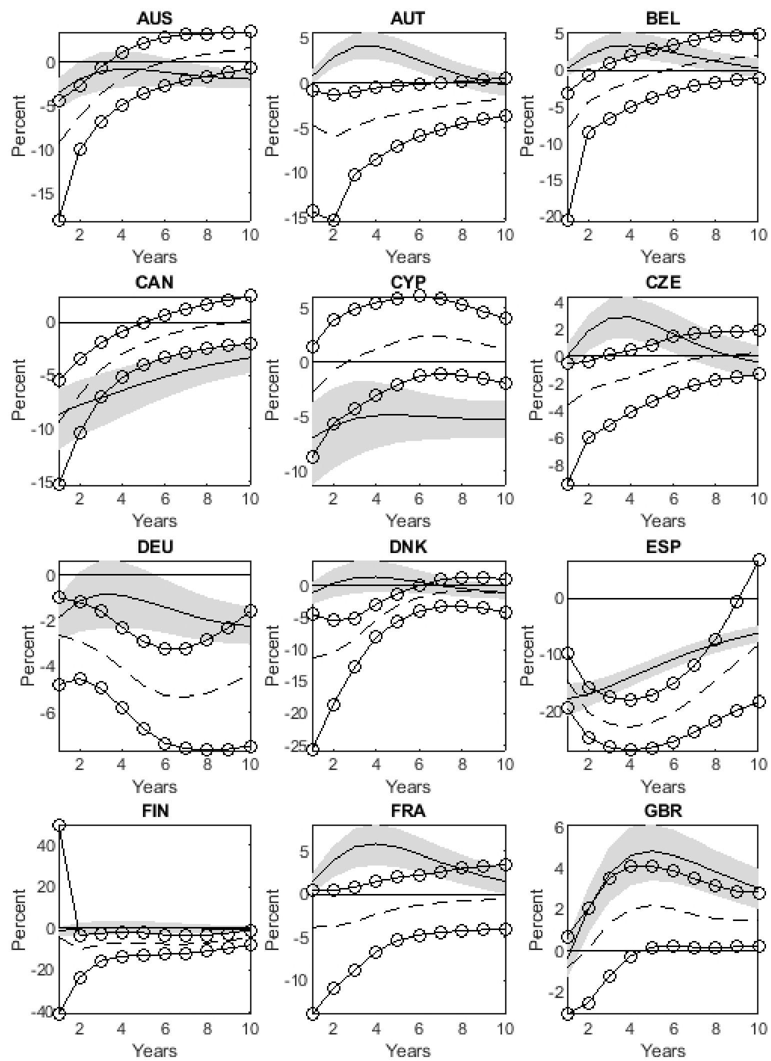

Figure 13: Advanced economy employment impacts

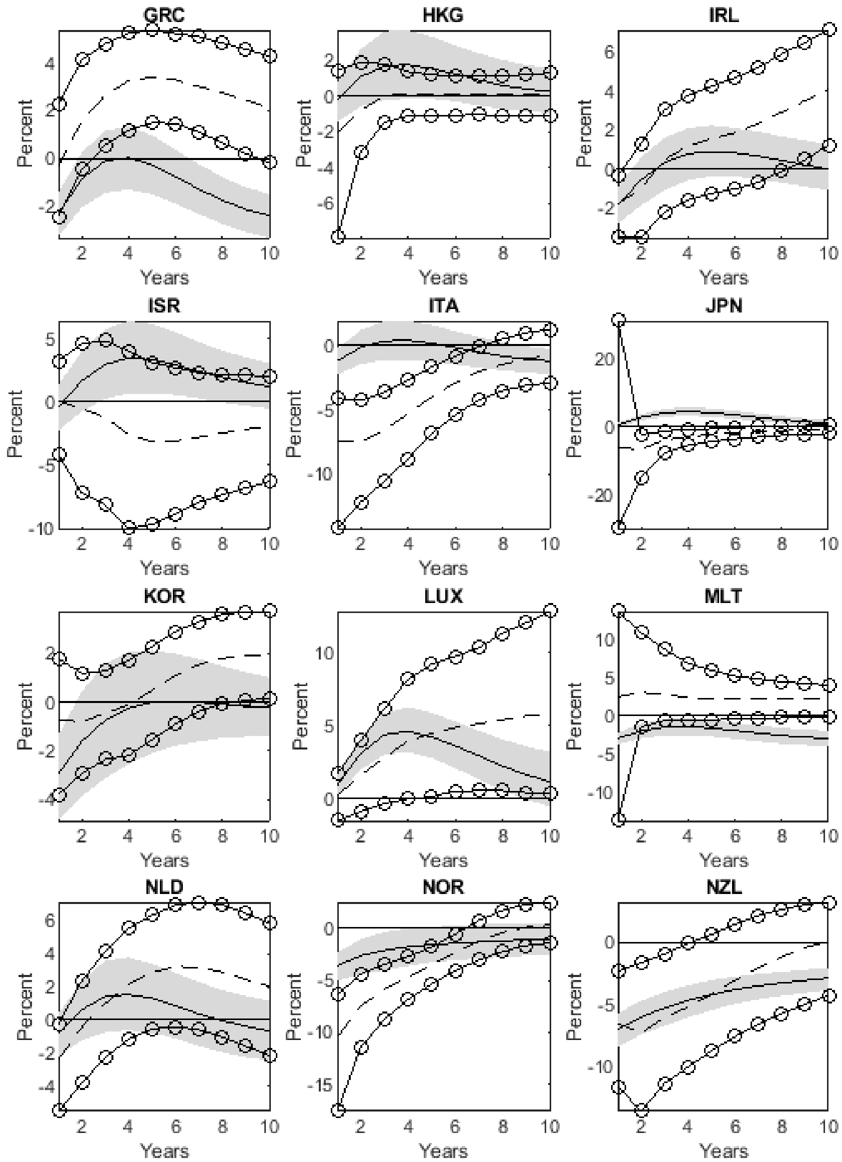

Figure 14: Advanced economy employment impacts

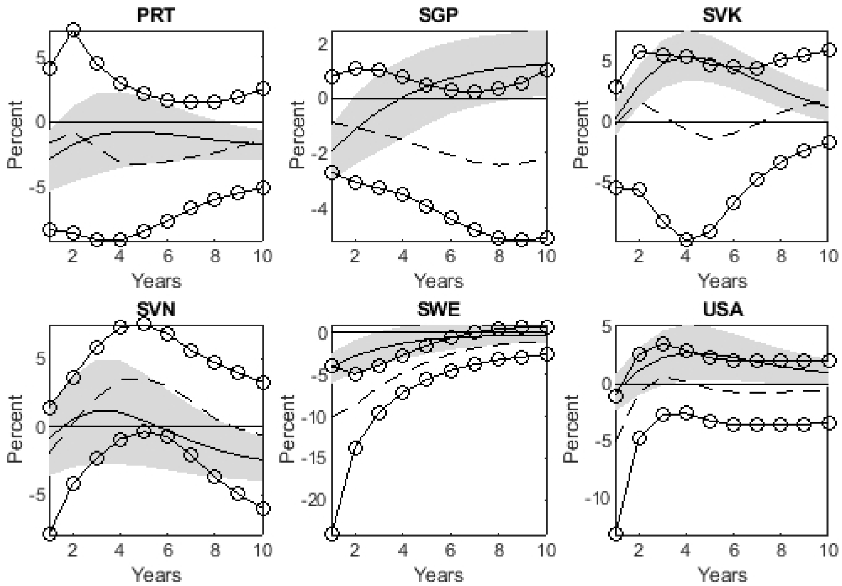

Figure 15: EMDE employment impacts

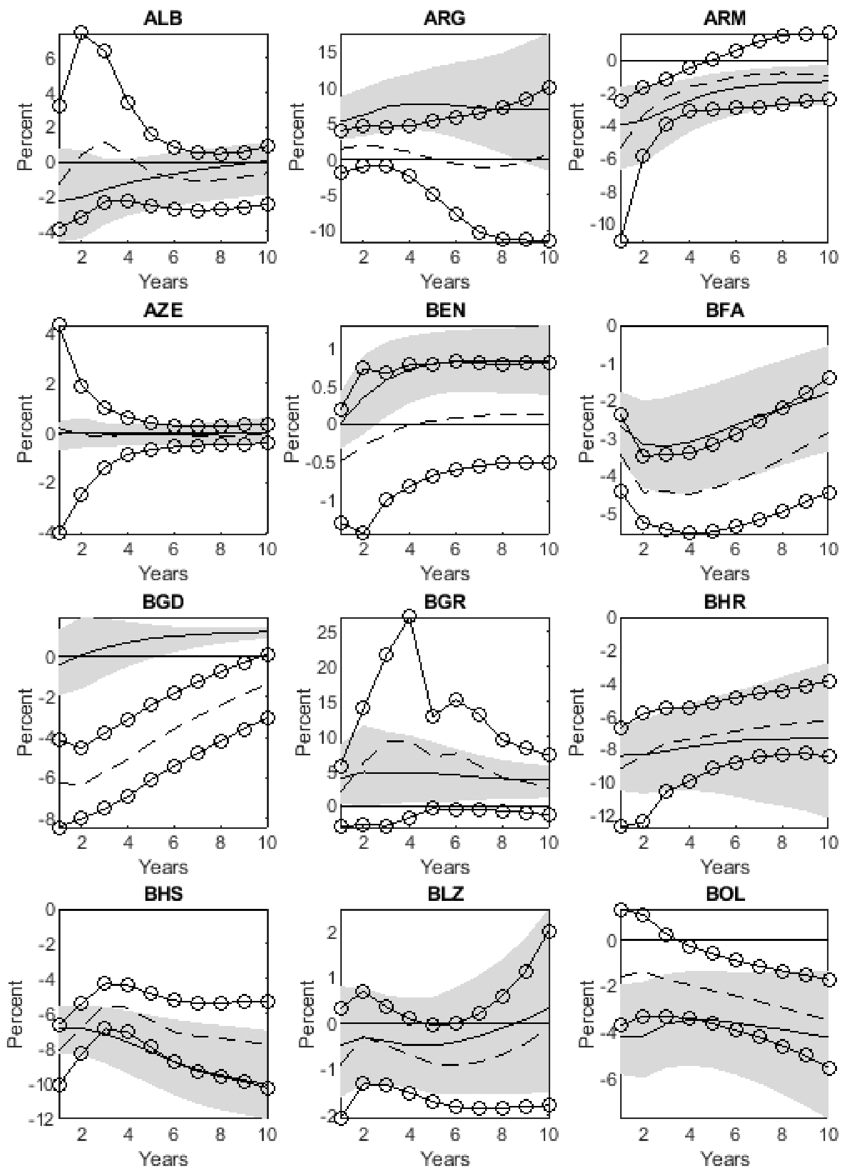

Figure 16: EMDE employment impacts

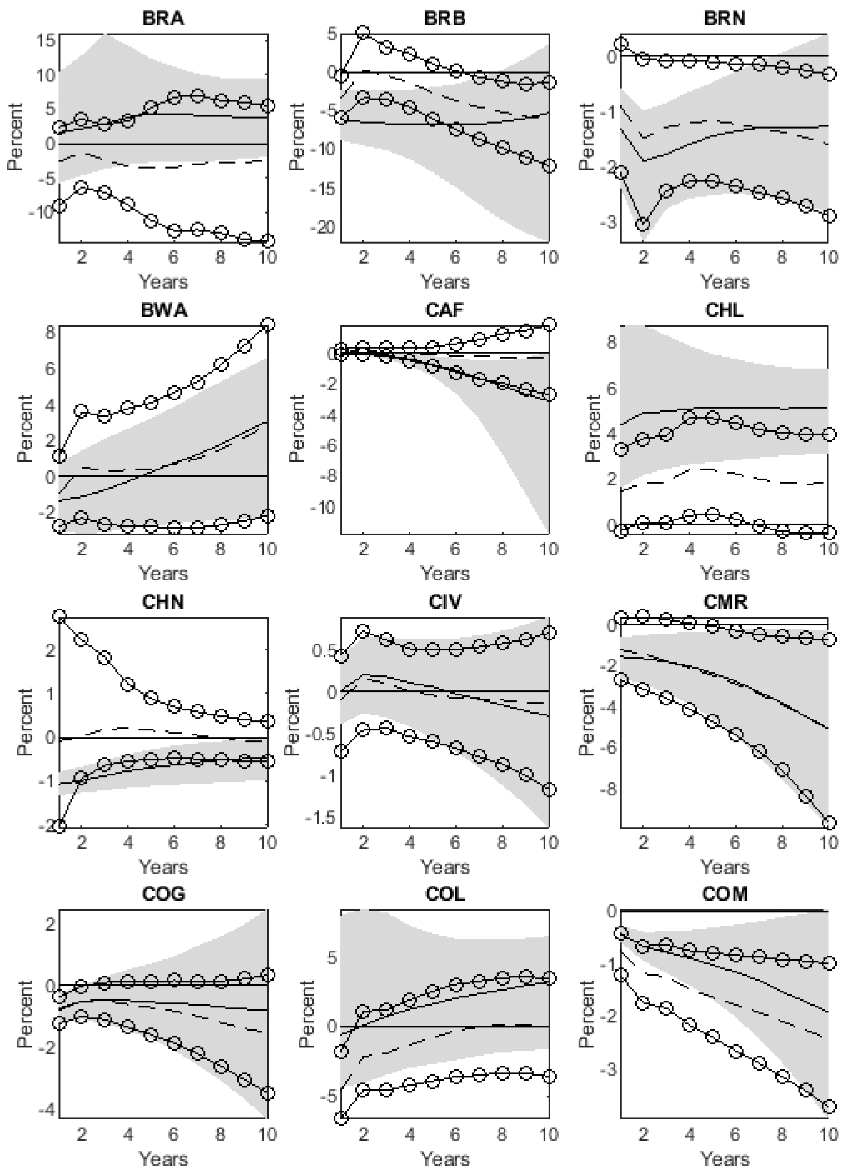

Figure 17: EMDE employment impacts

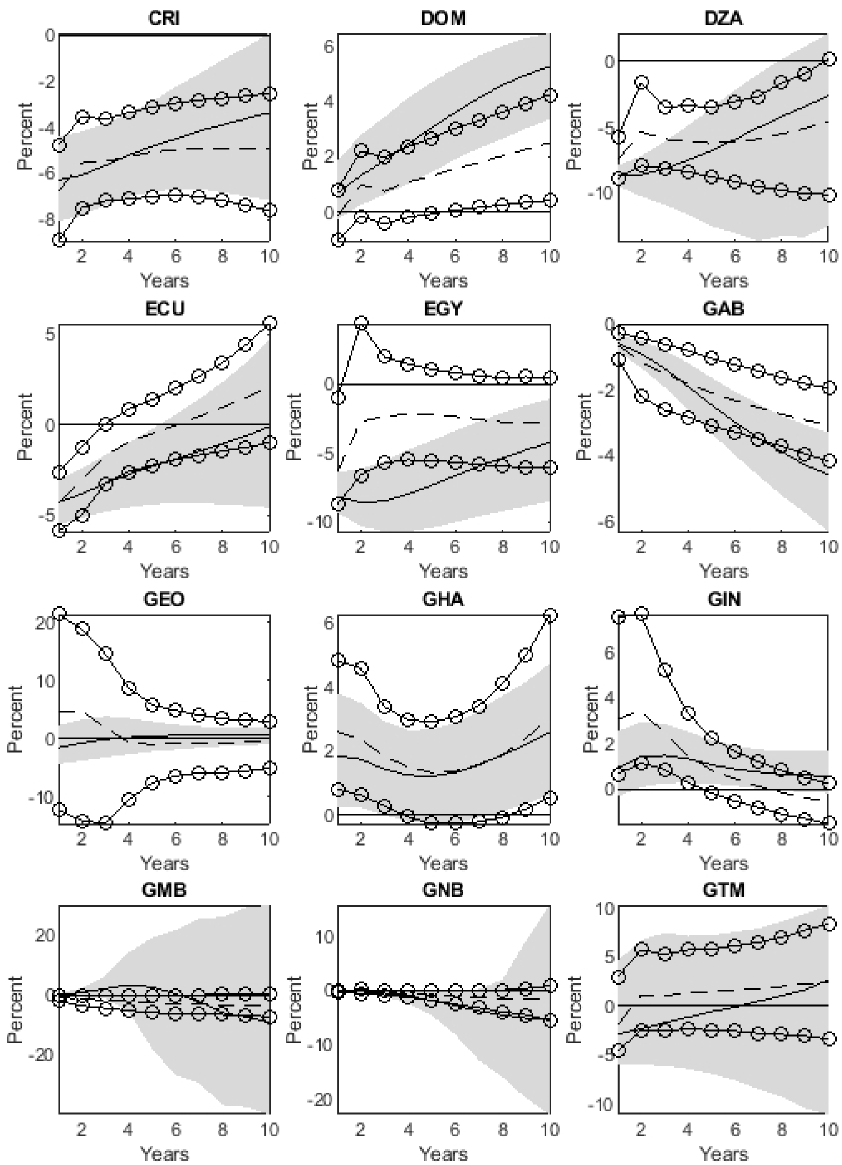

Figure 18: EMDE employment impacts

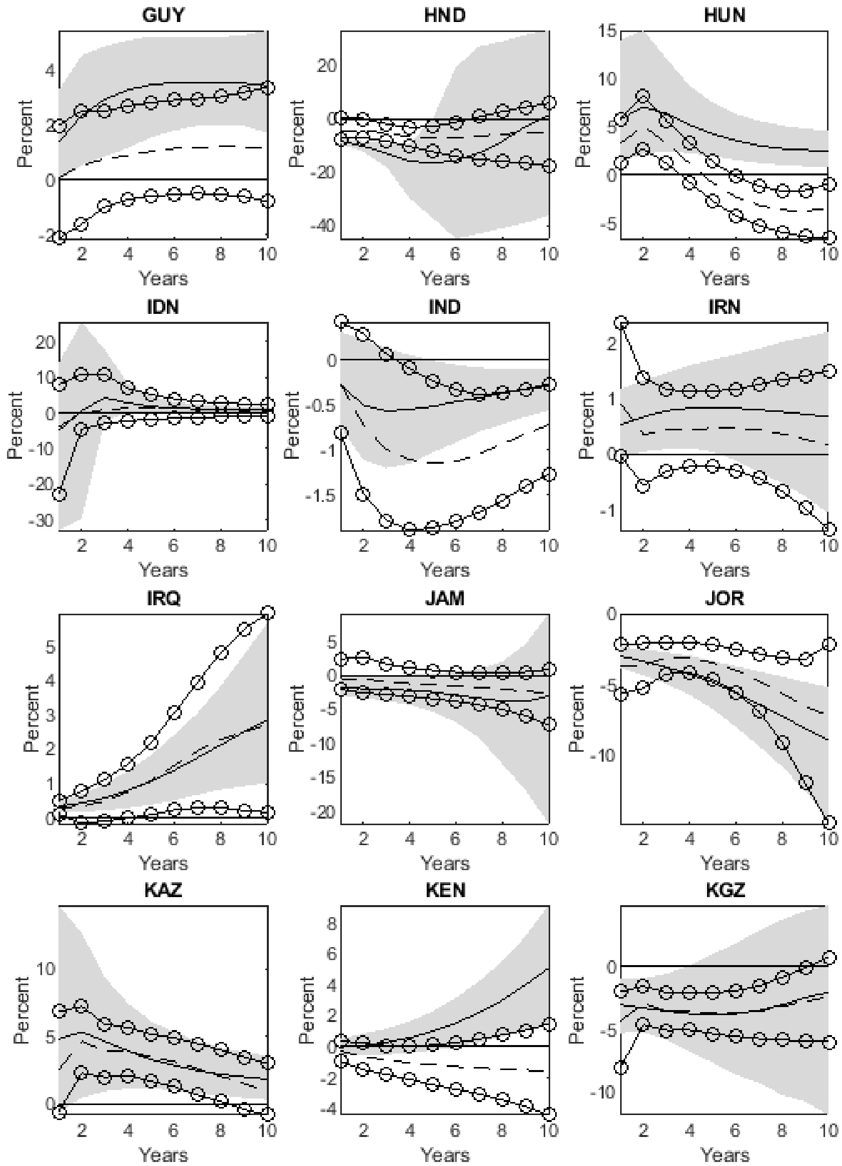

Figure 19: EMDE employment impacts

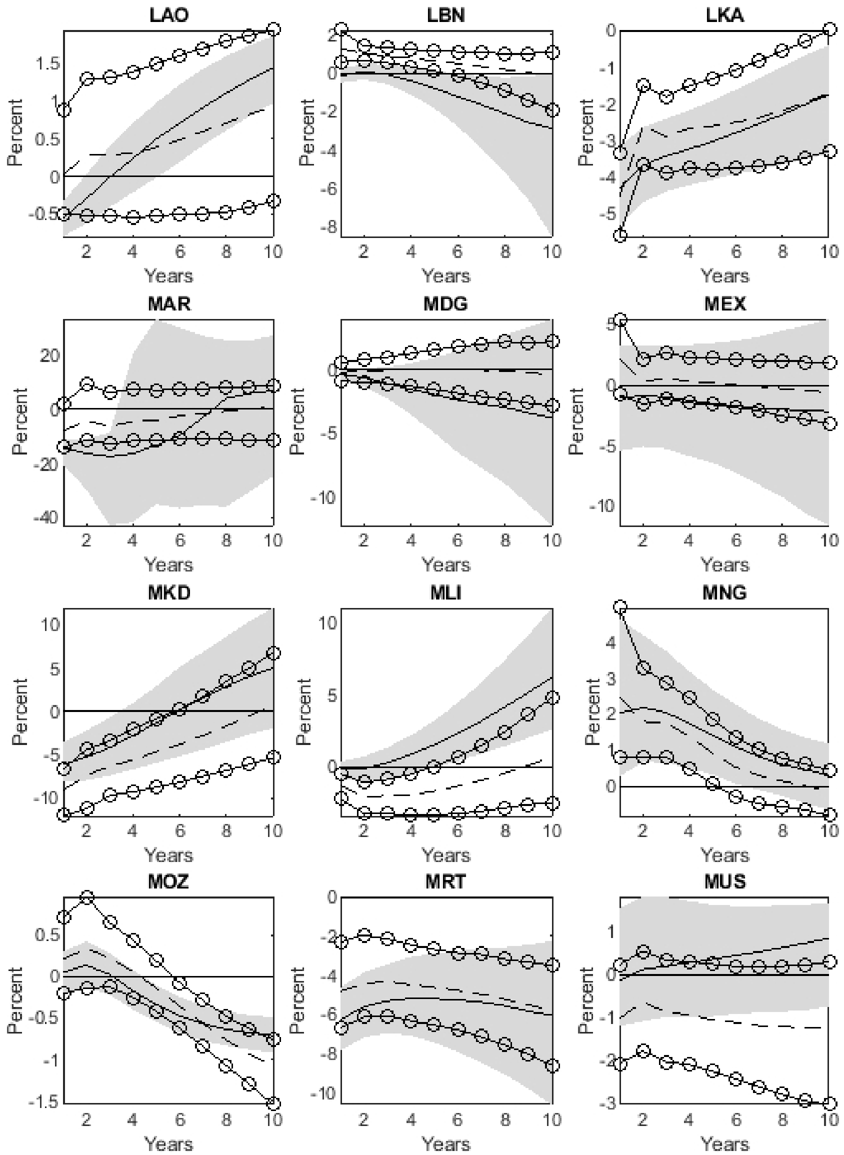

Figure 20: EMDE employment impacts

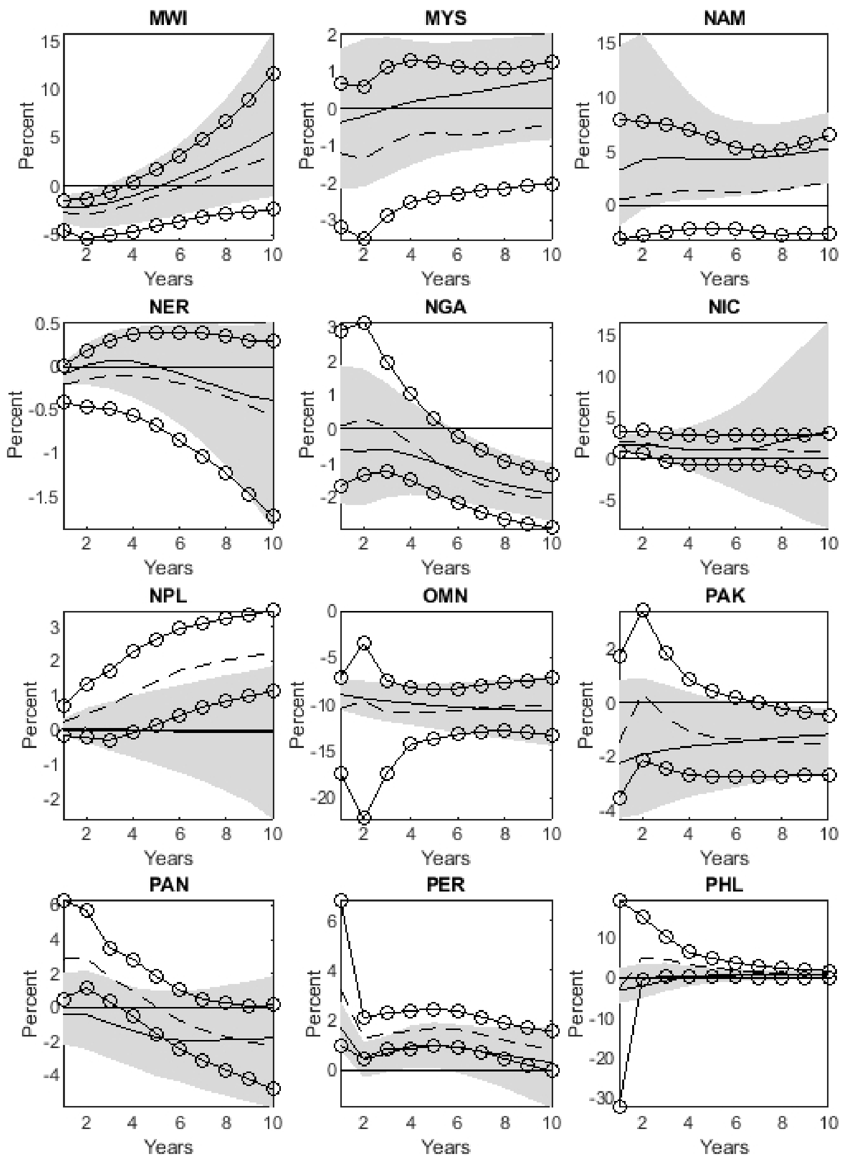

Figure 21: EMDE employment impacts

