Staff Working Paper ERSD-2017-15 **9 November 2017**

World Trade Organization

Economic Research and Statistics Division

Trade, Technology, and Prosperity: An Account of Evidence from a Labor-market Perspective

Marc-Andreas Muendler
UC San Diego, CESifo and NBER

Manuscript date: September 2017

Trade, Technology, and Prosperity

An Account of Evidence from a Labor-market Perspective*

Marc-Andreas Muendler¶

UC San Diego, CESifo and NBER

September 26, 2017

Abstract

Trade and technological change continually alter the workplace and labor-market outcomes, with consequences for economy-wide welfare and the distribution of real incomes. This report assesses the state of economic research into those areas, with a particular focus on empirical methodologies and their adequacy for an assessment of general-equilibrium outcomes. While difference-in-differences techniques and instrumental-variable approaches provide answers, they exhibit shortcomings that limit conclusiveness. Recent advances in structural estimation of multi-country and multi-sector models that allow for reallocation frictions in domestic labor markets hold promise to deliver more definite empirical answers. Interestingly, a conclusion from a two-decades old strand of literature seems to be vindicated by conclusions from a related recent literature: roughly one-quarter of changes in labor-market outcomes (wage inequality then and manufacturing job losses now) was predicted by trade integration and roughly one-third by technological change. The remainder of changes in labor-market outcomes remains unaccounted. The report offers candidate explanations, rooted in recent evidence, how interactions between globalization, technological progress, and structural change may account for that remainder.

JEL classification: F16, F32; J23, J24; L16; O14.

Keywords: Trade, current account balance, automation, choice of technology, industrial structure and structural change, labor-market outcomes, employment, jobs, wages, inequality

---

*I thank Prashant Bharadwaj, Gordon Dahl, Hartmut Egger, and especially Julie Cullen for many insightful discussions. Sam Bazzi, Sascha Becker, Gordon Hanson, Elhanan Helpman, James Rauch and Steve Redding provided insightful feedback on an earlier draft. I am grateful to Marc Bacchetta, José-Antonio Monteiro, Victor Kummritz, Roberta Piermartini and Stela Rubinova for numerous fruitful exchanges. This literature review was commissioned by the World Trade Organization, Geneva, in December 2016 as a background paper for the World Trade Report 2017 under a consultancy contract. Throughout this report, all empirical results that were originally reported in log points have been transformed into percentage changes: a log point change of x is restated as a percentage change of $(\exp\{x/100\} - 1) \cdot 100\%$.

¶muendler@ucsd.edu (*econ.ucsd.edu/muendler*).

1 Introduction

Manufacturing employment declined from 22 to 16 million workers in the United States between 1983 and 2008. After the 2008 financial crisis, the U.S. manufacturing sector lost an additional 2 million jobs (Fontagné and Harrison 2017). Over about the same period from 1980 to 2014, real pre-tax income per U.S. adult has increased by 60 percent on average (Piketty, Saez, and Zucman 2016). However, the average masks a considerable change in inequality. Income stagnated for the lower half of the distribution. For the upper middle class between the median and the 90th percentile, income increased, but less than average, by 40 percent. Among the top 10 percent of U.S. adults, most income increases were concentrated among the top 1 percent: in 1980 the top 1 percent adults earned on average 27 times more than the lower 50 percent of the adults, while they earn 81 times more today (Piketty, Saez, and Zucman 2016).[1] The top earners receive much attention, and a considerable share of their income is from returns on capital. Autor (2014) points out, however, that earnings inequality is also strongly affected by the changes that are consequential for the "other 99 percent of citizens," mainly because of the wage premium associated with higher education and cognitive ability. It is therefore important to look within the labor share in value added and analyze the inequality of wage earnings. Wage inequality is partly driven by the so-called college wage premium, that is the additional pay that a college graduate commands compared to a high-school graduate. The college wage premium was 48 percent in 1979 in the United States and has doubled to 96 percent by 2012, despite a continuous increase in the relative supply of college graduates.

The rise in inequality and the losses of manufacturing employment are common to many OECD countries, and to some developing countries. What factors drive these changes? Globalization in trade of final goods, offshoring of intermediate production stages, or multinational production within the network of plants operated by multinational enterprises? Technical change in the form of automation, computer adoption, or robot use? Organizational change that favors new combinations of workplace tasks within and across firms, and within and across borders? Or a common deindustrialization as productivity change saves manufacturing labor, demand for services expands with per-capita incomes, factoryless producers gain market shares, and manufacturing firms include services in the product mix?

International economic integration and technological progress in their varying forms have been a feature of modern life for centuries. Think of tea. What other product could be more British than tea? One might also wonder, what product could be less British than tea. Much of the global tea harvest used to be and still is from China, India, and Kenya. Depending on how you take the tea, the lemon is from Mediterranean climates, and for much of history the sugar used to come from the Caribbean. At the time of imperialism, these goods were luxuries. Agricultural productivity change, trade, and mass production and consumption, turned these consumer goods into affordable and common parts of our

[1] In the mid-1970s the top 1 percent of U.S. income earners received 8 percent of U.S. income, down from a share of 18 percent in 1913. Since the mid-1970s, however, the income share of the top 1 percent of U.S. earners started to climb again, reaching 13 percent in 1990 and returning to 18 percent by 2008, thus performing a U-turn to where it had been a century earlier (Helpman 2016).

diets. Coffee is widely associated with Italian brands and, with the advent of the retail company Starbucks, now also American. Just the milk likely continues to come from local suppliers. These examples quickly teach simple lessons. Globalization is older than we might think, and even old products continue to be subject to innovation. We receive many of the benefits from globalization and change through consumption that becomes more affordable, while quality improves too. Yet, the most visible consequences of international economic integration and technological progress are often those to workplaces and employment. Are lost jobs coming back, which new jobs are created? To what extent should we support declining activities to decelerate change rather than help the adoption of new activities to accelerate moves? Are globalization and automation today different?

In its most recent economic outlook, the International Monetary Fund (2017) considers one statistic related to both overall income inequality and employment shifts: the share of labor income in total income. The labor share—that is the fraction of national income paid in wages to workers and not in capital returns to investors—has declined in most industrialized countries since the 1980s and in many emerging economies since the 1990s. In a first-pass prediction exercise, the IMF uses broad measures for technological change, participation in global value chains and financial integration, and estimates the partial covariations of the labor share by country with those variables. In advanced economies, technology contributes to about half the variation in labor-share declines across countries, with global-value-chain participation and financial integration accounting for one quarter. In emerging markets, in contrast, global-value-chain participation is the dominant predictor of labor-share declines, offset by a positive covariation with financial integration, while technology plays a minor role. In their study of inequality trends in the United States, China, France, and the United Kingdom, Alvaredo, Chancel, Piketty, Saez, and Zucman (2017) observe rising top income and wealth shares in nearly all of these countries in recent decades. However, the magnitude of the changes in income and wealth concentration varies considerably between countries, even though they are arguably subject to similar global changes in economic integration and technology. Theses disparities in outcomes suggest that country-specific policies and institutions may matter substantively.

This report aims to extend the view, to consider labor-market and household-level outcomes in more detail, to assess real earnings inequality also among wage earners beyond the focus on capital income versus labor income, and to trace changes back to concrete and identifiable mechanisms that may vary from country to country. Most importantly, in the search for explanations, the public discourse as well as economic research tend to consider factors in isolation—attributing observed changes to trade in final goods and global competition, or to offshoring of intermediate production stages, or to labor saving technological change, or to the shift of economies towards services. These factors are not likely independent of each other. Would today's smartphones be viable consumer products in the absence of global value chains that integrate the production stages globally from design, to electronic component production, to assembly, to after-sales services? Would tea have turned from a luxury into a commonplace consumer good in the absence of global integration?

In academic research over the past years, ever more detailed micro data and novel statistical methods have enabled economic researchers to condition on aspects of economic effects that might otherwise confound the conclusiveness of results. For example, longitudinal data that track households over time allow researchers to remove common effects of households of specific socio-economic types. Data on individual firms enable researchers to condition away inherent capability differences between firms. Those data, and associated methods, have made certain empirical insights more definitive. Researchers at the forefront of those innovations have rightly claimed credit for this progress. However, we must not mistake this improvement in definiteness for generality. In fact, the opposite is the case. The more researchers condition away potentially confounding effects, the narrower the answers become. For example, a segment of recent empirical work compares regional outcomes within countries as a function of the local exposures to global trade or technological change, tracking local labor markets or individual households, or both. These studies help answer with much definiteness how disparities between regions and between households have changed. However, these studies are by design incapable of answering the question whether any of these regions or households have benefited or lost compared to their real economic status prior to changes. Did trade and technology lift or lower the boats? A specific segment of recent empirical work can clearly answer whether the boats are further apart from or nearer to each other after the shocks, but by design that research cannot answer whether the typical boat, or any individual boat, was in fact lifted or lowered. To be clear, this point of caution is not a call to a return to less sophisticated analysis. Instead, it is a reminder that, study by study, we need to carefully discern between the questions that the methods actually ask because, naturally, the answers depend on the questions.

In the next section, I outline a labor-market related view on prosperity and propose a measure to guide the literature discussion. The conclusiveness of empirical findings largely depends on the empirical methods that are used, so Section 3 discusses the advantages and limitations of a number of recent approaches. Several of the approaches take inspiration from empirical research in applied microeconomics and attempt to extend the analysis to global general equilibrium, with associated shortcomings. Section 4 turns to research investigating the impact of trade on labor-market outcomes, covering employment responses, changes to the nature of work, wage responses, and outcomes of the matching of diverse workers to heterogeneous firms. Section 5 covers the impact of technology on similar labor-market outcomes, where technology is studied in isolation of trade. Neither of the two literatures can explain more than a quarter to a third of the overall variation in measures of labor-market outcomes. Section 6 then explores the potential of interactions between trade and technology, and additional candidate explanations, that may help predict the observed labor-market outcomes. Section 7 concludes.

2 A Labor-market Lens on Prosperity

To gauge consequences of trade and technological change for individual households' prosperity, a labor-market lens is useful. Two-thirds to three-quarters of national income accrue to wage earners, the remainder to capital owners, and the vast majority of households derives most income from labor earnings.

2.1 Household prosperity and the real wage

A particular object of interest to measure a household's prosperity is the real wage

$$\frac{w_i}{P_i},$$

where w_i denotes nominal labor earnings, P_i denotes the average price of the household's typical consumption bundle, and i stands for a household type (such as the household's percentile in the income or wealth distribution). Theories emphasize different aspects of this ratio's responses to trade and technology shocks for households, and the ratio presents an opportunity to provide a brief preview of the ways in which empirical work addresses individual prosperity.

In a first pass, empirical evaluations of trade or technology shocks consider economy-wide average wage levels, $(1/N)\sum_{i=1}^{N} w_i$, and not individual household outcomes. The average wage is closely related to nominal per-capita income and, when standardized by an average price index P such as the GDP deflator, to overall real income. The average real wage contrasts with real capital earnings and re-weights to a per-worker basis from a per-capita normalization.

Another group of studies considers relative wages for two broad groups of workers such as production or nonproduction workers, or college and less educated workers. Such studies assume that households across the skill spectrum face similar consumption baskets so that there is a single denominator $P = P_i$. In the context of technological change, empirical studies emphasize the potential complementarity of skills with technology, such as automation and computer use, so that a simple ratio of the nominal wages $w_i/w_{i'}$ becomes the relevant object of study (because $P_i = P$ drops from the real-wage ratio), where w_i is the nominal wage of skilled and $w_{i'}$ the nominal wage of unskilled workers.

Similarly, studies of the skill premium can be inspired by classic trade theory, which is often cast in its most intuitive version of two factors of production, two industries, and two countries. The Stolper-Samuelson Theorem in the canonical two-factor, two-industry, two-country case (the classic Heckscher-Ohlin model) predicts that a country's relatively abundant factor will gain from free trade but the country's relatively scarce factor will lose, regardless of the industry that employs them. The reason is that, with free trade, a country will specialize in the industry that heavily draws on the country's abundant factor for production, so the abundant production factor will be in economy-wide higher demand and command higher earnings with trade. In the two-factor case, one country's relatively abundant factor is the other country's relatively scarce factor, so the Stolper-Samuelson Theorem predicts that wage inequality (the ratio $w_i/w_{i'}$) will widen in one country but narrow in the other. For example, if industrialized countries are relatively abundant in high-skilled labor, the skill premium should rise in these economies, whereas the skill premium is predicted to decline in developing countries. While the Stolper-Samuelson Theorem's prediction is an instructive benchmark, its applicability has proven limited.

A large body of empirical work has shown that the skill premium (the ratio $w_i/w_{i'}$), and wage inequality more broadly, has in fact risen in many industrialized and developing countries simultaneously—casting doubt on the adequacy of the canonical Heckscher-Ohlin model. (An influential literature review by Goldberg and Pavcnik (2007) compiles the evidence.) Feenstra and Hanson (1996) proposed to augment the canonical model with a set of production stages behind each final product to capture the idea of offshoring—that is the outsourcing of intermediate-input production stages to foreign suppliers. Production stages differ in the intensity with which they draw on skilled or unskilled workers and, as offshoring costs fall, industrialized countries will specialize in a narrower set of production stages that are relatively more skill intensive while developing countries will attract a wider set of production stages. In the developing country, the newly attracted production stages from industrialized countries will be relatively more skill intensive than the stages they used to host. As a consequence, relative demand for skilled labor is predicted to increase in every country, in industrialized countries because their relatively low-skill intensive production stages depart and in the developing countries because for them the newly hosted production stages are relatively skill intensive. Feenstra and Hanson (1999) documented for the United States that about one-third of the observed increase in the skill premium $w_i/w_{i'}$ in the 1970s and 1980s could be attributed to offshoring.[2]

For the 1990s and later, however, Feenstra (2017) argues that the outcomes and plausible explanations appear different. The skill premium in U.S. manufacturing continued to rise, but the relative employment of skilled workers has tended to fall. That finding is equally suggestive of technological change (Autor, Katz, and Kearney 2008) as it is suggestive of a new form of offshoring, sometimes called trade in tasks. The idea behind trade in tasks is that intermediate goods shipping across countries, and production services provided directly, contain different intermediate services activities, whereby more routine and less cognitively demanding tasks are assigned to offshore suppliers. Models of trade in tasks such as those by Grossman and Rossi-Hansberg (2008) and Rodríguez-Clare (2010) can account for related patterns. Grossman and Rossi-Hansberg (2008) in particular predict a rich set of possible outcomes from trade in tasks, including the possibility that the more offshorable factors of production in fact command relatively stronger wage gains.[3]

[2]More precisely, (Feenstra and Hanson 1999) presented a range of estimates for different trade and technological change predictors and found that between 15 and at most 40 percent of the observed change in the wage-bill share of nonproduction workers in the United States could be explained by offshoring, measured as imports of intermediate inputs.

[3]The possibility that factors of production with offshorable tasks gain from offshoring arises in general

Wage inequality is not generally well captured by a binary skill premium. Work by Autor, Levy, and Murnane (2003) and Goos and Manning (2007) documents that, in the United States and other industrialized countries, polarization has come to dominate wage inequality since at least the 1990s, by which a middle range of workers in the earnings distribution suffers a relative decline of income compared to both lower-income and higher-income groups, while the high-income group experiences faster earnings growth than both the low and middle groups. A richer account of the dispersion in w_i is therefore called for. Moreover, much of the observed wage dispersion in both industrialized and developing countries is driven by wage differentials between industries and firms for otherwise equally skilled workers (Card, Heining, and Kline 2013; Helpman, Itskhoki, Muendler, and Redding 2017). In particular, larger employers pay higher wages for workers with observably identical characteristics (Oi and Idson 1999). Larger firms or plants are globally more integrated and more productive (Bernard, Jensen, and Schott 2009), frequently import, export and adopt advanced technology, and therefore potentially drive the economy-wide dispersion in w_i *within* industries and occupations. For explanations of firm-size and industry-level wage premia, the processes by which employers post vacancies, workers search for jobs, and workers and employers then match, is therefore of particular importance. Recent theoretical advances in macroeconomics (e.g. Lise and Robin 2017), as well as in trade theory at the industry level (e.g. Costinot and Vogel 2010) and at the firm level (e.g. Helpman, Itskhoki, and Redding 2010; Sampson 2014), provide analytic frameworks for the empirical study of modern labor markets and their functioning in matching workers to employers.

A complete account of household prosperity needs to consider that household members can suffer spells of no earnings, $w_i = 0$, and experience wage cuts upon re-employment. Early evidence by Levinsohn (1999) for Chile's trade liberalization, for example, more systematic evidence on multiple trade liberalization episodes across countries by Wacziarg and Wallack (2004), and a comprehensive literature review by Goldberg and Pavcnik (2007) speak to the fact that labor reallocations in response to trade reform and other economic changes, such as structural shifts, are fraught with frictions. Motivated by this evidence, a literature has emerged to study departures from the frictionless labor-market environment and to account for the transitional dynamics in the reallocation of workers across sectors and markets or into temporary non-employment: Kambourov (2009), Artuç, Chaudhuri, and McLaren (2010), Dix-Carneiro (2014), Dix-Carneiro and Kovak (2015a), Menezes-Filho and Muendler (2011), Dix-Carneiro and Kovak (2015b), and Caliendo, Dvorkin, and Parro (2015) are examples. Even in arguably relatively flexible labor markets such as Denmark's, there workers face substantive switching costs

equilibrium because the quasi-rent from task-specific contracting of foreign labor accrues to the factor of production that can specialize in a narrower task range under offshoring. The reason is that the quasi-rent from trade in tasks is generated in the onshore economy, since foreign labor services are employed at domestic productivity but paid the lower foreign wage, while the industry incurs task-specific offshoring costs. Given simultaneous trade in final goods and perfect competition, the quasi-rent can only accrue to one factor in general equilibrium: the workers with offshorable tasks who remain fully employed in equilibrium but now specialize in tasks that are relatively more costly to offshore. The resulting wage gain for the workers with offshorable tasks is also called a "productivity effect" because the quasi-rent accrual is similar to a reduction in those workers' unit labor requirements in the home economy.

between sectors and especially occupations, partly because of large forgone returns to occupational tenure (Traiberman 2016).

The prosperity considerations so far are still predicated on the assumption that $P_i = P$ for all households. However, consumption choices differ considerably across income groups. There are many reasons, intuitive and theoretical, why consumption choices require attention of their own. Suppose consumption baskets did not differ across income groups, and richer households only consumed the same consumption baskets multiple times over so that $P_i = P$ were correct (as many empirical and theoretical approaches pre-suppose). Then a person like Warren Buffet, who is perhaps 10,000 times wealthier than the average American household member, would have to consume 10,000 times as many hamburgers as the average person. By their nature, economic changes such as structural shifts between sectors and technological progress move relative production costs and therefore substantively alter relative prices. Similarly, the gains from trade are first and foremost consumer gains. Even if there is no production—say fruits fell from trees and no labor needed to be allocated to any activity—an open economy will realize welfare gains from globalization because relative prices change when an economy opens to trade and capital markets, and consumers can re-optimize their consumption baskets to be better off after globalization. In short, it is an essential feature of technological change and globalization that relative prices move, and it is an essential fact about consumers that their behavior differs between income groups: $P_i = P$ is not a defensible assumption.

Despite the implied importance of price changes for individual prosperity, there is little empirical work to systematically address income-group specific price changes because of trade or technological progress. Feenstra (1994) and Broda and Weinstein (2006) document that the availability of additional varieties—through trade but potentially also through innovation—raises welfare for the average consumer (expressed in a fall of P for a properly measured consumer price index). Fajgelbaum and Khandelwal (2016) show that, especially in high-income countries, trade tends to be strongly pro-poor in that prices for their relatively preferred consumption goods fall relatively faster—reducing P_i more for low-income households and thus raising their prosperity more than for other households. In fairness to the trade and technology literatures, even pure studies of income inequality tend to set aside real-income changes by using a common price index P, such as the GDP deflator, for all households (e.g. Atkinson, Piketty, and Saez 2011; Piketty, Saez, and Zucman 2016). However, prosperity and the dispersion of prosperity between households crucially depend on the denominator in w_i/P_i.

2.2 Comparable longitudinal household data across countries

The evolution of the real wage w_i/P_i for individual households in turn depends on the underlying employment spells, their durations, and the transition times between employment spells. To assess the resulting employment and wage outcomes within countries, data from repeatedly surveyed household members are particularly useful. Such longitudinal household data exist for a set of countries. Longitudinal household data track household members, and workers in particular, over time. A remaining challenge is to make information from longitudinal household data consistent across countries over longer time horizons. One such project is the construction of the so-called Cross-national Equivalent File (CNEF), at Ohio State University, which strives to make definitions of economic activity and inactivity consistent over time between longitudinal household data from diverse countries.

The CNEF currently spans the time period 1973-2013 for the most comprehensive household survey from the United States and shorter time spans for other countries. The CNEF covers eight countries: the British Household Panel Study (BHPS), the Household Income and Labour Dynamics in Australia (HILDA), the Korea Labor and Income Panel Study (KLIPS), the Panel Study of Income Dynamics (PSID) of the United States, the Russia Longitudinal Monitoring Survey (RLMS-HSE), the Swiss Household Panel (SHP), the Canadian Survey of Labour and Income Dynamics (SLID), and the German Socio-Economic Panel (SOEP). In the present report, evidence is based on mostly five economies: the British Household Panel Study (BHPS), the Household Income and Labour Dynamics in Australia (HILDA), the Korea Labor and Income Panel Study (KLIPS), the U.S. Panel Study of Income Dynamics (PSID), and the German Socio-Economic Panel (SOEP).[4] The reported economic activity in CNEF is not fully consistent across countries at the two-digit level and requires further aggregation to 17 industries, which in this report are consistently defined across the countries (3 industries in the primary sector, 9 industries in manufacturing, and 5 industries in the services sector).[5]

As globalization and technology progress, a common perception is that workers have to transition between economic activities more frequently. This perception would imply that, beyond an assessment of the real wage levels w_i/P_i as outlined above, the expected real wage volatility might also be crucial. However, the perception that cross-activity transitions have become more frequent is not borne out by longitudinal household data for the group of CNEF countries which offer such data. To assess the frequency of lasting transitions between economic activities, we consider worker transitions over a period of four years. To make the household survey statistics nationally representative we base the counts of worker transitions on the respective household members' cross-sectional population weight in the survey's current year.[6]

Table 1 reports the share of workers in a given survey year who are trackable to employment four years prior and did not have a job in the same industry out of the 17 consistently defined industries four years earlier. From the 1980s to the early 2000s, around one-quarter of the trackable workers over four years changed economic activity (while three-quarters of workers held the same or a different job within the same industry). In all countries with comparable longitudinal household data at this level of economic activity, the frequency of cross-industry moves has gradually declined, from 23 to 19 percent in Germany between 2001 and 2013, from 27 to 13 percent in the United Kingdom and

[4]The Canadian Survey of Labour and Income Dynamics (SLID) is omitted from this report because of data access restrictions.

[5]The 17 industries are: Agriculture and forestry; Fisheries; Mining; Chemicals; Synthetics; Earth, clay and stone; Iron and steel; Mechanical equipment; Electrical equipment; Wood, paper and print; Clothing and textiles; Food industry; Construction; Commerce; Transport; Finance and insurance; Other services.

[6]We use the sampling weight in countries with unreported population weight.

Table 1: SHARES OF WORKER TRANSITIONS ACROSS 17 INDUSTRIES, LOOKING BACK FOUR YEARS

	Australia	Germany	Rep. of Korea	United Kingdom	United States
	(1)	(2)	(3)	(4)	(5)
1977					.204
1983					.237
1989		.248			.262
1995		.254		.242	.259
2001		.234		.271	.288
2007	.269	.182	.173	.210	.250
2013	.228	.194	.155	.129	.221

Source: Own calculations based on longitudinal household data with consistently defined variables across countries.
Notes: Share of four-year worker transitions between 17 consistently defined industries. Entries show the share of longitudinally trackable household members, with wage employment in the reported year and four years prior, who are currently employed in one of 17 industries but were employed in another one four years prior (one minus the reported share is the fraction of longitudinally trackable workers who are employed in the same industry in the current year and four years prior). Longitudinal household data are based on the Cross-national Equivalent File (CNEF), Ohio State University, extracted for 1973-2013 from the British Household Panel Study (BHPS), the Household Income and Labour Dynamics in Australia (HILDA), the Korea Labor and Income Panel Study (KLIPS), the U.S. Panel Study of Income Dynamics (PSID), and the German Socio-Economic Panel (SOEP). Reported economic activity In CNEF at the two-digit level was further aggregated to 17 industries consistently defined across the five countries (year 1995 replaced with 1994 in the United States): Agriculture and forestry; Fisheries; Mining; Chemicals; Synthetics; Earth, clay and stone; Iron and steel; Mechanical equipment; Electrical equipment; Wood, paper and print; Clothing and textiles; Food industry; Construction; Commerce; Transport; Finance and insurance; Other services. Worker transition counts based on household members' cross-sectional population weight in current year (sampling weight in countries with unreported population weight).

from 29 to 21 percent in the United States. In Australia, the frequency of four-year worker transitions fell from 27 to 23 percent between 2007 and 2013, and in the Republic of Korea from 17 to less than 16 percent. In the United States, where the longitudinal household data goes back the longest, only the 1970s show a slightly lower frequency of worker transitions than today. In short, workers today transition between economic activities less frequently than at any other time during the past three decades and the perceived increase in real-wage volatility from industry transitions is not supported by data.

Worker transitions between main economic sectors—primary production,[7] manufacturing, and services—are available for two more countries: in the Russia Longitudinal Monitoring Survey (RLMS-HSE) and the Swiss Household Panel (SHP). Table 2 reports the share of workers in a given survey year who are trackable to employment four years prior and did not have a job in the same sector out of the three consistently defined main economic sectors four years earlier. A similar pattern of evidence as for the more detailed 17 economic activities before emerges for the three main economic sectors: the share of

[7]The primary sector includes agriculture, energy and mining, consistent across all countries.

Table 2: Shares of Worker Transitions Across Three Sectors, Looking Back Four Years

	Australia	Germany	Rep. of Korea	Russian Federation	Switzer-land	United Kingdom	United States
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1977							.153
1983							.117
1989		.134					.116
1995		.145				.118	.115
2001		.108				.107	.121
2007	.116	.101	.076		.091	.089	.101
2013	.088	.102	.064	.746	.068	.045	.093

Source: Own calculations based on longitudinal household data with consistently defined variables across countries.
Notes: Share of four-year worker transitions between three consistently defined sectors (primary, manufacturing, services). Entries show the share of longitudinally trackable household members, with wage employment in the reported year and four years prior, who are currently employed in one sector but were employed in another sector four years prior (one minus the reported share is the fraction of longitudinally trackable workers who are employed in the same sector in the current year and four years prior). Longitudinal household data are based on the Cross-national Equivalent File (CNEF), Ohio State University, extracted for 1973-2013 from the British Household Panel Study (BHPS), the Household Income and Labour Dynamics in Australia (HILDA), the Korea Labor and Income Panel Study (KLIPS), the U.S. Panel Study of Income Dynamics (PSID), the Russia Longitudinal Monitoring Survey (RLMS-HSE), the Swiss Household Panel (SHP), and the German Socio-Economic Panel (SOEP). Reported economic activity In CNEF at the single-digit level was corrected according to the two-digit level information (year 1995 replaced with 1994 in the United States). Worker transition counts based on household members' cross-sectional population weight in current year (sampling weight in countries with unreported population weight).

cross-sector transitions has gradually declined in the economies that offer historic longitudinal household data. The frequency of cross-sector moves has fallen from 11 to 10 percent in Germany between 2001 and 2013, from 11 to 5 percent in the United Kingdom and from 12 to 9 percent in the United States. In Australia, the frequency of cross-sectoral worker moves dropped from 12 to 9 percent between 2007 and 2013, in the Republic of Korea from 9 to 6 percent, and in Switzerland from 9 to 7 percent. The only economy with a different experience, both in terms of the level of transitions frequencies and their changes over time, is the Russian Federation, which started to collect longitudinal household data in the mid-2000s. The reported frequency of workers' main-sector transitions at the four-year horizon gradually increased from 64 to 75 percent between 2008 and 2013 in the Russian Federation (not reported in the table). In summary, in all economies that offer historic longitudinal household data since 2003 or before, the frequency of workers' cross-sector transitions has gradually fallen and reached a share of about one in ten workers changing main economic sector at the four-year horizon. Much of this report therefore considers the different levels of real wages across households and sets aside questions about the expected volatility of real wages over time.

3 Methods

Answers depend on the questions. Empirical methods are perhaps best characterized by the exact questions that they ask. The less restrictive the set of assumptions behind a method, typically the less general the question that it addresses. Recent studies of trade and technology, and their impact on labor-market outcomes, adopt combinations of three types of methods: local labor-market studies that have much in common with the difference-in-differences (DD) method; an instrumental-variable (IV) method that can address economy-wide as well local outcomes for affected units of analysis; and theory-based quantification (TQ) exercises that potentially can, but need not, address economy-wide outcomes for all units of analysis. Methods, by design, can only answer the class of questions that they are setup for. It is perhaps ultimately the cumulation of evidence from multiple methods, or their combination, that helps research approximate conclusive answers.

3.1 Difference-in-differences and local labor-market outcomes

An economy can be thought of as a collection of (small open) regional economies, such as commuting zones, with segmented labor markets. Such local labor markets have been the units of observation for policy analysis in a number of contexts (see e.g. Pischke and Velling 1997; Moretti 2010; Kline and Moretti 2013). The local labor market approach has also become the method of choice for several studies of the economic impact of trade or technology. In a precursor study to many subsequent ones, Topalova (2010) analyzes the effect of India's trade liberalization in the early 1990s on the regional incidence of poverty within India. Kovak (2013) analyzes the effect of Brazil's trade reform in the early 1990s on wage and employment outcomes across regions of Brazil; Autor, Dorn, and Hanson (2013) relate earnings and employment outcomes in U.S. commuter zones to the exposure to imports from China; and Acemoglu and Restrepo (2017) relate labor-market outcomes in U.S. commuter zones to the adoption of robots. The design of these local labor market studies is methodologically similar to the difference-in-differences (DD) approach.

To elicit the precise empirical question that studies using DD pose, consider a community (local labor market) i at time t and an outcome Y_{it}. In the context of the so-called China shock, China's accession to the World Trade Organization (WTO) in 2001 and its rapid subsequent expansion of net exports, one can think of a binary treatment for simplicity: community i may or may not be exposed to the China shock D_{it}, so that $D_{i0} = 0$ in the initial period $t = 0$ for all i and $D_{i1} = 1$ in the late period $t = 1$ for at least some communities i. In the absence of a China shock, the expected outcome is $\mathbb{E}[Y_{it}|i,t] = \alpha_i + \gamma_t$, and in the presence of a China shock it becomes $\mathbb{E}[Y_{it}|i,t] = \alpha_i + \gamma_t + \rho$. We are interested in the effect of the China shock on the outcome, ρ. One way to estimate ρ is to specify an ordinary least-squares regression

$$Y_{it} = \bar{\alpha} + \alpha_i + \gamma_t + \rho D_{it} + \varepsilon_{it} \tag{1}$$

for some unrelated (mean independent) error term ε_{it}. The DD estimator of ρ is the linear regression coefficient.

What does ρ tell us? Suppose we happened to know, from conclusive independent analysis, that California ($i = \text{CA}$) is completely unaffected by the China shock.[8] Further suppose that, from conclusive independent analysis, we also know that the China shock adversely hits Michigan's labor market ($i = \text{MI}$). We therefore know the sign of ρ from conclusive independent analysis: $\rho < 0$. This suffices for inference. The difference of outcomes over time in Michigan is $\mathbb{E}[Y_{it}|i = \text{MI}, t = 1] - \mathbb{E}[Y_{it}|i = \text{MI}, t = 0] = \gamma_1 - \gamma_0 + \rho$. In contrast, the difference of outcomes over time in California is $\mathbb{E}[Y_{it}|i = \text{CA}, t = 1] - \mathbb{E}[Y_{it}|i = \text{CA}, t = 0] = \gamma_1 - \gamma_0$ because we know that California is immune to the China shock. Then the population difference-in-differences (DD) is the effect of the China shock ($\rho < 0$ for MI):

$$\begin{aligned}&\big\{\mathbb{E}[Y_{it}|i = \text{MI}, t = 1] - \mathbb{E}[Y_{it}|i = \text{MI}, t = 0]\big\}\\ &\quad -\big\{\mathbb{E}[Y_{it}|i = \text{CA}, t = 1] - \mathbb{E}[Y_{it}|i = \text{CA}, t = 0]\big\} = \rho < 0.\end{aligned}$$

Now suppose that the conclusive independent analysis had actually shown something quite different: Michigan ($i = \text{MI}$) is not affected by the China shock, and California ($i = \text{CA}$) is strictly positively affected by the China shock, so $\rho' > 0$ for CA.[9] We run the same regression as before, and the DD estimator of ρ' is the linear regression coefficient. Then the population difference-in-differences (DD) is the effect of the China shock ($\rho' > 0$ for CA):

$$\begin{aligned}&\big\{\mathbb{E}[Y_{it}|i = \text{MI}, t = 1] - \mathbb{E}[Y_{it}|i = \text{MI}, t = 0]\big\}\\ &\quad -\big\{\mathbb{E}[Y_{it}|i = \text{CA}, t = 1] - \mathbb{E}[Y_{it}|i = \text{CA}, t = 0]\big\} = -\rho' < 0.\end{aligned}$$

The DD estimate is exactly the same as before, but the interpretation is completely different. The estimated coefficient now shows the positive effect of the China shock on California, not the negative effect on Michigan. How can we obtain the conclusive independent analysis to help us understand which case we are dealing with? Well, whatever the conclusive independent analysis, it cannot be DD. By its design, the DD estimator can only identify differential effects between locations, not absolute effects.

Setting aside the mathematics, a common image to describe the consequences of trade, growth, or structural change is that of boats being lifted or lowered. Does the China shock lift or lower the boats in the United States? It is impossible to tell from DD estimation. The DD estimator shows conclusively whether the difference in altitude between two typical boats grew or shrank. For $\rho < 0$ or $\rho' > 0$, we can safely infer that the boats were pulled apart by an additional difference in altitude of $|\rho|$. In other words, regional disparities in the United States worsened. By its methodological design, however, the DD estimator does not allow us to infer whether all boats were lifted (just some less than others), or all boats were lowered (just some more than others), or some boats were lifted (little)

[8] In fact, structural estimation by Caliendo, Dvorkin, and Parro (2015) suggests to the contrary that employment losses in California were the most severe among U.S. states, but welfare still increased.

[9] As mentioned in the preceding footnote, structural estimation by Caliendo, Dvorkin, and Parro (2015) suggests that California's welfare increased (despite job losses).

and others lowered (much). In short, the DD estimator can precisely answer the question how regional disparities between communities in the United States changed. But the DD estimator is, by construction, incapable of showing how the U.S. economy as a whole (the average U.S. boat) was affected (unless we know from conclusive independent analysis of some local community that was immune).

Recent publications go one step further and do not address levels of outcomes but instead rates of change in outcomes. Rather than specifying a level equation such as (1) above, they consider

$$\Delta Y_i = \hat{\gamma} + \rho D_{i1} + \hat{\varepsilon}_i,$$

in the first differences $\Delta Y_i \equiv Y_{i1} - Y_{i0}$, $\hat{\gamma} \equiv \gamma_1 - \gamma_0$ and $\hat{\varepsilon}_i \equiv \varepsilon_{i1} - \varepsilon_{i0}$. The specification still provides a DD estimator of ρ, but level effects $(\bar{\alpha} + \alpha_i)$ are left out and only differential trends $\hat{\gamma} \equiv \gamma_1 - \gamma_0$ can be inferred. Examples for this estimation equation are: Autor, Dorn, and Hanson (2013, equation (5)), where a dependent variable is the decadal change in the manufacturing employment share of the working-age population in a U.S. commuting zone; Acemoglu, Autor, Dorn, Hanson, and Price (2016, equation (7)), who estimate stacked first-difference models for changes in commuting zone employment-to-population rates; and Autor, Dorn, and Hanson (2015, equation (6)). Whether or not the China shock affects the trend itself is, of course, still outside the scope of a DD estimator.

In practice, studies use a continuous variable of the local China shock *exposure* for D_{i1}, but the main logic is unchanged: only disparities between local communities are identified, not a national effect. Some DD studies augment the specification with variables that measure concomitant shocks. Dauth, Findeisen, and Suedekum (2014, equation (3)), for instance, include an Eastern-Europe shock D^+_{i1} (from the fall of the Iron Curtain) in addition to a China shock D_{i1} (from WTO accession) for the German economy in a specification such as

$$\Delta Y_i = \hat{\gamma} + \rho D_{i1} + \rho^+ D^+_{i1} + \hat{\varepsilon}_i. \tag{2}$$

While the DD estimator of ρ remains unchanged, additional coefficients ρ^+ provide for benchmarks that can be judged against each other.

