Introduction

Increasing inequality within countries and globalization are two of the defining economic trends of the last three decades. Thus, it is only natural to ask whether these phenomena are causally linked. This book contains a set of recent articles that attempt to address the questions of if and to what extent globalization has contributed to the increase in inequality.

Robert Lucas once famously remarked, "Of the tendencies that are harmful to sound economics, the most seductive, and in my opinion the most poisonous, is to focus on questions of distribution.[1]" Why then study the relationship between *inequality* and trade? Of course, one may want to study inequality motivated by moral considerations and not economics. But even if sound economics is one's goal, in the context of trade and globalization, inequality warrants attention because of its potential detrimental effects on economic efficiency. Specifically, it is the unequal effects of trade that often stand in the way of implementing desirable economic policies. For instance, the overwhelming majority of economists support free markets. Yet, the public often opposes trade reform and demands protection from global forces. The reason for this behavior is not that protectionists question the aggregate gains from trade that accrue to the economy as a whole. Rather, their concern arises from the distributional conflict brought about by trade reforms which generate winners and losers. Students of International Trade are taught that it is possible, in principle, to compensate the losers so that in the end everyone is better (or at least not worse) off, but in practice such compensation is rarely observed. Realizing the full gains from globalization thus requires a thorough understanding of its distributional consequences, uncovering of the mechanisms through which resources get reallocated in a liberalizing economy, and identifying the associated winners and losers. This book represents an effort in that direction.

In contrast to the recent public debate that has focused on the top 1% of the income distribution, the focus of the book is on the other 99%. While the growing disparity between the top 1% and the rest of the population is startling, empirical evidence suggests that the increasing incomes of the top 1% are not largely driven by globalization. The sources of the very top earnings tend to be domestic, and international markets matter only to the extent that they may contribute to the "superstar" syndrome. Kaplan and Rauh (2013) describe this "superstar" syndrome, arguing that one of the explanations for the rise of top incomes is that recent technological advances and globalization have enabled highly skilled individuals, so-called "superstars", to apply their talents on a substantially larger scale to markets of increasing size. However, given that the increase in inequality at the top is observed across a broad set of occupations (public company executives, private company executives, financial executives, lawyers, professional athletes) with varying degrees of exposure to globalization, the authors

[1] Lucas (2004), last paragraph.

ultimately conclude that while globalization does contribute to a larger market, it is unlikely to be the primary driver of growing inequality at the top.[2]

Why is the bottom 99% interesting? Well, for one, it represents 99% of the population... Moreover, a recent study by David Autor (2014) suggests that earnings inequality in the U.S. has substantially increased among the 99% during the past thirty years due to a pronounced increase in the skill premium. Autor estimates that since 1980, the earnings gap between workers with a high school degree and those with a college education has become four times greater than the gap between the 99% and the top 1% over the same period. In other work, the same author as well as many other economists[3] have documented increasing "job polarization" in several industrialized countries: growth of high- and low-skill jobs at the expense of middle-skill jobs, and a widening in the gap between jobs that pay the most and jobs that pay the least. These aspects of the inequality phenomenon are more relevant to the majority of the population.

Until recently, the consensus among economists had been that trade had a relatively small effect on inequality. This consensus emerged in the late 1990s after nearly a decade of studies by both labor and trade economists on the effects of trade on the U.S. and other developed countries' labor markets. While the methods used and underlying models differed across fields and specific papers, a common theme was the focus on explaining the increase in the wage gap between skilled and unskilled workers, the so-called skill premium, observed in many countries during the past few decades. The predictions of the workhorse model of International Trade, the Heckscher-Ohlin model, and the associated Stolper-Samuelson theorem were in principle consistent with the increase in the skill premium in the industrialized world. However, the quantitative evidence suggested that other mechanisms, in particular skill-biased technological change (SBTC), were more significant, and that to the extent that trade had any effects at all, these played out through interactions with technology. As a result of this evidence, interest in the relationship between trade and inequality in developed countries subsided. In the context of developing countries, the evidence revealed a similar increase in the skill premium that seemed to contradict Heckscher-Ohlin, even in qualitative terms. Motivated by this puzzle, research shifted towards developing countries where the wave of trade liberalizations caused by the accession of many of these countries into the World Trade Organization (WTO) and their integration in global markets provided an exciting setting for studying the effects of trade on inequality.

---

[2] Kaplan and Rauh (2013), p. 53.

[3] Acemoglu and Autor (2011), Autor and Dorn (2013), Autor, Katz and Kearney (2006), Firpo, Fortin and Lemieux (2011), Goos and Manning (2007), Goos, Manning and Salomons (2014).

The last five years have seen a revival of interest in the effects of trade on inequality in the U.S. and Europe. Two developments contributed to this renewed interest. First, there is the entry of China into the global trading system. While imports from low-wage developing countries have been steadily increasing since the early 1970s, the rapid increase of imports from China in the last decade has exposed the U.S to unprecedented import competition. This has sparked widespread concern that Chinese imports have displaced U.S. manufacturing employment with particularly dire consequences for the regions that specialize in manufacturing. Second, there is the increase in “offshoring” of production and services, which may have contributed to the job polarization in the U.S. and Europe. Recent research on the U.S. and Europe has focused on these two developments.