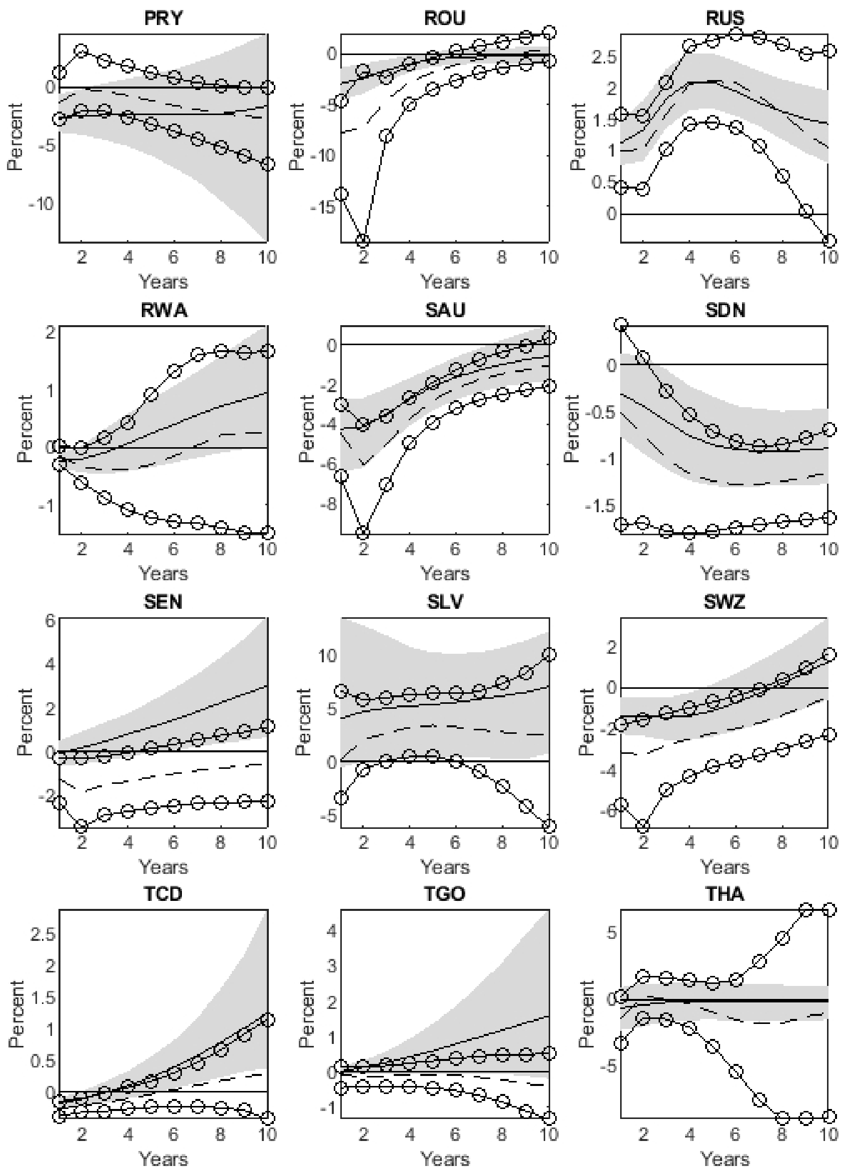

Figure 22: EMDE employment impacts

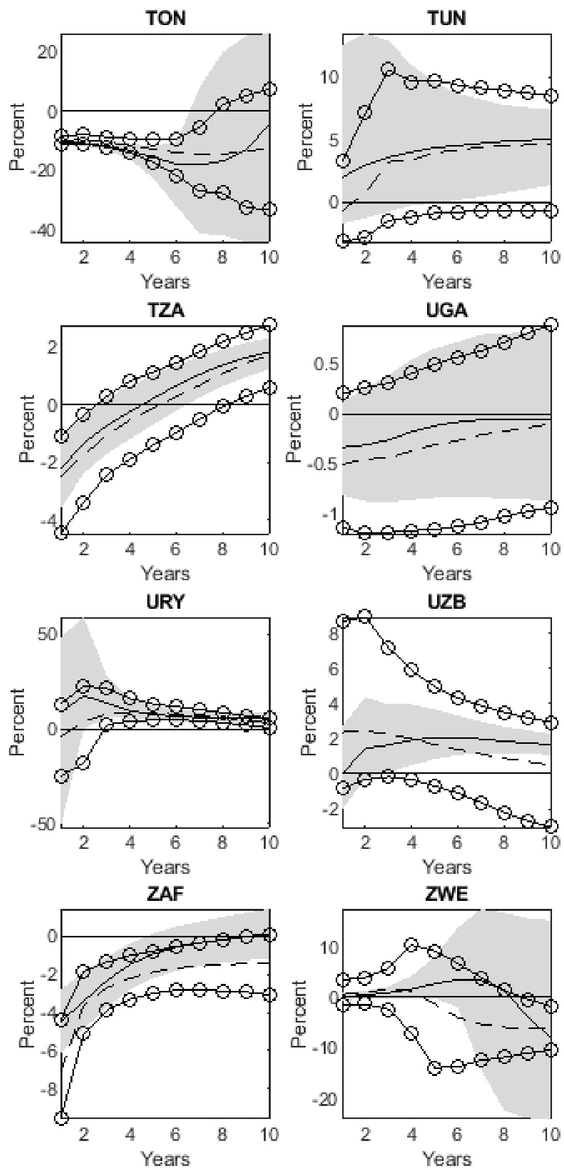

References

Acemoglu, D. (1999). Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence. *American Economic Review*, 89(5):1259–1278.

Acemoglu, D. (2003). Labor- and Capital-Augmenting Technical Change. *Journal of the European Economic Association*, 1(1):1–37.

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, pages 1043–1171.

Acemoglu, D., LeLarge, C., and Restrepo, P. (2020). Competing with Robots: Firm-Level Evidence from France. NBER Working Paper 26738, National Bureau of Economic Research, Cambridge, MA.