Early publications in the local labor market literature, such as Topalova (2010) and Kovak (2013), for example, carefully acknowledge the limited question that they address. Topalova (2010) finds that the local incidence of poverty in India increases in localities whose industries suffer heightened import competition after trade reform, but acknowledges explicitly and prominently that India's trade reform may have reduced India's economy-wide poverty—an outcome that the DD estimator cannot address. Perhaps because of the strong emphasis of limitations in earlier publications, later publications adopting DD have been more cavalier. In particular, much attention has been paid to summary statements such as one in the abstract of Autor, Dorn, and Hanson (2013, p. 2121) who state that the China shock "explains one-quarter of the contemporaneous aggregate decline in US manufacturing employment." Such an inference would only be possible if we knew, from independent analysis outside a DD approach, how exactly at least one benchmark region responded to the China shock. In the absence of such insight, a local labor-market estimator can only address regional disparities, not national outcomes. While those later research publications, including that by Autor, Dorn, and Hanson (2013,

p. 2121), present careful explanations of the restrictive conditions under which broader inferences from DD estimation would be permissible in the main body of subsequent text, the public reception of research results naturally rests on summary statements whose brevity precludes the enumeration of necessary conditions.[10]

In summary, the DD estimator conclusively identifies disparities between the units of analysis (such as local labor markets). Absent additional assumptions or insight, a DD estimator cannot identify the magnitude of effects that are common across the units of analysis. Even for cross-unit disparity, the DD estimator requires cautious interpretation. Because the common cross-unit effect is unidentified, there is a tendency to exaggerate regional adversity. It remains a possible outcome, consistent with the DD design, that no single U.S. community is negatively affected by the China shock, just some communities with better initial industrial composition are more strongly positively affected than others. In the presence of omitted variables, such as local industrial conditions in terms of productivity or a reduction in U.S. savings rates that drives a traded sector contraction, a DD estimator risks an exaggerated attribution to a single shock. A combination of DD analysis with instrumental variables can partly address omitted variable concerns, as discussed in the following subsection. A remaining question is whether alternative units of analysis, such as industries instead of local labor markets, are less prone to misattribution—a question to be taken up in subsection 3.4. The shortcomings of DD analysis for general-equilibrium considerations, and of local labor-market studies for economy-wide inference in particular, are also understood in macroeconomics. Beraja, Hurst, and Ospina (2016), for example, point out that regional variation alone does not permit the inference of macroeconomic outcomes because (i) the local and aggregate elasticities to the same type of shock are quantitatively different and (ii) purely aggregate shocks are differenced out when using cross-region variation.

What is the likelihood that regional differences, identified in local labor market studies, capture or misrepresent the national effect? The nature of manufacturing is undergoing changes. One example is what Bernard and Fort call factoryless goods producers: wholesale firms that design the goods they sell and then merely coordinate production

[10]Autor, Dorn, and Hanson (2013, p. 2139) acknowledge the precarious nature of their prominent inference of the contemporaneous aggregate decline: "One way to gauge the economic magnitude of these effects is to compare the estimated trade-induced reduction in manufacturing employment with the observed decline during 1990 to 2007. Such an exercise supposes that increased exposure to Chinese imports affects the absolute level of manufacturing employment in the United States and not just relative employment across US commuting zones." The exercise in fact presupposes even more: that increased exposure to Chinese imports affects the absolute level by exactly the same magnitude as relative employment. Autor, Dorn, and Hanson (2013, p. 2139) continue: "Given the magnitudes of the US trade deficit and China trade surplus (and the much larger increase in US imports from China than in US exports to China, as seen in Table 1), the possibility seems real that import competition from China has an absolute impact on US manufacturing (at least as long as trade imbalances persist)." It is indeed plausible that the China shock had a real impact on economy-wide labor market outcomes, but neither the DD estimator nor the simple time trend (in their Table 1) are suited to make the inference. In a related subsequent study that combines local labor market DD approaches and cross-industry estimation, Acemoglu, Autor, Dorn, Hanson, and Price (2016, p. S141) state in the abstract that their "central estimates suggest job losses from rising Chinese import competition over 1999-2011 in the range of 2.0-2.4 million." Similar caution is called for with that inference, given the reliance on local labor market DD estimation.

activities. Bernard and Fort (2015, p. 519): "The best-known example of a factoryless goods producer is Apple Inc. Apple designs, engineers, develops, and sells consumer electronics, software, and computers. For the vast majority of its products, including iPhones, iPads, and MacBooks, Apple does none of the production and the actual manufacturing is performed by other firms in China and elsewhere. While Apple is known for its goods and services and closely controls all aspects of a product, almost none of Apple's US establishments would be in the manufacturing sector." Bernard and Fort document empirically that factoryless goods producing firms are larger, pay higher wages, are active in more industries, and are more likely to be engaged in importing than the typical comparison firm. Given that factoryless goods producers are currently classified as wholesale companies in U.S. data, the China shock can statistically be associated with a manufacturing decline and a simultaneous manufacturing or wholesale expansion in other local labor markets.

The magnitude of factoryless goods production may be substantive. Fontagné and Harrison (2017) report that, in 2017, the U.S. Census Bureau plans to reclassify factoryless goods producing plants from the wholesale sector into the manufacturing sector. Bernard and Fort (2017) calculate that, had this reclassification of factoryless goods producers as manufacturers been conducted in 2007, it would have increased the number of manufacturing jobs in the United States by between 430,000 and 1.9 million—a magnitude comparable to estimates of manufacturing employment losses from the China shock. In that view, the apparent decline of manufacturing employment may merely be a statistical artefact from the past classification of factoryless producers in the wholesale sector. If the manufacturing decline occurs in other locations than the employment increase of wholesale activities, manufacturing outcomes predicted by DD methods can be particularly misleading. In general, even if the DD estimator of local labor markets properly identifies a substantive disparity between regional activities, the economy-wide China shock can be net positive if the build-up of (higher-end) manufacturing or wholesale activities elsewhere outpaces the (lower-end) manufacturing activities in the local labor markets that are most exposed to direct Chinese import competition.

A recent study by Magyari (2017) suggests that a conceivably favorable China shock for the economy as a whole, in the wake of heightened regional disparities, may indeed have been a possibility in the U.S. economy. Magyari (2017) uses U.S. census data on exports, imports and production at the level of plants and firms and argues that the employment of U.S. manufacturing firms rose in response to Chinese import competition in U.S. product markets. In her assessment, more exposed U.S. firms expanded employment in manufacturing, partly by shifting their workforce towards products less exposed to Chinese competition, and in non-manufacturing activities, by adding jobs in research and development, design, engineering, and headquarters services. The hypothesis that China caused a relative expansion of U.S. employment in manufacturing firms that experienced the largest growth in Chinese imports is entirely consistent with the evidence from DD (local labor market) studies, as long as the employment build-up occurs at the firms' plants that are outside the most adversely affected local labor markets.

3.2 Instrumental variable approaches to local or national outcomes

Instrumental variable (IV) estimation helps isolate simultaneous effects from each other and to establish a causal relationship. A valid instrumental variable is a causal variable that predicts only the effects from a single channel, so that the accordingly isolated effect through that single channel can be used to measure the ultimate (labor market) outcome of interest, separately from any potentially confounding effects. It is hard in practice to find a valid instrumental variable which satisfies the assumption that it is not associated with any other simultaneous effect. If available, a valid instrumental variable successfully isolates the effect through a single channel, but at the expense of making the inference context specific. Potential context dependence is an important caveat to keep in mind when pondering possible policy interventions.

It is important to note that the use of an IV approach neither alleviates any shortcomings nor compounds any benefits of other empirical frameworks, when IV estimation is combined with another framework. An IV approach, in addition to any other methodological convention adopted in the research design, brings with it its own advantages and weaknesses. For example, whether or not IV is used within a DD approach, the DD approach only permits a comparison across local labor markets and leaves the national (aggregate) effects unidentified.

Numerous recent studies in the local labor markets approach to the China shock have used Chinese shipments to other high-income countries, outside the region of study, as the instrumental variable—under the assumption that those shipments are not related to the direct shipments from China to the region of study, except for domestic Chinese reasons (e.g. Autor, Dorn, and Hanson 2013; Dauth, Findeisen, and Suedekum 2014). Typically, the instrument is then made location specific by using Chinese trade flows to other high-income countries at the industry level and mapping these flows into the region, using the initial local industry composition to weight the foreign flows by their local relevance. Concretely, industry-level exports from China to other high-income countries are attributed to the local labor market using the local employment shares of industries as weights when assigning the relevant Chinese export flows abroad. This localized instrumental variable is then used to predict the local exposure to Chinese import competition, which in turn is measured similar to the instrument: with economy-wide Chinese imports by industry times the baseline employment share of that industry in a labor market. This exposure measure then predicts local labor market outcomes.

While most of those studies allude to sectoral productivity change in China during the 2000s as the driving domestic Chinese change, they surprisingly choose not to use the observed Chinese productivity change itself as the basis for the instrumental variable, even though productivity measures for Chinese industries exist. Other studies have used the risk in tariff adjustments as an instrumental variable, in particular the uncertainty regarding potential import tariff reversions on Chinese imports to the United States and the removal of that uncertainty when China joined the WTO in 2001 (e.g. Pierce and Schott 2016; Handley and Limao 2013).[11] Studies in the context of unilateral trade reforms, such

[11] Handley and Limao (2013) quantify the impact of U.S. trade policy uncertainty toward China and show that increased trade policy uncertainty reduces trade flows and real income. They argue that, in the case

as in Brazil and India during the 1990s, use the politically implemented tariff changes as the instrumental variable under the assumption that imposed tariff paths are unrelated to the influence of domestic industries (e.g. Topalova 2010; Kovak 2013). In the context of automation, studies consider as an economic shock the local exposure to robots, measured with the economy-wide penetration of robots into each industry times the baseline employment share of that industry in a labor market, and as an instrumental variable the industry-level spread of robots in other advanced economies to predict the adoption of robots in U.S. local labor markets given their initial industry composition (e.g. Acemoglu and Restrepo 2017; Bloom, Draca, and Van Reenen 2016).

Consider an economic outcome Y_i, where the subscript i indexes the unit of study, such as a local labor market. Suppose the outcome is related to an underlying treatment with

$$Y_i = \alpha_0 + \rho_i D_i + \varepsilon_i, \tag{3}$$

where α_0 stands in for other controls (think of $\alpha_0 = \bar{\alpha} + X_i'\alpha$ where $\bar{\alpha}, \alpha$ are constant coefficients and X_i is a set of control variables for other effects), D_i is a binary treatment variable such as the China shock, and there is an error term ε_i. We are interested in the effect of the China shock on the outcome, ρ_i. A problem is that the sign or the magnitude of ρ_i, or both, can depend on effects that happen to coincide with the treatment. The IV method is designed to isolate the effect of just the treatment. The effect of the China shock on the outcome can in general vary across the units of study i, similar to a random effect. This variation is a crucial aspect to understanding the IV method. The following paragraphs explain why, so we can return subsequently to the question how context dependance afflicts inference.

Suppose there is a binary causal variable, Z_i, or an *instrument* for short. The following *treatment equation* (also called first-stage equation) shows how the instrument causally moves the binary treatment variable:

$$D_i = \delta_0 + \lambda_i Z_i + \xi_i, \tag{4}$$

where D_i is the binary treatment, δ_0 stands in for other controls (think of $\delta_0 = \bar{\delta} + X_i'\delta$ with constant coefficients δ_0, δ and the same set of control variables X_i as in the outcome equation (3) above), Z_i is the binary instrument, λ_i the instrument's causal impact on the treatment variable, and ξ_i an unrelated (mean independent) error term. Similar to before, think of the binary treatment D_i as an indicator that shows whether community i is or is not exposed to the China shock (in the late period $t = 1$) so that $D_i = 1$ for at least some communities i (timing is not necessarily important for IV so we can omit time subscripts). In the context of recent local labor market studies, the instrument Z_i is typically the trade flow from China to other high-income countries under the assumption that this trade flow stands in only for Chinese domestic conditions, such as China's industrial produc-

of the United States, the most important policy effect of China's WTO accession was a reduction in trade policy uncertainty: granting permanent normal trade relationship status and thus ending the annual threat to revert to Smoot-Hawley tariff levels. Their estimates imply that there was a welfare gain from removing this trade policy uncertainty for U.S. consumers, similar in magnitude to the U.S. gain from new imported varieties in 1990-2001.

tivity, weighted with the local importance of the exposed industries. For a binary causal variable, think of Chinese trade prowess in a local market; China can either possess trade prowess in a local market or not, depending on both China's prowess by industry and the initial local industry mix. The instrument's causal effect on the treatment is λ_i and can in principle vary across the units of study i, similar to a random effect. For the instrument to be relevant ("strong"), λ_i needs to be non-zero among sufficiently many units of study. For the instrument to be valid, Z_i must not be related to the outcome Y_i in any other way than through the treatment D_i. Concretely, Z_i must not be correlated with ε_i. Depending on a local market's initial conditions, λ_i can vary in magnitude; for a valid instrument we only require that λ_i never switch sign (satisfying monotonicity).

A causal effect answers the hypothetical question: what would have *counterfactually* happened to a local labor market in the absence of Chinese prowess that caused trade flows and hit the local industries? To understand the IV estimator, it is therefore useful to think through hypothetical treatments first, and potential outcomes next. Think of the Z_i as possible states of the world. It can either be the case that Z_i is equal to one (China has prowess in a local market) or Z_i is equal to zero (China does not). Denote a *hypothetical treatment* (hypothetical exposure to trade competition from China) with

$$D_i^0 \text{ if } Z_i = 0, \text{ and } D_i^1 \text{ if } Z_i = 1.$$

The definition of hypothetical treatments allows us to rewrite the actually observed treatment as

$$D_i = D_i^0 + (D_i^1 - D_i^0)Z_i = \delta_0 + \lambda_i Z_i + \xi_i,$$

using the treatment equation (4) from above and setting $\delta_0 = \mathbb{E}[D_i^0]$ and $\lambda_i = (D_i^1 - D_i^0) \geq 0$.[12] The first equality in the expression above simply restates: if $Z_i = 0$ then $D_i = D_i^0$, but if $Z_i = 1$ then $D_i = D_i^0 + D_i^1 - D_i^0 = D_i^1$.

Thinking through hypothetical treatments this way clarifies why there must be a subscript i on the λ_i: the instrument's causal effect on the treatment λ_i varies in general across the units of study i, similar to a random effect, because $\lambda_i = (D_i^1 - D_i^0)$. In the language of medical research, when patients may or may not take the prescribed treatment, there are compliers (patients who follow orders precisely), always takers (patients disobeying by always taking the treatment), and never takers (patients never accepting the treatment). In the context of the China shock on local labor markets, there are communities i that respond to the China shock depending on China's prowess in that community's market (D_i varies with Z_i), then there are always susceptible communities that take a hit by China's trade shock no matter whether China actually possesses trade prowess in that community ($D_i = 1$ regardless of Z_i and therefore $\lambda_i = 0$), and there are never susceptible communities that repel the Chinese trade shock no matter whether China actually possesses trade prowess in that community ($D_i = 0$ regardless of Z_i and therefore also $\lambda_i = 0$). Or think back to the image of boats. Which boats get lifted by the trade wave from China? There are always-taker boats out on the high seas; they get shaken by the China wave no matter

[12] $\mathbb{E}[\cdot]$ is the expectations operator and denotes the statistical mean of the variable inside. $\mathbb{E}[\cdot|\textit{condition}]$ is the conditional expectations operator and denotes the statistical mean of the variable if the condition is satisfied.

what China's actual prowess is like. There are never-taker boats moored in port; they do not move under the China wave no matter what China's actual prowess is like. And then there are complier boats close to shore but not in port; whether the China wave moves them (up or down) depends on China's actual prowess. As a consequence, whether or not the Chinese trade prowess afflicts a particular local market is not a given, D_i^1 is location specific (the same causal China shock may but need not result in a treatment), so $D_i^1 - D_i^0$ is location specific.

Let's turn to the economic outcomes, the objects of our ultimate interest. Denote a *potential outcome* with

$$Y_i(d, z) \quad \text{if } D_i = d \text{ and } Z_i = z.$$

In our context, there are two potential outcomes, $Y_i(1, Z_i)$ or $Y_i(0, Z_i)$ for a level of China's global trade prowess Z_i in a local market i. Given China's prowess Z_i in i, the local effect of Chinese trade competition is then $Y_i(1, Z_i) - Y_i(0, Z_i)$. The definition of potential outcomes $Y_i(1, Z_i)$ or $Y_i(0, Z_i)$ allows us to rewrite the actually observed outcome as

$$Y_i = Y_i(0, Z_i) + [Y_i(1, Z_i) - Y_i(0, Z_i)]\, D_i = \alpha_0 + \rho_i D_i + \varepsilon_i,$$

using the outcome equation (3) above and setting $\alpha_0 = \mathbb{E}[Y_i(0, Z_i)]$ and $\rho_i = Y_i(1, Z_i) - Y_i(0, Z_i)$. The first equality simply restates: if the locality is not exposed, so $D_i = 0$, then $Y_i = Y_i(0, Z_i)$; but if there is a treatment $D_i = 1$ then $Y_i = Y_i(0, Z_i) + Y_i(1, Z_i) - Y_i(0, Z_i) = Y_i(1, Z_i)$. Thinking through potential outcomes this way clarifies why there must be a subscript i on the ρ_i: the treatment's causal effect on the outcome ρ_i varies in general across the units of study i, similar to a random effect, because $\rho_i = Y_i(1, Z_i) - Y_i(0, Z_i)$.

How to estimate the common effect of the China shock on local labor market outcomes (a version of ρ above)? An IV estimator first uses the instrument Z_i (China's trade prowess in local market i) to predict treatment D_i (a local labor market's exposure to Chinese trade competition), and then uses the predicted treatment to measure the effect of the treatment on the outcome (in the local labor market). Under specific conditions on the validity of the instrument,[13] the IV estimator is equal to the average (expected value) of the individual causal effects ρ_i:

$$\bar{\rho} = \mathbb{E}[\rho_i |\, i = \text{ complier}\,]. \tag{5}$$

The IV estimator $\bar{\rho}$ is also called the *local average treatment effect* (LATE). Importantly, it is not equal to the average exposure effect (which would be $\mathbb{E}[Y_i(1, Z_i) - Y_i(0, Z_i)]$). The reason why the IV estimator $\bar{\rho}$ is merely a local (context-specific) average treatment effect, and not the general average effect, is that not all units of study respond to treatment in the same way: the expectations operator weights the individual responses with their

[13] The specific identifying conditions include, importantly, the exclusion restriction. The exclusion restriction postulates that the instrument is valid, that is the causal variable successfully isolates the effect on the outcome through a single channel. In mathematical terms: given a treatment $D_i = d$, there is no further effect of Z_i,

$$Y_i(d, 1) = Y_i(d, 0) \quad \text{for all } d = 0, 1.$$

Another important condition postulates that $\lambda_i > 0$ for all units of study i (or $\lambda_i < 0$ for all unites) but there cannot be individual reversals.

frequencies, and it excludes all localities that are never-takers ($D_i = 0$ regardless of Z_i) and all localities that are always-takers ($D_i = 1$ regardless of Z_i). The IV estimator $\bar{\rho}$ represents an average for a subset of the localities—those for whom their treatment is affected by the value of the instrument. The LATE is the average causal impact among those localities that are not exposed when $Z_i = 0$ and are exposed when $Z_i = 1$. The other localities do not help to identify the IV. Returning to the image of boats, only the complier boats close to shore but not in port are behind the estimated LATE, they get moved by the China wave and $\bar{\rho}$ measures their average response. Next time a China wave hits, the effect will depend on how many boats happen to be always-taker boats out on the high seas at that moment (whose response the LATE estimator does not capture), how many boats happen to be never-taker boats moored in port (whose response the LATE estimator also fails to capture), and how many boats happen to be close to shore but not in port (whose varying responses were measured by the LATE on average).

This understanding of the IV estimator is important, and it applies also to the case when China's prowess in a local market Z_i is not binary but continuous. The estimator profoundly intertwines the individual conditions in the local labor markets with the China trade shock itself. Concretely, the measured labor market response to a China shock will depend on the initial mix of local industries, from highly productive industries with quality products that do not directly compete with Chinese imports to less productive industries under direct competition. Similarly, the employment response to the China shock will vary with the initial capacity of the non-traded sectors active in the community, and less affected traded-goods industries, to absorb displaced labor. The communities with the most adverse industry mix and the least local labor market flexibility will arguably exhibit the most sluggish employment responses. In statistical terms, and for a binary treatment, the average effect of the China shock on exposed communities is a mixture of the effect on always-susceptible and responsive (LATE) communities ($\mathbb{E}[Y_i(1, Z_i) - Y_i(0, Z_i)|D_i = 1]$), whereas the average effect on non-exposed communities is a mixture of the effect on never-susceptible and responsive (LATE) communities $\mathbb{E}[Y_i(1, Z_i) - Y_i(0, Z_i)|D_i = 0]$. In contrast, the IV (LATE) estimator only measures the effect on responsive communities.

If the instrument is valid, that is uncorrelated with the error in the outcome equation and λ_i satisfies monotonicity, then studies based on IV estimators can authoritatively and conclusively isolate the causal effect of economic shocks *under the given historic conditions*. When conducted within a DD method, the IV estimates from recent local labor market studies can therefore definitively isolate the effect of the China shock on cross-regional disparities within a country. But IV approaches can lack the external validity to inform predictions outside the specific context, because a LATE IV estimator cannot dissociate the causal China effect from the dependence on local productivity conditions, local industry mixes, and local labor market flexibility. Outside the framework of local labor markets, IV approaches can be used to measure economy-wide outcomes by changing the unit of analysis to industries or firms and, instead of studying the disparities between regions, infer from the between-industry or between-firm variation a national effect, once a benchmark industry or firm is established (for a discussion see subsection 3.4 below).

Whether the instrument is valid cannot be tested directly. Plausibility checks are possible. A common approach is to conduct a falsification exercise and test whether future shocks can predict past outcomes in the spirit of what is known as Granger (1969) causality. Under certain conditions, including a lacking anticipation of future shocks by economic agents, the future should not cause the past. Neither a future causal variable (instrument) $Z_{i,t=1}$, nor the predicted $\hat{D}_{i,t=1}$ treatment, should predict a past outcome $Y_{i,t=0}$. Interestingly, those falsification tests fail in the context of recent local labor-market studies. In Autor, Dorn, and Hanson (2013, p. 2135), for instance, employment outcomes $Y_{i,t=-1}$ in the distant past (two decades prior to the China shock) are positively related to the China shock $\hat{D}_{i,t=1}$, employment outcomes $Y_{i,t=0}$ in the recent past (one decade prior) are not statistically significantly related to $\hat{D}_{i,t=1}$, while the current employment outcome $Y_{i,t=1}$ in local labor markets is negatively related to the China shock $\hat{D}_{i,t=1}$. This pattern is reminiscent of a mean reversion in local labor market conditions (to take a concrete example, an autoregressive productivity evolution at the decadal time horizon). In such a case, the IV (LATE) estimator will partly capture mean reversion in local communities over a two-decade horizon, while those communities that happen to decay towards the mean because of sluggish productivity performance are concomitantly hit by the China shock.[14]

In summary, if the instruments are valid, then IV based studies inform us with considerable historic precision about the causal impact of economic changes such as the China shock in the past. The use of IVs allows the researcher to make causal statements about the reasons for the measured changes in outcomes, such as regional disparities, beyond predictive statements. But the cost is that IVs turn the inference context specific. The LATE nature of the estimator hampers predictions for future outcomes and therefore obviates policy inference. Suppose the least productive local industries, the least prepared communities in terms of industry mix, and the most inflexible local labor markets respond to the China shock with the strongest manufacturing employment losses, as one might expect. Then we cannot infer how a mitigated exposure to Chinese competition now would affect disparities in regional outcomes, since local conditions may have changed. Neither can we assess to what extent a more resilient local industry mix and outright labor-market interventions would buffer adverse employment outcomes, because a LATE estimator cannot separate the local conditions from the economic shock. The LATE interpretation of IVs implies that the China shock causally altered regional disparities under the given industry mixes, innovative capacities and labor-market institutions during the 2000s; but the China shock would have played out differently for alternative industry

[14]Aware of such potential shortcomings in their instrument, Autor, Dorn, and Hanson (2013, p. 2130) write: "Industry-level or local US productivity shocks may be driving growth in imports from China. If, for instance, the United States has poor productivity growth in furniture, sales of US furniture may fall on both the US and European markets, leading each to import more from third countries, including China. While we cannot rule out this possibility, evidence suggests that productivity growth in China is likely to be an important driver of China's export surge. . . . [G]rowth in imports from China may reflect technology shocks common to high-income countries that adversely affect their labor-intensive industries, making them vulnerable to Chinese competition. In this story, rather than imports from China driving the move toward automation (as in Bloom, Draca, and Van Reenen 2011), automation drives imports from China. Again, we cannot categorically reject this possibility."

mixes, innovative capacities and labor-market institutions and the IV approach cannot tell how differently. Our interest in predictions under possible future shocks and the conduct of policy evaluations therefore calls for structural methods to enter into the mix of approaches.

3.3 Theory-based quantification

Computational general equilibrium (CGE) models have a long tradition in the assessment of trade, technology and other economic changes (Dixon and Parmenter 1996). Historically, an important application of computational general equilibrium models has been the analysis of regional trade agreements, such as NAFTA and the European Union, in which the counterfactual shocks of interest were not productivity shocks but rather changes in trade policy (for a survey see Baldwin and Venables 1995). Recently, the extension of the Ricardian trade model by Eaton and Kortum (2002) has opened new avenues for CGE analysis. The Eaton-Kortum framework accounts for both multiple industries and multiple countries (Dornbusch, Fischer, and Samuelson 1977 had only allowed for one or the other), thus accommodating rich geographic impediments to trade and a country-industry specific productivity dispersion. Solution and simulation algorithms such as those by Alvarez and Lucas (2007) and Dekle, Eaton, and Kortum (2007) permit equilibrium computations in relative terms despite the potential multiplicity of unknown parameters. Generalizations to a wider class of models, including with different factor endowments, are possible. One important extension of the quantitative general equilibrium models is to allow for rich input-output relationships between industries and across borders, using data on global input-output tables. Examples of such extensions include Di Giovanni, Levchenko, and Zhang (2014) and Caliendo, Dvorkin, and Parro (2015) focusing on China, Caliendo and Parro (2015) focusing on NAFTA, and Costinot and Rodríguez-Clare (2014) and Caliendo, Feenstra, Romalis, and Taylor (2015) focusing on global tariff cuts. Structural empirical studies have recently taken on the China shock—promising additional insight into the links between trade openness and labor market outcomes in general equilibrium and offering the possibility for overall welfare analysis.

The estimation of structural models can identify national effects under the assumption that the theoretical model adequately captures reality. Importantly, the subsequent simulation of the estimated model permits the prediction of how economic shocks affect outcomes at present or in the future, not just in the past. Simulation also permits the assessment of consequences from policy interventions. Not all structural models go equally far, however. Adão (2016), for instance, proposes and estimates a structural model that provides additional nonparametric identification conditions within the framework of local labor markets. However, the structural model itself remains confined to analyzing regional disparities between local labor markets, just like DD. Applied to the Brazilian economy, Adão nonparametrically identifies relevant parameters from the cross-regional variation responses of employment and wages to another aspect of the China shock: China's commodity demand shock that shifted relative commodity prices in global markets. Adão estimates that commodity demand shocks from China account for up to 10 percent of the fall in Brazilian wage inequality between 1991 and 2010. This approach is robust to the specific production structure of the economy, but it is not able to capture nationwide effects. More generally, any nonparametric identification results are asymptotic in nature and establish theoretical conditions under which one can precisely (point) identify each one of a potentially infinite set of parameters with a potentially infinitely large dataset. In practice, of course, the available data at any moment in time are finite, so actual estimation typically needs to compromise between parametric and non-parametric elements.

The conclusiveness of structural work depends on the completeness of the model. In the context of labor-markets outcomes, important aspects to account for include how internal migration decisions depend on local amenities and moving costs, how occupational choice depends on frictions to transitions and losses from deteriorating human capital, on the elasticity of labor supply responses to earnings, and on the matching process between workers and employers. The plausibility of estimates depends on the set of parameters that can and cannot be recovered. Recent studies include Coşar, Guner, and Tybout (2016), who allow for simultaneous trade liberalization and institutional labor-market reforms in the context of Colombia during the 1990s, Artuç, Chaudhuri, and McLaren (2010) who use a discrete-choice migration framework for the United States to track employment responses to globalization, Caliendo, Dvorkin, and Parro (2015) who develop and simulate a dynamic economic model of the United States while allowing for a rich set of costs of labor mobility and product transportation between local communities and goods trade across countries as well as geographic factors including local productivities, agglomeration forces and amenity benefits. On the side of research into technological change, Cortes, Jaimovich, and Siu (2016a) formulate a neoclassical model with endogenous labor supply and occupation choice to quantitatively assess the extent to which automation technology can account for increases in unemployment and employment in (low-wage) non-routine manual occupations, as observed over the past three decades in the United States. In the context of short-term labor-market responses to macroeconomic shocks, Lise and Robin (2017) develop a tractable equilibrium model of on-the-job search with diverse workers and employers, aggregate uncertainty, and vacancy creation. Many of these papers consider single shocks in isolation. A rare study that assesses multiple sources of change simultaneously is Burstein, Morales, and Vogel (2015).

The recent advance in structural methods appears ready to produce substantive and conclusive results within relatively short time. Adão, Costinot, and Donaldson (2017) establish the equivalence of a widely used family of trade models with simpler exchange models, in which countries directly exchange factor services. In other words, conventional trade models behave as if countries were directly trading factor services. The method is suitable to study technological shocks (in particular factor-augmenting productivity change) and changes in trade costs at the same time. Adão, Costinot, and Donaldson (2017) show that the quantification of counterfactual changes in important economic outcomes depends only on a reduced factor (labor and capital) demand system, and solution algorithms such as that by Alvarez and Lucas (2007) are applicable. The transformation of the quantification exercise to a factor demand system does require other sacrifices, however. The factor demand system needs to be invertible, which precludes the possibility of zeros in bilateral (implicit) *factor* trade—a circumstance that may require empirical scrutiny before the method is applied widely. In a sample quantification exercise to prove the method's applicability, Adão, Costinot, and Donaldson (2017) use China's recent integration into the world economy and ask how a country's real per-capita income would have differed had Chinese trade costs remained at the level of 1995 (a more limited exercise than employed in estimation approaches to the China shock such as those discussed above). Under a flexible factor demand system, the quantified model suggests that China itself gains relatively much (a small welfare increase by 1.5 percent), rich countries less (United States 0.71 percent, Germany 0.20 percent) while some poorer countries, such as Indonesia and Romania, are predicted to experience welfare losses.

This structural line of research, and the definitiveness of its inference, depend on the respective assumptions on economic primitives, which vary strongly from study to study. A discussion of those studies is therefore better relegated to the topical Sections 4 through 6.

3.4 Units of analysis

Local labor markets are a common unit of analysis (Moretti 2011). A comparison of the regional outcomes permits the precise and conclusive estimation of changing disparities between local labor markets. By design, however, economy-wide effects are removed from estimation and not identified, unless one local labor market with no response is separately known. Under what conditions can other units of analysis—such as the industry, occupation, firm, plant, or worker—overcome the shortcoming and potentially identify national effects?

Industries. Pierce and Schott (2016) apply a DD method to industries and test whether employment fell more strongly in manufacturing industries that were more exposed to China's WTO accession after 2001, as predicted by the removal of uncertainty (concretely the difference between the U.S. most-favored-nation (MFN) tariff and the U.S. non-MFN tariff to which the U.S. Congress could have reverted China prior to WTO membership). Pierce and Schott (2016) estimate that China's WTO accession reduced post-2001 manufacturing employment by 16 percent in exposed industries relative to non-exposed industries. This result reflects an even more pronounced predicted decline than what the cross-regional findings by Autor, Dorn, and Hanson (2013) imply; the latter authors estimate that, in the absence of the China shock, manufacturing employment would have been 5 percent higher. Acemoglu, Autor, Dorn, Hanson, and Price (2016) move between the local labor market and the industry as the unit of analysis to identify varying outcomes from the China shock. They estimate industry effects of the China shock, which can be direct or indirect through linkages, relying on cross-industry variation. They also estimate intra-regional reallocation effects and demand effects of the China shock from the cross section of local labor markets, and then add up the estimated effects. Whereas the industry-level estimates may capture national employment outcomes, local labor market estimates can only capture cross-regional disparities.

To identify the national outcome from a cross-industry approach with non-structural estimation, a benchmark industry is needed, for which there is a priori no predicted effect. Alternatively, an economy-wide constraint that imposes a zero response to the aggregate (average) of industries can provide the benchmark. Then the cross-industry effect relative to that benchmark industry, or the industry composite, is the national effect. In studies that use international trade flows or tariff changes to predict industry effects, for example, a candidate benchmark industry can be one with neither comparative advantage nor disadvantage compared to a trade partner, especially when trade is balanced (exports pay for imports). In the presence of unbalanced trade (such as between China and the United States during the 2000s), the associated alteration in the real exchange rate (or the relative price of non-traded to traded goods) can affect also non-traded industries such as domestic services and a correction may be called for. In practice, controlling for industry-specific effects in adequately specified regressions can establish the benchmark industry relative to (implicit) base-year comparative advantage or relative to a base-year net trade balance at the industry level. However, existing studies do not seem to make explicit efforts to establish either a benchmark industry or an average industry to firmly establish a national effect.

Occupations. Ebenstein, Harrison, and McMillan (2015) use the occupation as the unit of analysis and argue that important worker reallocations, which may occur between occupations across local labor markets, can explain why their results differ from findings confined to local labor markets. Ebenstein, Harrison, and McMillan find that globalization, as measured by either the occupation-level offshore employment at U.S. multinational enterprises or the U.S. product market penetration of industry-occupations with import flows, is not strongly associated with low labor force participation. To the contrary, Ebenstein, Harrison, and McMillan assemble evidence that import competition is associated with higher rates of labor force participation, higher employment, and lower unemployment (where unemployment is a pseudo-occupation). Similarly, Burstein, Morales, and Vogel (2015) use U.S. data at the occupation level in a structural model and find that computerization and changes in task demand jointly explain much of the change in the U.S. college premium. Under what conditions can a cross-occupation approach identify the national outcome? Trade theory can be understood as linking domestic economic activity to the pattern of foreign trade. It is common to think of economic activity in terms of industries, but occupations or combined industry-occupations are similarly natural categories of economic activity to which trade theory can apply. Comparable to the case of industries, the inference of national effects from occupation-level studies then requires a benchmark occupation, for which there is a priori no plausible effect, or a composite benchmark occupation under a national full-employment constraint.

Firms, plants, households, workers. Hummels, Jorgensen, Munch, and Xiang (2014) study employment responses at Danish firms to globalization choices. Muendler and Becker (2010) consider German multinational firms and find that employment at foreign affiliates is a strong substitute for domestic employment. Dix-Carneiro and Kovak

(2015a) move to the household level and find a modest effect of Brazil's trade liberalization on the college premium. Autor, Dorn, Hanson, and Song (2014) and Menezes-Filho and Muendler (2011) move to the worker level and trace worker reallocations, and their failures, across economic activities in response to trade shocks in the United States and Brazil. The finer the unit of analysis, the clearer it can become to define an a priori plausible benchmark unit with no direct effect, such as a domestically owned non-exporter in the case of firms, for example. However, indirect effects in general equilibrium can still shift the benchmark. Moreover, the finer the unit of analysis, the harder it can become to identify cross-unit effects. At the worker level, for example, a differential increase in the frequency of being employed under changed economic conditions can mean that the economic change raises the chance of employment for the directly affected worker type at the expense of a reduced chance of employment for other worker types, with no expected net employment benefit. Similarly, at the firm level a globalization-predicted gain in product market shares can come at the expense of a declining product market share for non-globalized firms, with no a priori grounds to sign the expected net employment effect. In general, units of analysis other than local labor markets can but need not facilitate inference of a national effect.

4 Trade, Labor-market Outcomes and Prosperity

Until recently, much of the empirical literature on the consequences of globalization for labor markets used to be predicated on the long-run perspective that the economy, after a period of adjustment, will return to full employment. Under this perspective, a main outcome of trade-induced labor reallocation is the relative earnings differential across industries and occupations for different groups of workers, such as skill groups, while full employment is restored. Even if there is no trade-induced long-run change to employment levels, a recent set of studies analyzes the durations of reallocation flows and the associated cost of non-employment spells to different worker groups and for the economy as a whole. In addition, recent theoretical advances investigate the dispersion of wages within industries and occupations and relate them to rich underpinnings of the labor-market functioning. Empirical studies based on such frameworks can ask, for example, to what extent the matching process between employers and, within employers, to jobs can explain aggregate wage inequality outcomes.