The articles included in this book cover for the most part newer research on the topic. Earlier articles are included mainly to provide background and context for the most recent work. The emphasis on recent work is motivated both by the questions this work asks which are more relevant today and by the broader availability of better data that contribute to more convincing empirical evidence.

The literature on trade and inequality is vast, and there are many possible organizing principles for existing research. A distinguishing feature of the organization offered in this book is that the empirical work is discussed first, followed by the recent theoretical advances. This sequence reflects the chronology of research development. In its early stages, empirical research on this topic was closely guided by traditional general equilibrium models of trade which offered stark predictions regarding the effects of trade on inequality. The evidence uncovered by these empirical papers seemed to provide little support for the traditional theory. As a result, new theories emerged that focus on different channels through which trade can impact inequality. Representative examples of these new theoretical directions are presented in the last part of the book.

The included articles are organized in four parts. The first part covers empirical work on trade and inequality focusing on developed countries. The second part covers a similar set of questions for developing countries. Both parts focus primarily on “evidence” in the sense of documenting “what actually happened”. These papers abstract for the most part from transitional dynamics following trade liberalizations, and interpret the variables and effects observed in the analysis as steady-state-outcomes corresponding to two different equilibria: a trade-restricted and an unrestricted (or less restricted) equilibrium. The third part is also empirical, but more methodologically oriented. The papers in that part are explicitly concerned with the transition process from a trade-restricted to a less-trade-restricted equilibrium and emphasize the importance of labor market and other frictions. As with the new theories featured in the last part of the book, this literature emerged in response to dissatisfaction with earlier studies that abstract from the adjustment process; for this reason, this work is discussed after the empirical evidence is presented in the first two parts.

The rest of this introduction offers a detailed guide for reading the articles included in the book, and outlines the main contribution of each article and the reason for its inclusion. Before proceeding, two notes are in order.

First, the discussion will abstain from providing a conceptual treatment of "inequality" or "globalization". For a detailed discussion of the challenges facing the definition and measurement of these concepts, I refer the reader to the *Journal of Economic Literature* article by Goldberg and Pavcnik (2007). Depending on the setting and question, the papers in the present volume use a variety of measures (e.g., the skill premium, wage, income or consumption inequality, wage gaps across different occupations and sectors, the wage gap between formal and informal workers). The limitations of each measure will be pointed out in the discussion of each article.

Second, from a methodological point of view, the papers span a variety of approaches, from reduced-form to structural, and from partial equilibrium to general equilibrium, and rely on different strategies to identify the causal impact of trade on inequality. The papers are organized based on questions and underlying mechanisms rather than the particular identification strategy, but I will discuss briefly the various identification approaches along with their advantages and disadvantages when introducing the papers.

Part I: Developed Countries

The book opens with an excellent article by Adrian Wood who provides a comprehensive overview and discussion of the empirical research on trade and inequality up to the mid-1990s. This research focused on the increase in the skill premium, the widening gap between wages of skilled and wages of unskilled workers that was observed starting in the early 1970s in the U.S. and other industrialized countries. Two main explanations were considered: trade and skill-biased-technological change. Wood presents the contrasting views and approaches of Labor economists who adopted a partial-equilibrium approach focusing on quantities, and Trade economists who adopted a general-equilibrium approach focusing on prices operating at the margin.

The first approach employed factor content calculations, i.e., calculations of how many workers it takes to produce imports from low-wage countries and hence such imports displace. These calculations suggested that import competition from low-wage countries could not be the driving force behind the increase in the skill premium because the numbers were simply too small: imports represented only a small fraction of economic activity in the U.S. and other developed countries. Wood discusses the measurement challenges that arise in the context of such factor content studies and suggests several modifications to the naïve approach employed in early papers showing that such calculations, when performed correctly, point to a larger role of imports. Nevertheless, in the end, he too concludes that no matter how one conducts the factor content calculations, the impact of trade *by itself* could not have been very large. However, one cannot rule out that increasing imports from low-wage countries contributed to the adoption of skill-biased technologies and that trade had larger effects in interaction with these technologies.

Trade economists viewed factor content calculations with suspicion and characterized them as inconsistent with the general equilibrium approach common in Trade. The latter approach suggested examining prices (rather than quantities of imports), and in particular relative prices of labor-intensive products imported from low-wage countries. If trade were the primary driver of the decline in the relative wages of unskilled workers, then the prices of these products should have declined in the last three decades. As Wood points out, in principle, the study of prices should have settled the debate on the role of trade in rising inequality. Yet, prices proved inconclusive. The reason: "The rising heterogeneity of goods within statistically defined sectors…which has become worse over time.[4]" This heterogeneity made comparisons of prices over time difficult if not impossible. This observation naturally leads to the next article by Khandelwal, which deals explicitly with the implications of product heterogeneity for the relationship between trade and inequality.