Acemoglu, D. and Restrepo, P. (2017). Robots and Jobs: Evidence from US Labor Markets. NBER Working Paper 23285, National Bureau of Economic Research, Cambridge, MA.

Acemoglu, D. and Restrepo, P. (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Aghion, P. and Howitt, P. (1994). Growth and Unemployment. *The Review of Economic Studies*, 61(3):477–494.

Alexopoulos, M. (2011). Read all about it!! What happens following a technology shock? *American Economic Review*, 101(4):1144–1179.

Arntz, M., Gregory, T., and Zierahn, U. (2016). The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis. OECD Social, Employment and Migration Working Papers 189, Organisation for Economic Co-operation and Development, Paris, France.

Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives*, 29(3):3–30.

Autor, D. H. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). The Fall of the Labor Share and the Rise of Superstar Firms. *The Quarterly Journal of Economics*, 135(2):645–709.

Autor, D. H. and Salomons, A. (2018). Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share. NBER Working Paper 24871, National Bureau of Economic Research, Cambridge, MA.

Bai, J. and Perron, P. (1998). Estimating and Testing Linear Models with Multiple Structural Changes. *Econometrica*, 66(1):47.

Barro, R. J. and Lee, J.-W. (2015). *Education Matters: Global Schooling Gains from the 19th to the 21st Century*. Oxford University Press, New York.

Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58:273–289.