On the theory side, approaches as early as Brecher (1974) considered effects of classic trade on wages in the presence of wage-setting frictions. More recently, Davidson, Martin, and Matusz (1988) included frictional unemployment in classic trade frameworks. The theoretical literature has since further evolved in two directions. One line of research continues to assume competitive labor markets and introduces assortative matching between heterogeneous workers and employers, with wages varying across employers as a result of differences in workforce composition (see for example Yeaple 2005, Verhoogen 2008, Bustos 2011, Monte 2011, Sampson 2014, and Burstein and Vogel 2016). Another line of research introduces labor market frictions so that workers with the same characteristics can be paid different wages by different firms. For example, efficiency or fair wages can result in wage variation across firms when the wage that induces worker effort, or is perceived to be fair, varies with the revenue of the firm (see for example Egger and Kreickemeier 2009, Davis and Harrigan 2011 and Amiti and Davis 2012). In addition, search and matching frictions and the resulting bargaining over the surplus from production can induce wages to vary across firms (see for example Davidson, Matusz, and Shevchenko 2008, Helpman, Itskhoki, and Redding 2010 and Helpman, Itskhoki, Muendler, and Redding 2017). Naturally, models with labor market frictions can also account for varying equilibrium levels of unemployment as globalization changes. In these models, predictions for trade-induced unemployment changes depend on specifics, and unemployment can rise or decline with globalization (see for example Helpman, Itskhoki, and Redding 2010).

4.1 Employment and Wages

When it comes to empirical analysis of globalization-induced employment consequences, there used to be few estimates to consider. Most research into the labor market implications of globalization used to address not the absolute employment effects of trade but the impact on relative wages and relative employment levels by skill. The paucity of trade studies analyzing absolute employment and absolute wage levels or poverty rates perhaps reflects modelling conventions that impose inelastic labor supply and full employment. There is now a broader empirical literature that examines the impact on labor market outcomes of heterogeneous exposure to import competition in terms of sector of employment—Menezes-Filho and Muendler (2011) and Autor, Dorn, Hanson, and Song (2014)—or region of residence—Topalova (2010), Kovak (2013), Autor, Dorn, and Hanson (2013), and Dix-Carneiro and Kovak (2015a). This subsection therefore pays particular attention to the novel findings on employment consequences of foreign trade.

Many studies that are concerned with employment, however, also investigate wages. This subsection therefore discusses both labor-market outcomes side by side to avert duplication. In recent years, earnings inequality, and especially wage inequality, has become a source of renewed concern. The McKinsey Global Institute (2016) reports that two-thirds of households in 25 industrialized countries were in income segments whose market incomes did not advance or were lower in 2014 than they had been in 2005, and that nearly one-third of those who felt they were not advancing thought that their children and the next generation would also advance more slowly in the future, while expressing negative opinions about trade and immigration. Turning to a prediction, the McKinsey Global Institute (2016) offers a calculation by which 40 percent of income segments may not experience market income gains in the next decade.[15] For France, Italy, and the United States they argue that today's younger generations are even at risk of ending up poorer than their parents.

Elsby, Hobijn, and Sahin (2013) report that, over the past quarter century, labor's share of income in the United States has trended downward, reaching its lowest level with the Great Recession. The labor share—that is the fraction of national income paid in wages to workers and not in capital returns to investors—has declined in most industrialized countries since the 1980s and in many emerging economies since the 1990s (International Monetary Fund 2017). However, the apparent stability of the aggregate labor share until the 1980s in the United States masks considerable volatility in the labor shares within U.S. industries during the period (Elsby, Hobijn, and Sahin 2013). The strong moves in the labor shares at the industry level just happened to offset each other in the aggregate until the 1980s. The recent decline in the labor share has been driven by manufacturing industries. Elsby, Hobijn, and Sahin (2013) argue that the substitution of capital for labor does not appear to be a strong explanation, even as technology embodied the capital goods evolves. In their view, offshoring of the labor-intensive production stages in U.S. manufacturing offers a leading potential explanation for the decline in the labor share over the past quarter century.

The labor share as an inequality measure does not account for the considerable wage dispersion within labor income. Over the past quarter century, another important source of the rise in inequality in industrialized has been the wage premium that college graduates command over high-school graduates. For a review of the literature until the turn of the century, see for example Katz and Autor (1999). Extending the data for the United States to 2012, Autor (2014) documents that the relative supply of college graduates,

[15]The McKinsey Global Institute (2016) predicates this prediction on the assumption that labor-market shifts such as workplace automation accelerate.

measured by the share of their hours in the aggregate number of hours worked by the working-age population, increased continuously from 1963 to 2012. Over the same period of time, the college wage premium initially followed a hump-shaped (inverted U) trajectory between 1963 and 1979, but then sharply increased since 1979. The college wage premium rose from 48 percent in 1979 to 96 percent by 2012.

4.1.1 Final goods trade in North America

Classic trade of final goods for final goods exerts an expansionary effect on the export industries and a contractionary effect on import-competing industries. These classic effects of trade are the foremost changes tracked in the context of the so-called China shock, China's accession to the World Trade Organization (WTO) in 2001 and its rapid subsequent expansion of net exports. Trade in final goods is typically also a main focus for studies of regional trade agreements as well as unilateral trade reforms.

Studying the U.S.-Canada free trade agreement (FTA) of 1989 in a DD analysis at both the industry and the plant-level for Canada, Trefler (2004) reports that manufacturing employment fell by 5 percent in the short run, for about three years. At the same time, the FTA raised labor productivity in the long run by 15 percent, with at least half of the productivity increase coming from the exit or contraction of low-productivity plants. The short-term employment losses are therefore tightly linked to some of the economy's long-term gains on the production side. When it comes to wages, Trefler (2004) estimates that the FTA inception also predicts that annual earnings rose modestly, by about 3 percent, while manufacturing employment fell by 5 percent in the short run because of the FTA. At the plant level, earnings rose for both production and nonproduction workers. Why did earnings rise slightly while employment was falling? Whereas union pressure and skill upgrading in impacted industries appear less plausible, Trefler (2004) argues that workers with lower earnings potential were those more affected by job loss, while the continuously employed workers commanded higher wages and drove up the average earnings in the remaining jobs.

Consequences of China's accession to the WTO are the focus of more recent empirical work. Using a DD approach that can assess changes in cross-regional disparities as a consequence of the China shock, Autor, Dorn, and Hanson (2013) estimate that Chinese import competition explains around one-quarter of the regional disparities in the decline of U.S. manufacturing employment. They define local labor markets along the lines of U.S. commuting zones. Commuting zones are exposed to import competition to different degrees because regions vary in the initial local industry mix. A stronger local exposure to the China shock also predicts higher local payments of government transfers and benefits for unemployment, disability, retirement, and health care. Autor, Dorn, and Hanson (2013), by adopting a DD approach, can investigate the change in regional disparities within the U.S. economy across local labor markets. Under the strong assumption that the estimated change in regional disparities is identical to the national average change (that is under the implicit and unverified assumption that the worst-hit boat is lowered while the most-favored boat is neither lifted nor lowered), the rise in Chinese imports between 2000 and 2007 would predict a net drop in manufacturing employment of up to 2 million jobs.[16] The inability of DD approaches to directly estimate a national economy-wide effect naturally draws attention to structural estimation alternatives.

Autor, Dorn, and Hanson (2013) consider not only employment but also wage effects. However, their estimates of the impact of Chinese import competition on wage differentials between local labor markets are somewhat surprising. Although Chinese import competition was concentrated in manufactured goods, Autor, Dorn, and Hanson find no statistically significant difference in wage outcomes in the manufacturing sector between local labor markets (commuting zones). Neither for college graduates nor for high-school graduates do manufacturing wages differ across regions with varying exposure to the China shock.[17] In explaining the absence of a wage effect from the China shock, despite adverse employment effects, Autor, Dorn, and Hanson (2013, p. 2147) argue that "manufacturing plants react to import competition by accelerating technological and organizational innovations that increase productivity and may raise wages." However, in a later study on U.S. industries and firms, Autor, Dorn, Hanson, Pisano, and Shu (2016) argue that the rate of patenting slows down when U.S. industries or firms are subject to a stronger exposure to Chinese import competition, apparently at variance with the earlier explanation that Chinese import competition spurs innovation. Turning from the local labor market as the unit of analysis to individual workers within the local labor markets, Autor, Dorn, Hanson, and Song (2014) find a differential wage impact of the China shock across local labor markets: Chinese import competition depresses wages but high-income workers were affected less adversely than low-income earners. Overall, the evidence to date on wage effects from the China shock does not appear strong.[18]

Caliendo, Dvorkin, and Parro (2015) build on the concept of local labor markets, extending the Artuç, Chaudhuri, and McLaren (2010) discrete-choice migration framework and combining it with Eaton and Kortum (2002) model for CGE analysis. Beyond more common static approaches, Caliendo, Dvorkin, and Parro develop a dynamic structural model, where production and consumption take place in spatially separate markets for labor and products, and allow for varying frictions impeding labor migration as well as

[16] In a related subsequent study that combines local labor market DD approaches and cross-industry estimation, Acemoglu, Autor, Dorn, Hanson, and Price (2016, p. S141) state in the abstract that their "central estimates suggest job losses from rising Chinese import competition over 1999-2011 in the range of 2.0-2.4 million." To put such net employment losses in perspective, note that U.S. gross job accessions and separations every three months are multiple times as large as this net job loss over roughly a decade, suggesting that a sufficiently flexible U.S. labor market may in principle have the absorptive capacity to accommodate an employment loss of several million jobs from a single sector. Lazear and Spletzer (2012) report that during the mid-2000s the U.S. private sector as a whole (the non-farm economy) hired and displaced approximately 12 million workers every quarter, roughly balancing private-sector job destruction and job creation every three months prior to the Great Recession. About two-thirds of these gross accessions and separations were "churn"—the hires and separations that offset each other within existing plants without a net employment change at the plant—and one-third were net hires or net displacements at the plant level (Lazear and Spletzer 2012, p. 577).

[17] Outside the manufacturing sector, Autor, Dorn, and Hanson (2013) do find larger wage declines for high-school graduates than for workers with a college degree in the wage of Chinese import competition.

[18] Helpman (2016, footnote 44, p. 40) stresses that "the size of these [wage] effects is just not large enough to explain the observed rise in inequality."

domestic and international goods trade.[19] There is a rich set of costs of labor mobility and product transportation, geographic factors including local productivities, agglomeration forces and amenity benefits, as well as input-output linkages between industries. Under the assumption that those costs and benefits remain constant, Caliendo, Dvorkin, and Parro show that the model equilibrium can be solved without the need to estimate productivities, migration frictions, or transport costs. When they feed the observed increase in China's imports into their calibrated model for 38 countries, the 50 U.S. states, and 22 industries, they find a loss of 0.8 million U.S. manufacturing jobs (only about one-quarter of the effect crudely extrapolated from regional disparities in Autor, Dorn, and Hanson 2013), whereas about half of the change in manufacturing employment is predicted by a secular trend. That dominant secular trend may include an employment decline because of labor-saving productivity change in the manufacturing sector, computerization, the adoption of automation technology, or potentially demand-driven deindustrialization.

Adão, Arkolakis, and Esposito (2017) go further and introduce elastic labor supply and agglomeration forces into a standard gravity model of international trade. As a result, the interaction between the labor supply decision and agglomeration forces magnifies the effect of international trade shocks on regional employment and real wages. The elasticity and agglomeration parameters require estimation. Adão, Arkolakis, and Esposito use U.S. regional data on employment, real wages and domestic trade shares for the years 1997-2012. Their results suggest, in contrast to Caliendo, Dvorkin, and Parro (2015), that increases in a U.S. state's trade openness lead to an *increase* in employment in addition to rising real wages (which Caliendo, Dvorkin, and Parro (2015) also found). The estimated endogenous employment expansion is a response to the rising real wage and suggests that identification of the economy-wide employment effect is highly sensitive to the exact specification of the labor supply elasticities. Compared to non-structural research, such as in Autor, Dorn, and Hanson (2013) and similar local labor market studies based on a DD approach, these structural papers are in principle capable of measuring the national U.S.-wide effects of a trade shock (beyond the mere change in cross-regional disparities to which estimation in DD approaches is confined by design).

Structural papers can also estimate welfare effects (changes in the average real wage given a common consumer price index across all income groups), and find those to be strictly positive (even in the absence of endogenous labor supply such as in Caliendo,

[19] At the same time, advances in the macroeconomic literature on labor markets offer the potential for important avenues to adopt realistic specifications of the labor market and its functioning. Lise and Robin (2017) propose a model that features aggregate uncertainty, diverse worker and employer types, and endogenous vacancy creation. Despite the stochastic nature of the model, they prove that surplus from filled jobs fully characterizes both the match value and the mobility decision of workers and, importantly, it does not depend on distributions, so that the model is tractable for estimation. The model and its estimation provide detailed insight into labor-market outcomes. As might be expected, unemployed persons are lower skilled relative to the population, and more skilled workers can match with more jobs than others. There is a positive complementarity between worker skills and firm technology, which becomes stronger in booms than in recessions. Workers are always offered their reservation value. The implication is that when an employer hires a worker out of unemployment, the employer receives the entire match surplus. However, when an employer poaches a worker away from another firm, the employer only receives the share of the surplus in excess of what is necessary to poach the worker.

Dvorkin, and Parro 2015). When the model allows households to endogenously decide their labor supply as in Adão, Arkolakis, and Esposito (2017), given real wages, employment outcomes appear to potentially differ. The papers by Caliendo, Dvorkin, and Parro (2015) and Adão, Arkolakis, and Esposito (2017) are unpublished work in progress but promise to further our insights into the important national outcomes, beyond regional disparities, in the years to come. In an independent non-structural estimation approach, using a nationally representative sample of U.S. households (the Current Population Survey 1983-2008) and the occupation as the unit of analysis, Ebenstein, Harrison, and McMillan (2015) find that globalization, as measured by either the offshore employment at U.S. multinational enterprises or the U.S. product market penetration with import flows, is not strongly associated with low labor force participation. To the contrary, Ebenstein, Harrison, and McMillan assemble evidence that import competition is associated with higher rates of labor force participation, higher employment, and lower unemployment—pointing in a similar direction as the structural work by Caliendo, Dvorkin, and Parro (2015) and Adão, Arkolakis, and Esposito (2017).

Caliendo, Dvorkin, and Parro (2015) find that U.S. aggregate welfare increased 0.35 percent due to China's import penetration growth and that the vast majority of U.S. labor markets experienced an increase in welfare (all local labor markets except the bottom five percent Caliendo, Dvorkin, and Parro 2015, Figure 5). In terms of manufacturing jobs, increased Chinese competition reduced the aggregate manufacturing employment share by 0.5 percentage points, which is equivalent to a loss of 0.8 million manufacturing jobs in the long run. However, Caliendo, Dvorkin, and Parro do not allow for endogenous labor supply, an extension that Adão, Arkolakis, and Esposito (2017) build into their structural model. With the real wage (welfare) increasing under the China shock, as in Caliendo, Dvorkin, and Parro (2015), Adão, Arkolakis, and Esposito (2017) find a manufacturing employment increase in the United States under their labor supply elasticity estimates.[20]

While inelastic labor supply in Caliendo, Dvorkin, and Parro (2015) may limit identification of the economy-wide manufacturing employment effect, their detailed regional and sectoral analysis nevertheless suggests interesting employment reallocation patterns. The computer and electronics and furniture industries contributed to about half of the decline in manufacturing employment predicted by the China shock, followed by metal and textiles. Even under the potentially overestimated economy-wide manufacturing employment decline, Caliendo, Dvorkin, and Parro find that some manufacturing industries, such as food and beverages, gained employment, because they were less exposed to Chinese import competition and benefited more from less expensive intermediate goods. In the regional dimension, Caliendo, Dvorkin, and Parro, Figure 4 estimate that the U.S. state suffering by far the largest predicted manufacturing employment reduction was Cal-

[20]In non-structural estimation, Ebenstein, Harrison, and McMillan (2015) use the U.S. household sample (CPS Merged Outgoing Rotation Groups) for 1983-2008 and measure exposure to globalization using both offshore employment at U.S. multinational enterprises and import flows penetrating U.S. goods markets. Ebenstein, Harrison, and McMillan study the impact of import competition on labor force participation rates. In their estimation, offshoring to China has a negative impact on U.S. labor force participation. (However, other factors such as increasing computer use and substitution of capital for labor are significantly more important determinants of U.S. employment rates across occupations.)

ifornia, accounting for about 12 percent of the economy-wide decline, because it hosted an initially larger share of industries that were subsequently strongly exposed to Chinese import competition (computer and electronics was the industry most exposed to China's import competition and much of Californian manufacturing employment was concentrated in that industry). Overall, globally more integrated local markets appear to fare better in terms of overall welfare: "[L]abor markets in states that trade more with the rest of the U.S. economy and purchase materials from sectors where Chinese productivity increases, tend to have larger welfare gains as they benefit from the access to cheaper inputs from China purchased from the rest of the U.S. economy" (Caliendo, Dvorkin, and Parro 2015, p. 6).

When it comes to unemployment, the DD approach in Autor, Dorn, and Hanson (2013) suggests that regional disparities in unemployment outcomes were such that higher exposure to Chinese imports in a local labor market caused a larger increase in unemployment in that market. Structural estimation by Caliendo, Dvorkin, and Parro (2015) suggests the opposite for the country as a whole. In their estimated model, unemployment falls in the U.S. economy in the wake of the China shock despite a predicted loss of 0.8 million manufacturing jobs. Caliendo, Dvorkin, and Parro find that all non-manufacturing sectors expand under the China shock, driven mostly by employment increases in health care, education, construction, and commerce. The economy-wide employment expansion, and unemployment reduction, is driven by access to lower-priced imports of Chinese intermediate inputs and a relative product demand increase in the expanding sectors. At the person level, Caliendo, Dvorkin, and Parro construct a measure of import changes per worker for every U.S. state over the period 2000-2007 and assess individual hazards of transitions to unemployment (Caliendo, Dvorkin, and Parro 2015, footnote 37). While unemployment falls economy-wide, they also find that states with stronger Chinese import penetration experience a smaller drop or a local increase in unemployment.

4.1.2 Final goods trade in Europe

Dauth, Findeisen, and Suedekum (2014) conduct DD estimation for local labor markets in Germany, similar to those by Autor, Dorn, and Hanson (2013). Dauth, Findeisen, and Suedekum study both a China shock and an Eastern Europe shock—the latter a result of the fall of the Iron Curtain and the accession of formerly socialist economies to the European Union a decade-and-a-half later.[21] While also confined to assessing only cross-regional disparities as a consequence of the China and Eastern-Europe shocks, their regional findings for the German economy exhibit important differences from the U.S. evidence. Most importantly, the net effect of the China and Eastern Europe shock on regional disparities in Germany is the converse of the predicted employment losses in the United States: German local labor markets that were specialized in export-oriented industries experienced strong employment gains and reductions in unemployment. On average across German regions, Dauth, Findeisen, and Suedekum estimate that trade integration

[21] Dauth, Findeisen, and Suedekum (2014, p. 1646): "The fall of the Iron Curtain in the late 1980s was arguably at least as sudden as the Chinese market openings, and the transformation of the former socialist block triggered substantial productivity gains in those economies (Burda and Severgnini 2009)."

has added 442,000 more jobs in more globalization-exposed regions than in less exposed ones and thus contributed to retaining the manufacturing sector. Moreover, in Germany the rise of Eastern Europe affected local labor markets more strongly than the rise of China, which played a relatively minor role.

Why was the adverse China shock much stronger in the United States than in Germany? Dauth, Findeisen, and Suedekum (2014) stress one hypothesis: the weaker effect of China's import competition on the cross-regional disparities in Germany seems to be mainly driven by the structure of imports: Dauth, Findeisen, and Suedekum find that Chinese import competition in the most penetrated sectors such as textiles, for example, had negligible labor market effects and argue that Germany already tended to import those labor-intensive goods in earlier decades. When China subsequently became the world's dominant supplier of products in those sectors, China largely displaced imports from other countries (such as from Italy or Greece) but had little effect on production and jobs in Germany. With the formation of global supply chains, the nature of trade flows changes and imports can differ in their value added content. Concretely, behind the value added of Chinese exports may be U.S. or German intermediate-input imports to different degrees, potentially explaining different outcomes from the China shock. Feenstra (2017, Figure 8) reports that the foreign valued added in Chinese exports was about 30 percent in 2013 and that this foreign share has nearly doubled from about 17 percent since China's entry to the WTO in 2001. However, both the United States and Germany are among the three most important countries that provide value added behind Chinese exports in 2013, following only Japan, suggesting that value added contributions to Chinese exports are not a discerning feature between the United States and Germany.

Germany's trade deficit with China is less pronounced than that in the United States, because German exports to China have also risen sharply with China's WTO accession. The China shock has therefore not just increased import competition (particularly in such sectors as textiles, toys, or office and computer equipment) in Germany, but also opened new market opportunities that German exporters seized to an apparently larger degree (notably in sectors such as automobiles, specialized machineries, and electronic and medical equipment) than U.S. industries did. In contrast, Eastern Europe had a more substantive impact on cross-regional disparities within Germany. Dauth, Findeisen, and Suedekum document that Germany initially tended to export goods where the subsequent Eastern European rise was particularly strong, causing substantial relative job displacements from rising Eastern European import penetration.

Another potential explanation for the stronger consequences of the China shock for U.S. outcomes than regional disparities in Germany arises in light of the LATE interpretation of IV estimation, which both Autor, Dorn, and Hanson (2013) and Dauth, Findeisen, and Suedekum (2014) adopt. If local labor markets in the United States are less well prepared to withstand economic shocks through within-regional labor reallocation, nonemployment spells will be longer and welfare losses from transitions more severe. Concretely, unfavorable local industry mixes and a high concentration in few activities will hamper reallocations, regardless of the type of economic shock that impacts on a commuting zone. Finally, one further possible reason for the more adverse regional disparities in labor-market outcomes in the United States, following the China shock, is that Europe is considerably further ahead in terms of robot penetration, so that at-risk middle-income jobs were already gone from European manufacturing industries because of automation by the time the China shock hit. As Acemoglu and Restrepo (2017) document, robot penetration (industrial robots per thousand workers) in the United States is only at around the 30th percentile of European penetration. In 2014, manufacturers in Europe operate almost 2.5 industrial robots per thousand workers but in the United States only 1.6.

4.1.3 Offshoring and global supply chains

Baldwin (2016) describes the formation of global supply chains as a second *unbundling* in the specialization patterns behind globalization. He points out that a first wave of globalization unbundled consumption from production, when the steam engine and other significant improvements in transportation boosted foreign trade in the early 19th century. Since the later half of the 20th century, the formation of global supply chains unbundles production stages from each other and spreads the steps in the production process across distant locations. Conversely, reintroducing trade restrictions today becomes tantamount to erecting walls in the middle of factories.[22] When classic trade models are augmented to allow for between-industry linkages and input-output relationships, access to lower priced intermediate inputs and capital goods, and additional varieties of intermediate inputs and capital goods, downstream client industries benefit in that a market-size effect allows these downstream industries to expand under reduced production costs. While offshoring may directly displace domestic workers, the resulting lower production costs allow domestic firms to increase efficiency, expand production and market shares at home and abroad, and thus create new jobs for domestic workers. The net employment effects are therefore theoretically ambiguous and the subject of empirical study.

As discussed in Section 2 before, when sourcing from foreign countries becomes less costly, a high-income country can be expected to stop hosting less skill-intensive production stages and instead to start sourcing the associated intermediate inputs from lower-income countries where the production stages relocate. An expansion in offshoring opportunities can come about because of a decline in transport costs, or improvements in technology (especially communication technology), or by falling costs of cross-border investments (including reductions in the political risk and costs to the repatriation of corporate earnings). A fall in the costs of cross-border capital flows makes additional offshoring possible within firm boundaries by multinational enterprises, as suggested by Feenstra and Hanson (1996). I will turn to evidence on the formation and operation of multinational enterprises, and associated employment outcomes, further below. Offshoring can also take place at arm's length, when domestic firms source their intermediate inputs from independent foreign suppliers (or domestic wholesalers that in turn source from independent foreign suppliers), as suggested by Trefler and Zhu (2005). Offshoring understood in this way predicts a rise in the relative employment of skilled workers within industries.

[22] On the measurement of traded components that may cross borders multiple times in the presence of global supply chains, see Johnson and Noguera (2012), Koopman, Wang, and Wei (2014), and Johnson (2014).

Acemoglu, Autor, Dorn, Hanson, and Price (2016) extend the approach in Autor, Dorn, and Hanson (2013), where import competition from China and final goods trade were the focus, to allow for offshoring through input-output linkages within and across borders. Acemoglu, Autor, Dorn, Hanson, and Price (2016) use a cross-industry DD approach to study the consequence of both import competition in the final-goods markets and input-output linkages. They conclude that import competition from China was a major force behind both recent reductions in U.S. manufacturing employment and—through input-output linkages and other general equilibrium channels—weak overall U.S. job growth. For the estimation of trade-induced reallocations to non-manufacturing sectors and potential changes in aggregate demand, however, Acemoglu, Autor, Dorn, Hanson, and Price (2016) resort again to a cross-regional (commuting zone based) DD approach, confining conclusiveness to cross-regional disparities as a consequence of the China shock and leaving economy-wide effects unaddressed.

When it comes to evidence on wages, prior to the 1990s offshoring and trade in final goods together were associated with modest increases in the wage premium for skilled manufacturing labor (Feenstra and Hanson 1999; Harrigan 2000). Feenstra and Hanson (1999) found that between 15 and at most 40 percent of the observed change in the wage-bill share of nonproduction workers in the United States could be explained by offshoring, measured as imports of intermediate inputs. Competing evidence pointed to other shocks, including skill biased technological change, and reviewers of the literature, including Katz and Autor (1999), suggested that technological change played a more important role in the evolution of the U.S. wage structure in that decade. Looking towards services offshoring, beyond the offshoring of imports of intermediate inputs, Amiti and Wei (2005, 2009) do not identify a significant impact of service offshoring on labor-market outcomes, possibly because they work with a single aggregate of labor. A study of white-collar workers in the United States by Crinò (2010), for both manufacturing and services, suggests that offshoring had an impact. Crinò finds that service offshoring raises labor demand for the high-skilled but lowers labor demand for the medium and low-skilled. Within each skill group, there is a differential response depending on whether the tasks being performed are classified as routine, intensive in personal communication, and requiring the use of a computer. Crinò combines those tasks to compute a "tradability index" for each occupation and finds service offshoring to depress wages in more tradable occupations while benefitting non-tradable occupations.

Ebenstein, Harrison, McMillan, and Phillips (2014) combine industry-level data on trade and offshoring with individual-level worker data from the Current Population Survey (CPS) from 1984 to 2002. They find that occupational exposure to globalization is associated with significant wage effects, while industry exposure has no significant impact. They present evidence that globalization has put downward pressure on worker wages through the reallocation of workers away from higher-wage manufacturing jobs into other sectors and other occupations. Using the panel structure of the data, they find that workers' transitions between occupations due to trade led to real wage losses of 12 to 17 percentage points. Liu and Trefler (2015) also use the CPS, but for a later period, 1996-2007, and find that offshoring predicts an increase in worker transitions out of high-paid into low-paid occupations. However, they argue that in their framework wage changes for occupational switchers are not identified. For workers who stay within their occupation, they find that wages fell by slightly more than 2 percent, a substantively smaller number than that found by Ebenstein, Harrison, McMillan, and Phillips (2014).

In a follow-up study, Ebenstein, Harrison, and McMillan (2015) expand the nationwide household sample from Ebenstein, Harrison, McMillan, and Phillips (2014) in coverage and time span, using the Current Population Survey Merged Outgoing Rotation Groups for 1983-2008. Ebenstein, Harrison, and McMillan (2015) measure exposure to globalization using offshore employment at U.S. multinational enterprises and import flows penetrating U.S. goods markets, thus accounting for both in-house and arm's length offshoring to foreign suppliers. Ebenstein, Harrison, and McMillan (2015) can arguably comprehensively account for the movement of workers out of manufacturing and into lower wage services within and between local labor markets in the United States. Ebenstein, Harrison, and McMillan report a significant downward pressure on U.S. wages at the occupation level (but not for exposure measures at the industry level). They find that offshoring to low-wage countries—including China, Mexico, India, and others—is associated with wage declines for U.S. workers, and that the workers most affected are those performing routine tasks. For workers in routine occupations, a ten-percent increase in low-income affiliate employment at U.S. MNEs abroad is associated with a 0.7 percent decline in U.S. wages, whereas workers in less routine occupations were largely unaffected by offshoring.

Beyond the evidence on trade in final or intermediate goods, Liu and Trefler (2015) consider trade in services, which mostly takes the form of offshoring since only a small share of traded services goes to final consumers. Liu and Trefler (2015) use a model of trade in tasks within a general equilibrium framework of occupational choice to motivate their empirical approach. They turn to U.S. household data, the economy-wide Current Population Survey (CPS) for 1996-2007, and conduct their analysis at the worker-occupation level. Given a shift of the unit of analysis from local labor markets to individual household members, results in Liu and Trefler (2015) offer a partial corroboration of earlier findings in Autor, Dorn, and Hanson (2013). Liu and Trefler characterize occupations by their average pay and track workers across occupations. In terms of employment changes, transitions of workers "down' from high-paid to lower-paid occupations became 17 percent more frequent, whereas transitions "up" from lower- to higher-paid jobs rose by only 4 percent, and transitions to unemployment increased by almost 1 percentage point during 1996-2007. They argue that, in their framework, wage changes for occupational switchers are not identified but find that wages of occupational stayers fell by slightly more than 2 percent. Ebenstein, Harrison, and McMillan (2015) use the U.S. household sample (CPS Merged Outgoing Rotation Groups) for 1983-2008 and consider trade in final or intermediate goods. They measure exposure to globalization using offshore employment at U.S. multinational enterprises and import flows penetrating U.S. goods markets, thus accounting for both in-house and arm's length offshoring to foreign suppliers, but not for services trade. Ebenstein, Harrison, and McMillan estimate that offshoring in goods markets has led to the reallocation of workers away from high-wage manufacturing jobs into other sectors and other occupations, with apparently larger declines in wages among workers who switch than Liu and Trefler (2015) estimate for services.

4.1.4 Multinational production

Multinational enterprises (MNEs) are important mediators of world trade. Transactions data for the United States show that around ninety percent of U.S. trade are conducted by MNEs (when defined as firms with ownership of at least one related party abroad, see Bernard, Jensen, and Schott 2009). While it is hard to track trade flows across borders but within multinational enterprises, estimates suggest that around 40 percent to one-half of U.S. imports cross the country border but do not leave the boundaries of the trading firm (an early precise estimate is by Zeile 1997). The multinational enterprise is an informative unit of study because its own offshore operations are directly observed, in the form of its own employment, capital and sales abroad. The operation of a multinational enterprise is commonly called *multinational production*. The precision of information on multinational production itself, however, comes at the expense of two important limitations. First, an alternative to multinational production is to arrange cross-border supplier and client relationship at arm's length in the market place through trade—in other words, the ownership of production network members is an outside decision margin that involves a cost-benefit assessment that may need to be quantified. Multinational production is only one aspect of offshoring. Second, trade and multinational production can, a priori, either complement or substitute each other. If the foreign location of a production facility serves the network of MNE plants with access to relatively low-cost materials and components, then final-goods trade and multinational production can be complements in that the low-cost sourcing of intermediate inputs reduces production cost and thus likely raises global product-market shares. Conversely, if the foreign location of a production facility serves the client market so as to avoid transportation, tariff and other transaction costs related to the geographic proximity to clients, then trade and multinational production are likely substitutes. Multinational production therefore involves more than just offshoring.

Perhaps partly owed to these complexities, for a long time empirical research failed to find that the operation of MNEs affected relative employment outcomes across locations (studies that found no significant relative labor demand or wage effects include e.g. Slaughter 2000; Konings and Murphy 2006, for U.S. and European MNEs). Muendler and Becker (2010) argue that one reason for absent employment findings at MNEs was that it is important to discern the formation stage of the MNE and its subsequent operation. An MNE's labor demand responds to international wage differentials at the extensive margin, when the MNE first enters a foreign market with a foreign affiliate, and at the intensive margin, when the MNE operates its existing affiliates. Using a structural model of the related labor demand decisions within the MNE, Muendler and Becker show for German MNEs that the dominant employment response in Germany occurs at the extensive margin when an MNE starts in-house production abroad, and at distant low-income locations only the extensive margin matters. Home and foreign employment are substitutes within MNEs, so MNEs exploit labor-cost differentials across locations. However, inasmuch labor productivity in foreign locations advances and the cross-border wage differentials close with regional integration, the estimates also imply that a shrinking wage gap brings manufacturing jobs back to the high-income locations. As one example, Muendler and Becker predict that a hypothetical five-fold increase in Central and Eastern European wages in the year 2000, so as to narrow the wage gap vis-à-vis Germany by half, would have brought almost 200,000 counterfactual manufacturing jobs to Germany—around an eighth of the estimated home employment at German manufacturing MNEs in 2000. In practice, cross-country income differentials narrow only gradually.[23]

Harrison and McMillan (2011) estimate the relationship between wages at U.S. MNEs' foreign affiliates and U.S. manufacturing employment and find that in-house offshoring to low-wage countries substitutes for domestic employment but, when the MNEs' domestic and foreign plants report different economic activity, foreign and domestic employment at MNEs are complements. Combining their estimates, Harrison and McMillan (2011) argue that foreign employment at U.S. MNEs is associated with a quantitatively smaller decline in U.S. manufacturing employment than what Muendler and Becker (2010) found for Germany. Head and Ries (2002) study the composition of MNE employment in Japan. Considering relative skill demand, they investigate the covariation between offshore production of Japanese MNEs and the Japanese onshore skill intensity. Head and Ries report that additional foreign affiliate employment in low-income countries predicts more skill intensive workforces at Japanese MNEs' domestic plants, while this positive covariation diminishes or is absent for Japanese MNE employment in foreign high-income countries.

The study of well-defined effects at the micro level of firms requires a gradual build-up towards economy-wide outcomes. Labor-demand estimation within MNEs fails to account for important additional employment consequences. MNEs with labor-cost or market-access advantages after a foreign expansion can gain in product-market shares and consequently raise employment (the market-size effect). Their expansion in market shares can come at the expense of foreign competitors in the global product market, but also at the expense of domestic competitors who might consequently lose employment. Using an estimation approach similar to DD, Becker and Muendler (2008) compare job stability at German MNEs, which happened to achieve a foreign expansion of multinational activity, to job stability at German firms that did not have the chance to expand abroad. Becker and Muendler document that there is a separation rate of 14 percent among workers at MNEs, compared to less stable employment with an 18-percent separation frequency at German firms with no foreign affiliate, and argue that about half of the 4-percentage point difference is explained by the MNEs expansion abroad itself (whereas the other half arguably stems from the fact that more capable firms are more likely to be MNEs with foreign expansions). One interpretation of their result is that preventing firms from exploiting cross-country wage differences in-house would threaten additional jobs in the home economy. However, as with any approach similar in identification assumptions to DD, the 2-percentage point lower separation rates at MNEs are relative to firms with no foreign affiliate, so the foreign expansion of MNEs could either cause more employment stability at MNEs, or cause less employment stability at their

[23]Barro and Sala-i-Martin (1992) estimated that, among regionally integrated economies, per-capita GDP converges to steady state at a half time of around 35 years.

domestic competitors, or both, and the estimator cannot tell.

Boehm, Flaaen, and Pandalai-Naya (2015) construct comprehensive data at the firm level for trade transactions within and between U.S. MNEs, combined with detailed employment records, and empirically discern the substitution effect between imported foreign inputs and local employment at U.S. MNEs from the market-size effect, by which production cost reductions from offshoring boost firm size and employment. Their empirical approach rests on a quantifiable model, and allows them to overcome the limitations of a DD approach. The net employment impact of the counterbalancing substitution and market-size effects depends on the elasticity of firm size to production efficiency, which Boehm, Flaaen, and Pandalai-Naya quantify. Putting together the individual channels empirically, Boehm, Flaaen, and Pandalai-Naya estimate that manufacturing employment at U.S. MNEs, which account for two-fifth of the U.S. manufacturing decline, fell by 13 percent in the wake of rising imported foreign inputs. However, the firm-level approach still leaves unanswered the question to what extent the gross market-size effect of offshoring by U.S. MNEs came at the expense of foreign competitors in the global product market, or at the expense of other U.S. firms who might lose additional employment beyond the directly attributable employment decline at the U.S. MNEs.

In addition to potential adverse employment effects at domestic competitors of MNEs, there can be further employment impacts at domestic upstream suppliers of the MNEs that successfully expand multinational production abroad. If the MNE used to resolve the make-or-buy decision by sourcing intermediate input from domestic suppliers, and now sources the same intermediate inputs from its own foreign affiliates, there is an adverse employment consequence at domestic suppliers. In contrast, if the MNE used to source from arm's length foreign suppliers before, there is no employment consequence at domestic suppliers. To arrive at an overall quantification of employment consequences from the expansion of multinational production, additional evidence is called for. Studies such as Becker and Muendler (2008), Muendler and Becker (2010), Harrison and McMillan (2011) and Boehm, Flaaen, and Pandalai-Naya (2015) estimate the responsiveness of employment in multinational companies to labor cost conditions in foreign labor markets. By design, work at the level of individual MNEs tends to emphasize the elasticity of employment with respect to changes in globalization and labor cost conditions, and has difficulty producing estimates of aggregate impacts of foreign competition on employment.