Khandelwal hypothesizes that firms in developed countries, like the U.S., can escape low-wage import competition by differentiating their products along the quality dimension. The idea goes back to two influential papers by Peter Schott (2003 and 2004) that explored product heterogeneity in trade flows. Schott (2003) showed that even when countries seem to trade the same (in terms of statistical definition) "products" (for example electronics), in reality they specialize in different goods (e.g., radios and satellites) that are produced with different factor intensities. Schott (2004) takes this observation to a further level of disaggregation and shows that even when it appears that countries trade the exact same product (e.g., shoes or shirts), in reality they trade different *qualities* of this product. Product quality is of course not observed. But under the assumption of a vertical differential model, price is a sufficient statistic for quality. Accordingly, Schott uses price (unit value, to be more precise) data to demonstrate that there is substantial heterogeneity in the quality of products imported into the U.S. and that this quality differentiation is systematically related to the country of origin: imports from low-wage countries are cheaper (and hence of lower quality) than imports of the same product from high-wage countries (for example, shirts from Taiwan are cheaper than shirts from Italy). While Schott does not explore the implications of his findings for inequality, the connection is

[4] Wood (1995), p. 73

immediately apparent: if high- and low-wage countries specialize in different varieties of the same product and if further these varieties require different combinations of inputs that correspond to countries' different factor endowments, then countries can escape the Stolper-Samuelson type effects of trade, even when they trade freely with each other[5].

This is the hypothesis that Khandelwal puts to the test. But before doing so, he makes two modifications to Schott's approach. First, rather than using price as a proxy for quality, which is valid only under vertical differentiation, he considers a more general setup characterized by both vertical and horizontal differentiation. He uses insights from the modern Industrial Organization literature to compute a measure of quality, which relies on both price and quantity data. Intuitively, a product is of "higher quality", if it commands a higher market share than another product conditional on prices. Second, Khandelwal hypothesizes that sectors differ in their potential for quality differentiation – in some sectors the potential quality ladder is long, while in others quality specialization may be infeasible. Based on this observation, Khandelwal examines whether low-wage import competition on U.S. has had differential impacts on U.S. manufacturing depending on the length of industries' quality ladders. His results confirm his hypothesis: Competition from low-wage countries has had a much smaller impact on output and employment in industries where it is feasible to differentiate the products along the quality dimension. This observation may also explain why earlier work that ignored this quality differentiation failed to find a quantitatively significant effect of import competition on U.S. employment and wages.

As noted earlier, examination of the "quantities" associated with trade in the U.S. and other developed economies made hard to believe that trade was the primary culprit in the rise of inequality in the 1980s and 1990s. But this changed with the entry of China in the world trading system. The paper by Autor, Dorn and Hanson examines the implications of this entry for local labor markets in the U.S. Figure 1 and Table 1 in their paper describe the changing landscape: In 1991, 2.9% of all manufactured imports in the U.S. came from low-income countries, 0.6% of all U.S. expenditures were on Chinese Imports, and 12.6% of all American workers were employed in manufacturing. In 2007, 11.7% of all manufactured imports in the U.S. came from low-income countries, 4.6% of all U.S. expenditures were on Chinese Imports, and the share of American workers employed in manufacturing had declined to 8.4%. To link casually the rise of imports from China to labor outcomes in the U.S., the authors first assume that China's emergence as a major exporting power was spurred by an exogenous productivity shock in China coupled with China's accession to the WTO (rather than demand changes in the U.S.). Even with this assumption, they face the usual identification challenge faced by trade economists in work on trade and inequality, namely the need to generate variation in the

[5] The reason is that factor price equalization is broken in this case.

explanatory and outcome variables that they can exploit in a regression framework. To this end, they adopt an approach popularized by Topalova (2010), whose paper will be discussed later under work on developing countries, and focus on the differential exposure of local economies within the U.S. to Chinese import competition. What are the sources of this differential exposure? The underlying assumption is that regional economies are not perfectly integrated due to barriers to capital and labor mobility. Hence, each of these local economies can be treated as a separate "country". For reasons that are outside the scope of the analysis, different industries are concentrated in different regions. In this setting, one can employ an approach similar to the one first used by Bartik (1991) to develop a measure of regional exposure to import competition by interacting a measure of pre-sample industrial composition with the increase in the import competition from China experienced by each industry at the national level. The so-constructed measure of regional exposure exhibits thus two sources of variation: time variation that is entirely due to the rise of imports from China in each industry, and cross-sectional variation across regions that is due to the initial industrial composition. With this identification strategy in place, Autor et al examine several outcomes, such as manufacturing employment, wages and transfer payments.

The results are startling in several respects. First, there is the sheer magnitude of the effects. If one interprets the estimates as capturing the effect of Chinese imports on the *absolute* level of U.S. manufacturing employment (and not only the *relative* employment across regions), then the main specification in Autor et al suggests that approximately one quarter of the aggregate decline of U.S manufacturing is due to rising import competition. Areas that are more exposed to import competition from China experience higher unemployment, lower labor force participation, lower wages and significant increase in transfer payments (for unemployment, disability, retirement, healthcare). Second, these effects play out over nearly a decade (2000 to 2007[6]); hence, they cannot be dismissed as short-term transitional effects that would die out in the long run. But the perhaps most surprising aspect of the study is the unit of the analysis at which these large effects are documented. A "local" economy in the Autor et al study is a commuting zone. The finding of such large effects of import competition at the commuting zone level suggests substantial barriers to mobility even across geographically proximate locations and calls into question the assumption of perfect integration and costless relocation of production factors in response to trade and other shocks. I will return to this theme later in the discussion of the third part of the book.