Basu, S., Fernald, J. G., and Kimball, M. S. (2006). Are Technology Improvements Contractionary? *American Economic Review*, 96(5):1418–1448.

Beaudry, P., Green, D. A., and Sand, B. M. (2016). The Great Reversal in the Demand for Skill and Cognitive Tasks. *Journal of Labor Economics*, 34(S1):S199–S247.

Beaudry, P. and Portier, F. (2006). Stock Prices, News, and Economic Fluctuations. *American Economic Review*, 96(4):1293–1307.

Canova, F., Lopez-Salido, D., and Michelacci, C. (2010). The effects of technology shocks on hours and output: a robustness analysis. *Journal of Applied Econometrics*, 25(5):755–773.

Canova, F., LopezâSalido, D., and Michelacci, C. (2013). The Ins and Outs of Unemployment: An Analysis Conditional on Technology Shocks. *The Economic Journal*, 123(569):515–539.

Chang, Y. and Hong, J. H. (2006). Do Technological Improvements in the Manufacturing Sector Raise or Lower Employment? *American Economic Review*, 96(1):352–368.

Chari, V. V., Kehoe, P. J., and Mcgrattan, E. R. (2009). New Keynesian Models: Not Yet Useful for Policy Analysis. *American Economic Journal: Macroeconomics*, 1(1):242–266.

Charles, K. K., Hurst, E., and Notowidigdo, M. J. (2016). The masking of the decline in manufacturing employment by the housing bubble. *Journal of Economic Perspectives*, 30(2):179–200.

Christiano, L. J., Eichenbaum, M., and Vigfusson, R. (2004). The response of hours to a technology shock: Evidence based on direct measures of technology. *Journal of the European Economic Association*, 2(2-3):381–395.

Collard, F. and Dellas, H. (2007). Technology shocks and employment. *Economic Journal*, 117(523):1436–1459.

Comin, D. A., Lashkari, D., and Mestieri, M. (2015). Structural Change with Long-run Income and Price Effects. NBER Working Paper 21595, National Bureau of Economic Research, Cambridge, MA.

Cortes, G. M., Nekarda, C. J., Jaimovich, N., and Siu, H. E. (2020). The dynamics of disappearing routine jobs: A flows approach. *Labour Economics*, 65:101823.

De Loecker, J., Eeckout, J., and Unger, G. (2020). The Rise of Market Power and the Macroeconomic Implications. *The Quarterly Journal of Economics*, 135(2):561–644.