4.1.5 Firm-level global integration

Firms can globalize by turning themselves into MNEs with multinational production networks that they own. Firms can also globalize by importing intermediate inputs and capital goods, either directly from foreign suppliers or through domestic wholesalers, and they can globalize by becoming exporters. Most firms that export also import (Bernard, Jensen, and Schott 2009).

At the heart of the firm-level approaches to globalization is a framework by Melitz (2003), in which heterogeneous firms sort into exporting. The framework's logic can easily be extended to sorting into additional forms of globalization beyond exporter status, including importer status and MNE status. A firm learns its capability—such as the productivity of the production technology or the appeal of its product variety to consumers—after the firm pays a sunk entry cost for sampling from a known distribution of capabilities. Once the firm has learned its capability, there are two more stages of decisions. First, the firm needs to decide which product markets it wants to enter, including the domestic market and foreign destinations. Second, the firm needs to operate production and manage sales to those markets under fixed costs of operation that differ between product markets. Exporting entails an additional fixed cost of establishing a beachhead in every destination country, so that fixed costs of exporting exceed the fixed costs of access to the domestic market. On the second operation stage, the firm sets the optimal prices for its variety in each of the product markets (under so-called monopolistic competition). In the distribution of capabilities there is exactly one level of capability at which a firm just breaks even—so that profits from operation just cover the fixed costs of operation. The more capable firms, with a capability above the threshold, earn strictly positive net profits (profits from operation less the fixed cost). Firms below the capability threshold for a given market do not enter that destination, so that there are non-exporters with capabilities below those of the exporters, and some firms do not enter any market. The expected value of the net profits by firms that do enter the market then cover the initial sunk cost in equilibrium. The Melitz (2003) framework matches crucial features of the empirically well established firm-level behavior: exporters and non-exporters coexist within both import-competing and export industries (Bernard and Jensen 1995), and innately more capable firms sort into exporting (Clerides, Lach, and Tybout 1998), so that larger firms are exporters, smaller firms are non-exporters.

By the Stolper-Samuelson theorem, classic trade in final goods raises the skill premium in skill-abundant (high-income) countries and reduces the skill premium in skill-scarce (low-income) countries. As stated before, this implication is at odds with evidence (Goldberg and Pavcnik 2007). Feenstra and Hanson (1999) preserved a notion of classic trade according to factor proportions but introduced offshoring to generate the possibility that the skill premium also rises in skill-scarce countries; they showed the success of the offshoring model in explaining the U.S. wage inequality evolution in the 1980s. Burstein and Vogel (2016) preserve classic trade but introduce heterogeneous firms within industries to generate the possibility that the skill premium also rises in skill-scarce countries; they use the enriched classic trade model with embedded firms to assess world-wide trade and wage inequality patterns in 2005-2007.[24]

Helpman, Itskhoki, and Redding (2010) develop a model of firms with heterogeneous capabilities, based on the heterogeneous-firm framework by Melitz (2003), and examine within-industry labor allocation across firms in the presence of labor market frictions. In the model, more productive firms pay higher wages, and exporting increases the wage

[24]For a theoretical exposition of the enriched classic trade model with embedded monopolistically competitive firms following Melitz (2003), see Bernard, Redding, and Schott (2007). In a variation of that model, Burstein and Vogel (2016) allow for oligopolistic competition with variable price-cost markups following Bernard, Eaton, Jensen, and Kortum (2003). Under variable markups, a trade opening can result in a simultaneous rise in the skill premium, as workers transition to more skill intensive exporters, and a fall in the relative price of skill-intensive goods—a potential magnifying effect.

paid by a firm with a given capability. Depending on labor market tightness (the proportion of vacancies to unemployed workers) and the elasticities of screening costs for worker abilities, in their model unemployment can increase or fall with further globalization. Felbermayr, Prat, and Schmerer (2011) consider a version of the Melitz trade model with search unemployment and show that, in the long run, trade openness lowers the rate of unemployment. In terms of relative employment, Helpman, Itskhoki, and Redding (2010) show that labor demand for more able workers is relatively higher at more capable firms, which are also exporters. I return to a fuller account of the labor-market functioning in subsection 4.3 on matching below.

When it comes to employment of skills in terms of observed educational attainment, Berman, Bound, and Griliches (1994) document that, during the final decades of the past century, the relative employment of skilled workers increased in all manufacturing industries, and these within-sector changes accounted for the majority of the rise in the aggregate employment of skilled relative to unskilled workers in manufacturing. Behind these within-industry changes can be firms and their plants. In a study of trade impacts on U.S. manufacturing plants, Bernard, Jensen, and Schott (2006) estimate that import penetration from low-income countries including China accounted for 14 percent of the total decline in manufacturing employment of 675,000 workers that occurred during the final decades of the past century (between 1977 and 1997).

That period of study precedes the 2000s, when China exhibited the strongest increase in net exports. However, Bernard, Jensen, and Schott (2006) find a relatively high sensitivity of employment to import competition. Whereas during the final two decades of the last century import penetration from low-income countries increased by less than 10 percent of 1 percent (less than 0.1 percent), during the early 2000s the annual increase in import penetration from China alone amounted to one-half of a percent (0.50 percent Autor, Dorn, and Hanson 2013). Had the elasticity of employment declines with respect to import been the same in the 2000s as it was in the two preceding decades, the job losses estimated by Bernard, Jensen, and Schott (2006) in the cross-plant sample would have exceeded those found in the cross-regional analysis by Autor, Dorn, and Hanson (2013). One reason why the cross-plant estimation may produce higher estimates of the impact of imports on employment is that workers may separate from declining plants and transition to other plants in the same local labor market and the same industry. The isolated study of smaller units, such as plants, requires an account for cross-plant outcomes to arrive at economy-wide inference. Similar to the case of MNEs, to achieve an overall quantification of employment consequences from plant-level responses, additional evidence is needed.

Burstein and Vogel (2016) consider two skill groups: high skilled (college educated) workers and low skilled workers (with less schooling). Burstein and Vogel specify that more capable firms use more skill intensive production technologies within each industry, industries vary in skill intensity so that the average firm differs across industries, and countries differ in skill abundance. As countries become more open to trade in final goods, workers of both skill levels reallocate to the country's comparative advantage industries, raising the skill premium in skill-abundant countries and reducing the skill premium elsewhere. Workers of both skill levels also reallocate to more capable and skill intensive firms within industries, and to more skill-intensive industries in all countries, raising the skill premium in all countries. For skill-scare countries, the cross-firm reallocation channel can overturn the classic wage inequality reduction. Burstein and Vogel parameterize the model for 60 countries using data in 2005-2007 on firms and industries. Holding constant endowments and technologies, they use the model to carry out counterfactual experiments. In one experiment, countries are driven to autarky by raising trade costs. A move from autarky back to the 2005-2007 trade levels raises the real wage for both skilled and unskilled workers in all countries. At the same time, the skill premium increases in almost all countries, including the skill-scarce countries (the Russian Federation is the only exception). This means that the wage effects of the within-industry reallocations, including firm selection into exporting and worker reallocations to exporting firms, overwhelm the Stolper-Samuelson effects. The increases in the skill premium vary across countries, being larger in more open economies and in countries that export skill-intensive products. The rise in the skill premium is largest in Lithuania, amounting to 12 percent, compared to an average of 5.1 percent. It rises by 2 percent in the United States and by 0.5 percent in Brazil. One important consequence is that globalization can raise the relative employment of skilled workers within industries even when it raises the skill premium.

In the Burstein and Vogel model, more capable firms are assumed to employ a larger proportion of workers with observably higher skills, such as a college education. Therefore, larger firms pay higher wages. However, larger firms also pay higher residual wages: identically skilled workers of the same gender and with the same observable level of education and labor market experience earn higher wages at larger firms. In the context of firm-level globalization, it is a common finding that MNEs, importers, and exporters are larger, pay higher wages and pay higher residual wages than domestically owned firms, non-importers and non-exporters. An explanation of firm-level difference in residual wages requires an account of labor market frictions. I therefore return to those frictions in subsection 4.3 on matching.

4.1.6 Labor-market conditions

To what extent are inflexible labor markets or import competition to blame for the observed declines in employment? The frequent use of IV estimation in the context of trade shocks, and the intricate relationship of outcomes with local conditions under the LATE interpretation, suggest that labor-market conditions are an important aspect of the predicted outcomes. A long list of empirical studies has found labor movement across U.S. cities and states in the aftermath of changes in regional labor demand to be slow and incomplete (Blanchard and Katz 1992; Glaeser and Gyourko 2005). In the context of adjustment to trade shocks, structural estimates of mobility costs across both regions and occupations suggest substantive losses from moving (Artuç, Chaudhuri, and McLaren 2010; Dix-Carneiro 2014), partly ascribable to the depreciation of human capital upon industry and occupation transitions. Using U.S. household level data from the CPS for 1975-2001, Artuç, Chaudhuri, and McLaren (2010) estimate that workers' moving costs between industries can amount to 13 times the average annual wage earnings, with a standard deviation about seven times the average wage.

Subsequent research has further looked into specific components of the moving costs. Feenstra, Ma, and Xu (2017) posit that the U.S. response to the China shock was sluggish because house values declined with adverse economic conditions in local labor markets, as did employment prospects, and may have hampered worker mobility. However, in a study using detailed credit reports for individual households, combined with loan-level mortgage data, Demyanyk, Hryshko, Luengo-Prado, and Sørensen (2017) find that not even negative home equity was a significant barrier to job-related mobility because, in 2005-2009, the benefits of accepting a job elsewhere outweighed the moving costs in a dynamic model of housing, consumption, employment, and location choice when calibrated to the estimated parameters. Beyond moving costs, Caliendo, Dvorkin, and Parro (2015) allow for a menu of product transportation costs between regions and countries as well as geographic forces including local productivities, agglomeration forces and amenity benefits. Under this broad set of parameters, moving costs decline in importance. Caliendo, Dvorkin, and Parro (2015) estimate that the costs from labor reallocation frictions reduce the long-term aggregate welfare gains by only 2.5 percent (and welfare gains are only on the order of a few percentage points of GDP).

It is hard to realistically ponder how trade shocks would play out differently in labor markets with varying degrees of flexibility, even in structural models with multiple parameters that capture elasticities of responses but not necessarily the primitives behind workers' employment and location choices. The context of labor-market and trade reforms in Colombia, which followed each other in short time during the late 1980s and early 1990s, provides a setting that can illuminate interactions between trade shocks and changing labor-market institutions. Coşar, Guner, and Tybout (2016) use Colombian plant data to estimate an open economy dynamic model that links trade to job flows in the presence of firing costs that alter plant-level dynamics and job turnover. They conduct counterfactual experiments and find that Colombia's trade opening increased welfare but at the expense of higher unemployment and heightened plant-level employment volatility. However, Colombia's labor market reforms dampened the rise in unemployment and aggregate job turnover.

Just as with employment, wage outcomes in response to economic shocks also depend on labor-market conditions. In IV estimation, those labor-market conditions closely relate to the measured effect of the shock. Krolikowski (2017) documents that U.S. workers who suffer a job separation experience surprisingly large and persistent earnings losses, regardless of the nature of the shock that caused their displacement. Krolikowski (2017) proposes a model for this empirical regularity, combining the common search framework with a "job ladder" and increased separation rates for recently hired workers. The notion of a job ladder captures the idea that workers are better suited to some jobs than others, and it takes time for workers to find the jobs for which they are well suited. When calibrated to U.S. labor-market data, Krolikowski's model captures the magnitude and the duration of earnings losses for displaced worker. (The model also matches employment-to-nonemployment and employer-to-employer transition frequencies by tenure, the empirical decomposition of earnings losses into wage cuts and employment effects, the observed wage dispersion, and the distribution of wage changes around a nonemployment event.)

In related empirical work, but with a focus on both globalization and technological change, Ebenstein, Harrison, and McMillan (2015) use the U.S. household sample (CPS Merged Outgoing Rotation Groups 1983-2008) to compare workers who change occupations to workers who switch sectors but remain in the same occupation within manufacturing. For industry-switchers (from textiles to steel, for example) who keep their occupation they find no significant wage change from displacement. The highest real wage declines occur for workers who leave manufacturing and also switch occupations. On average, workers who leave manufacturing and change occupation experience a real wage decline of four percent. Those results are also consistent with findings in Kambourov and Manovskii (2009), who report that wage declines are particularly large among workers who switch occupations. This body of evidence suggests an important role for occupation-specific human capital in workers' wage profiles and welfare changes as a consequence of economic shocks. Kambourov and Manovskii (2009) elaborate that frictions in occupation switching may contribute to rising U.S. wage inequality, in particular if younger workers have to switch more frequently and thus forego benefits from occupational tenure more than workers in previous decades.

4.2 Nature of work

As globalization progresses, the nature of work changes. Novel data on occupational details allow researchers to characterize work by the tasks that a worker performs at the workplace.[25] One set of task definitions characterizes *how* a worker performs the job. Much empirical research considers dichotomous categories of these task types, including routineness (the incidence with which work steps are repeated), manual work intensity (the extent to which work steps are performed by hand), and codifiability (the frequency with which work steps can be summarized in explicit rather than tacit instructions). Some data also provide information on task types that characterize *what* a worker does at the workplace, including whether or not she manufactures or produces goods, conducts research, organizes or plans own and others' work, controls machines or technical processes, or tends to others as an interpersonal service.

Workplace tasks are an important, employer-driven characteristic of the labor market. Most prominently, a decline of employment in routine occupations has been documented for many industrialized countries since the late 1980s (Goos and Manning 2007; Goos, Manning, and Salomons 2009; Acemoglu and Autor 2011; Cortes, Jaimovich, Nekarda, and Siu 2014). Shifts in the task composition of work have also been shown to relate closely to other economic changes, including the relative drop in demand for workers in the middle of the wage distribution (Autor, Katz, and Kearney 2006; Goos, Manning, and Salomons 2009), technological change (Cortes, Jaimovich, and Siu 2016a), the employment of immigrants (Ottaviano, Peri, and Wright 2013), and the offshorability of jobs (Leamer and Storper 2001; Levy and Murnane 2004; Blinder 2006). The assignment of tasks in an open economy, and the implications for welfare and wage inequality,

[25]Data sources include the U.S. Occupational Information Network (O*NET), a database of occupational definitions, and the German Qualifications and Career Survey (BIBB-BAuA), an occupational survey of about one percent of the German workforce conducted in seven-year intervals.

have been studied from a theoretical perspective in industry-level models, including the Heckscher-Ohlin (Grossman and Rossi-Hansberg 2008, 2010) and the Ricardian framework (Rodríguez-Clare 2010; Acemoglu and Autor 2011) and, more recently, at the plant level (Becker, Egger, Koch, and Muendler 2017).

Becker and Muendler (2015) use the German Qualifications and Career Survey and describe how tasks at German workplaces evolved between 1979 and 2006. The authors relate the task evolution to the globalization of the German economy over the period. While task categories of *how* workers perform their jobs—such as the routiness and codifiability of their work steps—retain a similar frequency over the three decades, there are considerable changes in *what* tasks workers perform. Most notably, the workplace activities as to *what* workers do on the job are naturally cumulative and do cumulate markedly over time. Multi-tasking progresses fast. By 2006, German workers execute six times more activities than they used to perform in 1979 (activities include tasks such as manufacturing goods, conducting research, organizing others' work, performing legal work, or controlling equipment). Importantly, the German workplace survey queries workers' tasks and their time variation within industries and occupations, and the statistically dominant change in task compositions occurs indeed within industries and occupations. In other words, those industries and occupations expand most that are initially less intensive in the faster growing tasks. Imputing the potential task content of import flows, which compete with German domestic products, Becker and Muendler describe how the task content of imported tasks shifts over time. They document that pure production activities become relatively more dominant in German imports, while in German workplaces coordination related tasks gain importance, including tasks such as organizing and planning, consulting, and computer programming as well as work under deadlines and adaptation to changing situations. These task changes are more pronounced in regional labor markets within Germany that exhibit less labor-market tightness (a smaller proportion of vacancies to unemployed workers).

Beyond such descriptive findings, theory-based estimation has established clear links between workplace tasks and employment outcomes. Ottaviano, Peri, and Wright (2013) implement an estimable version of the Grossman and Rossi-Hansberg (2008) model and study empirically how tasks are assigned to workers with heterogeneous ability. Ottaviano, Peri, and Wright use employment records from the American Community Survey (ACS) 2000-2007, combined with task information from the U.S. Occupational Information Network (O*NET), and offshore employment in affiliates of U.S. multinational enterprises from the U.S. Bureau of Economic Analysis (BEA), which they scale proportionally by the overall imports shipping to these U.S. multinational companies so as to account for arm's length offshoring in addition to in-house offshoring. Ottaviano, Peri, and Wright present industry-level evidence consistent with the idea that tasks in the middle of the complexity range are assigned to offshore workers, while the least complex tasks are performed by immigrant workers onshore and the most complex tasks by native-born U.S. workers onshore. In the Grossman and Rossi-Hansberg (2008) framework, the home economy captures quasi-rents from offshoring tasks to low-wage workers abroad while employing their remote labor services at the productivity level of the home economy—a productivity effect. Ottaviano, Peri, and Wright's results suggest that offshoring, as well as immigration, do improve industry efficiency similar to the theoretical productivity effect.

Wright (2014) revisits the productivity effect in a trade-only framework and discerns occupations by worker skill intensity: low-skilled production jobs and high-skilled nonproduction jobs. His quantitative objective is to evaluate whether the magnitude of the productivity effect compares favorably to the magnitude of welfare losses from worker displacement, or not. As regards employment, Wright's version of the Grossman and Rossi-Hansberg (2008) model implies that offshoring tasks has a direct worker displacement effect (reducing employment), a product-market expansion effect under the productivity gains from offshoring (expanding domestic employment), and a substitution effect between factors and tasks (with an ambiguous employment effect). Wright (2014) constructs an index of offshorability from the U.S. Occupational Information Network (O*NET) data and uses U.S. manufacturing data for 2001-2007 to test whether low-skilled production and high-skilled nonproduction job losses respond differently in industries that exhibit greater offshorability. More low-skilled jobs are lost when the tasks performed in an industry are predominantly routine, the converse occurs for high-skilled jobs—a result consistent with theory. The gain of high-skilled jobs more than offsets the loss of low-skilled jobs, and Wright (2014, p. 76) concludes that "the aggregate effect of offshoring is estimated to have led to a cumulative increase in aggregate employment of 2.6 percent over the period 2001 to 2007, a relatively minor effect."

Rather than overall employment effects, Becker, Ekholm, and Muendler (2013) study relative demand for different worker groups. They use data for German multinational enterprises, German social security records on workers matched to those MNE data, and task information from the German Qualifications and Career Survey. They show that multinational production (in-house offshoring) is associated with a statistically significant shift towards more non-routine and more interactive tasks in Germany, and with a shift towards more skilled workers. Importantly, the shift towards more skilled workers is in excess of what is explained by changes in the task composition of work. This finding suggests that, while changes in tasks are closely associated with globalization and while offshorable tasks tend to be carried out by low-skilled rather than high-skilled workers, multinational production affects the domestic matching of workers to employers by skill independently of task assignments. At German plants, expansions of multinational production explain between 10 and 15 percent of the observed changes in employment shares of highly educated workers and measures of non-routine and interactive tasks (as measured by wage-bill shares).

Becker, Egger, Koch, and Muendler (2017) investigate workplace arrangements as a function of employer size. They document that larger plants and exporters organize production into more occupations, and that workers at larger plants and exporters perform fewer tasks within the occupations. Becker, Egger, Koch, and Muendler propose a workplace model, in which the plant bundles tasks into occupations and workers match to occupations. By splitting the task range into more occupations, the plant can assign workers to a narrower task range per occupation, reducing worker mismatch and raising worker efficiency. Embedding the workplace model into a Melitz (2003) framework, where employers face fixed span-of-control costs that increase with the occupation count, generates an economy in which inherently more capable (exporter) plants exhibit higher workforce efficiency, but at the expense of a more unequal within-plant wage dispersion. Those theoretical predictions are borne out in the German plant-worker data.

4.3 Matching

Harrison, McLaren, and McMillan (2011, p. 261) write in the abstract of their literature review: "The 1990s dealt a blow to traditional Heckscher-Ohlin analysis of the relationship between trade and income inequality, as it became clear that rising inequality in low-income countries and other features of the data were inconsistent with that model. As a result, economists moved away from trade as a plausible explanation for rising income inequality. In recent years, however, a number of new mechanisms have been explored through which trade can affect (and usually increase) income inequality. These include within-industry effects due to heterogeneous firms, the effects of offshoring of tasks, effects on incomplete contracting, and the effects of labor-market frictions. A number of these mechanisms have received substantial empirical support."

Common to many of the new mechanisms that link trade to labor-market outcomes is the idea of matching in the labor market—the sorting of diverse workers to heterogeneous firms. Importantly, the matching process in the labor market can unlock efficiency gains for the economy and therefore also plays an important economic role beyond the consequences for earnings inequality.

Matching can be based on unobserved worker attributes, such as unknown worker characteristics not reflected in years of schooling or years of labor market experience, and on unobserved employer attributes, such as innate but unknown firm capabilities. Matching is therefore also a natural process to relate to residual wages and residual wage inequality. Skills are predictors of individual earnings and include both unobserved components, such as ability or talent, and observed components, such as schooling and experience. Since at least the work of Mincer (1974), log wages have been decomposed into their observable part—based on the commonly observed predictors education, experience and gender—and an unobservable residual part, which arguably reflects unknown ability or talent. The observed component in the log wage explains only a fraction of the variation in log wages, while a large residual component drives what is termed "residual inequality" (Katz and Murphy 1992). Residual inequality has increased over time in the United States, and the sources of this increase are subject to debate. Lemieux (2006) attributes the rise in residual inequality to episodic events and compositional changes. In contrast, Autor, Katz, and Kearney (2008) conclude from their data work that the rise in residual inequality was not transitory. They find that between 1973 and 2005 the 90/10 percentile ratio of wages of males increased by 41 percent, of which they ascribe 14 percent to residual inequality (about two-thirds of the residual inequality increase is attributable to the rise of the 90/50 percentile ratio). Rising residual wage inequality played a major role in rising wage dispersion in other countries, too. In Sweden, for example, residual wage inequality accounted for 70 percent of wage dispersion in 2001 (measured as the standard deviation of log wages) and it contributed 87 percent to the rise in wage inequality between 2001 and 2007 (Akerman, Helpman, Itskhoki, Muendler, and Redding 2013, Table 3). Similarly, residual wage inequality accounted for 59 percent of wage dispersion in 1994 in Brazil, where it also explained 49 percent of the rise in wage inequality between 1986 and 1995 (Helpman, Itskhoki, Muendler, and Redding 2017).

The idea of matching has a long tradition in economics. Matching occurs in the assignment of sellers to buyers, creditors to borrowers, managers to projects, and workers to employers. In the context of marriage, which matches spouses to each other, Becker (1973) first formulated conditions for *positive assortative matching.* Positive assortative matching in the context of marriages means that the additional gain from one spouse's positive characteristic increases with the other spouse's positive characteristic, a property also known as supermodularity. Becker (1973) then showed that the realized marriages in a competitive marriage market exhibit positive assortative matching. The logic carries over from marriage markets to other contexts of matching, including the matching of workers to firms. Think of a worker's positive characteristic as her ability. Think of the employer's positive characteristic as the firm's capability. Then, in a competitive labor market, the realized matches between workers and firms exhibit positive assortative matching. In an extension of the marriage-market idea, a realistic matching process in the labor market needs to allow many-to-one matches, so that multiple workers can each have one job at the same firm.

Matching in the labor market can take many forms. Workers can be matched with coworkers, with managers, with capital equipment, with firms, or with industries, and the consequences of trade for wages can depend on these particular features. Most applied theory models with matching assume that the natural logarithm of the value of a match exhibits complementarity. This property is also known as log supermodularity. Log supermodularity implies that a marginal increase in the characteristic of one party raises the marginal value of the other party's characteristic proportionately more than the value of the match.

Changes in matching can worsen or alleviate wage inequality. In particular, if a change in product market conditions leads to higher-ability workers matching with more capable firms, then the relative wage gap between any pair of workers with different ability is larger than before, raising wage inequality. Stronger positive assortative matching brings out the complementarity forces between worker abilities and firm capabilities. Beyond its effect on wage inequality, matching generates surplus (a match value), so that stronger positive assortative matching can raise an economy's welfare. Globalization can therefore impact both earnings inequality and economy-wide welfare through its influence on the assortative matching of workers to firms.

A large body of literature considers the employer-employee matching process, both theoretically (see e.g. Legros and Newman 2002; Eeckhout and Kircher 2011) and empirically (Burdett and Mortensen 1998; Postel-Vinay and Robin 2002; Cahuc, Postel-Vinay, and Robin 2006; Postel-Vinay and Turon 2010). The potential efficiency gains from improved assortative matching have received according attention in the trade literature (Costinot and Vogel 2010; Sampson 2014). Several studies highlight trade-induced change of match quality as a key aspect of trade in terms of welfare, employment and wage inequality (Amiti and Pissarides 2005; Davidson, Matusz, and Shevchenko 2008; Davidson, Heyman, Matusz, Sjöholm, and Zhu 2014). More recent studies have started to complement the analysis of cross-industry and cross-firm matches with an analysis of within-firm matches.[26] Larch and Lechthaler (2011) study the assignment of workers across plants within multinational firms, and Bombardini, Orefice, and Tito (2015) investigate the permissible ability ranges of workers at firms when worker-firm matches are formed. Becker, Egger, Koch, and Muendler (2017) point out that an additional source of efficiency gains for employers is to improve match quality by narrowly assigning tasks to workers with the best fit to those tasks (a core ability within the occupation's task range).[27]

4.3.1 Final goods trade and matching

Ohnsorge and Trefler (2007) propose a model to study trade and labor-market outcomes in an assignment setting, in which workers sort across industries. Costinot and Vogel (2010) specify a set of general primitives and develop a variant of the Heckscher-Ohlin model that features multiple industries, ranked by their capability, and multiple types of workers, ranked by their ability. Workers then match to industries in a competitive market. One interpretation of their model is that industries produce intermediate inputs that are traded internationally, and every country uses these intermediate inputs to assemble its own final consumer goods. As Helpman (2016) points out, another interpretation is that workers are assigned tasks (in this view an industry is relabelled to be a task), and tasks are combined to produce the final consumer goods. The worker's ability and the industry's capability are complementary in a production function that is log supermodular. As a consequence, the competitive market outcome results in positive assortative matching: more able workers are matched to more capable industries in every country.

International trade has an unambiguous effect on wage inequality when the countries differ in factor endowments with a clear ranking of relative ability abundance. Suppose that, for any two ability levels, one country has relatively more workers with the higher ability (the so-called monotone likelihood ratio property). Costinot and Vogel show for the case when high ability workers are strictly more abundant in one country compared to the other that trade opening improves the positive assortative matching for all workers in the high-ability abundant country but worsens the matches for all workers in the low-ability abundant country. In other words, in the skill-abundant country trade leads to a reallocation of workers between industries so that every worker is employed in a more-capable industry afterwards, and the opposite happens in the low-skill abundant country. As described at the outset of this section, when workers are matched with more capable employers the wage gap between any higher-ability and any lower-ability worker widens, resulting in more wage inequality. Trade therefore raises wage inequality in the high-ability abundant country and reduces wage inequality in the low-ability abundant country. This finding may be viewed as a generalization of the Stolper-Samuelson Theo-

[26] In a review of the literature on the structure of wages within and across firms, Lazear and Shaw (2009) conclude that the wage structure appears to be more dependent on firm- or within-plant sorting of workers to occupations than on sorting of workers to firms or plants.

[27] An interpretation related to the core ability of workers, most suitable for specific tasks, is that human capital is occupation specific. Kambourov and Manovskii (2009) and Sullivan (2010) provide empirical evidence on occupation-specific human capital.

rem, which applies to the two-factor and two-industry case under classic labor markets. As the quote by Harrison, McLaren, and McMillan (2011) at the outset of this section suggests, "the 1990s dealt a blow to traditional Heckscher-Ohlin analysis," so the baseline results in Costinot and Vogel are arguably of limited empirical relevance. However, it is conceptually reassuring that a shift of primitives from classic labor markets to matching preserves the well established implications of relative endowments (factor proportions) for the trade pattern.

The Costinot and Vogel (2010) framework provides analytic tools to investigate the implications of higher moments of the worker ability distribution, beyond the classic view on factor proportions (which relate to the first moments of the ability worker distribution). Countries can differ in the dispersion of factor endowments, that is in the variance of worker ability while identical in factor proportions, leading to a trade pattern that Costinot and Vogel call North-North trade. Under specific conditions, North-North trade worsens the matches of low-ability workers (workers with low ability get matched to less capable industries) and improves the matches of high-ability workers (workers with high ability get matched to more capable industries) in the country with a more diverse factor endowment (a more dispersed worker ability distribution), and the opposite occurs in the country with the less diverse factor endowment.

Grossman, Helpman, and Kircher (2013) consider a variation of the matching idea where firms bring together workers and managers, each group with their own ability distribution. Grossman, Helpman, and Kircher show how matching of workers to managers within firms determines the inequality of both workers' wages and managers' earnings. Log supermodularity in production ensures positive assortative matching within industries. In this generalization, sorting across industries does not necessarily satisfy positive assortative matching anymore. An economy's opening to trade can improve matching for one factor of production, either workers or managers, in both industries, and worsen matching in both industries for the other factor of production. This multitude of possibilities awaits empirical quantification. Importantly, an improvement in positive assortative matching can enhance an economy's welfare.

An empirical study by Davidson, Heyman, Matusz, Sjöholm, and Zhu (2014) is one of the first to analyze the response of matching to progressing globalization. Davidson, Heyman, Matusz, Sjöholm, and Zhu combine multiple data sources from Statistics Sweden to form linked firm-worker data for Swedish industries over the period 1995-2005. It is a challenge in this literature to elicit measures of unobserved worker ability beyond the observed worker skills (such as education and labor-market experience). One widely used empirical approach is to extract both a time invariant worker fixed effect, to capture unobserved worker ability, and a time invariant employer fixed effect, to capture unobserved firm capability, from a Mincer (1974) regression of log wages on observed worker skills (following Abowd, Kramarz, and Margolis (1999) and the corrected method in Abowd, Creecy, and Kramarz (2002)). However, separate identification of both a worker fixed and an employer fixed effect needs to rest on the assumption that matching between workers and firms is random, conditional on the observed worker characteristics. Yet, matching theory itself, which the empirical work aims to test and quantify, implies that competitive markets lead to non-random (positive) assortative matching. In other words, the Abowd, Creecy, and Kramarz (2002) method is inconsistent with the matching theory it is used to address, suggesting that results on measures of unobserved worker ability may be less reliable than results on the observed worker characteristics.[28]

Davidson, Heyman, Matusz, Sjöholm, and Zhu nevertheless argue that their preferred measures of matching and trade openness imply that globalization improves the matching between workers and firms in industries with greater comparative advantage. In their baseline empirical results, lower import tariffs in Sweden result in better worker-firm matching in more than four out of five industries and the welfare gains can be substantial. For the industries with the largest comparative advantage, matching improves on average by 1.15 standard deviations when tariffs fall by one standard deviation. Davidson, Heyman, Matusz, Sjöholm, and Zhu conclude that there may be significant gains from globalization that have not been identified in previous work, as globalization can improve the efficiency of the matching process in the labor market. Their results remain largely unaltered after controlling for technological change, suggesting that globalization exerts a direct impact on the matching process.[29]

4.3.2 Offshoring, multinational production and matching

A natural extension of trade theory with matching in the labor market is to consider offshoring and multinational production. Costinot, Vogel, and Wang (2012) show how the introduction of offshoring (the formation of global supply chains) can alter basic insights from the Costinot and Vogel (2010) model of matching under final goods trade. Costinot, Vogel, and Wang (2012) extend the Costinot, Vogel, and Wang (2013) model of global supply chains, which has one factor of production, to the case of multiple worker types. It is instructive to consider wage inequality in low-income countries (the South). There is a parallel in the matching literature to implications of the Stolper-Samuelson Theorem in the Heckscher-Ohlin model for the South, where wage inequality falls with trade. Similarly, Costinot and Vogel (2010) show in the context of matching that trade reduces wage inequality in the low-ability abundant country. Introducing intermediate production stages and offshoring into the Heckscher-Ohlin model, Feenstra and Hanson (1996) point out for the South that wage inequality no longer falls but increases with trade. Similarly, Costinot, Vogel, and Wang (2012) introduce global supply chains into the Costinot and Vogel (2010) model with final goods trade only and show for the South that wage inequality need not fall but may increase with trade along the global value chain. Costinot, Vogel, and Wang (2013) consider a continuum of intermediate inputs and one final good

[28]Card, Cardoso, Heining, and Kline (2016) provide a review of the state of evidence on the separate identification of time invariant employee and employer fixed effects in Mincer regressions. Card, Heining, and Kline (2013) argue for German plant data, under a specific test, that endogenous worker mobility between jobs may not adversely affect estimation. Krishna, Poole, and Senses (2014) implement a more general testing procedure—based on the correlation of residual wages between jobs as proposed by Abowd and Schmutte (2010)—and show for Brazilian linked employer-employee data that the independence of wage outcomes from worker mobility is strongly rejected.

[29]Note that the observed directed improvement of worker-firm matches based on unobserved characteristics can also be viewed as invalidating the identification of unobserved worker ability from Mincer regressions in the first place.

and assume that production of the final good is sequential and subject to mistakes. With free trade in intermediate inputs, countries with lower probabilities of making mistakes specialize in later stages of the supply chain. Within that framework of global supply chains, Costinot, Vogel, and Wang (2012) show that offshoring has opposite effects on wage inequality among workers employed at the bottom and the top of the global supply chain.

The impact on wage inequality at the bottom of the South's income distribution is reminiscent of the Stolper-Samuelson effect. Low-skill workers are relatively more abundant in the South than on world average, which tends to raise their relative wages in the South. However, when global supply chains form, trade in intermediate inputs leads to an increase in inequality at the top of the Southern income distribution. Why do global supply chains change the inequality predictions? In the absence of offshoring (global supply chains), changes in wages reflect changes in the prices of the goods produced by different workers. Free trade makes the prices of the goods produced by high-skill workers relatively cheaper in the South compared to autarky, so wage inequality must fall in the South. In the presence of offshoring (global supply chains), by contrast, changes in wages also reflect changes in the prices of the intermediate goods used by these workers. Now, if free trade reduces the prices of the intermediate goods used by high-skilled workers in the South compared to autarky, then wage inequality in the South tends to increase.

Larch and Lechthaler (2011) study the assignment of workers across plants within multinational firms. A multinational enterprise (MNE) has an advantage in workers' assignment choices over national competitors because an MNE can post workers to plants in different countries. An MNE's multinational internal labor market ("labor pool") reduces the mismatch of its workforce—leading to lower mismatch, higher average productivity of workers, lower prices, higher output, and higher employment at MNE plants compared with plants at a national firm. These theoretical predictions for matching under multinational production, as well as those for matching under offshoring more broadly, still await empirical scrutiny.

4.3.3 Firm-level global integration and matching

A particularly promising extension of the standard trade approach to more realistic labor-market settings is related to firm-level globalization patterns and worker-firm matching. First, looking to the firm as the unit of analysis permits an account of firm heterogeneity within industries. Second, looking to the worker as the complementary unit of analysis in addition allows us to take seriously worker heterogeneity beyond the dichotomy of two groups with low-skilled or high-skilled individuals. Third, matching in the labor market can address labor market frictions such as unemployment, wage bargaining, and costly mobility.

A number of studies have combined the Melitz (2003) framework with frictions in the labor-market, based on different primitives, to address the potential impact of globalization on residual wage inequality. For a detailed description of the Melitz (2003) framework in general, recall subsection 4.1.5 above. The labor-market studies based on the

Melitz framework emphasize alternative mechanisms of wage determination, such as labor market frictions in the form of fair wages (Egger and Kreickemeier 2009; Amiti and Davis 2012), efficiency wages (Davis and Harrigan 2011), and search and matching (Davidson, Matusz, and Shevchenko 2008; Helpman, Itskhoki, and Redding 2010; Helpman, Itskhoki, Muendler, and Redding 2017). Combined with Melitz-style firm heterogeneity, each one of these mechanisms generates a wage distribution among workers with similar characteristics, because these models predict that within the same industry firms with better capabilities pay a higher wage to the same type of workers. The positive correlation between firm size and wages, called the wage-size premium, has been known for a long time (Oi and Idson 1999). Different flavors of the Melitz logic allow for variations in the globalization mode of the firms, so that more capable firms ship their final product to more export destinations, or import intermediate inputs from more source countries, or both.