The Autor et al paper is one of the first studies documenting a substantial impact of trade on U.S. workers, and in this case, the size of the impact is due to the size of the culprit: China. A different line of research explores an alternative mechanism, the increasing importance of

[6] The study actually covers a longer period, 1990-2007, but imports from China do not surge until 2000.

offshoring, also known as "outsourcing" or "global production sharing" (these terms are used interchangeably in the literature). This mechanism is the focus of the last three papers in Part I of the book. The importance of offshoring for U.S. wages was highlighted as early as in 1999 in the influential contribution of Feenstra and Hanson. Their work was the first to show that trade *can* – via outsourcing -- account for the growing relative demand for educated workers *within* industries. Feenstra and Hanson measure outsourcing by the estimated imports of intermediate inputs into each industry. In their main specification, this measure of outsourcing accounts for approximately 15% of the increase in the relative wages of skilled workers (defined here as nonproduction workers), while computer use contributes around 35%. When gauging the magnitude of these effects, one should keep in mind that the Feenstra and Hanson study examines outsourcing only up to the mid-1990s, so it does not include the last two decades during which the role of developing countries has become more prominent. While their work demonstrates the significance of outsourcing for U.S. wages, it also highlights the main measurement and identification challenges that empirical studies face in this literature. Their measure based on imports of intermediate inputs misses other aspects of outsourcing (e.g., outsourcing of computer programming or secretarial activities to lower wage countries) and relies on aggregate industry data. The study does not exploit an exogenous trade shock, and the industry data used in the analysis do not report workers' education so that the skill premium needs to be proxied by the difference in the wages between production and nonproduction workers.

A very recent study by Hummels et al improves on several (though not all) of these issues by relying on substantially better data: matched employer-employee data for the universe of private-sector firms in Denmark. The data include detailed information on worker and firm characteristics as well as trade flows *at the firm-level* that are broken down by product and origin and destination country. Such data are becoming increasingly available and permit an in-depth look into firms' and workers' adjustment to globalization. Hummels et al examine the impact of an increase in imports of intermediate inputs on workers of different skill and occupation. Their identification strategy relies on the observation that the input-output structure is highly specific to firms. Based on this observation, the authors consider various instruments for offshoring, the most successful of which turns out to be firm-specific transport costs. These are predicted primarily by oil prices, distance to trading partners and mode of transportation for different products, weighted by pre-sample firm-specific imports. The main identification assumption is that the pre-sample outsourcing structure (i.e., set of inputs that outsourced and identity of trading partners) is uncorrelated with within-sample shocks to wages. The authors document significant effects of offshoring on wages, with the wages of skilled workers increasing and the wages of unskilled workers decreasing as a result of offshoring. They also find that workers who were displaced from offshoring firms suffer greater earnings losses than other displaced workers, and that the earnings losses are greater and more persistent for low-skill than high-skill workers.

An interesting aspect of the Hummels et al paper is that they also document that the wage effects of offshoring differ by "task" characteristics: Conditional on skill, workers whose occupations involve routine tasks experience larger wage cuts while occupations that utilize knowledge from mathematics, social sciences and languages gain from offshoring. Moreover, occupations that are related to the natural sciences and engineering do not. These results speak to the literature on job polarization that emphasizes the role of specific job characteristics, and in particular routine tasks, in explaining recent changes in the wage distribution.

The last paper in Part I examines the potential role of offshoring in explaining job polarization. In contrast to Hummels et al who examine data only from Denmark, Goos et al utilize data from 16 Western European countries over the period 1993-2010. They document pervasive job polarization in all these countries. Then they set out to assess the relative importance of technology versus trade. The question is reminiscent of the earlier debate on trade versus technology in the 1990s that focused on the skill premium, but with job polarization having replaced skill premium. As with the earlier literature on the skill premium, identification of the relative contributions of technology and offshoring is an uphill battle because the same characteristics that make a job easy to replace by a machine make it also easy to offshore this job to a foreign country. Goos et al address this concern by employing two separate measures, one for "technology" (the "Routine Task Intensity (RTI) index used by Autor and Dorn (2013)), and one for "offshorability" (a measure proposed by Blinder and Kruger (2013) that is based on professional coders' assessment of the ease with which each occupation can be offshored). While the two measures are positively correlated, the correlation is only 0.46, so it is in principle possible to identify separately the effects of technology and offshorability. The authors find that technology, "routine-based-technological-change" as they call it, is substantially more important than offshoring.

At first, this message seems to contrast with the results of Hummels et al, but one should keep in mind that the two papers focus on different aspects of "offshoring". In Hummels et al, the primary focus is on the *imports of intermediate inputs* that may replace domestic unskilled labor. In Goos et al, the focus is on the direct offshoring of *jobs* to other countries. The Goos et al results notwithstanding, the papers by Autor et al and Hummels et al point to a new trend in the literature, where trade and globalization are assigned a great role in explaining changes in the income distribution. As emphasized earlier, in the Autor et al paper, this is because of the China phenomenon. It is possible that the significance attributed to offshoring in the Hummels et al paper is due to the fact that Denmark is a small, open economy, so trade and offshoring matter more for its economy and workers' incomes that they do for larger countries. We need many more studies on other countries before we can say with confidence whether or not offshoring has had a quantitatively significant impact on countries' income distributions.