Dedola, L. and Neri, S. (2007). What does a technology shock do? A VAR analysis with model-based sign restrictions. *Journal of Monetary Economics*, 54(2):512–549.

Dieppe, A., Francis, N., and Kindberg-Hanlon, G. (2019). New Approaches to the Identification of Low-Frequency Drivers: An Application to Technology Shocks. Policy Research Working Paper, World Bank, Washington, DC.

Dieppe, A., Francis, N., and Kindberg-Hanlon, G. (2021). The Identification of Dominant Macroeconomic Drivers: Coping with Confounding Shocks. ECB Working Paper, European Central Bank, Frankfurt, Germany.

European Commission (2020). Financing the Green Transition: The European Green Deal Investment Plan and Just Transition Mechanism. Technical report, European Commission, Brussels.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The Next Generation of the Penn World Table. *American Economic Review*, 105(10):3150–3182.

Fernald, J. (2007). Trend breaks, long-run restrictions, and contractionary technology improvements. *Journal of Monetary Economics*, 54(8):2467–2485.

Francis, N., Owyang, M. T., Roush, J. E., and Dicecio, R. (2014). A Flexible Finite-Horizon Alternative to Long-run restrictions with an application to technology shocks. *Review of Economic Statistics*, 96(4):638–647.

Francis, N. and Ramey, V. A. (2005). Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited. *Journal of Monetary Economics*, 52:1379–1399.

Galí, J. (1999). Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations? *The American Economic Review*, 89(1):249–271.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*, 104(8):2509–2526.

Hale, G. and Xu, M. (2016). FDI effects on the labor market of host countries. FRBSF Working Paper, Federal Reserve Bank of San Francisco, San Francisco.

Ilzetzki, E., Mendoza, E. G., and Végh, C. A. (2013). How big (small?) are fiscal multipliers? *Journal of Monetary Economics*, 60:239–254.

Jaimovich, N. and Siu, H. E. (2019). Job Polarization and Jobless Recoveries. *The Review of Economics and Statistics*, pages 1–19.

Jarociński, M. (2010). Responses to monetary policy shocks in the east and the west of Europe: a comparison. *Journal of Applied Econometrics*, 25(5):833–868.

Jarocinski, M. and Marcet, A. (2010). Autoregressions in Small Samples, Priors about Observables and Initial Conditions. ECB Working Paper 1263, European Central Bank, Frankfurt.

Khan, H. and Tsoukalas, J. (2013). Effects of productivity shocks on hours worked: UK evidence. *The B.E. Journal of Macroeconomics*, 13(1):549–579.

Kurmann, A. and Sims, E. (2017). Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks. NBER Working Paper 23142, National Bureau of Economic Research, Cambridge, MA.

Maloney, W. F. and Molina, C. (2016). Are Automation and Trade Polarizing Developing Country Labor Markets, Too? Policy Research Working Papers, World Bank, Washington, DC.

Michelacci, C. and Lopez-Salido, D. (2007). Technology Shocks and Job Flows. *The Review of Economic Studies*, 74:1195–1227.

Miranda-Agrippino, S., Hacioglu Hoke, S., and Bluwstein, K. (2019). When creativity strikes: news shocks and business cycle fluctuations. Staff Working Paper, Bank of England, London.

Mortensen, D. T. and Pissarides, C. A. (1998). Technological progress, job creation, and job destruction. *Review of Economic Dynamics*, 1(4):733–753.

Mortensen, D. T. and Pissarides, C. A. (1999). Unemployment Responses to 'Skill-biased' Technology Shocks: the Role of Labour Market Policy. *The Economic Journal*, 109(455):242–265.

Nedelkoska, L. and Quintini, G. (2018). Automation, skills use and training. OECD Social, Employment and Migration Working Papers, Organisation for Economic Co-operation and Development, Paris.

Organisation for Economic Co-operation and Development (2019). *OECD Employment Outlook 2019*. Organisation for Economic Co-operation and Development, Paris.

Park, K. (2012). Employment responses to aggregate and sectoral technology shocks. *Journal of Macroeconomics*, 34(3):801–821.

Pesaran, M. and Smith, R. (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, 68(1):79–113.

Ramey, V. (2016). Macroeconomic Shocks and Their Propagation. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier.

Restrepo, P. (2015). Skill Mismatch and Structural Unemployment. Job market paper, Massachusetts Institute of Technology.

Rodrik, D. (2016). Premature deindustrialization. *Journal of Economic Growth*, 21(1):1–33.

Rozenberg, J. and Fay, M. (2019). *Beyond the Gap : How Countries Can Afford the Infrastructure They Need while Protecting the Planet*. World Bank, Washington, DC.

Rujin, S. (2019). What are the effects of technology shocks on international labor markets? Ruhr Economic Papers 806, RWI, Essen.

Schumpeter, J. (1942). *Capitalism, Socialism, and Democracy*. Harper & Bros., New York.

Uhlig, H. (2003). What moves real GNP? *mimeo HU Berlin*.

Waldkirch, A., Nunnenkamp, P., and Bremont, J. E. A. (2009). Employment effects of FDI in Mexico's non-maquiladora manufacturing. *Journal of Development Studies*, 45(7):1165–1183.

World Bank (2018). *World Development Report 2018: Learning to Realize Education's Promise*. World Bank, Washington, DC.

World Bank (2019). *World Development Report: The Changing Nature of Work*. World Bank, Washington, DC.