Studying Indonesia's trade liberalization during the 1990s, Amiti and Davis (2012, Table 1A) document a clear ranking of firms from exporters, to importers, to firms that globalize in both dimensions: exporters paid 27.5 percent higher wages than non-exporters, importers of intermediate inputs paid 46.8 percent higher wages than non-importers, and firms that both imported intermediates and exported their final product paid 66.4 percent higher wages, consistent with firm-level trade theory. They found that reductions in Indonesian tariffs on product imports increased wages at exporters and reduced wages at non-exporters, as expected in theory because non-exporters suffer shrinking market shares under foreign competition. In contrast, Indonesian cuts in tariffs on intermediate inputs raised wages at large importers and reduced wages at non-importing firms, as firm-level trade theory predicts.

Going beyond effects on mean wages by firm type, Sampson (2014) studies wage dispersions by firm type. To generate a wage distribution, Sampson considers workers with diverse abilities and shows that this distribution differs between globalized and non-globalized economies. A firm's realized total factor productivity depends on its technology and its workers' abilities, while production exhibits log supermodularity. As a result there is positive assortative matching between firms and workers: more capable firms with more productive technologies match with more able workers and wage inequality results. Inequality of wages within an ability interval is higher when the workers are matched with more capable firms. Trade raises inequality among workers employed by exporters (in the upper tail of the overall wage distribution). If trade raises productivity of the least productive incumbents, then trade also raises inequality at the lower end of the wage distribution, but otherwise reduces low-end inequality.

Helpman, Itskhoki, and Redding (2010) developed a firm-level theory of trade and wages that incorporates search and matching in the labor market, where a firm can screen job applicants for their ability. A worker's ability to perform a specific job is a random draw from a known distribution. Firms post vacancies and unemployed workers search for jobs. A fraction of the unemployed workers are matched with a fraction of the vacancies, and at a cost firms can screen applicants in order to improve the ability mix in their workforce. While a firm can improve the mean ability of its workers by screening to a tougher standard, testing whether or not a worker exceeds an ability threshold, the firm cannot infer the exact ability level of a worker. Firms differ in their elemental capability as in the Melitz (2003) model, so more capable firm with a larger anticipated product market share have a stronger incentive to engage in screening and therefore end up with a higher labor productivity because of both their elemental capability and their more able workforce. As a consequence, equally able workers are paid different wages, with wage premia for those workers who are matched to more capable firms, which command larger market shares in the domestic product market and are more likely to export.

The Helpman, Itskhoki, and Redding (2010) model has important implications for the behavior of economy-wide wage inequality as globalization progresses. Suppose trade costs are so high that no single firm is an exporter, the economy is completely closed. Then the wage dispersion in the economy is driven by the underlying ability distribution of workers and its combination with the firm capability distribution through assortative matching, but there is no globalization-related wage differential between firms because all firms are non-exporters. Now think of the other extreme and suppose trade costs are so low that every firm is an exporter, the economy is fully globalized. Economy-wide wage inequality is then essentially the same as when the economy was completely closed. The reason is that the wage dispersion in the economy is again driven only by the underlying ability distribution of workers and its combination with the firm capability distribution through assortative matching, but there is no globalization-related wage differential between firms because all firms are exporters.[30] In a partly globalized economy, however, where some firms are non-exporters and others exporters, wage inequality is strictly higher than in the two extremes because now there is also a globalization-related wage differential between non-exporters and exporters. As a result, wage inequality follows a hump-shaped curve as the economy moves from completely closed to fully open. At some level of globalization, where a sufficient fraction of firms has globalized (has turned into exporters), the economy reaches peak inequality. From that point on the economy experiences a decline in wage inequality as globalization progresses and even more firms globalize. Whether an economy has or has not yet reached peak inequality is an empirical question.

The Helpman, Itskhoki, and Redding (2010) model yields compact log-linear estimation equations that describe an industry's cross-firm variation of wages, employment, and selection into exporting. Helpman, Itskhoki, Muendler, and Redding (2017) extend this theoretical model to allow firms to differ also in the fixed cost of exporting and in the ability to screen workers, so as to prepare the model for estimation with Brazilian linked firm-worker data. The matching and screening process in the model implies that average worker ability at a firm is conceptually not discernable from the elemental capability of a firm, so Helpman, Itskhoki, Muendler, and Redding (2017) do not use the Abowd, Creecy, and Kramarz (2002) method but instead only extract the firm fixed effect from Mincer (1974) regressions across workers (not separating out any fixed worker abilities).

[30] In the parametrization of Helpman, Itskhoki, and Redding (2010), where both worker abilities and firm capabilities are Pareto distributed, the wage inequality in the completely closed and the fully globalized economy is exactly the same. In the parametrization of Helpman, Itskhoki, Muendler, and Redding (2017), where worker abilities are Pareto but firm capabilities log normal, simulations show that the wage inequality in the completely closed and the fully globalized economy is quantitatively similar.

In Brazil, the firm-specific component of the log wage accounted for 39 percent of residual log wage inequality in 1994, and for 86 percent of the change in residual log wage inequality between 1986 and 1995. After estimating the model structurally for Brazil in 1994, Helpman, Itskhoki, Muendler, and Redding (2017) simulate the impacts of reducing fixed export-market entry costs and variable trade costs on wage inequality. They show that peak inequality occurs at a later point for variable trade costs, than for fixed costs, because variable trade costs (such as tariffs) affect both a firm's entry decision on the first stage and the operation decision on the second stage. Helpman, Itskhoki, Muendler, and Redding (2017) use the implied trade cost reductions from observed increases in the fraction of globalized firms and calculate that the model predicts a rise in inequality of about 2 percent from 1990 to 1994, when the Brazilian economy reached peak inequality, followed by a decline in wage inequality of about 1.5 percent thereafter. Compared to the data, these simulated trade-cost reductions account for about two-fifths of the inequality increase between 1990 and 1994 and about a quarter of the inequality reduction thereafter.

Further extensions of the Melitz (2003) model can also address within-employer wage inequality, which Becker, Egger, Koch, and Muendler (2017) find to account for 71 percent of residual log wage inequality in Germany in 1996-2014. They also show that more than half of residual wage inequality is within German plants' occupations (in fact within the plant hierarchies' occupations)—a common fact in many economies that lead Lazear and Shaw (2009) to argue that the wage structure appears to be more dependent on employer-internal sorting of workers to occupations than on worker sorting to employers. Becker, Egger, Koch, and Muendler embed a model of the internal labor market into the Melitz (2003) framework. In the internal labor market, a plant can bundle tasks into occupations and workers then match to occupations. By splitting the task range into more occupations, the plant can assign workers to a narrower task range per occupation, thus reducing worker mismatch and raising worker efficiency (reminiscent of specialization in Adam Smith's pin factory) and hence increasing the within-plant within-occupation dispersion of wages under an efficiency wage schedule to elicit worker effort. Placing this model of the internal labor market into the Melitz framework, where fixed span-of-control costs increase with occupation counts, Becker, Egger, Koch, and Muendler show that plants with higher elemental capability exhibit higher worker efficiency and a higher within-plant within-occupation wage inequality, and that economy-wide wage inequality is higher in the open economy for an empirically confirmed parametrization. Estimation of their model shows that a worker's average number of tasks in her occupation is inversely related to plant size, while the within-plant within-occupation wage dispersion is positively related to plant size. The principal selection of more capable firms into exporting remains a basic force in their version of the Melitz model (as documented empirically by Clerides, Lach, and Tybout 1998, e.g., and others). However, the feedback of exporting into worker efficiency through narrower task ranges at exporters is akin to a learning-by-exporting effect (for direct evidence on learning-by-exporting see, e.g., Crespi, Criscuolo, and Haskel 2008).

4.4 Consumer prices

The most basic gains from trade are gains to consumers, irrespective of any changes in production. Think of a simple endowment economy where fruits fall from trees, ready for consumption, with no labor or other input needed. If the relative endowments with fruit trees of different kinds differ between countries, then the trade of fruits between countries will move relative prices compared to when the economies where closed, and consumer theory shows that a change in prices allows all consumers to reshuffle their consumption baskets and raise their well-being. In addition, consumers who value the consumption of multiple varieties of products, gain from the availability of additional varieties with trade (Feenstra 1994; Broda and Weinstein 2006). These insights also affect the welfare assessment of technological change. We tend to characterize technological progress, such as the use of computers, information technology and the internet, as resource saving (in the workplace). However, the use of computers, information technology and the internet also allow consumers a more convenient and efficient shopping experience, raising living standards beyond real GDP growth (Hulten and Nakamura 2017).

Conventional trade theory considers so-called homothetic preferences, by which consumers maintain expenditures on the goods in their consumption basket in fixed proportions as they grow richer. These preferences mostly serve a conceptual purpose: they allow trade models to elicit production-side explanations of trade flows without confounding influences from demand distortions. Of course, in reality a person whose income doubles will not typically double the expenditures on all its consumer goods. For example, expenditures on more basic consumer goods, such as food items and beverages, typically decline in relative importance as persons become richer. Restated more formally, the income elasticity is relatively smaller for basic goods whose expenditure share in the consumption basket drops as incomes rise. In short, realistic preferences are non-homothetic. Fajgelbaum, Grossman, and Helpman (2011) analyze in theory rich interactions between inequality, shifting consumption choices within countries across income groups and between countries by per-capita incomes, and their feedback into prosperity in the presence of trade.

Fajgelbaum and Khandelwal (2016) conduct a global empirical analysis of trade-induced relative price changes and their consequences for persons at different levels of income. They estimate an extended model of country-to-country trade flows, called the gravity equation (which relates trade between country pairs positively to their respective sizes and negatively to their distance). Fajgelbaum and Khandelwal first derive additional terms to be included in the gravity equation so as to account for non-homothetic preferences, using a tractable demand system (the so-called Almost-Ideal Demand System) that offers a first-order approximation to any demand system. Non-homothetic preferences mean that consumers at different income levels within an economy spend different shares of their expenditure on goods from different source countries and different industries. Empirically, low-income consumers spend relatively more on goods that are more traded, whereas high-income individuals consume relatively more services, which are among the least traded sectors. Additionally, low-income consumers happen to concentrate spending on goods with a lower cross-price elasticity of substitution between source countries (such as agricultural products, food and beverages) and, under the specified demand system, price declines with trade are relatively faster in the goods with low cross-price elasticities given perfect competition. Taken together, larger expenditures in more tradeable sectors and a lower rate of substitution between imports and domestic goods lead to larger gains from trade for the poor than the rich (Fajgelbaum and Khandelwal 2016, p. 1117). In their most general specification of world trade, with multiple industries, they find a pro-poor bias of trade in every country: poor consumers are predicted to gain more from trade than rich consumers everywhere. On average across the countries in their sample, the real-wage gains from trade are 63 percent at the 10th percentile of the income distribution and 28 percent at the 90th percentile.

The importance of varying prices across consumer baskets at different quantiles of the income distribution suggest that the real wage w_i/P_i can vary substantively for different income groups. Only few studies consider the interaction between nominal wage outcomes and prices. There is research into select countries. Porto (2006) studies the distributional effects of MERCOSUR (a regional trade agreement among Argentina, Brazil, Paraguay and Uruguay) during the 1990s; Nicita (2009) extends Porto's approach by adding a link from trade policy to domestic prices and studies Mexico's trade liberalization 1990-2000; Marchand (2012) studies India's trade reforms between 1988 and 2000. In recent unpublished work, He and Zhang (2017) use a unified framework of matching in the labor and non-homothetic demand in the product market and estimate it for 40 countries using industry-level trade and production data. They find that trade cost reductions decrease the relative nominal wage of the poor and the relative price index for the poor in all countries. Similar to Costinot and Vogel (2010), the income channel benefits the poor more than the rich in low-income countries and the rich more than the poor in high-income countries. Similar to Fajgelbaum and Khandelwal (2016), lower trade costs increase real incomes and, therefore, decrease the relative demand for and the relative price of low-income elastic goods. On net, real-wage inequality falls everywhere. The findings for individual countries, and worldwide outcomes, depend on how tariff cuts and global price changes translate into local prices that are relevant for consumers. As Goldberg and Pavcnik (2007) argue, predictions crucially depend on the degree of pass-through from trade policy changes to product prices as well as the wage-price elasticities.

5 Technological Change, Labor Markets, and Prosperity

Today we feed 7 billion people on Earth with considerably fewer farmers than centuries or just decades ago. Agriculture may have undergone one of the most impressive labor productivity changes of all sectors of the economy over the past decades and centuries. As mechanization, more efficient crop cycles, and labor-saving innovations progressed in agriculture, where did all the farmers go? Today, technological change in the manufacturing sector is a foremost concern for employment, and a presumption seems to prevail that displaced manufacturing workers now may face inferior prospects than workers displaced from technological process in the past.

Ford (2015), a Silicon Valley entrepreneur, argues that today's industrial revolution may not unfold like earlier similar episodes. Yet, such concerns may in fact be far older than recent publications suggest. Acemoglu and Restrepo (2017, p. 1) cite Keynes in 1930 and Leontief in 1952: "'We are being afflicted with a new disease of which some readers may not have heard the name, but of which they will hear a great deal in the years to come—namely, technological unemployment' (Keynes, 1930). More than two decades later, Wassily Leontief would foretell similar problems for workers writing 'Labor will become less and less important. More and more workers will be replaced by machines. I do not see that new industries can employ everybody who wants a job' (Leontief, 1952)." In light of those dated warnings, it remains somewhat unclear why technological changes today would exert more severe adverse consequences, while similar predictions were not borne out in the past. If Say's law of markets broadly applies, then labor supply creates its own product demand and employment will result. Ford's main argument, why labor-market consequences of technological change may be more severe today than in the past, is based on the perception that past solutions to technological disruption, such as more training and education, may no longer be applicable. Paralegals, journalists, office workers, and even computer programmers, Ford points out, are poised to be replaced by smart software and robots may dislocate both blue- and white-collar jobs. Bill Gates, the founder of Microsoft Inc. and now a philanthropist, has proposed a tax on robots to fund adjustment programs for displaced workers. Despite these concerns, we have little systematic evidence on the equilibrium impact of new technologies.

Conceptually, it is perhaps the principle of comparative advantage itself that can give us the clearest guidance as to how humans and machines will interact in the near future. Machines, especially prospective machines with artificial intelligence, may outperform humans in every single activity—from physical manufacturing tasks, to production planning and marketing, to playing Go and Chess. But it is merely comparative advantage, not absolute advantage, that matters for specialization in exchanges and economic transactions.[31] When employers organize the interaction between workers and workplace equipment in the future, there will remain at least one task at which a human worker is *relatively* better than a machine, even if the machine is more productive at every task in an absolute sense. It will be at least that one task for which an employer will hire a hu-

[31]The idea that comparative advantage will govern human-machine interactions, in a coming economy with machines that command artificial intelligence, is a leading tenet behind a current book project by Ronald Davies of the University College Dublin, to whom I owe this insight.

man, and the wage that this human worker will command for performing the task is going to reflect the stupendous marginal labor productivity that a highly specialized and machine-empowered human worker exerts when generating surplus for the firm. It is therefore arguably the principle of comparative advantage, first formulated in the context of foreign trade by David Ricardo (1821) two centuries ago, that may offer the profound explanation as to why technological change through automation has not displaced aggregate employment in the past and will not do so in the future.

Katz and Murphy (1992) investigated whether the rise in the college wage premium during the 1980s was driven by labor supply, such as varying rates of schooling, or labor demand factors, such as trade, technological change, and structural shifts. Katz and Murphy concluded that the dominant cause was an increase in the relative demand for skilled workers, and especially for those with a college degree: "Although much of this shift in relative demand can be accounted for by observed shifts in the industrial and occupational composition of employment toward relatively skill-intensive sectors, the majority reflects shifts in relative labor demand occurring within detailed sectors. These within-sector shifts are likely to reflect skill-biased technological change" (Katz and Murphy 1992, p. 37). Subsequent research during the 1990s and 2000s continued to support their inference that skill-biased technological change was the main cause behind the widening wage gap between skilled and unskilled workers. Other factors, such as the decline in unionization rates, changes in the minimum wage, or deregulation of labor and product markets, were not found to contribute much to the rising college wage premium (see Bourguignon (2015, Ch. 3) for a review of evidence concerning these factors in several countries). In light of the firm-level literature and firm-worker matching (see the preceding section), an alternative explanation of within industry-occupation shifts in employment and wage outcomes is the between-firm dispersion of worker skills and wages.

In practice, much of the literature on demand causes behind wage inequality or other employment outcomes uses observed labor demand shifters, including trade in final goods or offshoring, and then considers the residual from the estimation to be exclusively driven by unobserved technological change. Such approaches, of course, ignore other potential demand shifters including changing expenditure patterns as incomes grow (Fajgelbaum and Khandelwal 2016). They ignore substitution to other types of labor, including non-native labor supply from migration. They ignore structural change, which has led to a significant deindustrialization trend in recent decades beyond the advanced, post-industrial economies (Rodrik 2016). Instead, all of the residual if often assigned to productivity change.[32] The idea of this section is to turn the burden of proof around, to review economic research that uses measures of observed technological change, and to see how far they themselves can push back the variation in the residual by attributing labor-market outcomes directly to variables of technological change.

[32] Examples of such residual arguments in the earlier literature on changes in relative wages include Katz and Murphy (1992) or Feenstra and Hanson (1999), a more recent example is Hicks and Devaraj (2015).

5.1 Technical change in isolation and the nature of work

Recent shifts in the nature of work include a strong decline of middle-skilled occupations that are intensive in routine work steps in the United States. This decline in middle-skilled occupations is also thought of as job polarization, by which low- and high-skilled occupations are faring relatively better. Cortes, Jaimovich, Nekarda, and Siu (2014) use matched individual-level CPS data and trace the worker flows responsible for the decline in U.S. middle-skilled jobs during the last 30 years. At the macroeconomic level, the driving work flow is not the separation from the declining occupations but, compared to earlier years, the main driver is a paucity of transitions from out of the labor force and from unemployment into routine employment. In other words, the middle-skilled routine jobs used to serve as entry occupations for transitions of persons out of work back into the workforce. At the same time as routine employment declines in the United States, non-routine manual employment expands (Autor and Dorn 2013; Mazzolari and Ragusa 2013). Autor, Levy, and Murnane (2003) argue that this process of job polarization is best explained by automation technologies that substitute for labor in routine tasks.

Michaels, Natraj, and Van Reenen (2014) test the hypothesis that information and communication technologies (ICT) polarize labor markets by increasing demand for the high skilled at the expense of the middle skilled, with little effect on low-skilled workers. Using data for the United States, Japan, and nine European countries from 1980 to 2004, they document that industries with faster ICT growth shifted demand from middle-educated workers to highly educated workers, consistent with ICT-induced polarization.[33] In their predictions, ICT can account for up to a quarter of the cross-country variation in demand growth for high skilled workers.

Acemoglu and Restrepo (2017) consider the use of robots as a particularly salient aspect of automation and show that exposure to robots is only weakly correlated with other aspects of automation—such as the use of software to process information, the installation of other types of information technology equipment, or the total capital stock—and final goods trade and offshoring.[34] Acemoglu and Restrepo report that between 1993 and 2007 the stock of robots in the United States and Western Europe increased fourfold. Acemoglu and Restrepo (2017) adopt a DD design for local labor markets, similar in spirit to the trade literature on the China shock (e.g. Autor, Dorn, and Hanson 2013), and compute the local exposure to robots, which is in turn predicted with the industry-level spread of robots in other advanced economies as an instrument. Acemoglu and Restrepo study the effect of increased industrial robot usage between 1990 and 2007 on U.S. local labor market outcomes. In their preferred specification, they estimate that one new robot in a local labor market reduces regional employment by 5.6 workers. Employment effects of

[33] Michaels, Natraj, and Van Reenen (2014) also find that trade openness is associated with polarization in some specifications, but the correlation of trade with high skilled demand is not robust to controlling for cross-border investment flows (mergers and acquisitions).

[34] An industrial robot is an automatically controlled, reprogrammable multipurpose machine. There are currently between 1.5 and 1.75 million industrial robots in operation in the world, and their number might increase to 4 to 6 million by 2025 (Boston Consulting Group 2015). Brynjolfsson and McAfee (2015) and Ford (2015) describe how industrial robots have impacted the labor market so far and how they can be expected to transform it further.

robots are most pronounced in routine manual, blue collar, assembly and related occupations and for workers with less than college education. Those results are from weighted regressions, after the top one percent of labor markets with the highest exposure to robots is excluded; otherwise effects are considerably less pronounced (Acemoglu and Restrepo 2017, Figure 8).

The cross-country evidence in Michaels, Natraj, and Van Reenen (2014) and the documentation of cross-regional disparities in robot adoption by Acemoglu and Restrepo (2017) are suggestive of potential nationwide changes. Cortes, Jaimovich, and Siu (2016a) provide a quantification. Cortes, Jaimovich, and Siu (2016a) document that the decline of middle-skilled occupations is primarily driven by the disappearance of routine jobs among workers in specific demographic groups, mostly young and prime-aged men with low levels of education who both become less likely to work in those occupations and become fewer compared to other demographic groups in the workforce.[35] On the labor supply side, increasing educational attainment and population aging in the United States reduce the fraction of workers with these demographic attributes. Cortes, Jaimovich, and Siu's attention lies on labor demand effects, however, and in particular on technological change as an explanatory variable. Importantly, labor supply alone cannot account for the labor-market experience of young and middle-aged male workers as these demographic groups experience a substantial increase both in non-employment and in employment in low-wage, non-routine manual occupations observed during the same time period.

To quantify the contribution of technological change to the decline in middle-skilled occupations, Cortes, Jaimovich, and Siu (2016a) adopt a neoclassic macroeconomic model and extend it to allow for endogenous workforce participation and occupation choices on the worker side and for endogenous technology adoption on the firm side. They use data on automation technology, so an interpretation of the model is that technology investments are for automation purposes. The model demonstrates analytically that advances in automation cause workers to leave routine occupations and sort more frequently into non-routine manual jobs and into non-employment. The reallocation from routine to non-routine manual work happens because adoption of automation technology drives down the marginal product of routine labor relative to the manual occupation, so there must be movement of labor out of the routine occupation until the marginal worker is indifferent between routine and the non-routine manual jobs. The framework clarifies an important tradeoff for society (Cortes, Jaimovich, and Siu 2016a, Proposition 1): while there are welfare gains from the automation-induced reallocation of workers from routine to non-routine manual occupations, an economic cost to society is the automation-induced failure of reallocations as some workers are displaced from employment to non-employment. Note, however, that the occupational reallocation is welfare enhancing for the economy as a whole but for individual workers the occupational reallocation amounts to a down-

[35] Cortes, Jaimovich, and Siu (2016b) consider the gender aspect more closely. They report that, since 1980, the probability that a college-educated man was employed in a cognitive/high-wage occupation fell, whereas for college-educated women the probability of working in these occupations rose, despite a much larger increase in the supply of educated women relative to men during this period. Cortes, Jaimovich, and Siu speculate that certain social skills that women possess more frequently may have become more important over time within these occupations.

grade (from routine to low-wage non-routine manual occupations) with potential real earnings losses unless the automation-induced price reductions raise real incomes more than the wage loss. When workers can endogenously choose their workforce participation in addition, then the tradeoff between welfare-enhancing occupational reallocation (from routine to non-routine manual occupations) and welfare-reducing displacement (from employment to non-employment) changes: the occupational downgrading from the routine to the manual occupation is less frequent, because some of the adjustment to automation results in fewer workers selecting into employment (Cortes, Jaimovich, and Siu 2016a, Proposition 2).

Cortes, Jaimovich, and Siu calibrate the model to the U.S. economy for 1979-2014 (or just 1989-2014), using matched individual-level CPS data, and feed in an automation shock under variations in parametrization. To measure the magnitude of the automation shock, Cortes, Jaimovich, and Siu infer capital-embodied automation technology from information and communication technology (ICT) equipment and account for the fact that, along a balanced growth path, capital expectedly grows and therefore measure the automation shock as the deviation of ICT capital from a balanced growth trend. Cortes, Jaimovich, and Siu ask much of the quantified model, namely that it deliver both the observed decline in middle-skilled occupations (the occupational reallocation to non-routine manual jobs) and displacement into non-employment at the same time (employment changes). The authors conclude that advances in automation technology on its own are unable to jointly generate changes in occupational shares and employment propensities that are quantitatively similar to those observed in the data for the relevant demographic groups. At most one-third of employment changes are explained by automation. But it is a tall order for any model to match two hard-to-model macroeconomic facts. In the dimension of occupational reallocation, the model can account for at most half of the observed transitions under a certain parametrization; under the same parametrization the model then accounts for just less than one-third of the fall in the employment rate.

The findings in Cortes, Jaimovich, and Siu (2016a), who consider technological change in isolation, suggest that automation explains one-third of the decline in middle-skilled employment, while Autor, Dorn, and Hanson (2013), who considered trade in isolation, argued that one-quarter of the decline in U.S. manufacturing employment is explained by final goods trade and import competition largely from China. Taking these two estimates together crudely, under the observation that middle-skilled occupations are concentrated in manufacturing, the two isolated studies jointly account for just over one-half of the salient employment decline. What may account for the other, dominant part of the employment decline that remains to be explained? Mazzolari and Ragusa (2013) point to shifting demand for low-skill intensive services: with rising inequality, high-income earners demand more than proportionally low-skill intensive childcare and gastronomy services, maintenance and contractor work, and the like. As a consequence, non-routine manual employment expands, but Mazzolari and Ragusa's explanation does not address the causes of the decline in routine employment. Charles, Hurst, and Notowidigdo (2016) point to a long-term structural decline in manufacturing employment, where middle-skilled routine occupations are concentrated, and emphasize the fact that a housing boom during the early 2000s tended to provide a temporary reprieve for middle-skilled routine work (a "mask" of the underlying manufacturing decline) that ended with the U.S. housing market collapse in 2007. Moreover, fluctuations in skill demand can result in a rematching of worker skills to occupations with lower skill intensities—a process to be discussed in subsection 5.2 below. Aspects of the structural decline in manufacturing employment are coming up in subsection 6.2.

When it comes to wages, common evidence across both industrialized and developing countries is often considered as suggestive of common technological change across the world. Berman, Bound, and Machin (1998) found for the 1980s that, across industrialized countries, most industries increased the proportion of high-skilled (nonproduction) wage bill to the low-skilled (production) wage bill despite rising or stable relative wages for high-skilled workers. The same manufacturing industries simultaneously increased relative wage bills for the high-skilled in different countries, while many developing countries showed increased skill premia—a global pattern not predicted by classic industry-level trade theories such as the Heckscher-Ohlin model (or for that matter by matching-based trade theories such as Costinot and Vogel 2010). The pattern is, however, consistent with common skill-biased technological change (SBTC) in industrialized and developing countries. Berman and Machin (2000) followed up on the evidence with a large country sample and documented that relative wage bills of high-skilled workers jointly increased in the manufacturing industries of 37 high-, middle-, and low-income countries during the 1980s. Concretely, within-industry skill upgrading (the increase in the wage-bill share of nonproduction workers) was pervasive across countries: the within-industry contribution to skill upgrading was large in 12 high-income countries (except Sweden), in 18 middle-income countries (except the Republic of Korea), and in all low-income countries (except Bangladesh). Moreover, the industry-level skill upgrading in all countries was positively correlated with skill upgrading in the same U.S. industries. Finally, industry-level skill upgrading in all countries—rich, middle-income and poor—was positively correlated with U.S. computer usage and OECD R&D intensity. In summary, changes in skill intensity were similar and wide spread across countries at different income levels and closely related to technology usage in industrialized countries—consistent with simultaneous global skill-biased technological change.

Relative wage and employment patterns for nonproduction and production workers changed in the 1990s and 2000s, with an increase in the relative wage but a decrease in the relative employment of skilled labor (Feenstra 2017). Autor, Katz, and Kearney (2008) point to a pattern of polarization, by which workers in the high-skilled and low-skilled occupations command rising relative earnings, while workers in the middle-skilled occupations fell behind. Concretely, Autor, Katz, and Kearney document with CPS data for 1963-2005 that more moderate increases in overall wage inequality since the 1990s, compared to before, conceals different paths for upper-tail (90/50) inequality–which has increased steadily since 1980—and lower-tail (50/10) inequality—which rose sharply in the first half of the 1980s and contracted thereafter. Autor, Katz, and Kearney point to a deceleration in relative demand growth for college workers in the early 1990s and the polarization of skill demands favoring high-wage and low-wage work at the expense of middle-wage jobs. Autor, Katz, and Kearney propose a revised skill-biased technological change hypothesis, by which information technology complements high-skilled workers engaged in abstract tasks, substitutes for middle skilled workers who perform routine tasks, and has less impact on low-skilled workers who are assigned to manual tasks.[36]

Similar findings regarding a polarization of the wage distribution have been found in several other countries. Goos and Manning (2007), for example, show for the United Kingdom since 1975 that a pattern of job polarization has occurred with rises in employment shares in the highest- and lowest-wage occupations. They predict that job polarization can explain one-half of the rise in the log(90/50) differential and one-third of the rise in the log(50/10) wage differential. Skill-biased technological change, at given factor supplies, tends to increase wage inequality (e.g. Autor, Levy, and Murnane 2003; Autor, Katz, and Kearney 2008; Acemoglu and Autor 2011). Studying the wage level, rather than wage inequality, and its relationship to robot use across local labor markets, Acemoglu and Restrepo (2017) estimate that one more robot per thousand workers in the United States reduces the regional wages by about 0.5 percent.

These studies of wage changes and their relation to employment patterns suggest that skill-based technological change, in a task-related reinterpretation, has possibly played a considerable role for earnings across skill groups. However, until very recently, there has been little evidence on the relation to variables that directly measure technology use. In many cases, technological change is merely the default explanation when alternatives appear less compelling (Helpman 2016, p. 13). Recent exceptions are Acemoglu and Restrepo (2017) and Cortes, Jaimovich, and Siu (2016a) as discussed.

5.2 Technical change in isolation and matching

Charles, Hurst, and Notowidigdo (2016) revisit technological change but propose an intriguing reverse explanation for employment and wage outcomes in the presence of matching worker skills to occupations. Charles, Hurst, and Notowidigdo's basic tenet is the idea that skilled-biased technological change can cause an initial boom and a subsequent partial bust in the demand for cognitive tasks. Charles, Hurst, and Notowidigdo argue that during an initial adoption phase, when a general purpose technology (GPT) makes deep inroads into the workplace, demand for cognitive tasks grows fast because the associated machinery and equipment need to be built and installed. However, once the general purpose technology has been widely adopted, demand for cognitive tasks partially drops because at maturity of the technology those activities are mainly needed for maintenance and occasional replacement of the technology, but no longer for its adoption. In absolute levels, demand for cognitive tasks at the technology's maturity still exceeds the demand before the general purpose technology introduction but is no longer as high as during the initial adoption phase. Charles, Hurst, and Notowidigdo document that demand for cognitive tasks in the United States reached its maximum around

[36] As discussed in Section 2 above (see footnote 3 in particular), an alternative explanation for the increase in the skill premium and a simultaneous decrease in the relative skill employment during the 1990s and later is that offshoring of tasks (trade in tasks) reduces the employment of lower-paid nonproduction workers but the quasi-rent from offshoring benefits these nonproduction workers through increased earnings in general equilibrium (see Grossman and Rossi-Hansberg 2008, who call this outcome a "productivity effect").

the year 2000 and has since been declining. As the relative demand for cognitive tasks falls, the displaced high-skilled workers are matching to less skill intensive occupations ("down the occupational ladder"). As a consequence, positive assortative matching between workers and occupations worsens, and less skill intensive occupations are being filled with higher proportions of skilled workers (a "de-skilling process"). This process is reminiscent of shifting skill proportions in the classic Heckscher-Ohlin model (or for that matter the matching-based model by Costinot and Vogel 2010), when a decline in relative demand for the skilled factor leads to an increasing skill proportion in employment across all industries (while the relative wage for skill comes down). To reconcile their argument with job polarization, Charles, Hurst, and Notowidigdo (2016) need to consider cognitive tasks as concentrated in middle-range occupations with intermediate wage levels.

5.3 Organizational change

Management practices have been found to be an important determinant of the variation in plant and firm efficiency within and across countries (Bloom and Van Reenen 2011), also in the absence of firm-level globalization strategies (in the context of firm-level trade integration, see the discussion of organizational change in subsection 6.6.4). Given the positive association between firm size and wages, called the wage-size premium (Oi and Idson 1999, see also subsection 4.3.3 above), a wider dispersion of management practices across employers can also be expected to lead to a wider wage dispersion. Moreover, managerial practices can shift the relative demand for skills. Caroli and Van Reenen (2001), for example, argue that organizational change and skills are complements and document for a panel of British and French plants that organizational change is indeed skill biased. Among the organizational changes that they consider is the flattening of hierarchies, the decentralization of authority, and increased multitasking. Caroli and Van Reenen (2001) show that these organizational changes reduce the demand for unskilled workers in both countries. Organizational change also leads to greater productivity increases at plants with larger initial skill intensities, suggesting one more link between employer capabilities and the cross-employer wage differentials. As product market conditions change with trade and import competition, incentives for management practices and organizational change shift—one instance of trade-induced technology change to be taken up in subsection 6.6.4 below.

5.4 Trade studies that control for technical change

Several trade studies control for potentially confounding effects through technological change and check the trade-related estimates for their sensitivity to varying technology measures. The objective of the studies collected and discussed in this subsection is to arrive at a careful measure of the trade impact on labor markets, conditioning on technological change. These studies do not necessarily aim to explore the nature of a potential covariation between trade and technology. The broader question as to how trade and technological change may interact, so that they determine labor markets not only in isolation but also exert a potential joint effect, is the subject of Section 6 further below.

In their study of relative wage and employment outcomes among U.S. workers in the 1980s, Feenstra and Hanson (1999) showed that their estimates of the impact of offshoring were sensitive to an industry-level measure for the use of high-tech equipment in the total capital stock. When they used this technology measure in the analysis, about one-quarter of the wage differential paid to nonproduction workers during the 1980s was attributable to offshoring and around 30 percent to technology. When Feenstra and Hanson (1999) altered their measure of high-tech equipment to place more weight on more recently installed (and thus arguably more advanced) equipment, the contribution of offshoring to wage outcomes fell by half while the contribution of technology more than tripled.[37] Their study has remained one of only a few that undertake a direct and explicit comparison between a trade-related variable and a technology-related variable in predictions of labor-market outcomes.[38]

Technological progress within manufacturing has been most rapid in recent decades in computer and skill-intensive sectors (Doms, Dunne, and Troske 1997; Autor, Katz, and Krueger 1998). Acemoglu, Autor, Dorn, Hanson, and Price (2016) extend the approach in Autor, Dorn, and Hanson (2013) to allow for offshoring through input-output linkages within and across borders and to control for technological change. Acemoglu, Autor, Dorn, Hanson, and Price add variables from the NBER-CES database to the specifications, including the workforce share of nonproduction workers, the average wage, the ratio of capital to value added, as well as computer and high-tech equipment investment. However, Acemoglu, Autor, Dorn, Hanson, and Price do not report estimates for these skill and technology variables nor do they perform comparisons of the relative importance of the China shock versus technological change.

Harrison and McMillan (2011) and Ebenstein, Harrison, and McMillan (2015) find only minor effects of globalization on labor force participation. But they report large effects of computer use and the price of investment goods on employment. Falling investment goods prices are associated with an increasing use of capital. Ebenstein, Harrison, and McMillan (2015) use a U.S. household sample (CPS Merged Outgoing Rotation Groups 1983-2008) and make the occupation the unit of analysis. Their estimation approach suggests that greater use of computers and capital equipment is associated with lower employment, higher unemployment, and lower labor force participation, whereas globalization is not strongly associated with low labor force participation. In contrast, import competition appears to be associated with higher rates of labor force participation, higher employment, and lower unemployment in their estimation. While the import-competition induced employment consequences are small in magnitude, general (and in an extension occupation-specific) measures of total factor productivity, the price of in-

[37] As mentioned earlier, Feenstra and Hanson (1999) documented for the United States that about one-third of the observed increase in the skill premium in the 1970s and 1980s could be attributed to offshoring. (Feenstra and Hanson 1999) presented a range of estimates for different trade and technological change predictors and found that between 15 and at most 40 percent of the observed change in the wage-bill share of nonproduction workers in the United States could be explained by offshoring.

[38] Autor, Dorn, and Hanson (2015) undertake a comparison of sorts between the China shock and technology but they equate technological change with routine task specialization, which other authors relate to offshoring (Blinder 2006; Becker, Ekholm, and Muendler 2013; Ottaviano, Peri, and Wright 2013).

vestment goods, computer use rates, and the capital to labor ratio predict substantive labor-market responses.

A rare study that assesses multiple sources of economic change simultaneously, and structurally, is Burstein, Morales, and Vogel (2015). Similar to the reduced-form evidence in Ebenstein, Harrison, and McMillan (2015), Burstein, Morales, and Vogel find closely related effects and a dominant role of computerization for relative wage outcomes. Burstein, Morales, and Vogel use a matching model to decompose changes in between-group wage inequality into changes in the composition of the workforce, the demand for tasks, computerization, and labor productivity. The model incorporates comparative advantage between groups of workers, types of equipment, and tasks, but is sufficiently parsimonious to be estimated. Burstein, Morales, and Vogel use U.S. data from the Combined CPS in 1984-2003 to track the allocation of workers to occupations and computer usage as well as changes in average wages across worker groups between 1984 and 2003 to parameterize their model.