Part II: Developing Countries

Since the mid-1990s, a substantial amount of research has been devoted to analyzing the impact of globalization on inequality in developing countries. Part of the reason for this shift in focus towards developing countries is that the recent trade liberalization episodes in these countries provide natural quasi-experiments that help identify the effect of trade on inequality. Of course, most developing countries enacted several other reforms at the same time, so a simple before-to-after comparison of outcomes would be subject to the usual identification challenges. However, there are three features of developing country liberalizations that help pin down the distributional effects of trade.

First, trade protection in developing countries still makes heavy use of tariffs, which are easy to measure. Countries use non-tariff barriers as well, but work to date has shown that these are positively correlated with tariffs, so that tariffs provide a good proxy of the overall level of protection. Second, most trade liberalizations were enacted either as a condition for countries' entry into the WTO or as a condition for aid from international organizations (e.g., IMF's aid package to India in 1991). These organizations often dictated the particular magnitude of the trade barrier reductions in each sector of the economy, so the usual political economy concerns of trade protection are mitigated. Third, a particular feature of these trade reforms was that the aim was to reduce not only the average level of protection, but also the dispersion across sectors. Therefore, sectors with initially larger levels of trade protection experienced larger trade barrier reductions relative to sectors with lower initial levels of protection. The *differential* changes of trade barriers across sectors can be exploited to overcome the limitations of before-to-after comparisons and to identify the distributional effects of trade liberalization episodes.

From an empirical point of view, these liberalizations offer a distinct advantage relative to work on developed countries which does not have such clean sources of identification (since all major developed economies liberalized their trade systems early as a result of GATT negotiations). Remaining trade barriers in developed countries primarily take the form of non-trade barriers, which are substantially more difficult to measure relative to tariffs. But in addition to this empirical advantage, there are also purely intellectual reasons that work on developing countries has been plentiful: Evidence has shown that inequality in developing countries has increased in the last three decades, mirroring the trend in the industrialized world. At first, this increase appeared to contradict the results of the traditional trade theory that predicted a decrease in inequality in countries abundant with less skilled labor. Hence, the empirical evidence led to a puzzle: How could the observed inequality trends be reconciled with the predictions of trade theory? Efforts to answer this question led to many interesting empirical studies as well as to new theoretical approaches that are discussed at the end of this volume.

There are nine papers included in this part of the book. The first six focus on import liberalizations in developing countries and employ the identification strategy described above. These papers all exploit quasi-experimental variation induced by the trade reforms in various countries in order to identify the effects of trade on inequality. The last three papers focus on exports rather than imports and explore channels through which a rise in exports can lead to changes in a country's income distribution. The sources of identification in these papers involve plausibly exogenous shocks (e.g., exchange rates, bilateral trade agreements) that affected exports in each case.

The first paper in this part by Attanasio et al focuses on Colombia that unilaterally liberalized its trade in the early 1990s. At the same time, the country experienced an increase in wage inequality. The paper lays out three potential channels through which trade may have impacted inequality: changes in the economy-wide skill premium, changes in industry wages that adversely affected workers with initially lower wages, and a shift towards the informal sector. The evidence suggests that the rise of the skill premium played only a modest role in increasing inequality in Colombia, as a significant part of the rise in aggregate inequality is accounted for by a rise in inequality *within* educational groups. Though the paper ultimately does not find strong support for the hypothesis that the change in inequality was brought about by trade, it provides a useful framework for thinking about the mechanisms through which such an impact could have taken place. These mechanisms are investigated in depth in the subsequent papers of this part of the book.

The paper by Porto provides an excellent illustration of the mechanism highlighted in the traditional factor proportions theory. Trade affects inequality through its direct impact on traded goods prices and its indirect impact on non-traded goods prices. Changes in prices in turn imply changes in the returns to factors used in the production of goods. Hence, trade liberalization affects inequality through two channels: it changes the prices facing consumers, and it changes the relative wages of skilled and unskilled workers. Porto explores these mechanisms in the context of the trade liberalization that Argentina experienced as a consequence of Mercosur. His analysis offers several advantages: (a) there is a close connection between the underlying theory and the empirics, so the results are easy to interpret; (b) the study investigates the effects of the liberalization on both prices and wages, which is an improvement to other studies that focus only on wages; (c) by using data from a Consumer Expenditure Survey, he is able to analyze the impact of the liberalization on *consumption* rather than wage inequality (the former is arguably a more informative measure of inequality). The main challenge for the paper is one that is encountered by every paper that attempts to analyze the effects of a liberalization employing a framework that assumes away all frictions (as is the case with the traditional model of comparative advantage that Porto employs). Because in Porto's framework there are no frictions, either in labor or product markets, and because both factor and product markets are perfectly competitive, there is only one economy-wide return to each factor of production, and only one price for each product in the country. As a result, any empirical investigation of the *actual* effects of a liberalization episode would have to rely exclusively on time-variation of prices and factor returns (e.g., variation of the economy-wide "skilled wage" or "unskilled wage" over several years). But such variation is hardly sufficient for identifying the effects of trade; apart from the fact that there simply very few annual observations one can use in the estimation, inference would be contaminated by the fact that many other changes took place over the same period. So rather than *estimate* the effects of Mercosur, Porto does something much more sensible in this context: he uses simulation to assess the impact of the actual tariff changes on the income distribution. But these simulations do not tell us what actually happened in Argentina. They simply tell us what would have happened if the model that Porto uses for the simulations were valid. Hence the conclusions are as credible as the underlying model. In this case, the conclusions are counterintuitive at times. Indicatively, Porto asserts that Mercosur raised the price of food products in Argentina (due to the increase on the external tariff on food products which dominated the decrease in the internal tariff). This price increase in food is shown in the simulations to ultimately benefit the poor; the reason is that the food sector employs a relatively higher share of unskilled workers, who according to the underlying model and key parameter estimates would see their wages increase when the price of food is higher. Because the wage effects dominate the price effects, these unskilled workers end up better off. While the result is consistent with the underlying framework, historical experience makes hard to believe that an increase in the price of food benefits the poor. Such counterintuitive predictions illustrate the limitations of simulations at times. More importantly, they show the limitations of the frictionless general equilibrium theory of comparative advantage from an empirical point of view: the theory predicts *no variation* in factor returns across sectors, so that if one takes the theory seriously, as Porto does, then one is left with the challenge of identifying an effect without sufficient variation.

These limitations are overcome by the next few papers that are all based on trade models that assume some type of friction. Frictions are useful not only because they are a-priori more plausible (there is substantial evidence pointing towards frictions and other imperfections in both product and factor markets), but also because they help generate variation in the outcome variables. Two types of friction are commonly explored in the literature: frictions to factor mobility across sectors/industries that give rise to variation in factor returns (wages most importantly) across industries, and frictions to factor mobility across regions within a country that give rise to regional variation in factor returns. In either case, the studies control for time effects, which capture among other things the potential economy-wide effects of a policy change. Accordingly, the effects identified by exploiting industry or regional variation are only *relative* effects. Interpreting them as *absolute* effects requires additional assumptions which may be more plausible in some settings than others.

The paper by Goldberg and Pavcnik is based on the specific factors model of trade, which assumes that factors of production are specific to the industries they are employed in, and cannot reallocate in response to a trade shock. This is a reasonable assumption in the short- and medium-run. In this case, factor returns vary across industries. Theory predicts that a differential reduction of tariffs across industries will benefit workers in those industries in which tariffs go up, and will hurt workers in the industries in which tariffs go down. Goldberg and Pavcnik take this prediction to the data. They focus on industry wage premia, i.e., the proportion of wages in each industry that cannot be accounted for by observable worker characteristics (for example, age, gender, education). They exploit the Colombian trade liberalization of the 1990s that offers all the aforementioned advantages: large and economically exogenous tariff reductions that differed across industries. The results are consistent with the theoretical predictions: wages do indeed go down by more in those industries in which tariffs were reduced by more. Moreover, the sectors that experienced the largest tariff reductions (e.g., textiles and apparel, footwear) had the lowest wages to start with, a pattern that contributes to a rise in inequality. While the results are statistically significant, the magnitudes of the trade effects are relatively small. Therefore, it is hard to argue based on the evidence presented in that paper that trade significantly contributed to a worsening of wage inequality via its effect on industry wages.

It is important to note at this point that relying on some form of friction for identification does not imply that the empirical approach is invalid if the friction is not present. In the absence of the presumed friction, the study will simply fail to find any effects. In the context of the industry wage premia discussed above for example, if production factors can move across industries so that their returns are equalized, then a regression of industry wage premia on tariffs will produce a zero coefficient on tariffs. The fact that the empirical results point to a tariff effect that is significantly differently from zero offers indirect evidence that a friction exists. Hence, the approach is internally consistent.

Goldberg and Pavcnik exploit frictions in the movement of labor and capital across sectors within an economy. A different type of friction is exploited in Topalova's work on the effects of the Indian trade liberalization on poverty and inequality. Topalova focuses on frictions to factor mobility across regions within a country. Under the assumption that capital and labor are not perfectly mobile, each regional economy can be considered as a separate country. Because different regional economies specialize in the production and trade of different goods, regions will differ in their exposure to trade. Topalova uses this differential exposure of regions to trade to devise an identification strategy that helps her pin down the effect of India's trade liberalization on poverty. The empirical strategy is the same as in Autor et al that was discussed earlier, but Topalova uses changes in tariffs, plausibly exogenous in the case of India, compared to Autor et al's use of potentially endogenous trade volumes[7]. Topalova finds large effects of the tariff reductions on rural poverty, with areas that were exposed more to liberalization experiencing an increase in poverty. Topalova is careful to point out that the documented effect is an increase *relative* to the economy-wide average. Because India experienced a large decrease in poverty during that time, Topalova's results, properly interpreted, suggest a slow-down of the poverty decline due to trade reforms, rather than an absolute increase in poverty.