Burstein, Morales, and Vogel (2015) find that computerization and changes in task demand, which are both measured without directly using data on changes in wages, jointly explain the majority of the rise in the skill premium and more disaggregated measures of between-education group inequality. Their results imply that computerization is the central force driving changes in between-education group inequality, whereas labor productivity plays a relatively minor role. Computerization alone accounts for most of the observed rise in between-education-group inequality over this period: 60 percent of the rise in the skill premium. The combination of computerization and occupation demand shifters explain roughly 80 percent of the rise in the skill premium, and almost all of the rise in inequality across more disaggregated education groups.

To summarize the evidence on employment and wage outcomes up to here, consider two benchmark studies from the respective literatures: Autor, Dorn, and Hanson (2013) have studied trade in isolation and Cortes, Jaimovich, and Siu (2016a) who consider technological change in isolation. Autor, Dorn, and Hanson (2013) argued that one-quarter of the decline in U.S. manufacturing employment is explained by final goods trade and import competition largely from China—even under the strong and unverified assumption that their estimated change in regional disparity is equal to the national change (as discussed in subsection 4.1.3 above). Cortes, Jaimovich, and Siu (2016a) consider technological change in isolation and find in a demanding structural model that automation explains one-third of the decline in middle-skilled employment (when the model is also parametrized to match occupational reallocations). Taking these two estimates together crudely for employment, under the observation that middle-skilled occupations are concentrated in manufacturing, the two isolated studies jointly account for just over one-half of the salient employment decline. What accounts for the remaining and perhaps dominant part of the employment decline that remains to be explained? When it comes to wages, we might follow Burstein, Morales, and Vogel (2015), who studied technological change in isolation and find that it may account for three-fifth of observed changes in the skill wage premium, but might also need to recognize that a benchmark study by Autor, Dorn, and Hanson (2013) finds hardly any predictable wage changes. What accounts for the remaining substantive part of relative wage changes that remains to be explained?

Burstein, Morales, and Vogel give one answer. They argue that the strength of the main two forces—computerization and changes in task demand—is likely linked in turn to the volume of international trade, which they control for separately. For example, they show that the large rise in computer imports relative to non-computer equipment reduces the relative price of computers to non-computer equipment, and this relative price change affects relative wages in exactly the same way as an increase in computer productivity in the closed economy. In other words, studying technical change and globalization in isolation may be a misleading exercise. The so-far unaccounted and potentially important interaction between trade and technology is the subject of the next section.

6 Interactions of Trade, Technology, and Prosperity

Globalization and technological change were discussed in isolation so far, but technology adoption and trade in merchandize and services can interact in important ways. Moreover, globalization and technological change together raise prosperity, which in turn alters the content of our consumer baskets that include ever more services, from health care to entertainment. Our changing consumer baskets may drive a large part of the lasting structural changes common to many economies. Finally, the disparities in prosperity across households influence our savings behavior, which in turn is a key determinant of trade and current account imbalances. These connections, and the potential magnitudes of their economic relevance, are touched upon in several studies, which this report partly covered so far. However, research into several potential connections is arguably at early stages, so that this section presents an exploratory collection of evidence from diverse strands of research. To start, it appears useful to elicit key facts from consistent longitudinal household across countries. Those facts can then help guide the discussion.

6.1 Evidence on labor-market outcomes and global structural change

Using consistent longitudinal household data across countries, based on CNEF and additional consistency treatments for economic activities, we documented above a steady decline in the frequency of workers' cross-industry and cross-sector moves (subsection 2.2, Tables 1 and 2). This steady decline in transition frequencies between economic activities is broad based across all economies in the sample and has lasted for decades. During business cycle contractions, especially in the recent Great Recession, workers also transition into unemployment or out of the labor force, rather than between economic activities. However, the long-lasting and broad decline in worker transitions between economic activities across diverse countries suggests that the reduced labor-market turnover reflects a longer-term phenomenon beyond business-cycle adjustments in the labor markets. Countries that were little affected by the Great Recession, such as Germany and Switzerland, or strongly affected, such as the United Kingdom and the United States, show similar declines in the frequency of workers' cross-industry transitions. Moreover, in Germany, the United Kingdom and the United States, where longitudinal household data are available for long time spans, a gradual decline in cross-industry transition frequencies has been observable since the mid 1990s. One potential explanation for the common but wrong perception that workers transition between economic activities more frequently today than in the past may have to do with the specific transition experiences: the performance of the manufacturing sector, where technical change can be disruptive and globalization act fast, the direction of worker flows, and the earnings experience associated with workers' moves between economic activities.

To shed light on the performance of the manufacturing sector, Figure 1 shows the share of manufacturing employment in a country's total employment for the period 1977-2013. In the seven economies, for which consistent longitudinal household data on this aspect are available, manufacturing employment exhibits a steady decline since the late 1970s. In Germany, reported manufacturing employment dropped from a share of 31 percent

Figure 1: **Share of Manufacturing Employment in Total Employment**

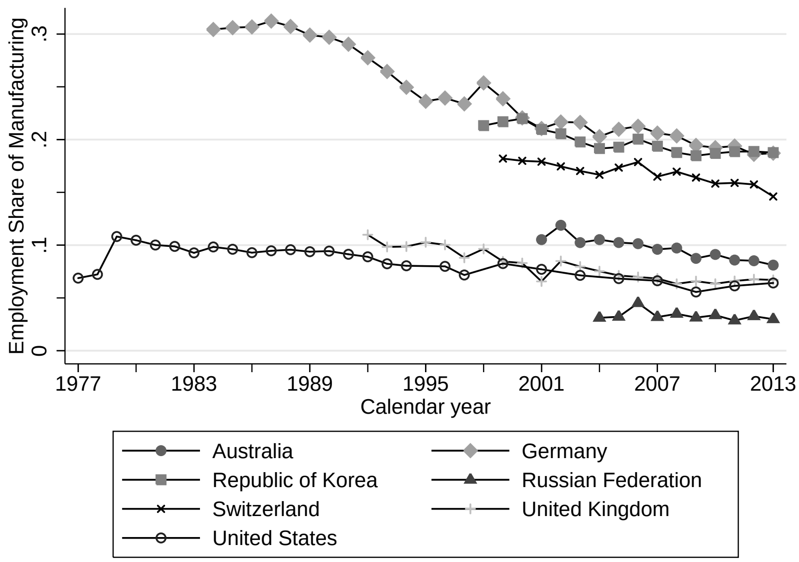

Source: Own calculations based on longitudinal household data with consistently defined variables across countries.
Notes: Share of workers employed in manufacturing in total number of employed workers. Entries show the share of household members, with wage employment in the reported year, who are currently employed in the manufacturing sector (one minus the reported share is the fraction of workers who are employed in the primary sector or the services sector). Longitudinal household data are based on the Cross-national Equivalent File (CNEF), Ohio State University, extracted for 1973-2013 from the British Household Panel Study (BHPS), the Household Income and Labour Dynamics in Australia (HILDA), the Korea Labor and Income Panel Study (KLIPS), the U.S. Panel Study of Income Dynamics (PSID), the Russia Longitudinal Monitoring Survey (RLMS-HSE), the Swiss Household Panel (SHP), and the German Socio-Economic Panel (SOEP). Reported economic activity In CNEF at the single-digit level was corrected according to the two-digit level information (year 1995 omitted in the United States). Worker counts based on household members' cross-sectional population weight in current year (sampling weight in countries with unreported population weight).

in economy-wide employment in 1987 to a share of 19 percent by 2013, matching the reported manufacturing share in the Republic of Korea by 2013, where it had peaked at 22 percent in 2000. In Switzerland, the manufacturing employment share dropped from 18 percent in 1999 to less than 15 percent by 2013. In the Russian Federation, the manufacturing share of employment remained steady at around 3 percent between 2004 and 2013. The reported levels of the manufacturing share in total employment are historically lower in the more services-intensive Anglo-Saxon economies of Australia, the United Kingdom and the United States, where a higher degree of outsourcing manufacturing-related services from manufacturing employers to independent companies is arguably more common. However, declines in the manufacturing share are still marked in the Anglo-Saxon economies: in Australia the manufacturing employment share fell from 12 percent in 2002 to 8 percent in 2013, in the United Kingdom from 11 percent in 1992 to less than 7 percent in 2013, and in the United States from 11 percent in 1979 to 6 percent in 2013.

The common decline in the manufacturing share across these diverse economies, with different patterns of participation in global trade, suggests that economic forces common to those economies are at work—such as labor-saving technological change or consumption-driven increases in services demand. As per-capita incomes rise and consumption baskets become lighter, filled with more services and fewer goods, deindustrialization rather than globalization can be expected to be a common source of global change across countries. Moreover, the frequently investigated consequences of China's accelerating insertion into the global networks of trade during the 2000s does not appear to have markedly altered the gradual evolution of manufacturing employment: the rate of decline looks relatively smooth over the three-and-a-half decades since 1977, and in fact similar across all economies except the Russian Federation, until the late 2000s. In two countries, the decline in the manufacturing employment share in fact stopped falling by the late 2000s. In the United Kingdom and the United States the manufacturing employment share reached a minimum in 2008/09, dropping to 6.3 percent in the UK and to 5.6 percent in the U.S., and slightly rebounded since, rising back to 6.7 percent in the UK and 6.4 percent in the United States by 2013.

This long-lasting relative employment decline of the manufacturing sector, similar across the countries in our longitudinal household data, can be expected to lead to worker transitions. Table 3 reports the employment share across the three main economic sectors for longitudinally trackable household members who were employed in the manufacturing sector four years prior and have wage employment in the current year. The manufacturing employment share (as reported in Figure 1 before), of course, also depends on first entrants into the workforce after schooling, transitions from and to unemployment, and exits from the workforce because of retirement or discouragement from job search. Those caveats notwithstanding, the frequencies of manufacturing worker continuations in the manufacturing sector, and the transitions out of manufacturing for workers who continue in the active labor force over four-year periods, provide a picture of the reallocation process for a large part of the economy-wide workforce. Worker continuations in the manufacturing sector include both retained jobs at the same employer and new jobs at another manufacturing employer.

At one extreme in our country sample, the economy with the highest frequency of

Table 3: SHARES OF WORKER CONTINUATIONS OR TRANSITIONS ACROSS THREE SECTORS, CONDITIONAL ON MANUFACTURING EMPLOYMENT FOUR YEARS BEFORE

From Manufacturing To:	Australia (1)	Germany (2)	Rep. of Korea (3)	Russian Fed. (4)	Switzer-land (5)	United Kingdom (6)	United States (7)
1989 Primary		.006					.017
Manufacturing		.787					.775
Services		.207					.208
1995 Primary		.013				.012	.017
Manufacturing		.721				.693	.725
Services		.266				.296	.257
2001 Primary		.010				.007	.011
Manufacturing		.774				.679	.694
Services		.217				.314	.295
2007 Primary	.052	.006	.010		.006	.024	.012
Manufacturing	.612	.782	.817		.765	.705	.688
Services	.336	.211	.173		.229	.271	.301
2013 Primary	.040	.026	.007	.033	.011	.009	.018
Manufacturing	.630	.757	.862	.550	.801	.818	.757
Services	.331	.217	.130	.417	.188	.173	.225

Source: Own calculations based on longitudinal household data with consistently defined variables across countries.

Notes: Share of continuations or transitions to three consistently defined sectors (primary, manufacturing, services) for workers with manufacturing employment four years prior. Entries show the share of longitudinally trackable household members, with wage employment in the reported year and four years prior, who were employed in the manufacturing sector four years prior. Longitudinal household data are based on the Cross-national Equivalent File (CNEF), Ohio State University, extracted for 1985-2013 from the British Household Panel Study (BHPS), the Household Income and Labour Dynamics in Australia (HILDA), the Korea Labor and Income Panel Study (KLIPS), the U.S. Panel Study of Income Dynamics (PSID), the Russia Longitudinal Monitoring Survey (RLMS-HSE), the Swiss Household Panel (SHP), and the German Socio-Economic Panel (SOEP). Reported economic activity In CNEF at the single-digit level was corrected according to the two-digit level information (year 1995 replaced with 1994 in the United States). Worker transition counts based on household members' cross-sectional population weight in current year (sampling weight in countries with unreported population weight).

worker continuations in manufacturing is the Republic of Korea, where 82 to 86 percent of manufacturing workers stay in manufacturing over a four-year period. The countries with the lowest continuation rates in the manufacturing sector are the Russian Federation, with only 55 percent of manufacturing workers continuing in manufacturing over four years leading up to 2013, and Australia, with 61 to 63 percent of manufacturing workers staying in manufacturing over the four years leading up to 2007 and 2013. Both countries exhibit relatively high transition rates to the primary sector (mostly to mining and energy but not much to agriculture), with 4 to 5 percent in Australia and 3 percent in the Russian Federation, compared to other sample countries where typically less than 2 and often less than 1 percent of manufacturing workers transition to the primary sector. However, even in the primary-sector intensive economies of Australia and the Russian Federation, the bulk of workers who transition out of manufacturing move into services.

Germany and Switzerland exhibit intermediate levels of worker continuations in manufacturing, where between 76 and 80 percent of manufacturing workers stay in manufacturing between 1989 and 2013 (except for one outlier period in 1995 when Germany after reunification only shows a 72 percent continuation rate of manufacturing workers in manufacturing). The United Kingdom and the United States exhibit similar but somewhat smaller manufacturing continuation rates that can fall as low as 68 percent in some years but also reach as high as Korean manufacturing continuation rates of 82 percent. Interestingly, the times of low manufacturing continuation rates seem to be a matter of the past, as both the United Kingdom and the United States have seen rebounds in the frequency of worker continuations in manufacturing to 76 percent in 2013 in the United States and the aforementioned 82 percent in the United Kingdom in 2013.

In summary, the speed of gross worker outflows from manufacturing—among workers who continue in the workforce—seems to be steady in most economies, while some countries such as Australia or the Russian Federation experience fast gross worker moves out of manufacturing and others such as Germany or Switzerland undergo more gradual transitions out of manufacturing, and yet other economies such as the United Kingdom and the United States are seeing a slowdown in gross worker moves out of manufacturing. From Figure 1 we know, however, that net worker flows must be directed away from manufacturing as the manufacturing share in total employment continually declines.

As workers continue in or move between sectors, their earnings change. If worker transitions out of a sector are driven by a pull from stronger labor demand in other sectors, one might expect moves to be associated with wage gains. In contrast, if workers move out of a sector because of a push from weakening labor demand in their initial sector of employment, one might expect the sector changes to be accompanied by wage losses. Table 4 reports mean four-year difference of log wages for workers with initial employment in manufacturing, four years earlier, discerning continuations in manufacturing and transitions to the primary sector and services. The reported figures are the mean differences of log gross annual wages, after deflating wages with the country-specific CPI (using 2014 as base year).[39] Workers who stay in the manufacturing sector over a four-year period,

[39] These mean log differences can be transformed into percentage changes: a log point change of x can be restated as a percentage change of $(\exp\{x/100\} - 1) \cdot 100\%$.

Table 4: ANNUAL WAGE CHANGE IN PERCENT OVER FOUR YEARS FOR WORKER CONTINUATIONS OR TRANSITIONS ACROSS THREE SECTORS, CONDITIONAL ON MANUFACTURING EMPLOYMENT FOUR YEARS BEFORE

From Manufacturing To: (in %)	Australia (1)	Germany (2)	Rep. of Korea (3)	Russian Fed, (4)	Switzer-land (5)	United Kingdom (6)	United States (7)
1989 Primary		.311					.466
Manufacturing		.200					.105
Services		.175					-.039
1995 Primary		.384				-.585	-.246
Manufacturing		.100				.159	.051
Services		.090				.071	-.050
2001 Primary		.150				.353	.100
Manufacturing		.147				.184	.194
Services		.213				.067	.075
2007 Primary	.406	.195	-1.099		-.285	-.111	.355
Manufacturing	.159	.034	.173		.055	.065	.072
Services	.250	-.030	.145		-.175	-.127	-.098
2013 Primary	.253	.060	.139		.625	.019	-.374
Manufacturing	.125	.108	.101		.103	-.006	.0003
Services	.160	-.004	.146		.153	.011	-.141

Source: Own calculations based on longitudinal household data with consistently defined variables across countries.

Notes: Four-year difference of wages for workers with continuations or transitions to three consistently defined sectors (primary, manufacturing, services), conditional on their manufacturing employment four years prior. Entries show the mean differences of log gross annual wages, after deflating wages with the country-specific CPI (base year 2014), for longitudinally trackable household members, with wage employment in the reported year and four years prior, who were employed in the manufacturing sector four years prior. Longitudinal household data are based on the Cross-national Equivalent File (CNEF), Ohio State University, extracted for 1985-2013 from the British Household Panel Study (BHPS), the Household Income and Labour Dynamics in Australia (HILDA), the Korea Labor and Income Panel Study (KLIPS), the U.S. Panel Study of Income Dynamics (PSID), the Russia Longitudinal Monitoring Survey (RLMS-HSE), the Swiss Household Panel (SHP), and the German Socio-Economic Panel (SOEP). Reported economic activity In CNEF at the single-digit level was corrected according to the two-digit level information (year 1995 replaced with 1994 in the United States). Means of log wage differences are based on household members' cross-sectional population weight in current year (sampling weight in countries with unreported population weight). Log points transformed into percentage changes: a log point change of x is restated as a percentage change of $(\exp\{x/100\} - 1) \cdot 100\%$

either retaining their job or moving to a different manufacturing job, command real wage increases in all reported periods, but for one exception: a real wage loss of 0.6 log points among continuing manufacturing workers in the United Kingdom in 2013. Experiences with wage real wage changes for workers who move to the primary or services sector are diverse over time and across countries.

In the most recent four-year period leading up to 2013, in Australia, the Republic of Korea, Switzerland and the United Kingdom, manufacturing workers in 2009 who moved to the primary or the services sector by 2013 commanded stronger real wage gains on average than workers who continued in manufacturing. The pattern of real wage changes for the same workers in these countries is consistent with a pull of faster increasing labor demand in the primary and services sector than in manufacturing. In Germany and especially the United States, in contrast, workers who stayed employed in the manufacturing sector had a stronger real wage performance than those workers who switched to employment in the other two sectors. A consistent explanation of that pattern for Germany and the United States might be that workers in surviving manufacturing jobs saw their labor productivity increase. However, the loss of underperforming manufacturing jobs in Germany and the United States may have pushed workers out of manufacturing into other sectors where labor demand for their (perhaps manufacturing specific) skills was too weak in Germany to allow the switching workers to command similar real wage gains as continuing manufacturing workers or so adverse in the United States that switching workers suffered real wage losses.

Interestingly, in the United States there was no period when switchers out of manufacturing and into services commanded stronger real wage increases than continuing manufacturing workers in the surviving manufacturing jobs, and in most years (except 1989 and 2007) switchers to the primary sector also performed worse than continuing manufacturing workers. This pattern in the United States is an outlier experience in our cross country data, suggesting that the heavy emphasis on U.S. evidence in studies of labor-market experiences from globalization and technological change may not be representative for other economies. Only the Republic of Korea and Switzerland had a U.S.-like experience in 2007, when they saw continuing manufacturing workers outperform switchers. But in most countries (except Germany) and most time periods (except 2007) switchers out of manufacturing do either not perform clearly worse or in fact better than continuing manufacturing workers. The singularly adverse wage performance of switchers out of manufacturing in the United States may explain why discontent with labor-market outcomes in the United States is particularly strong. The four-year time period of analysis could conceal a possible subsequent catch-up in real wage gains among switchers, but periods of more than four years can be a long time horizon for households faced with workplace disruption.

Figure 1 above documented a broad based and common relative manufacturing employment decline across all countries in our sample, which has arguably been underway for three or more decades now. Despite the similarity in the overall manufacturing employment evolution, however, the individual labor-market outcomes across the countries in our sample are diverse. As Table 4 showed, workers who transition to non-manufacturing sectors in the United States have suffered slower real-wage changes or outright real wage declines from their switches, whereas economies such as Australia or more recently the Republic of Korea, Switzerland and the United Kingdom offered workers who switched out of manufacturing faster real wage gains than those workers who stayed behind in manufacturing.

6.2 Deindustrialization

Deindustrialization is the decline in manufacturing employment and value added as a share of economy-wide activity. It has been lasting for decades and is pervasive across countries. Acemoglu, Autor, Dorn, Hanson, and Price (2016) show that U.S. manufacturing as a share of employment has been declining since the 1950s, and the number of manufacturing employees has also trended downward since the 1980s. Imbs (2017) documents main features of structural change in fifteen OECD countries since 1970. By that account, an accelerated deindustrialization process began in the OECD in the 1980's, but only in terms of changes in the allocation of labor, not the allocation of value added. The share of manufacturing value-added remained roughly constant until the turn of the century, when deindustrialization further accelerated. Those findings closely resemble the evolution of the share of manufacturing employment for the period 1977-2013 in countries with longitudinal household data, as plotted in Figure 1 above.

Turning to a larger group of countries, Rodrik (2016) documents a significant deindustrialization trend in recent decades that goes beyond the early industrialized economies and afflicts emerging markets as well. Rodrik documents a hump-shaped relationship between industrialization (measured by employment or output shares of manufacturing in the economy as a whole) and per-capita incomes. He argues that this hump-shaped curve has shifted downwards and closer to the origin, so that deindustrialization now begins at lower levels of per-capita incomes than in earlier periods. Rodrik (2013) pointed out that, unlike economies as a whole, manufacturing industries exhibit strong unconditional convergence in labor productivity, with a coefficient of unconditional convergence in the magnitude of an annual growth rate in industrialized countries (between 2 and 3 percent). This convergence in manufacturing labor productivity reinforces the idea that deindustrialization is pervasive and faster than in earlier periods. However, as deindustrialization in terms of manufacturing employment progresses globally, leading firms, especially in industrialized countries, will likely continue to act as global coordinators of production networks.

At the same time, boundaries between manufacturing and services are becoming increasingly blurred. First, manufacturing firms move into non-manufacturing activities and offer more services than in the past. Crozet and Milet (2017) show for French firms that, in 2007, 83 percent of manufacturing firms sold some services, 40 percent sold more services than goods, and 26 percent did not even produce goods. The McKinsey Global Institute (2012) reports that after-sales services and customer care accounted for 39 percent of Sweden's manufacturing employment in 2007. In other words, deindustrialization can take place not just between sectors but also within firms. Second, some services firms act like manufacturers. Bernard and Fort (2015) document the prevalence of factoryless goods producers in the United States: wholesale-sector firms that design the goods they sell and that coordinate global production activities, but operate no domestic manufacturing plant. Bernard and Fort (2017) calculate that reclassifying these factoryless goods producers from the wholesale sector to manufacturing would raise the number of manufacturing jobs in the United States by between 430,000 and 1.9 million in 2007. Third, manufacturers build global supply chains that combine in-house assembly with outside service suppliers. The McKinsey Global Institute (2012) reports for Germany, for instance, that service suppliers contribute 34 percent of the total domestic value added in manufacturing exports in 2007.

Those issues of sectoral classification for precise measurement notwithstanding, the decline in relative manufacturing employment and value added—documented across industrialized and emerging economies for decades in the studies cited above—suggests that deindustrialization is a lasting and pervasive economic process. To what extent do globalization, technological change, and advances in prosperity accelerate deindustrialization, directly or indirectly through their interactions? Inasmuch as foreign trade at the sectoral level is driven by comparative advantage, the reallocation of labor between manufacturing and services should show reverse patterns, with services-export countries exhibiting deindustrialization and manufactured- goods exporting countries experiencing a re-industrialization. The lasting and pervasive nature of deindustrialization therefore suggests that not trade but technological change and advances in prosperity are the likely *direct* causes of the relative decline in manufacturing, but globalization in its interaction with those processes could well have a mediated effect.

The reallocation of employment away from manufacturing is consistent with Baumol's (1967) view that sectors with relatively fast labor productivity growth lose employment. The reallocation of employment away from manufacturing is also consistent with the observation that our consumer baskets are getting lighter. From the cross-country evidence in Fajgelbaum and Khandelwal (2016) we know that households with lower incomes consume a larger proportion of basic manufactured goods. Both technological change and shifting demand are therefore the common suspects behind deindustrialization.

Leaving out globalization, Comin, Lashkari, and Mestieri (2015) model productivity growth and changing consumer baskets. The technology-oriented view by Baumol (1967), that sectors with fast productivity growth shrink in terms of employment, requires one key demand-side component to fall in place: the elasticity of substitution between manufactured goods and services must be less than unity. The consumption-oriented view, that our consumption baskets contain more services as we grow more prosperous, requires another set of key demand-side components to fall in place: the income elasticity for services must exceed that for manufactured goods. Comin, Lashkari, and Mestieri (2015) propose a new system of tractable consumer preferences, which have the convenient feature for quantification that the income elasticity parameters are separate from the elasticity of substitution parameters between sectors. Comin, Lashkari, and Mestieri (2015) document for a sample of 25 countries in Africa, Asia, Europe, Latin America and North America, and for 10 sectors since 1947, that estimates of the elasticity of substitution and the income elasticities across sectors are robust across countries, time periods and economic measures of sectoral activity. Their baseline estimate of the elasticity of substitution between goods and services is around 0.7. Comin, Lashkari, and Mestieri (2015) show that their model generates income-dependent consumption patterns at any level of economic development and a decline in agriculture since 1947, a hump-shaped evolution of manufacturing for developing and emerging economies (such as South Africa with a peak manufacturing employment share in the early 1980s or the Republic of Korea with a peak manufacturing employment share in the mid 1990s), a continual deindustrialization since 1947 for many industrialized countries (such as France, Sweden, the United Kingdom and the United States) with the notable exception of Japan (whose manufacturing employment share peaked in the early 1970s), and a continual expansion of services employment shares over time in every sample country since 1947. A main finding is that realistic preferences with (non-homothetic) income elasticities play a dominant role in generating structural change. In that sense, as both globalization and technological change advance prosperity, there is scope for globalization and technological change to be jointly behind deindustrialization, though indirectly. The exact connections, however, remain largely unexplored in empirical economic research.

A shortcoming of the structural change literature to date is that it cannot address some of the labor-market outcomes that are associated with the worker transitions out of manufacturing, while deindustrialization progresses. As documented above (Table 4 in subsection 6.1), in Germany and the United States in recent years workers who stayed employed in the manufacturing sector had a stronger real wage performance than those who moved to services, whereas in Australia, the Republic of Korea, Switzerland and the United Kingdom, for instance, the opposite was the case. The structural change literature so far mostly relies on competitive factor markets, especially competitive labor markets, often lacks a separate consideration of skill groups, and does not allow for rents in competitive product markets, so it cannot explain wage changes from sector switching. Arguably, the job opportunities outside manufacturing are a crucial cause of discontent with deindustrialization and the potential contribution of globalization, so real-wage changes from employment transitions seem to be an important aspect to address within the structural-change framework.

6.3 Market structure

Recent models of trade and technology adoption, both at the industry level and at the firm level, have moved beyond the competitive market paradigm when considering the labor market and have embedded rigorous theories of worker-employer matching and wage bargaining (see subsections 4.3 and 5.2). This move beyond the competitive paradigm has so far rarely been extended to the product market. Most trade theories and many macroeconomic models are set in frameworks of monopolistic or even perfect competition, and so are the empirical studies. The market structure, however, regulates in important ways how trade and technology and how trade and the dispersion of prosperity interact. On the theory side, progress has been made in modelling oligopolistic firms that interact strategically with their competitors in their own product markets but act as price takers with regard to the aggregate economy, which consists of a continuum of industries. Neary and Tharakan (2012), for instance, allow for Cournot and Bertrand oligopoly and endogenous alternations between the two market structures, and trace ways in which shocks such as globalization or technological change lead to switches between market structures and affect the skilled wage premium. Neary (2016) allows for Cournot oligopoly by industry and explores the distribution of income between wages and profits, as markups of prices over production costs change.[40] To date, there is still little systematic empirical evidence as to how the distribution of incomes between skilled and unskilled workers, and the distribution of wages and profits in the presence of few dominant and highly profitable globalized firms, depend on trade-induced or technology-induced changes to the market structure.

In an early investigation for the 1970s through the 1980s, Borjas and Ramey (1995) argued that the earnings impact of the U.S. trade deficits dependent crucially on the domestic market structure of affected U.S. industries. They documented that the adversely affected import competing U.S. industries, mostly in durable-goods manufacturing, used to earn substantive rents, because they were highly concentrated, and used to share those rents with their workers. Those industries were also low-skill intensive industries. Heightened competition eroded the rents of these oligopolistic industries, so the U.S. trade deficits resulted in a decline of the relative wages of the mostly low-skilled workers employed in those industries. In other words, rising inequality from this type of foreign competition is the result of falling low-skilled wages (rather than increasing high-skilled wages), and it is more likely to occur in industries such as automobiles, which used to command oligopoly rents, than in industries such as apparel or textiles, where rents were low under tough domestic competition to begin with. (The evidence in Borjas and Ramey (1995) also serves as a reminder of the earlier discussion of LATE findings in the methodology subsection 3.2, which emphasized that domestic market conditions interact with foreign shocks in important ways that make findings specific to the exact historic context.) To our knowledge, this line of reasoning has not been taken up in investigations of more recent episodes of globalization.

Economists hold conflicting beliefs about the benefits of competition for innovation. A common tenet has it that fierce product market competition instills innovative behavior, as firms seek to leap ahead of competitors or pursue other first-mover advantages. Managers arguably respond to heightened product-market competition with more effort (Schmidt 1997), and lower entry costs for competitors can induce firms to adopt more efficient performance-based labor contracts (Raith 2003). That is the competition tenet. But then there is the argument for patents. Economists also widely support a patent system, so as to award monopoly power to innovators for decades—under the premise that monopoly profits are necessary for innovators to earn back the sunk costs of innovations.

[40]Not all economic outcomes that depend on imperfect competition necessarily require oligopoly. Feenstra and Weinstein (2010) document, for instance, that monopolistic competition combined with translog preferences on the demand side allows for rich structural estimation of the impact of globalization on price markups over production costs. They estimate that, between 1992 and 2005, import competition led to a fall in U.S. markups. Compared to a less flexible demand system with constant elasticities of substitution, Feenstra and Weinstein (2010) argue that gains from additional varieties are reduced by one-third, but the newly measured reduction in U.S. price markups over production costs results in a comparable welfare gain overall as under the simpler demand system.

How can perfect competition and its opposite, monopoly power, both be good for innovation?

One resolution of the open question is offered by Aghion, Bloom, Blundell, Griffith, and Howitt (2005): the impact of global competition on innovation depends on the market structure. Firms far from the global technological frontier give up innovation in the face of foreign competition, whereas firms close to the technological frontier can have stronger incentives to innovate under tougher competition. Aghion, Bloom, Blundell, Griffith, and Howitt (2005) document an according hump-shaped relationship using panel data for India. In their model, competition discourages laggard firms from innovating but encourages what they call "neck-and-neck" firms to innovate. Together with the effect of competition on the equilibrium market structure, these opposing forces generate an inverted U. Two additional predictions of the model—namely that the average technological distance between leaders and followers increases with competition, and that the hump-shaped relationship has a steeper slope when industries are more neck-and-neck—are also supported by the data. Beyond the context of India, case studies such as Freeman and Kleiner (2005) on U.S. footwear manufacturers, Bartel, Ichniowski, and Shaw (2007) on U.S. valve producers, or Bugamelli, Schivardi, and Zizza (2010) on Italian manufacturers show firms innovating in response to foreign import competition.

Acemoglu, Aghion, and Zilibotti (2006) underpin those empirical relationships with additional theory. In their model, firms can undertake innovation or adopt technologies from the world technology frontier, or both. The selection of high-skill workers and firms is more important for innovation than for adoption. As an economy approaches the global technological frontier, selection becomes more important. Countries at early stages of development therefore might choose to pursue an investment-based strategy, which relies on existing firms and managers to maximize investment but foregoes selection. Closer to the world technology frontier, economies switch to an innovation-based strategy with younger firms and stronger selection of firms and their workers. The importance of these empirical and theoretical findings notwithstanding, globalization, technological change and advances in prosperity arguably alter the market structure themselves, so that the way in which market structure mediates the impact of globalization or technological change on innovation shifts.

6.4 Savings and trade imbalances

The strong focus on China's productivity advances in manufacturing (Autor, Dorn, and Hanson 2013) and tariff certainty after WTO accession (Pierce and Schott 2016) in the existing literature are hard to reconcile with canonical macroeconomic explanations. Fast productivity increases in the traded sector are commonly associated with real exchange rate appreciations, as the relative price of local non-traded goods increases (akin to the so-called Balassa-Samuelson effect, see also Baumol 1967).[41] Moreover, faster productiv-

[41] The assumption that Chinese productivity change explains the observed rise in Chinese export shipments is a common assertion but so far largely untested. The recent structural approach by Caliendo, Dvorkin, and Parro (2015) may offer a suitable framework to carry out this important empirical check. As mentioned before, Caliendo, Dvorkin, and Parro calibrate a trade-and-internal-migration model to 38 coun-

ity increases in one economy's traded goods sector, compared to the rest of the world, are typically expected to attract net capital inflows and an associated *current account deficit* in an economy with a fast-growing traded-goods sector such as China during the 2000s (for microfoundations of open-economy macroeconomic models with representative households and this prediction, see canonical textbook models such as in Obstfeld and Rogoff 1995). The converse occurred. China's real exchange rate appears to have remained undervalued compared to its main trading partners throughout the 2000s (see e.g. Campbell 2016) and China commanded a sizeable current account *surplus*, with associated net exports mainly to the United States. Over the same period of time, China's private household consumption declined from 47 percent of GDP to merely 36 percent, the lowest consumption among the world's major economies (private household consumption usually accounts for between two-thirds and three-quarters of a country's GDP). China's manufacturing productivity change does therefore not necessarily appear to be the prevalent macroeconomic factor.[42]

By the logic of economic definitions, a country's current account balance equals its national savings less its domestic investment. Empirically, the trade balance and the current account balance are similar, with differences of only a few percentage points for most countries (oil exporters such as Kuwait or foreign-investment driven economies such as Ireland are among the handful exceptions). A key factor in China's expanding net exports to the rest of the world during the 2000s is therefore likely its national savings behavior. The average Chinese savings rate has been rising over time, with a marked acceleration in the 2000s. Economists at the Bank for International Settlements (2015) calculate that the Chinese aggregate savings rate exceeded a globally unprecedented 50 percent during the period (China's ratio of savings per GDP was 51.2 percent in 2005, and 54.3 percent by 2008)—far above every OECD economy's savings rate during the 2000s and overtaking Singapore, which has traditionally been among the highest global savers. Other countries with persistent trade surpluses such as Germany or Japan had savings rates of 22.2 and 26.8 percent in 2005. The high Chinese savings rate is also unusual in historic perspective, especially compared to East Asian economies that rapidly developed during the 20th century: Japan's aggregate savings rate increased by 15 percentage points during the period

tries, all 50 U.S. states, and 22 industries and (similar to the IV approach in Autor, Dorn, and Hanson 2013) predict changes in U.S. imports from China using as instrument the change in imports from China by other high-income countries between 2000-2007. To maintain consistency with the structural model, they compute the implied changes in sectoral total factor productivity (TFP) in China between 2000 and 2007 that match the predicted imports in the model. These predicted productivity changes could easily be contrasted with the actually measured productivity changes so as to assess whether indeed productivity, or perhaps other macroeconomic factors, best explain the rise in Chinese imports.

[42]Interestingly, major studies such as Autor, Dorn, and Hanson (2013), Acemoglu, Autor, Dorn, Hanson, and Price (2016) or Dauth, Findeisen, and Suedekum (2014) never use Chinese manufacturing productivity as a first-stage predictor of global trade flows from China to high-income countries, despite their emphasis on productivity change for explanations of the identification strategy. Instead, the studies use realized trade flows, which reflect the economy-wide trade balance that results from overall macroeconomic conditions beyond merely productivity. To make empirical headway, the vast majority of empirical papers, when they undertake quantification, proceed under the assumption that trade imbalances remain constant at their observed level in terms of the factor price of the reference country. Autor, Dorn, and Hanson (2013) assume, in addition, that the share of the trade imbalance in total expenditure is the same across U.S. regions.

1955-70, and the Republic of Korea's saving rate rose from 16 percent to 40 percent between 1983 and 2000. In the 2000s, India's savings rate climbed by 10 percentage points, reaching 38 percent of India's GDP in 2008. But none came close to China's savings rate.