The advantage of Topalova's reduced form approach is that the results could be in principle compatible with several alternative theoretical explanations. In contrast to Porto, Topalova does not take a stand on the right theoretical model underlying her analysis. At the same time, the results on the (relative) poverty increase are thought-provoking, and one would like to understand the mechanisms leading to these effects. Kovak's work offers insight into the mechanisms by providing a theoretical model that rationalizes Topalova's approach. The model is a Ricardo-Viner model of trade featuring one factor of production (capital) that is immobile across sectors and one factor of production (labor) that is mobile within a region, coupled with Topalova's premise that neither capital nor labor can move across regions.

Apart from providing a rationalization for Topalova's approach, Kovak's theoretical formulation offers another advantage. One of the challenges facing work that exploits regional variation in exposure to trade policy is that researchers need to construct an appropriate measure of *regional* trade liberalization. The data report tariffs and other forms of trade protection at the product or sectoral level, but these measures need to be aggregated to the regional level to provide a measure of a region's exposure to trade protection. How should one aggregate? The approach Topalova employs is sensible: it weighs the tariff in each sector by the fraction of a region's workers employed in this sector in the period prior to trade liberalization, and assigns a zero tariff to the non-traded sector. This measure is intuitively appealing as regions that have large non-traded sectors will have only a small fraction of the population employed in each tradable sector. Hence, the weights assigned to each tradable sector will be small, and the aggregate measure of exposure to trade will also be small. Topalova realizes that this measure

[7] However, Autor et al use an instrumental variables strategy to address the potential endogeneity of trade volumes.

may induce spurious correlation in the analysis, and instruments it using a different measure of aggregate regional protection: an alternative weighted average of tariffs across sectors, in which each tariff is weighted by the proportion of workers working in this sector as a fraction of all workers employed in the traded sector. Compared to the original measure, this second measure will tend to produce larger measures of aggregate exposure to liberalization in regions with large non-traded sectors, because the large employment in the non-traded sector will not appear in the denominator of the weights.

The main insight from Kovak's work is that this second measure is in fact a more appropriate measure of regional exposure to trade protection, and that empirical work should employ this measure as an explanatory variable (and not use it just as an instrument to the alternatively constructed measure Topalova proposes). The reason is that in general equilibrium, the prices of non-traded products adjust (even though these prices are not directly affected by trade liberalization). Kovak shows that the measure Topalova used implicitly holds the prices of non-traded products constant. But if one lets the prices of non-traded products adjust, then the second measure provides a better proxy for the "correct", that is theoretically consistent, measure of liberalization. Why does this matter? The particular weights used in the measure of trade liberalization do not matter for the finding of a statistically significant effect of trade liberalization on various outcomes; the results are robust in the statistical sense. However, the particular measure used does matter for the magnitudes of the effects. Using the latter measure suggested by Kovak generates smaller estimates of the effect of regional tariffs on wages (as shown in Kovak's tables). The overall effect of the trade reform is less clear as this effect depends on both the regression estimates and the particular measure of regional trade liberalization used, and the two measures suggested by Topalova and Kovak respectively have different means and variances.

The above three papers raise another question that is relevant to interpreting the results from existing work: Are empirical studies with identification strategies premised on the existence of a friction partial or general equilibrium in nature? The discussion of Topalova's and Kovak's papers shows that "frictions" is not synonymous with "partial equilibrium". The approach pioneered by Topalova is perfectly consistent with a general equilibrium model of trade assuming that the proper measure of trade liberalization is used. Similarly, the specific factors model that motivates Goldberg and Pavcnik's approach is a general equilibrium model, albeit a general equilibrium model that is more relevant in the short and medium run. Along the same lines, "reduced form" is not the same as "partial equilibrium". These terms are sometimes used interchangeably in the literature, and mistakenly so. Porto's approach is structural and general equilibrium in nature. Topalova's approach is reduced form, but (with the appropriate choice of regional tariff) consistent with a general equilibrium model of trade. The difference is that the equilibrium considered in Porto is one that is more relevant as a long run concept, while the model underlying Topalova's analysis is more applicable to the short or medium run.

A related question is whether it is valid to interpret the results of studies relying on industry or regional variation as revealing the *absolute* effects of a policy change under certain conditions. Again, this is a question that cannot be answered without a model that lays out the mechanisms through which a reduction of trade barriers affects industry, regional and economy-wide outcomes. The relevant question here is whether the trade reform generates economy-wide changes that are absorbed by the time effects. To use a specific example, if one interprets Topalova's results through the prism of Kovak's model, which features perfectly integrated product markets with prices of tradable goods equalized across regions, then the time effects in Topalova's regressions capture (among other things) the effect of India's trade liberalization on economy-wide prices of tradable products. The decline of such prices contributes to a decline in poverty, hence the aggregate decline in poverty should be in part attributed to the trade liberalization. In this case, Topalova's results capture only the effects that operate through the labor market channel and changing prices in the non-tradable sector, and one can only interpret them as relative effects (as she does). But if one assumes segmented product markets so that prices of tradable products are not equalized across regions, and if trade is not balanced, then one might interpret the estimated effects as capturing the absolute effects of the liberalization. Similarly, in Autor et al, the estimated effects of the Chinese import competition on local labor markets are strictly speaking relative effects. However, if one additionally imposes the assumption that the surge of imports from China did not affect U.S. exports to China or other destinations, then one can interpret their results as revealing the absolute impact of Chinese imports on U.S. manufacturing employment. The assumption seems reasonable in the medium run, especially given the size of the U.S. trade deficit and China's trade surplus. In general, making the leap from a relative to an absolute effect interpretation requires: (a) specifying the exact mechanisms through which trade impacts the outcomes of interest which in turns requires a model and (b) evaluating whether the assumptions underlying the model specified under (a) are plausible in the particular setting.