In the view of economists at the Bank for International Settlements (2015), China's rising aggregate savings is not necessarily attributable to a single source. They document that savings have simultaneously risen in the corporate sector, among private households (the flip side of low consumption), and in government. While the Chinese household sector is the largest saver in levels, the corporate and government sectors have been the main contributors to the savings increase since the mid-1990s, explaining more than four-fifths of the 17 percentage point rise in China's savings rate. China's corporate savings doubled from 12 percent of GDP in 1992 to a peak of 24 percent in 2004, but fell back to 19 percent in 2008, while government savings rose faster. In the 2000s, half of the buildup of Chinese savings has come from the government sector. The contribution to national savings from all three sectors casts some doubt on a prevalent hypothesis that economic distortions and subsidies alone can account for China's high savings rate. The Bank for International Settlements (2015) posits that tough corporate restructuring (including pension and home ownership reforms), a marked rural-to-urban reallocation process (where the average wage exceeded the marginal product of labor in the rural area), and a rapid ageing process have played more important roles. However, the prolonged transition of Chinese rural labor to urban and more industrialized localities may have both demographic and institutional roots. In capping wage growth, the rural-urban migration boosted corporate profits in the process. For dominant state-owned enterprises are not required to pay dividends, corporate savings can well have risen with profits. Dollar and Jones (2013) discuss interactions between migratory and capital controls and their impact on China's past development path, while a number of China scholars suggest that recent real wage growth may indicate an end of China's experience with surplus labor (e.g. Zhang, Yang, and Wang 2011; Das and N'Diaye 2013). China's current account surplus has meanwhile declined to about one-quarter of its peak magnitude.

A flip side of the Chinese current account surplus is the U.S. current account deficit, but the interplay of international and domestic factors can be hard to disentangle. As discussed in the preceding subsection 6.3, Borjas and Ramey (1995) argued that the domestic market structure of the import-competing U.S. industries was a crucial factor explaining the earnings impact of the U.S. trade deficits in the 1970s and 1980s. Low-skill intensive oligopoly industries, which shared surplus with their workers, suffered an erosion of their rents from global competition and their contraction lead to a decline in the relative wage of low-skilled workers. Classic trade theory points in a similar direction: low-skilled labor, a relatively scarce production factor in advanced economies, tends to suffer real-wage losses from deeper global integration in final goods trade (Deardorff and Staiger 1988).

Among the domestic causes of the U.S. current account deficit is the low U.S. savings rate, which might be depressed itself because of earnings losses concentrated among low-skilled workers. In its annual financial survey of U.S. households, the Board of Governors of the Federal Reserve System (2016) asks respondents how they would pay for a hypothetical emergency expense of just $400. Only a little more than half (54 percent) of U.S.

households reported that they could handle such a small emergency expense out of their pocket (paying for it using cash, checking or savings account balances, or a credit card that they would pay in full at their next statement). Almost half of U.S. households (the remaining 46 percent) answer that such a small emergency expense would be a financial challenge and that they either could not pay the expense or would borrow or sell something to do so. This large share is similar to that seen in 2014, when 47 percent answered that they would not be able to pay for a $400 expense out of pocket (Board of Governors of the Federal Reserve System 2016). Among the respondents with financial difficulty, 38 percent would use a credit card that they pay off over time and 31 percent simply could not cover the expense. Over a quarter would borrow from friends or family, and smaller fractions would either sell something, use a payday loan, bank overdraft, or bank loan. In short, a substantive fraction of low-income households have such low amounts of savings that they cannot cover unexpected expenses of just $400. The low U.S. savings rate in turn is a possible reason for the economic uncertainty that households face, especially in their preparedness to address shocks to income or expenses.

A low U.S. savings rate, compared to foreign trade partners, in turn is also a cause of U.S. current account deficits, which potentially further erode the real incomes of low skilled workers. This interplay of possibly self-aggravating factors offers a candidate explanation for the unusually adverse real-wage changes in the United States when workers move out of manufacturing (Table 4 in subsection 6.1): the contraction of low-skill intensive and rent-sharing manufacturing industries depresses not only low-skilled wages but also savings rates, which in turn lead to a further contraction of traded-goods industries in the country with a current-account deficit, while the expanding services sector can hire low-skilled switchers out of manufacturing at depressed wages. In countries with high households savings rates such as the Republic of Korea or Japan, this potentially self-aggravating process is arguably less severe (but the import-competition-to-savings channel would need to operate differently in a surplus-economy like Germany where there are neither real-wage losses nor gains for sector switchers).

There is little direct and systematic empirical evidence of a nexus from import competition to household savings rates. Coeurdacier, Guibaud, and Jin (2015) document a similar channel as the one described but for demographic differences in the economy's age structure and differences in household credit constraints. Coeurdacier, Guibaud, and Jin show that important observable global imbalances follow from the savings channel if household credit constraints bind more severely in fast-growing emerging economics: household saving rates diverge between industrialized and emerging economies, and capital flows from the emerging to the advanced economies result.[43] Their model explains about one-third of the divergence in aggregate saving rates between advanced and emerging economies. The heterogeneity in savings rates between income groups, beyond age groups, might further add to the observed divergence in aggregate saving rates between economies.

[43]Savings are, of course, not the only cause of capital flows. Jin (2012) shows on the investment side that capital tends to flow toward countries that specialize in capital-intensive industries, when a financial market is embedded into a classic trade model, and can also generate substantive current account imbalances between industrialized and emerging economies.

6.5 Technology diffusion through trade in capital goods

Only a few countries produce much of the world's capital equipment and, as a consequence, many countries import a large share of their machinery. Poorer countries are more dependent on capital goods (Mutreja, Ravikumar, and Sposi 2014 report a correlation coefficient of -0.27 between capital goods imports and per-capita incomes): Malawi imports 47 times as many capital goods as it produces, Argentina imports twice as many as it makes, while the United States imports only half as many as it manufactures. At the same time, research suggests that capital-skill complementarity is an important feature of the technology embedded in capital equipment, as discussed before in subsection 4.1 above (for an early review see Katz and Autor 1999; for more recent findings consider Autor, Katz, and Kearney 2008 and Acemoglu and Autor 2011). Burstein, Cravino, and Vogel (2013) compute for the year 2000 that four-fifth of the world's capital equipment manufacturing took place in just eight countries: the United States, Japan, Germany, China, France, the Republic of Korea, the United Kingdom, and Italy. Conversely, the share of domestic absorption imported from abroad in the equipment sector in the year 2000 was 73 percent in the United Kingdom, 81 percent in Australia, 84 percent in Chile, and 96 percent in Cameroon.

Burstein, Cravino, and Vogel (2013) specify a model to quantify the effect of international trade on capital equipment stocks and, through capital-skill complementarity, the skill premium—combining a canonical classic trade model with an aggregate production function that exhibits capital-skill complementarity. Given the lasting economic debate on the magnitude, or existence, of capital-skill complementarity, Burstein, Cravino, and Vogel are cautious to emphasize that their quantitative analysis should only be taken as suggestive. In a main counterfactual exercise, Burstein, Cravino, and Vogel (2013) take technologies and factor endowments as given and putatively raise all trade costs to infinity, so they can quantify how much each country's real wages and skill premia would change when the economy is hypothetically moved to autarky. Real wages fall for both skilled and unskilled workers in all countries, and fall relatively more in countries that used to be more open. Second, the losses from moving to autarky are unevenly distributed within countries: skilled workers lose disproportionately, suffering a 2.4 times stronger decline in the skill wage in the median country. The real-wage loss is larger for countries that are highly dependent on imports of capital equipment, such as Cameroon and the Czech Republic, whereas the decline in the skill premium is negligibly small for Japan and the United States (but their average real-wage loss is large).

While inequality within countries matters, in particular for attitudes towards globalization and the formulation of trade policy, the cross-country dispersion of real incomes remains a concern for development. Mutreja, Ravikumar, and Sposi (2014) consider how international trade in capital goods affect per-capita incomes through capital formation and labor productivity. They propose, and quantify, a model of the global economy, combining productivity differences and Ricardian trade with production from two factors, capital and labor. Trade barriers and sectoral productivities affect how much of a country's capital stock originates from domestic equipment manufacturing and how much from trade, which in turn affects the amount of capital per worker in a country. They calibrate the model to match bilateral trade (in capital goods and intermediate goods), the observed relative prices of capital goods, and the global dispersion of per-capita incomes, and capture main facts of the global economy: the price of capital goods is roughly the same across countries but the relative price of capital is higher in poorer countries because the price of the nontradable consumption good is lower in poorer countries, and the investment rate measured in domestic prices is uncorrelated with per-capita income (Restuccia and Urrutia 2001) while the investment rate measured in international prices is positively correlated with per-capita incomes (Hsieh and Klenow 2007). In Mutreja, Ravikumar, and Sposi's main counterfactual exercise, the cross-country income differences decline by more than half when trade frictions are eliminated, with four-fifth of the change in each country's income attributable to change in its capital stock. If the world hypothetically moved to autarky in capital goods, poor countries would suffer a real income loss of 16 percent. Anecdotes can shed light on the prevalence of the equipment-imports channel for development. Nam (1995) tracks the Republic of Korea's imports of capital goods over a period of 40 years since the 1950s: while the relative price of capital in the Republic of Korea fell by roughly half, its investment rate more than quadrupled. Hsieh (2008) describes the more recent case of India during the 1990s: when India reduced its barriers to capital goods imports, the relative price of capital between 1990 and 2005 fell by one-fifth and the Indian investment rate increased 1.5-fold.

Trade in capital goods is also an area where the global trade pattern is undergoing marked changes since the Great Recession. Weak demand for capital goods and investment helps explain the global trade slowdown and strongly contributed to the contraction of global merchandize exports by one-tenth between 2011 and 2015.[44] The low rate of robot penetration (industrial robots per thousand workers) in the United States at around the 30th percentile of European penetration (Acemoglu and Restrepo 2017) has a counterpart in the U.S. trade pattern of advanced machinery. The United States runs a trade surplus in factory equipment and machinery, but its exports are mainly components and less advanced machinery, shipped to developing countries, while U.S. imports are concentrated in more advanced capital goods.[45]

6.6 Trade-induced technical change and technology-induced trade

Economy-wide connections between structural change, market structure, the savings behavior, and global technology diffusion arguably remain aspects of economic change for which evidence is relatively scant, especially when it comes to their mediating effects as globalization, technological progress and advances in prosperity interact. In contrast, at

[44] Wall Street Journal, March 30, 2017 ("Globalization Backers Face End of an Era").

[45] Wall Street Journal, March 27, 2017 ("Driving U.S. Factories: Foreign Robotics"). Tesla Inc., for example, a U.S. electric car and solar equipment manufacturer, has equipped its factory in Fremont, California, with robotic machinery from Kuka in Germany, now Chinese owned, and has bought the German factory-automation supplier Grohmann GmbH to equip its battery factory in Nevada. Some U.S. machinery suppliers, in turn, are acquiring foreign producers. The U.S. automation-testing equipment manufacturer Teradyne Inc. bought Denmark's Universal Robots A/S in 2015 and General Electric Co bought 3-D metal printing companies in Sweden and Germany in 2016.

the microeconomic level of firms and their workforces or industry-occupations and their workers, empirical economic research has built up a relatively large body of evidence on the interplay between technological change and trade.[46]

Would the iPhone exist in the absence of the deep trade integration between East Asia and the rest of the world? The bulk of assembly and production tasks for Apple products—iPhones, iPods and iPads—are offshored to East Asia (Linden, Dedrick, and Kraemer 2011). Without offshoring opportunities, it may not have been profitable for Apple to introduce the current varieties of iPhones and iPods because higher labor costs elsewhere might have raised production costs, and final product prices, so that consumer demand would not have sufficed for large-scale production. In terms of labor-market consequences, a lack of access to parts manufacturers in East Asia would therefore likely also have reduced the demand for high-skill engineering and design jobs at Apple. In this regard, globalization can lead to a positive relationship between offshoring and skill-biased technological change (technological change that favors demand for highly skilled workers). As Acemoglu and Restrepo (2017) point out in a rigorous model of offshoring and (Ricardian) trade, a complete account of the implications when offshoring opportunities are absent also needs to consider that Apple may have designed iPhones, iPods and iPods differently when faced with higher production costs at suppliers elsewhere, in order to reduce its dependence on more expensive unskilled labor outside East Asia. Somewhat paradoxically, the presence of offshoring opportunities therefore can imply a *higher* demand for low-skilled workers in the United States. This additional channel opens the possibility of a novel link between offshoring and skill-biased technological change—available cost savings from offshoring can raise an industry's product-market size and thus enable demand for domestic low-skilled labor which, in the absence of viable consumer goods varieties, would go away.

In a comment on Linden, Dedrick, and Kraemer (2011), Alberro (2012) pursues some instructive back-of-the-envelope computations along those lines for the Apple iPod. Linden, Dedrick, and Kraemer assessed the direct effects of the global value chain "that designs, builds, and brings iPods to consumers." They considered the jobs and wages that an iPod sustains, concluding that "in 2006, the iPod supported nearly twice as many jobs

[46]The questions about the relationship between new forms of trade and productivity change are reminiscent of an older, now less active, literature that explored the elusive link between openness and growth. Influential publications include Rivera-Batiz and Romer (1991), Grossman and Helpman (1991), Young (1991), Grossman and Helpman (1993), Lucas (1993), Frankel and Romer (1999), and Rodríguez and Rodrik (2000). Over the past half century, economies such as the Republic of Korea, Hong Kong (China), Chinese Taipei (the Separate Customs Territory of Taiwan, Penghu, Kinmen and Matsu) and Singapore have had sustained high growth rates and have undergone substantial economic transformations. Lucas (1993) called the Republic of Korea a miracle, where between 1960 and 1988 annual real per-capita income growth averaged 6.2 percent. That growth contrasted with a world average growth of just 1.8 percent over the same period. The South East Asian economies were also characteristically open to the global economy. In its 1987 classification of developing countries according to trade orientation, the World Bank listed only the Republic of Korea, Hong Kong (China) and Singapore as strongly outward oriented economies over the 1963-1985 period. More recently, annual output growth in China averaged 9.7 percent between 1989 and 2017, accompanied by a dramatic simultaneous reduction in poverty. China's openness to global trade and capital flows, though clearly politically managed, contributed to what is called the China shock elsewhere in this report.

offshore as in the United States. Yet the total wages paid in the United States amounted to more than twice as much as those paid overseas." Alberro (2012) adds additional considerations. The direct computations by Linden, Dedrick, and Kraemer (2011) incorporate neither the effect of the iPod's introduction on demand for parts and services in upstream industries (that supply inputs) nor in downstream industries (that provide sales and after-sales services). While far from fully conclusive, the illustrative calculations by Alberro (2012) suggest that the total number of jobs in the United States sustained by the introduction of the iPod increases significantly: once accounting for cross-industry linkages, by Alberro's measure the iPod supports two and one-half times as many jobs in the United States as Linden, Dedrick, and Kraemer calculated.

Those back-of-the-envelope computations cannot account for the far more complex, and arguably crucial, counterfactual question how technology adoption and relative demand for skills would differ in the absence of offshoring. Are offshoring and technology adoption complements or substitutes? As falling trade costs permit firms to perform some production tasks offshore, production factors employed onshore (in the domestic economy) become more productive. Reduced trade barriers may thus cause simultaneous growth in productivity and trade, so that technology adoption and offshoring are complements. However, an alternative strand of reasoning that links technology and trade recognizes that many jobs, especially in occupations with routine and codifiable tasks, are suitable for automation, technological upgrading and organizational change but, at the same time, also highly susceptible to offshoring (Blinder 2009). In that alternative view, trade and technology act on the labor-market outcome as if they were substitutes, each potentially reducing jobs that are intensive in routine and codifiable tasks.

6.6.1 Final goods trade and technology

Just as imports of consumer goods reduced retail prices in developed economies, prices of intermediate inputs for industrial uses have come down and new varieties become available. The less expensive imported intermediate and capital goods reduce production cost for incumbents on the upside, but heightened import competition in final product markets tend to reduce domestic incumbents' market shares on the downside. Models that link international trade to directed technological change with a skill bias include Acemoglu (2003) and Thoenig and Verdier (2003), for example. These papers show how import competition or export opportunities, or both, can bring about innovations that raise demand for skill, thus amplifying the direct impact of international trade on the wage structure through induced technological change.[47]

China's trade growth, and its WTO accession in 2001, constitute a shift in import competition that several studies consider exogenous to domestic economic conditions. In this context, Bloom, Draca, and Van Reenen (2016) document that increased Chinese trade has induced faster technological change among European incumbent firms from both innovation and the adoption of new technologies, contributing to productivity growth

[47] A related channel occurs under increasing returns to scale: global trade integration increases firms' market sizes and can lead to higher returns to education, and thus raise wage inequality, if the returns from additional economies of scale accrue to high-skilled workers (Epifani and Gancia 2008).

beyond mere production cost savings. Their results suggest that trade stimulates technical progress in Europe, which in turn raises demand for skilled labor. Bloom, Draca, and Van Reenen study the impact of imports on a broad range of innovation-related activities: patent applications, investments in information technology, the conduct of research and development, changes in total factor productivity and the adoption of new management practices. They infer that, over the period 2000-2007, import competition from low wage countries like China has a greater effect on innovation than imports from high wage countries and that China accounts for almost 15 percent of the increase in European patenting and 12 percent of productivity change. These effects are not uniform, however. The average firm suffers a strong negative employment effect of Chinese imports (a 10 percentage point increase in Chinese imports is associated with a 3.5 percent drop in employment), but the high-tech firms that innovate grow in workforce size. In contrast to the findings for European firms, Autor, Dorn, Hanson, Pisano, and Shu (2016), who focus their study on patenting, find for U.S. industries and firms that their rate of patenting slows down when they are subject to faster increases in exposure to Chinese import competition.

6.6.2 Offshoring, multinational production and technology

Burstein, Morales, and Vogel (2015) explore the interactions between offshoring, in the form of capital and occupational services imports, and technological change in a number of the empirical exercises. Evidence from the Burstein, Morales, and Vogel (2015) study was also discussed above in subsection 5.4 as a contribution to trade research that explicitly conditions on technological change. Burstein, Morales, and Vogel (2015, Table 3) report that the wage premium for college educated workers increased by 16 percent between 1984 and 2003. In the absence of labor demand changes, the additional supply of college educated workers would have reduced the wage premium for college educated workers. Instead, labor demand shifts raised the skill premium by 30 percent. As mentioned before, Burstein, Morales, and Vogel (2015) estimate that computer use explains almost three-fifth of this demand effect behind the skill premium (about 60 percent), while almost one-fifth (18.5 percent) of the demand effect was due to occupational demand shifts. Labor productivity and other sources account for the remaining (residual) change in the demand effect. In Burstein, Morales, and Vogel's structural model, the skill premium responds to computer use through two channels: the comparative advantage of educated workers in computer use, and the comparative advantage of educated workers in computer-intensive occupations.

Burstein, Morales, and Vogel ask the counterfactual question how the absence of trade may have altered computer use and the occupational shift and therefore labor demand. They estimate that if the U.S. economy had not been integrated into the global markets for equipment (including computers), then between 1984 and 2003 the rise in the skill premium would have been lower by 2.1 percentage points, and if the U.S. economy had not traded occupation services, then the skill premium would have been lower by 1.3 percentage points (Tables 7 and 8). Each of these estimates represents only a fraction, 13 percent and 27 percent, of the contribution of computers and occupations to the rise in the skill premium. Helpman (2016) concludes from this evidence that international trade had only a modest impact on U.S. wage inequality.

In theory, the effect of offshoring on the skill premium through technological change is ambiguous. Acemoglu, Gancia, and Zilibotti (2015) show in a Ricardian trade model with directed technological change and with offshoring that the effect of a reduction in offshoring costs on wages in industrialized countries depends on the varying effects that offshoring can have on technological change. In particular, Acemoglu, Gancia, and Zilibotti derive theoretically that, when the offshoring cost is high initially, then an increase in offshoring opportunities causes a fall in the real wages of unskilled workers in industrial countries, skill-biased technological change and a rising skill premium. However, when the offshoring cost becomes sufficiently low, offshoring induces technological change biased in favor of the unskilled workers, because offshoring expands the market size of technologies complementary to unskilled workers. Taking the model to data from the United States and China in 2000, the U.S. economy does not appear to be a low offshoring cost case yet. Acemoglu, Gancia, and Zilibotti (2015) evaluate varying scenarios for offshoring. Compared to a global economy with no offshoring, the trade-induced skill biased change raises the skill wage premium one-and-a-half times more than it would have counterfactually been in the absence of offshoring. In another scenario, Acemoglu, Gancia, and Zilibotti compute that a reduced offshoring cost that raises U.S. offshoring imports by one-fifth can boost the U.S. growth rate, raising welfare for workers across all skill groups, but resulting in an unequal distribution of those gains in terms of consumption variations: while Chinese workers can afford a more than 26 percent higher consumption, U.S. low-skilled workers gain only one percent, and U.S. high-skilled workers 8 percent. These welfare gains mostly emanated from a simulated boost to economy-wide innovation and growth, brought about by offshoring-induced skill biased change.

There are more aspects to interactions between trade and technology adoption or innovation—beyond intermediate goods, capital and occupational services imports. Product market access abroad can raise a firm's profitability while foreign product market competition in the domestic market can reduce a firm's profitability—altering incentives to install technology and adopt innovations. Depending on whether global competition weakens incentives for innovation (the patent argument) or strengthens the incentives (the competition tenet), trade in final goods will alter technology and skill demand beyond capital imports. This caveat also applies to other studies of the interaction between offshoring and technological change. As discussed before in subsections 4.1.3 and 4.1.4, multinational production can take the form of in-house offshoring when a multinational enterprise owns its foreign supplier. This type of multinational production is also called vertical foreign direct investment, because the foreign supplier is considered vertically related to the subsequent production stages in the MNE's value chain. Multinational production can also take the form of production replication, as an MNE may seek proximity to foreign clients by locating a copy of its production unit abroad, giving up on economies of scale from concentration of its production in a single central location in lieu of better market access for its products. This form of multinational production is also called horizontal foreign direct investment, because the foreign affiliate can be considered a horizontal competitor to the MNE's own domestic production unit. Horizontal foreign direct investment substitutes trade, since it displaces the exports that the central unit would otherwise ship to the foreign market. As a consequence, the interaction between multinational production, trade and technology can potentially exhibit complex patterns, which to our knowledge have not been the subject of theoretical or empirical investigation.

6.6.3 Firm-level global integration and technology

In research on matching in the labor market, discussed above (Section 4.3), employers did not have the choice to invest in technologies. However, incentives to trade and incentives to adopt technology can relate in important ways. Naturally, an exporter can benefit from a cost-saving or quality-improving technology investment in more product markets than a non-exporter, so exporting raises the returns to technology adoption. Conversely, a more capable firm with cost-saving or quality-improving technology in place will experience higher returns to exporting since the technology permits a deeper penetration of any product market, so technology adoption raises the returns to exporting.[48]

Yeaple (2005) analyzes theoretically the matching of workers with different abilities to firms with different capabilities. Firms are ex ante identical but can choose one of two capability levels. A firm's capability originates in one of two production technologies: one technology comes with a high fixed cost but low variable cost of production (high capability), and the other technology is associated with a low fixed cost but high variable cost of production (low capability). There are two sectors in the economy. A traditional sector supplies a homogeneous good, for which firms have no technology choice; labor productivity simply rises with worker ability. In the advanced sector, in contrast, firms can adopt one of the two technologies to produce varieties of a differentiated product; labor productivity in the advanced sector rises with worker ability for each technology. Yeaple (2005) shows that, in a competitive equilibrium, there is sorting of firms into capability choices and positive assortative matching of workers by ability to firms. If some but not all firms export, then the firms that adopt low capability must serve only the domestic market, and the firms that select into exporting must be adopters of high capability. Selection into both exporting and capability results in a sharp distinction between exporters and non-exporters in two dimensions: exporters have high capability and are matched with more able workers. In this model, opening to trade results in a wage outcome that exhibits polarization: the wage rises for high-ability workers who match with highly capable firms, wages decline for middle-ability workers who match with little capable firms, and wages do not change for workers in the traditional sector.

Bustos (2011) takes this rationale to the data, but she omits the homogeneous sector from her model, so there is no potential polarization (just a skill premium), and she makes firms ex ante heterogeneous in their elemental productivity. To simplify matching in the theory, and data work in practice, Bustos considers only two types of workers, highly skilled and low-skilled workers. Both types of workers are indispensable in production.

[48] An example of mutually reinforcing returns is in fact present in Section 4.3: in the Helpman, Itskhoki, and Redding (2010) model investments in the screening technology for higher ability workers raise the returns to exporting, and vice versa. An example more closely related to conventional technology adoption is the model by Aw, Roberts, and Xu (2011), where R&D investments raise the returns to exporting, and vice versa.

Technologies differ in factor intensities: the high-capability technology uses more skilled relative to unskilled workers. With technology adoption, a firm can scale its elemental productivity, but firms with any elemental productivity level gain the same proportional variable-cost advantage if they adopt the technology with low production costs.

Bustos (2011) uses this framework to study the improved market access for Argentine firms to Brazil with the formation of MERCOSUR in 1991, a customs union with the initial members Argentina, Brazil, Paraguay, and Uruguay. A reduction in variable trade costs leads firms with the lowest elemental productivity to exit, non-exporters with relatively high elemental productivity to become exporters, and induces those exporters that used to use the low-capability technology to adopt the high-capability technology. As a consequence of these within-industry reallocations, the relative demand for skilled workers increases and the skill premium rises, both because a larger fraction of firms adopts the high-capability technology and because firms with the high-capability technology expand while those with the low-capability technology shrink. The Brazilian tariff cuts on Argentine imports with the formation of MERCOSUR varied across industries, affecting initially different firms differentially. Bustos finds that among firms above median size 76 percent exported, while only 38 percent of firms among those below median size were exporters. Her estimates suggest that the average reduction of the Brazilian tariffs on Argentine imports of around 23 percent brought about an 8-percent reduction in the share of skilled workers in firms below the median and an increase of 6 percent in the share of skilled workers in firms above the median—a marked shift in labor composition across firms.

6.6.4 Firm-level global integration and organizational change

Aspects of the internal labor market and residual wage inequality are difficult to observe directly. Recent studies of the firm's internal labor market have turned to the importance of hierarchical layers (Caliendo and Rossi-Hansberg 2012; Caliendo, Monte, and Rossi-Hansberg 2015) and their response to firm-level trade. The internal organization of plants and firms also involves the motivation of workers to exert effort. Related studies analyze the response of employers' incentives for workers, and observable incentive pay in particular, when global competition changes (Guadalupe 2007; Cunat and Guadalupe 2009).[49]

[49] Antràs, Garicano, and Rossi-Hansberg (2006) provide a theoretical rationale for changes in the matching between workers and managers when offshoring changes. In the managerial hierarchy, every manager has a team of workers, and workers differ in ability. Worker ability determines the range of production problems that the worker can solve. An unresolved problem results in zero production. A more able worker can solve all the problems of a less able worker, and more. If a problem occurs that is outside a worker's range of problem solving, the worker hands the problem over to the manager and the manager solves the problem if it is within in her range. A consequence of the hierarchical setup is that a worker's ability and a manager's ability are complements, so in a competitive market there is positive assortative matching. Final goods trade leads to a rematching of workers and managers around the globe, which changes the distribution of income. Antràs, Garicano, and Rossi-Hansberg show that trade raises wage inequality among Southern workers. Trade can reduce wage inequality among Northern workers if management consumes much time and the skill gap between North and South is small.

7 Conclusions

Does globalization or does technological change better explain the labor-market outcomes that we are witnessing? It seems, the largely separate lines of economic literature on international trade and on technological change keep coming back to numbers that suggest neither one explanation can account for more than one-quarter to one-third of the variation in labor-market outcomes.

When studying U.S. labor-market outcomes in the 1980s, Feenstra and Hanson (1999) found that at most two-fifth of the observed change in the wage-bill share of nonproduction workers in the United States could be explained by offshoring alone, measured as imports of intermediate inputs. But their estimates of the impact of offshoring were sensitive to a measure of high-tech equipment use. Once technology adoption was included, about one-quarter of the wage differential paid to nonproduction workers during the 1980s could be attributed to offshoring and around one-third to technology. When Feenstra and Hanson (1999) altered their measure of high-tech equipment to place more weight on more recently installed (more advanced) equipment, the contribution of offshoring to wage outcomes dropped by half and the contribution of technology more than tripled. Their study has remained one of only a few that undertake an explicit comparison between a trade-related variable and a technology-related variable in predictions of labor-market outcomes. Most of the contributions in the two literatures on trade and on technological change measure one economic shock in isolation and treat the other as part of the unexplained residual.

For the 1990s and 2000s, two globalization trends have become prominent objects of study: trade in tasks (the offshoring of routine and less cognitively demanding tasks) and its possible association with job polarization (a relative decline in labor-market prospects for workers in the middle range of the earnings distribution compared to both lower- and higher-income groups), and the China shock (China's accession to the WTO in 2001 followed by a substantive Chinese trade surplus) and its possible association with manufacturing employment declines elsewhere. Using evidence from comparisons between U.S. regions, and treating them as indicative of nationwide outcomes, Autor, Dorn, and Hanson (2013) argue that the China shock explains one-quarter of the decline in U.S. manufacturing employment. In light of their methodological shortcomings and recent additional evidence, the number of one-quarter is likely too large. In a structural estimation approach for 38 countries, 50 U.S. states, and 22 industries, Caliendo, Dvorkin, and Parro (2015) find that the China shock can only account for one-quarter of the one-quarter in Autor, Dorn, and Hanson (2013), or one-sixteenth of the decline in U.S. manufacturing employment. Ebenstein, Harrison, and McMillan (2015) turn their attention away from U.S. regions to nationwide worker transitions between occupations and find only minor effects of globalization on labor force participation, but large effects of computer use and other equipment investment on employment. Coming at the problem from the other side, that of technological change alone, Cortes, Jaimovich, and Siu (2016a) calibrate a demanding structural model of automation to the U.S. economy for 1979-2014 and find that automation explains at most one-third of the decline in middle-skilled employment (when requiring the model to match worker transitions between occupations). Different times, different methods, different economic shocks, different outcomes, and we still find that one-quarter or one-third of the labor-market outcomes go to each of the two candidate explanations, globalization and technological change. Summing one-quarter and one-third adds to just over one-half, meaning that the dominant part of the variation in labor-market outcomes remains unexplained by either technology or trade.

What accounts for the other, dominant part of the variation in labor-market outcomes that remains to be explained? One hypothesis is that treating trade and technological change in isolation may miss important interactions between the two forces and their potential mutual reinforcement. A rare study that assesses multiple sources of change simultaneously is Burstein, Morales, and Vogel (2015). They use U.S. data at the occupation level in a structural model and find that computerization and changes in task demand jointly explain four-fifth of the rise in the U.S. skill premium of college educated workers' wages over less educated workers' wages. While computerization and changes in task demand are arguably technological changes in nature, they might still be attributable in part to globalization because computerized capital goods are traded and task inputs can be sourced offshore. However, Burstein, Morales, and Vogel find that those forms of trade represent only between one-eighth and one-quarter of the contribution of computers and occupations to the rise in the U.S. skill premium. Should we now count out globalization? Probably not. Trade in intermediate inputs (computerized capital) and intermediate services (tasks) is perhaps more recent than classic trade in final goods, but final goods trade continues to matter and can relate to technological change in complex ways. Depending on whether global product-market competition weakens incentives for innovation (the patent argument) or strengthens the incentives (the competition tenet), trade in final goods will alter technology and skill demand.

How plausible is it that rich interactions between globalization and technological change might account for the unexplained part of the variation in labor-market outcomes? Perhaps that part, which accounts for most of the decline in manufacturing employment and middle-range occupations, originates simply in deindustrialization. Labor saving productivity change in agriculture has cut the agricultural workforce over the past decades and centuries, while today we feed a larger world population with richer diets than ever before. It may now be the turn of manufacturing. As per-capita incomes rise and our consumption baskets become lighter, filled with more services and fewer goods, deindustrialization can be expected to remain a global source of change. In this light, another hypothesis that relates the so far unexplained part of the variation in labor-market outcomes to globalization and technological change is that both globalization and technological change raise prosperity, and prosperity in turn makes our consumption baskets lighter. Like with any form of economic change, inevitably there will be winners and losers. We may therefore need to look into redistributive measures in new ways. There is no reason, however, why our solutions for compensating those who would otherwise lose should decelerate technological change or globalization, because both technological change and globalization make our diets, consumption goods, and services more affordable to everyone.

References

ABOWD, J. M., R. H. CREECY, AND F. KRAMARZ (2002): "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," *LEHD Technical Paper*, 2002-06.

ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67(2), 251–333.

ABOWD, J. M., AND I. M. SCHMUTTE (2010): "Endogenous Mobility," Cornell University, unpublished manuscript.

ACEMOGLU, D. (2003): "Patterns of Skill Premia," *Review of Economic Studies*, 70(2), 199–230.

ACEMOGLU, D., P. AGHION, AND F. ZILIBOTTI (2006): "Distance to Frontier, Selection, and Economic Growth," *Journal of the European Economic Association*, 4(1), 37–74.

ACEMOGLU, D., AND D. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of labor economics*, ed. by D. Card, and O. Ashenfelter, vol. 4B, chap. 12, pp. 1043–1171. Elsevier Science, North-Holland, Amsterdam, New York and Oxford.

ACEMOGLU, D., D. AUTOR, D. DORN, G. H. HANSON, AND B. PRICE (2016): "Import Competition and the Great U.S. Employment Sag of the 2000s," *Journal of Labor Economics*, 34(1 Part 2), S141–98.

ACEMOGLU, D., G. GANCIA, AND F. ZILIBOTTI (2015): "Offshoring and Directed Technical Change," *American Economic Journal: Macroeconomics*, 7(3), 84–122.

ACEMOGLU, D., AND P. RESTREPO (2017): "Robots and Jobs: Evidence from U.S. Labor Markets," *NBER Working Paper*, 23285.

ADÃO, R. (2016): "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil," Princeton University and University of Chicago, unpublished manuscript.

ADÃO, R., C. ARKOLAKIS, AND F. ESPOSITO (2017): "Trade Openness, Agglomeration Effects, and Endogenous Labor Supply: Theory and Evidence," Yale University, unpublished manuscript.

ADÃO, R., A. COSTINOT, AND D. DONALDSON (2017): "Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade," *American Economic Review*, 107(3), 633–89.

AGHION, P., N. BLOOM, R. BLUNDELL, R. GRIFFITH, AND P. HOWITT (2005): "Competition and Innovation: An Inverted-U Relationship," *Quarterly Journal of Economics*, 120(2), 701–28.

AKERMAN, A., E. HELPMAN, O. ITSKHOKI, M.-A. MUENDLER, AND S. J. REDDING (2013): "Sources of Wage Inequality," *American Economic Review: Papers & Proceedings*, 103(3), 214–19.

ALBERRO, J. (2012): "Innovation and Job Creation in a Global Economy: The Case of Apple's iPod: Comment," *Journal of International Commerce and Economics*, 4(0), 89–95.

ALVAREDO, F., L. CHANCEL, T. PIKETTY, E. SAEZ, AND G. ZUCMAN (2017): "Global Inequality Dynamics: New Findings from WID.world," *NBER Working Paper*, 23119.

ALVAREZ, F., AND R. E. J. LUCAS (2007): "General Equilibrium Analysis of the Eaton-Kortum Model of International Trade," *Journal of Monetary Economics*, 54(6), 1726–68.

AMITI, M., AND D. R. DAVIS (2012): "Trade, Firms, and Wages: Theory and Evidence," *Review of Economic Studies*, 79(1), 1–36.

AMITI, M., AND C. A. PISSARIDES (2005): "Trade and Industrial Location with Heterogeneous Labor," *Journal of International Economics*, 67(2), 392–412.

AMITI, M., AND S. J. WEI (2005): "Fear of Service Outsourcing: Is It Justified?," *Economic Policy*, 42, 307–39.

——— (2009): "Service Offshoring and Productivity: Evidence from the U.S.," *World Economy*, 32(2), 203–20.

ANTRÀS, P., L. GARICANO, AND E. ROSSI-HANSBERG (2006): "Offshoring in a Knowledge Economy," *Quarterly Journal of Economics*, 121(1), 31–77.

ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade Shocks and Labor Adjustment: A Structural Empirical Approach," *American Economic Review*, 100(3), 1008–1045.

ATKINSON, A. B., T. PIKETTY, AND E. SAEZ (2011): "Top Incomes in the Long Run of History," *Journal of Economic Literature*, 49(1), 3–71.

AUTOR, D., D. DORN, G. HANSON, G. PISANO, AND P. SHU (2016): "Foreign Competition and Domestic Innovation: Evidence from U.S. Patents," *CEPR Discussion Paper*, 11664.

AUTOR, D. H. (2014): "Skills, Education, and the Rise of Earnings Inequality Among the 'Other 99 Percent'," *Science*, 344(6186), 843–851.

AUTOR, D. H., AND D. DORN (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market," *American Economic Review*, 103(5), 1553–97.

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103(6), 2121–68.

——— (2015): "Untangling Trade and Technology: Evidence from Local Labour Markets," *Economic Journal*, 125(584), 621–46.

AUTOR, D. H., D. DORN, G. H. HANSON, AND J. SONG (2014): "Trade Adjustment: Worker Level Evidence," *Quarterly Journal of Economics*, 129(4), 1799–1860.

AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): "The Polarization of the U.S. Labor Market," *American Economic Review: Papers and Proceedings*, 96(2), 189–194.

——— (2008): "Trends in U.S. Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics*, 90(2), 300–323.

AUTOR, D. H., L. F. KATZ, AND A. B. KRUEGER (1998): "Computing Inequality: Have Computers Changed the Labor Market?," *Quarterly Journal of Economics*, 113(4), 1169–1213.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118(4), 1279–1333.

AW, B. Y., M. J. ROBERTS, AND D. Y. XU (2011): "R&D Investment, Exporting, and Productivity Dynamics," *American Economic Review*, 101(4), 1312–1344.

BALDWIN, R. E. (2016): *The great convergence: Information technology and the new globalization*. Harvard University Press, Cambridge, Mass.

BALDWIN, R. E., AND A. J. VENABLES (1995): "Regional Economic Integration," in *Handbook of international economics*, ed. by G. M. Grossman, and K. Rogoff, vol. 3 of *Handbooks in Economics*, chap. 31, pp. 1597–1644. Elsevier, North-Holland, Amsterdam, New York and Oxford.

BANK FOR INTERNATIONAL SETTLEMENTS (2015): *China's High Saving Rate: Myth and Reality*, vol. 312 of *BIS Working Papers*. Bank for International Settlements, Monetary and Economic Department, Basel, Switzerland, authored by Ma, Guonan and Wang Yi.

BARRO, R. J., AND X. SALA-I-MARTIN (1992): "Convergence," *Journal of Political Economy*, 100(2), 223–51.

BARTEL, A., C. ICHNIOWSKI, AND K. SHAW (2007): "How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills," *Quarterly Journal of Economics*, 122(4), 1721–58.

BAUMOL, W. J. (1967): "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crises," *American Economic Review*, 57(3), 415–26.

BECKER, G. S. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81(4), 813–46.

BECKER, S. O., H. EGGER, M. KOCH, AND M.-A. MUENDLER (2017): "Tasks, Occupations, and Wage Inequality in an Open Economy," University of California, San Diego, unpublished manuscript.

BECKER, S. O., K. EKHOLM, AND M.-A. MUENDLER (2013): "Offshoring and the Onshore Composition of Tasks and Skills," *Journal of International Economics*, 90(1), 91–106.

BECKER, S. O., AND M.-A. MUENDLER (2008): "The Effect of FDI on Job Security," *The B.E. Journal of Economic Analysis & Policy: Advances*, 8(1), Article 8, pp. 1–45.

——— (2015): "Trade and Tasks: An Exploration Over Three Decades in Germany," *Economic Policy*, 30(84), 589–641.

BERAJA, M., E. HURST, AND J. OSPINA (2016): "The Aggregate Implications of Regional Business Cycles," *NBER Working Paper*, 21956.

BERMAN, E., J. BOUND, AND Z. GRILICHES (1994): "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures," *Quarterly Journal of Economics*, 109(2), 367–97.

BERMAN, E., J. BOUND, AND S. MACHIN (1998): "Implications of Skill-Biased Technological Change: International Evidence," *Quarterly Journal of Economics*, 113(4), 1245–79.

BERMAN, E., AND S. MACHIN (2000): "Skill-Based Technology Transfer around the World," *Oxford Review of Economic Policy*, 16(3), 12–22.

BERNARD, A. B., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): "Plants and Productivity in International Trade," *American Economic Review*, 93(4), 1268–90.

BERNARD, A. B., AND T. C. FORT (2015): "Factoryless Goods Producing Firms," *American Economic Review*, 105(5), 518–23.

——— (2017): "Factoryless Goods Producers in the USA," in *The factory-free economy: Outsourcing, Servitization, and the Future of Industry*, ed. by L. Fontagné, and A. E. Harrison, chap. 5, pp. 136–168. Oxford Univeristy Press, Oxford.

BERNARD, A. B., AND J. B. JENSEN (1995): "Exporters, Jobs, and Wages in U.S. Manufacturing: 1976-1987," *Brookings Papers on Economic Activity: Microeconomics*, 1995(1), 67–112.

BERNARD, A. B., J. B. JENSEN, AND P. K. SCHOTT (2006): "Survival of the Best Fit: Exposure to Low-Wage Countries and the (Uneven) Growth of U.S. Manufacturing Plants," *Journal of International Economics*, 68(1), 219–37.

——— (2009): "Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. that Trade Goods," in *Producer Dynamics: New Evidence from Micro Data*, ed. by T. Dunne, J. B. Jensen, and M. J. Roberts, vol. 68 of *Studies in Income and Wealth*, chap. 14, pp. 513–552. University of Chicago Press, Chicago.

BERNARD, A. B., S. J. REDDING, AND P. K. SCHOTT (2007): "Comparative Advantage and Heterogeneous Firms," *Review of Economic Studies*, 74(1), 31–66.

BLANCHARD, O. J., AND L. F. KATZ (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992(1), 1–61.

BLINDER, A. S. (2006): "Offshoring: The Next Industrial Revolution?," *Foreign Affairs*, 85(2), 113–28.

——— (2009): "How Many U.S. Jobs Might Be Offshorable?," *World Economics*, 10(2), 41–78.

BLOOM, N., M. DRACA, AND J. VAN REENEN (2016): "Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity," *Review of Economic Studies*, 83(1), 87–117.

BLOOM, N., AND J. VAN REENEN (2011): "Human Resource Management and Productivity," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4B of *Handbooks in Economics*, chap. 19, pp. 1697–1767. Elsevier, Amsterdam.

BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM (2016): *Report on the Economic Well-Being of U.S. Households in 2015*. Board of Governors of the Federal Reserve System, Consumer and Community Development Research Section, Washington, DC, authored by Larrimore, Jeff, Sam Dodini and Logan Thomas.

BOEHM, C. E., A. FLAAEN, AND N. PANDALAI-NAYA (2015): "Multinationals, Offshoring, and the Decline of U.S. Manufacturing," Princeton University, unpublished manuscript.

BOMBARDINI, M., G. OREFICE, AND M. D. TITO (2015): "Does Exporting Improve Matching? Evidence from French Employer-Employee Data," *NBER Working Paper*, 21225.

BORJAS, G. J., AND V. A. RAMEY (1995): "Foreign Competition, Market Power, and Wage Inequality," *Quarterly Journal of Economics*, 110(4), 1075–1110.

BOSTON CONSULTING GROUP (2015): *The Robotics Revolution: The Next Great Leap in Manufacturing*. Boston Consulting Group, Boston, Mass., authored by Sirkin, Harold L., Michael Zinser and Justin Rose.

BOURGUIGNON, F. (2015): *The globalization of inequality*. Princeton University Press, Princeton and Oxford.

BRECHER, R. A. (1974): "Minimum Wage Rates and the Pure Theory of International Trade," *Quarterly Journal of Economics*, 88(1), 98–116.

BRODA, C., AND D. E. WEINSTEIN (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2), 541–85.

BRYNJOLFSSON, E., AND A. MCAFEE (2015): "Will Humans Go the Way of Horses? Labor in the Second Machine Age," *Foreign Affairs*, 94(4), 8–14.

BUGAMELLI, M., F. SCHIVARDI, AND R. ZIZZA (2010): "The Euro and Firm Restructuring," in *Europe and the Euro*, ed. by A. Alesina, and F. Giavazzi, NBER Conference Report series, chap. 3, pp. 99–138. University of Chicago Press, Chicago and London.

BURDA, M. C., AND B. SEVERGNINI (2009): "TFP Growth in Old and New Europe," *Comparative Economic Studies*, 51(4), 447–66.

BURDETT, K., AND D. T. MORTENSEN (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39(2), 257–73.

BURSTEIN, A. T., J. CRAVINO, AND J. VOGEL (2013): "Importing Skill-Biased Technology," *American Economic Journal: Macroeconomics*, 5(2), 32–71.

BURSTEIN, A. T., E. MORALES, AND J. VOGEL (2015): "Accounting for Changes in Between-Group Inequality," *NBER Working Paper*, 20855.

BURSTEIN, A. T., AND J. VOGEL (2016): "International Trade, Technology, and the Skill Premium," Columbia University, unpublished manuscript.

BUSTOS, P. (2011): "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *American Economic Review*, 101(1), 304–40.

CAHUC, P., F. POSTEL-VINAY, AND J. M. ROBIN (2006): "Wage Bargaining with On-the-Job Search: Theory and Evidence," *Econometrica*, 74(2), 323–64.

CALIENDO, L., M. DVORKIN, AND F. PARRO (2015): "The Impact of Trade on Labor Market Dynamics," *NBER Working Paper*, 21149.

Caliendo, L., R. C. Feenstra, J. Romalis, and A. M. Taylor (2015): "Tariff Reductions, Entry, and Welfare: Theory and Evidence for the Last Two Decades," *NBER Working Paper*, 21768.

Caliendo, L., F. Monte, and E. Rossi-Hansberg (2015): "The Anatomy of French Production Hierarchies," *Journal of Political Economy*, 123(4), 809–52.

Caliendo, L., and F. Parro (2015): "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies*, 82(1), 1–44.

Caliendo, L., and E. Rossi-Hansberg (2012): "The Impact of Trade on Organization and Productivity," *Quarterly Journal of Economics*, 127(3), 1393–1467.

Campbell, D. L. (2016): "Relative Prices, Hysteresis, and the Decline of American Manufacturing," New Economic School, Moskow, unpublished manuscript.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2016): "Firms and Labor Market Inequality: Evidence and Some Theory," University of California, Berkeley, unpublished manuscript.

Card, D., J. Heining, and P. Kline (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 128(3), 967–1015.

Caroli, E., and J. Van Reenen (2001): "Skill-Biased Organizational Change? Evidence from a Panel of British and French Establishments," *Quarterly Journal of Economics*, 116(4), 1449–92.

Charles, K. K., E. Hurst, and M. J. Notowidigdo (2016): "The Masking of the Decline in Manufacturing Employment by the Housing Bubble," *Journal of Economic Perspectives*, 30(2), 179–200.

Clerides, S. K., S. Lach, and J. R. Tybout (1998): "Is Learning by Exporting Important? Micro-dynamic Evidence from Colombia, Mexico, and Morocco," *Quarterly Journal of Economics*, 113(3), 903–47.

Coşar, A. K., N. Guner, and J. Tybout (2016): "Firm Dynamics, Job Turnover, and Wage Distributions in an Open Economy," *American Economic Review*, 106(3), 625–63.

Coeurdacier, N., S. Guibaud, and K. Jin (2015): "Credit Constraints and Growth in a Global Economy," *American Economic Review*, 105(9), 2838–81.

Comin, D. A., D. Lashkari, and M. Mestieri (2015): "Structural Change with Long-run Income and Price Effects," *NBER Working Paper*, 21595.

Cortes, G., N. Jaimovich, C. J. Nekarda, and H. E. Siu (2014): "The Micro and Macro of Disappearing Routine Jobs: A Flows Approach," *NBER Working Paper*, 20307.

Cortes, G., N. Jaimovich, and H. E. Siu (2016a): "Disappearing Routine Jobs: Who, How, and Why?," *NBER Working Paper*, 22918.

——— (2016b): "The End of Men and Rise of Women in the High-Skilled Labor Market," University of Southern California,, unpublished manuscript.

COSTINOT, A., AND A. RODRÍGUEZ-CLARE (2014): "Trade Theory with Numbers: Quantifying the Consequences of Globalization," in *Handbook of International Economics*, ed. by E. Helpman, K. Rogoff, and G. Gopinath, vol. 4, chap. 4, pp. 197–261. Elsevier, Amsterdam.

COSTINOT, A., AND J. VOGEL (2010): "Matching and Inequality in the World Economy," *Journal of Political Economy*, 118(4), 747–86.

COSTINOT, A., J. VOGEL, AND S. WANG (2012): "Global Supply Chains and Wage Inequality," *American Economic Review: Papers and Proceedings*, 102(3), 396–401.

——— (2013): "An Elementary Theory of Global Supply Chains," *Review of Economic Studies*, 80(1), 109–44.

CRESPI, G., C. CRISCUOLO, AND J. HASKEL (2008): "Productivity, Exporting and the Learning-by-Exporting Hypothesis: Direct Evidence from U.K. Firms," *Canadian Journal of Economics*, 41(2), 619–638.

CRINÒ, R. (2010): "Service Offshoring and White-Collar Employment," *Review of Economic Studies*, 77(2), 595–632.

CROZET, M., AND E. MILET (2017): "The Servitization of French Manufacturing Firms," in *The factory-free economy: Outsourcing, Servitization, and the Future of Industry*, ed. by L. Fontagné, and A. Harrison, chap. 4, pp. 111–135. Oxford Univeristy Press, Oxford.

CUNAT, V., AND M. GUADALUPE (2009): "Globalization and the Provision of Incentives inside the Firm: The Effect of Foreign Competition," *Journal of Labor Economics*, 27(2), 179–212.

DAS, M., AND P. N'DIAYE (2013): "Chronicle of a Decline Foretold: Has China Reached the Lewis Turning Point?," in *China's Economy in Transition: From External to Internal Rebalancing*, ed. by A. Singh, M. Nabar, and P. N'Diaye, chap. 9, pp. 157–79. International Monetary Fund, Washington, D.C.

DAUTH, W., S. FINDEISEN, AND J. SUEDEKUM (2014): "The Rise of the East and the Far East: German Labor Markets and Trade Integration," *Journal of the European Economic Association*, 12(6), 1643–75.

DAVIDSON, C., F. HEYMAN, S. MATUSZ, F. SJÖHOLM, AND S. C. ZHU (2014): "Globalization and Imperfect Labor Market Sorting," *Journal of International Economics*, 94(2), 177–94.

DAVIDSON, C., L. MARTIN, AND S. MATUSZ (1988): "The Structure of Simple General Equilibrium Models with Frictional Unemployment," *Journal of Political Economy*, 96(6), 1267–93.

DAVIDSON, C., S. J. MATUSZ, AND A. SHEVCHENKO (2008): "Globalization and Firm Level Adjustment with Imperfect Labor Markets," *Journal of International Economics*, 75(2), 295–309.

DAVIS, D. R., AND J. HARRIGAN (2011): "Good Jobs, Bad Jobs, and Trade Liberalization," *Journal of International Economics*, 84(1), 26–36.

DEARDORFF, A. ., AND R. W. STAIGER (1988): "An Interpretation of the Factor Content of Trade," *Journal of International Economics*, 24(1/2), 93–107.

DEKLE, R., J. EATON, AND S. KORTUM (2007): "Unbalanced Trade," *American Economic Review: Papers and Proceedings*, 97(2), 351–55.

DEMYANYK, Y., D. HRYSHKO, M. J. LUENGO-PRADO, AND B. E. SØRENSEN (2017): "Moving to a Job: The Role of Home Equity, Debt, and Access to Credit," *American Economic Journal: Macroeconomics*, 9(2), 149–181.

DI GIOVANNI, J., A. A. LEVCHENKO, AND J. ZHANG (2014): "The Global Welfare Impact of China: Trade Integration and Technological Change," *American Economic Journal: Macroeconomics*, 6(3), 153–83.

DIX-CARNEIRO, R. (2014): "Trade Liberalization and Labor Market Dynamics," *Econometrica*, 82(3), 825–885.

DIX-CARNEIRO, R., AND B. K. KOVAK (2015a): "Trade Liberalization and the Skill Premium: A Local Labor Markets Approach," *American Economic Review*, 105(5), 551–57.

——— (2015b): "Trade Reform and Regional Dynamics: Evidence From 25 Years of Brazilian Matched Employer-Employee Data," *NBER Working Paper*, 20908.

DIXON, P. B., AND B. R. PARMENTER (1996): "Computable General Equilibrium Modelling for Policy Analysis and Forecasting," in *Handbook of computational economics*, ed. by H. M. Amman, D. A. Kendrick, and J. Rust, vol. 1 of *Handbooks in Economics*, chap. 1, pp. 3–85. Elsevier Science, North-Holland, Amsterdam, New York and Oxford.

DOLLAR, D., AND B. F. JONES (2013): "China: An Institutional View of an Unusual Macroeconomy," *NBER Working Paper*, 19662.

DOMS, M., T. DUNNE, AND K. R. TROSKE (1997): "Workers, Wages, and Technology," *Quarterly Journal of Economics*, 112(1), 253–90.

DORNBUSCH, R., S. FISCHER, AND P. A. SAMUELSON (1977): "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods," *American Economic Review*, 67(5), 823–39.

EATON, J., AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–79.

EBENSTEIN, A., A. E. HARRISON, AND M. MCMILLAN (2015): "Why are American Workers Getting Poorer? China, Trade and Offshoring," *NBER Working Paper*, 21027.

EBENSTEIN, A., A. E. HARRISON, M. MCMILLAN, AND S. PHILLIPS (2014): "Estimating the Impact of Trade and Offshoring on American Workers using the Current Population Surveys," *Review of Economics and Statistics*, 96(4), 581–95.

EECKHOUT, J., AND P. KIRCHER (2011): "Identifying Sorting–In Theory," *Review of Economic Studies*, 78(3), 872–906.

EGGER, H., AND U. KREICKEMEIER (2009): "Firm Heterogeneity and the Labor Market Effects of Trade Liberalization," *International Economic Review*, 50(1), 187–216.

ELSBY, M. W. L., B. HOBIJN, AND A. SAHIN (2013): "The Decline of the U.S. Labor Share," *Brookings Papers on Economic Activity*, 2013(2), 1–52.

EPIFANI, P., AND G. GANCIA (2008): "The Skill Bias of World Trade," *Economic Journal*, 118(530), 927–60.

FAJGELBAUM, P., G. M. GROSSMAN, AND E. HELPMAN (2011): "Income Distribution, Product Quality, and International Trade," *Journal of Political Economy*, 119(4), 721–65.

FAJGELBAUM, P. D., AND A. K. KHANDELWAL (2016): "Measuring the Unequal Gains from Trade," *Quarterly Journal of Economics*, 131(3), 1113–1180.

FEENSTRA, R. C. (1994): "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84(1), 157–77.

——— (2017): "Statistics to Measure Offshoring and its Impact," *NBER Working Paper*, 23067.

FEENSTRA, R. C., AND G. H. HANSON (1996): "Foreign Investment, Outsourcing, and Relative Wages," in *The political economy of trade policy: Papers in honor of Jagdish Bhagwati*, ed. by R. C. Feenstra, G. M. Grossman, and D. A. Irwin, chap. 6, pp. 89–127. MIT Press, Cambridge and London.

——— (1999): "The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979-1990," *Quarterly Journal of Economics*, 114(3), 907–40.

FEENSTRA, R. C., H. MA, AND Y. XU (2017): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States: Comment," University of California, Davis, unpublished manuscript.

FEENSTRA, R. C., AND D. E. WEINSTEIN (2010): "Globalization, Markups, and the U.S. Price Level," *NBER Working Paper*, 15749.

FELBERMAYR, G., J. PRAT, AND H.-J. SCHMERER (2011): "Globalization and Labor Market Outcomes: Wage Bargaining, Search Frictions, and Firm Heterogeneity," *Journal of Economic Theory*, 146(1), 39–73.

FONTAGNÉ, L., AND A. E. HARRISON (2017): "The Factory-Free Economy: Outsourcing, Servitization and the Future of Industry," in *The factory-free economy: Outsourcing, Servitization, and the Future of Industry*, ed. by L. Fontagné, and A. E. Harrison, pp. 1–23. Oxford Univeristy Press, Oxford.

FORD, M. R. (2015): *Rise of the robots: Technology and the threat of a jobless future*. Basic Books, New York, NY.

FRANKEL, J. A., AND D. ROMER (1999): "Does Trade Cause Growth?," *American Economic Review*, 89(3), 379–99.

FREEMAN, R. B., AND M. M. KLEINER (2005): "The Last American Shoe Manufacturers: Decreasing Productivity and Increasing Profits in the Shift from Piece Rates to Continuous Flow Production," *Industrial Relations*, 44(2), 307–30.

GLAESER, E. L., AND J. GYOURKO (2005): "Urban Decline and Durable Housing," *Journal of Political Economy*, 113(2), 345–75.

GOLDBERG, P. K., AND N. PAVCNIK (2007): "Distributional Effects of Globalization in Developing Countries," *Journal of Economic Literature*, 45(1), 39–82.

GOOS, M., AND A. MANNING (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics*, 89(1), 118–33.

GOOS, M., A. MANNING, AND A. SALOMONS (2009): "Job Polarization in Europe," *American Economic Review*, 99(2), 58–63.

GRANGER, C. W. J. (1969): "Investigating Causal Relations by Econometric Models and Cross-Spectral Methods," *Econometrica*, 37(3), 424–38.

GROSSMAN, G. M., AND E. HELPMAN (1991): "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 58(1), 43–61.

——— (1993): *Innovation and growth in the global economy*. MIT Press, Cambridge, Mass.

GROSSMAN, G. M., E. HELPMAN, AND P. KIRCHER (2013): "Matching and Sorting in a Global Economy," *NBER Working Paper*, 19513.

GROSSMAN, G. M., AND E. ROSSI-HANSBERG (2008): "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98(5), 1978–97.

——— (2010): "External Economies and International Trade Redux," *Quarterly Journal of Economics*, 125(2), 829–58.

GUADALUPE, M. (2007): "Product Market Competition, Returns to Skill, and Wage Inequality," *Journal of Labor Economics*, 25(3), 439–74.

HANDLEY, K., AND N. LIMAO (2013): "Policy Uncertainty, Trade and Welfare: Theory and Evidence for China and the U.S.," *NBER Working Paper*, 19376.

HARRIGAN, J. (2000): "International Trade and American Wages in General Equilibrium, 1967-1995," in *The impact of international trade on wages*, ed. by R. C. Feenstra, NBER Conference Report series, chap. 5, pp. 171–93. University of Chicago Press, Chicago and London.

HARRISON, A., AND M. MCMILLAN (2011): "Offshoring Jobs? Multinationals and U.S. Manufacturing Employment," *Review of Economics and Statistics*, 93(3), 857–75.

HARRISON, A. E., J. MCLAREN, AND M. MCMILLAN (2011): "Recent Perspectives on Trade and Inequality," *Annual Review of Economics*, 3(1), 261–89.

HE, Z., AND F. ZHANG (2017): "Trade and Real Wages of the Rich and Poor: Cross-Country Evidence," Columbia University, unpublished manuscript.

HEAD, K., AND J. RIES (2002): "Offshore Production and Skill Upgrading by Japanese Manufacturing Firms," *Journal of International Economics*, 58(1), 81–105.

HELPMAN, E. (2016): "Globalization and Wage Inequality," *NBER Working Paper No. 22944*, 22944.

HELPMAN, E., O. ITSKHOKI, M.-A. MUENDLER, AND S. J. REDDING (2017): "Trade and Inequality: From Theory to Estimation," *Review of Economic Studies*, 84(1), 357–405.

HELPMAN, E., O. ITSKHOKI, AND S. J. REDDING (2010): "Inequality and Unemployment in a Global Economy," *Econometrica*, 78(4), 1239–1283.

HICKS, M. J., AND S. DEVARAJ (2015): *The myth and the reality of manufacturing in America*. Ball State University, Muncie, Indiana.

HSIEH, C. T. (2008): "Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence: Comment," in *Globalization and Poverty*, ed. by P. Collier, and J. W. Gunning, vol. 1 of *Globalization of the World Economy series*, chap. 5, pp. 159–64. Elgar, Cheltenham, U.K. and Northampton, Mass.

HSIEH, C. T., AND P. J. KLENOW (2007): "Relative Prices and Relative Prosperity," *American Economic Review*, 97(3), 562–85.

HULTEN, C., AND L. NAKAMURA (2017): "Accounting for Growth in the Age of the Internet: The Importance of Output-Saving Technical Change," *NBER Working Paper*, 23315.

HUMMELS, D., R. JORGENSEN, J. MUNCH, AND C. XIANG (2014): "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, 104(6), 1597–1629.

IMBS, J. (2017): "Structural Change in the OECD: Some Facts," in *The factory-free economy: Outsourcing, Servitization, and the Future of Industry*, ed. by L. Fontagné, and A. Harrison, chap. 3, pp. 86–110. Oxford Univeristy Press, Oxford.

INTERNATIONAL MONETARY FUND (2017): "Understanding the Downward Trend in Labor Income Shares," in *World Economic Outlook: Gaining Momentum*, vol. 1 of *Handbooks in Economics*, chap. 3, pp. 1–52. International Monetary Fund, Washington, DC, authored by Dao, Mai Chi, Mitali Das, Zsoka Koczan and Weicheng Lian.

JIN, K. (2012): "Industrial Structure and Capital Flows," *American Economic Review*, 102(5), 2111–46.

JOHNSON, R. C. (2014): "Five Facts about Value-Added Exports and Implications for Macroeconomics and Trade Research," *Journal of Economic Perspectives*, 28(2), 119–42.

JOHNSON, R. C., AND G. NOGUERA (2012): "Accounting for Intermediates: Production Sharing and Trade in Value Added," *Journal of International Economics*, 86(2), 224–36.

KAMBOUROV, G. (2009): "Labour Market Regulations and the Sectoral Reallocation of Workers: The Case of Trade Reforms," *Review of Economic Studies*, 76(4), 1321–1358.

KAMBOUROV, G., AND I. MANOVSKII (2009): "Occupational Specificity of Human Capital," *International Economic Review*, 50(1), 63–115.

KATZ, L. F., AND D. H. AUTOR (1999): "Changes in the Wage Structure and Earnings Inequality," in *Handbook of labor economics*, ed. by O. Ashenfelter, and D. Card, vol. 3 of *Handbooks in Economics*, chap. 26, pp. 1463–1555. Elsevier Science, North-Holland, Amsterdam, New York and Oxford.

KATZ, L. F., AND K. M. MURPHY (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107(1), 35–78.

KLINE, P., AND E. MORETTI (2013): "Place Based Policies with Unemployment," *American Economic Review*, 103(3), 238–43.

KONINGS, J., AND A. P. MURPHY (2006): "Do Multinational Enterprises Relocate Employment to Low-Wage Regions? Evidence from European Multinationals," *Review of World Economics/Weltwirtschaftliches Archiv*, 142(2), 267–286.

KOOPMAN, R., Z. WANG, AND S. J. WEI (2014): "Tracing Value-Added and Double Counting in Gross Exports," *American Economic Review*, 104(2), 459–94.

KOVAK, B. K. (2013): "Regional Effects of Trade Reform: What Is the Correct Measure of Liberalization?," *American Economic Review*, 103(5), 1960–76.

KRISHNA, P., J. P. POOLE, AND M. Z. SENSES (2014): "Wage Effects of Trade Reform with Endogenous Worker Mobility," *Journal of International Economics*, 93(2), 239–52.

KROLIKOWSKI, P. (2017): "Job Ladders and Earnings of Displaced Workers," *American Economic Journal: Macroeconomics*, 9(2), 1–31.

LARCH, M., AND W. LECHTHALER (2011): "Multinational Firms and Labor Market Pooling," *Review of International Economics*, 19(4), 728–49.

LAZEAR, E. P., AND K. L. SHAW (2009): "Wage Structure, Raises, and Mobility: An Introduction to International Comparisons of the Structure of Wages Within and Across Firms," in *The Structure of Wages: An International Comparison*, ed. by E. P. Lazear, and K. L. Shaw, chap. 1, pp. 1–57. University of Chicago Press, Chicago and London.

LAZEAR, E. P., AND J. R. SPLETZER (2012): "Hiring, Churn, and the Business Cycle," *American Economic Review: Papers & Proceedings*, 102(3), 575–79.

LEAMER, E. E., AND M. STORPER (2001): "The Economic Geography of the Internet Age," *Journal of International Business Studies*, 32(4), 641–65.

LEGROS, P., AND A. F. NEWMAN (2002): "Monotone Matching in Perfect and Imperfect Worlds," *Review of Economic Studies*, 69(4), 925–42.

LEMIEUX, T. (2006): "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?," *American Economic Review*, 96(3), 461–498.

LEVINSOHN, J. (1999): "Employment Responses to International Liberalization in Chile," *Journal of International Economics*, 47(2), 321–44.

LEVY, F., AND R. J. MURNANE (2004): *The New Division of Labor*. Princeton University Press, Princeton.

LINDEN, G., J. DEDRICK, AND K. L. KRAEMER (2011): "Innovation and Job Creation in a Global Economy: The Case of Apple's iPod," *Journal of International Commerce and Economics*, 3(0), 223–40.

LISE, J., AND J.-M. ROBIN (2017): "The Macrodynamics of Sorting between Workers and Firms," *American Economic Review*, 107(4), 1104–1135.

LIU, R., AND D. TREFLER (2015): "A Sorted Tale of Globalization: White Collar Jobs and the Rise of Service Outsourcing," Rotman School of Management, unpublished manuscript.

LUCAS, R. E. J. (1993): "Making a Miracle," *Econometrica*, 61(2), 251–72.

MAGYARI, I. (2017): "Firm Reorganization, Chinese Imports, and U.S. Manufacturing Employment," Columbia University, unpublished manuscript.

MARCHAND, B. U. (2012): "Tariff Pass-Through and the Distributional Effects of Trade Liberalization," *Journal of Development Economics*, 99(2), 265–81.

MAZZOLARI, F., AND G. RAGUSA (2013): "Spillovers from High-Skill Consumption to Low-Skill Labor Markets," *Review of Economics and Statistics*, 95(1), 74–86.

MCKINSEY GLOBAL INSTITUTE (2012): *Trading myths: Addressing misconceptions about trade, jobs, and competitiveness*. McKinsey & Company, San Francisco, authored by Roxburgh, Charles, James Manyika, Richard Dobbs and Jan Mischke.

——— (2016): *Poorer than their parents? Flat or falling incomes in advanced economies*. McKinsey & Company, San Francisco, authored by Bughin, Jacques, James Manyika and Jonathan Woetzel.

MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.

MENEZES-FILHO, N. A., AND M.-A. MUENDLER (2011): "Labor Reallocation in Response to Trade Reform," *NBER Working Paper*, 17372.

MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2014): "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years," *Review of Economics and Statistics*, 96(1), 60–77.

MINCER, J. (1974): *Schooling, experience, and earnings*. Columbia University Press for the National Bureau of Economic Research, New York.

MONTE, F. (2011): "Skill Bias, Trade, and Wage Dispersion," *Journal of International Economics*, 83(2), 202–218.

MORETTI, E. (2010): "Local Multipliers," *American Economic Review*, 100(2), 373–77.

——— (2011): "Local Labor Markets," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4B of *Handbooks in Economics*, chap. 14, pp. 1237–1313. Elsevier, Amsterdam.

MUENDLER, M.-A., AND S. O. BECKER (2010): "Margins of Multinational Labor Substitution," *American Economic Review*, 100(5), 1999–2030.

MUTREJA, P., B. RAVIKUMAR, AND M. J. SPOSI (2014): "Capital goods trade and economic development," *Federal Reserve Bank of St. Louis Working Paper*, 2014-12.

NAM, C. H. (1995): "The Role of Trade and Exchange Rate Policy in Korea's Growth," in *Growth Theories in Light of the East Asian Experience*, ed. by T. Ito, and A. O. Krueger, no. 4 in NBER-East Asia Seminar on Economics series, chap. 6, pp. 153–77. University of Chicago Press, Chicago and London.

NEARY, J. P. (2016): "International Trade in General Oligopolistic Equilibrium," *Review of International Economics*, 24(4), 669–98.

NEARY, J. P., AND J. THARAKAN (2012): "International Trade with Endogenous Mode of Competition in General Equilibrium," *Journal of International Economics*, 86(1), 118–32.

NICITA, A. (2009): "The Price Effect of Tariff Liberalization: Measuring the Impact on Household Welfare," *Journal of Development Economics*, 89(1), 19–27.

OBSTFELD, M., AND K. ROGOFF (1995): "Exchange Rate Dynamics Redux," *Journal of Political Economy*, 103(3), 624–60.

OHNSORGE, F., AND D. TREFLER (2007): "Sorting It Out: International Trade with Heterogeneous Workers," *Journal of Political Economy*, 115(5), 868–92.

OI, W. Y., AND T. L. IDSON (1999): "Firm Size and Wages," in *Handbook of labor economics*, ed. by O. Ashenfelter, and D. Card, vol. 3B, chap. 33, pp. 2165–2214. Elsevier Science, North-Holland, Amsterdam, New York and Oxford.

OTTAVIANO, G. I. P., G. PERI, AND G. C. WRIGHT (2013): "Immigration, Offshoring, and American Jobs," *American Economic Review*, 103(5), 1925–59.

PIERCE, J. R., AND P. K. SCHOTT (2016): "The Surprisingly Swift Decline of U.S. Manufacturing Employment," *American Economic Review*, 106(7), 1632–62.

PIKETTY, T., E. SAEZ, AND G. ZUCMAN (2016): "Distributional National Accounts: Methods and Estimates for the United States," *NBER Working Paper*, 22945.

PISCHKE, J. S., AND J. VELLING (1997): "Employment Effects of Immigration to Germany: An Analysis Based on Local Labor Markets," *Review of Economics and Statistics*, 79(4), 594–604.

PORTO, G. G. (2006): "Using Survey Data to Assess the Distributional Effects of Trade Policy," *Journal of International Economics*, 70(1), 140–60.

POSTEL-VINAY, F., AND J. M. ROBIN (2002): "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 70(6), 2295–2350.

POSTEL-VINAY, F., AND H. TURON (2010): "On-the-Job Search, Productivity Shocks, and the Individual Earnings Process," *International Economic Review*, 51(3), 599–629.

RAITH, M. (2003): "Competition, Risk, and Managerial Incentives," *American Economic Review*, 93(4), 1425–36.

RESTUCCIA, D., AND C. URRUTIA (2001): "Relative Prices and Investment Rates," *Journal of Monetary Economics*, 47(1), 93–121.

RICARDO, D. (1821): *On the Principles of Political Economy and Taxation*. John Murray, London, 3rd edn.

RIVERA-BATIZ, L. A., AND P. M. ROMER (1991): "International Trade with Endogenous Technological Change," *European Economic Review*, 35(4), 971–1001.

RODRÍGUEZ, F., AND D. RODRIK (2000): "Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence," in *NBER Macroeconomics Annual 2000*, ed. by B. Bernanke, and K. Rogoff, vol. 15. MIT Press, Cambridge, Massachusetts.

RODRÍGUEZ-CLARE, A. (2010): "Offshoring in a Ricardian World," *American Economic Journal: Macroeconomics*, 2(2), 227–58.

RODRIK, D. (2013): "Unconditional Convergence in Manufacturing," *Quarterly Journal of Economics*, 128(1), 165–204.

——— (2016): "Premature Deindustrialization," *Journal of Economic Growth*, 21(1), 1–33.

SAMPSON, T. (2014): "Selection into Trade and Wage Inequality," *American Economic Journal: Microeconomics*, 6(3), 157–202.

SCHMIDT, K. M. (1997): "Managerial Incentives and Product Market Competition," *Review of Economic Studies*, 64(2), 191–213.

SLAUGHTER, M. J. (2000): "Production Transfer within Multinational Enterprises and American Wages," *Journal of International Economics*, 50(2), 449–472.

STOLPER, W. F., AND P. A. SAMUELSON (1941): "Protection and Real Wages," *Review of Economic Studies*, 9(1), 58–73.

SULLIVAN, P. (2010): "Empirical Evidence on Occupation and Industry Specific Human Capital," *Labour Economics*, 17(3), 567–80.

THOENIG, M., AND T. VERDIER (2003): "A Theory of Defensive Skill-Biased Innovation and Globalization," *American Economic Review*, 93(3), 709–28.

TOPALOVA, P. (2010): "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India," *American Economic Journal: Applied Economics*, 2(4), 1–41.

TRAIBERMAN, S. (2016): "Occupations and Import Competition: Evidence from Denmark," Yale University, unpublished manuscript.

TREFLER, D. (2004): "The Long and Short of the Canada-U.S. Free Trade Agreement," *American Economic Review*, 94(4), 870–95.

TREFLER, D., AND S. C. ZHU (2005): "Trade and Inequality in Developing Countries: A General Equilibrium Analysis," *Journal of International Economics*, 65(1), 21–48.

VERHOOGEN, E. A. (2008): "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *Quarterly Journal of Economics*, 123(2), 489–530.

WACZIARG, R., AND J. S. WALLACK (2004): "Trade Liberalization and Intersectoral Labor Movements," *Journal of International Economics*, 64(2), 411–39.

WRIGHT, G. C. (2014): "Revisiting the Employment Impact of Offshoring," *European Economic Review*, 66(0), 63–83.

YEAPLE, S. R. (2005): "A Simple Model of Firm Heterogeneity, International Trade, and Wages," *Journal of International Economics*, 65(1), 1–20.

YOUNG, A. (1991): "Learning by Doing and the Dynamic Effects of International Trade," *Quarterly Journal of Economics*, 106(2), 369–406.

ZEILE, W. J. (1997): "U.S. Intrafirm Trade in Goods," *Survey of Current Business*, 77(2), 23–38.

ZHANG, X., J. YANG, AND S. WANG (2011): "China Has Reached the Lewis Turning Point," *China Economic Review*, 22(4), 542–54.