The final study focusing on the effects of unilateral trade liberalizations and import competition on inequality is by Goldberg and Pavcnik on informality. The study investigates a channel that has been absent from formal models of trade, yet important in the public debate focusing on trade liberalizations in developing countries: the possibility that globalization leads to an expansion of the informal sector. Because this sector is associated with lower wages, lower job security and overall lower job satisfaction, an increase in informality would worsen inequality within society. The paper exploits trade liberalizations in Brazil and Colombia to investigate the impact of these reforms on informality. The identification strategy is the same as in earlier work by these authors on these two countries: the differential changes of tariffs across industries are used to investigate whether industries exposed to larger tariff reductions experienced an increase in informality. The authors do not find any strong effects. Only in Colombia does the evidence suggest that tariff reductions led to increases in informality, and only in a period that precedes a major labor market reform in this country. A limitation of this study is that the identification strategy permits one to infer the trade reform's effect only *within* sectors, but misses changes that may have been accompanied by movement of workers to different industries or to the non-traded sector. The data suggest that mobility across industries was limited, but transitions to the non-traded sector have been significant. It is possible that these transitions were associated with transitions to informality, but the data and identification strategy used by Goldberg and Pavcnik do not permit an analysis of this margin.

The last three papers in the second part of the book focus on the role of exports in the rise of inequality. The first two papers by Verhoogen and Brambilla et al argue that it is not exporting per se that matters for the income distribution, but the destination to which countries export. Verhoogen develops this argument in the context of Mexico. He argues that a major peso devaluation in 1994 increased Mexico's exports to high-income countries, particularly the U.S. Consumers in high-income countries demand higher quality products than consumers in mid- or lower-income countries. Hence, in order for firms from developing countries to become successful exporters to high-income destinations, they need to upgrade the quality of their products. But higher quality of output requires higher quality inputs. Hence, firms also need to upgrade their labor force. This leads to higher wages in the firms that export, with skilled workers experiencing a larger increase than low-skill workers. This "skill-upgrading" of the labor employed in exporting firms contributes to rising inequality. Brambilla et al develop this thesis further by arguing that firms exporting to high-income destinations not only sell higher quality products to these destinations, but also provide additional services (marketing, advertising, distribution) that are highly skill-intensive and associated with higher wages. Hence, as with Verhoogen, exporting to high-income destinations implies a rise in inequality in developing countries.

Both papers develop theoretical models that formalize the above arguments, and provide strong empirical support for the mechanisms suggested by the respective models, Verhoogen for Mexico and Brambilla et al for Argentina. In both cases, the arguments emphasize the role of quality differentiation in the rising demand for skill in developing countries. It is instructive to compare these arguments to Khandelwal's analysis of the impact of quality differentiation on the U.S. wage distribution. In all these papers, the ability of firms to differentiate in the quality dimension of their product lines plays a key role. In Khandelwal's analysis of the U.S., quality differentiation on the part of U.S. firms eases the pressure low-wage exporting countries exert on the U.S. labor market. In Verhoogen and Brambilla et al, it is firms in developing countries that engage in quality differentiation in order to penetrate high-income markets, and this differentiation worsens the relative position of low-skill workers in developing countries. Hence, quality differentiation tends to decrease inequality in developed countries, but increase inequality in developing countries.

Another noteworthy feature of the above two papers is that though they both document statistically significant effects of exporting to high-income countries on wages, it is not clear that these effects are large enough to explain the rise in inequality observed in these countries in the past few decades. A distinguishing feature of these two studies is that both sets of authors rely on currency devaluations in order to identify the causal effect of exports. It has been argued that because exchange rate fluctuations are temporary shocks, their effects on economic outcomes may be more modest than the effects of other shocks, such as trade policy changes or trade agreements, that are considered more permanent. However, exchange rates exhibit high persistence (the peso devaluation Verhoogen exploits, for example, lasted for several years), so it is unlikely that economic agents do not respond to exchange rate shocks because they view them as short-lived. An alternative hypothesis is that exporting is still a small part of economic activity in countries such as Mexico or Argentina, so mechanisms that operate via changes in exports are unlikely to have a first-order impact on the income distribution.

Against this background, it is interesting to contrast the results of the above studies to those of McCaig, the last paper in this part of the book. McCaig differs from the previous two studies in that he focuses on a poor country, Vietnam. Further, he uses a bilateral trade agreement between Vietnam and the U.S. as a source of identification and focuses on exporting per se rather than the destination of exports. His approach is similar to Topalova's; he focuses on regional differences in the trade exposure to identify the effects of the increased exporting induced by the trade agreement on regional wages and poverty.

What is striking about McCaig's paper is the magnitude of the reported effects. McCaig finds very large reductions in regional poverty rates relative to the economy-wide average as a result of Vietnam's trade liberalization. What explains such large effects? A plausible explanation is that the economic significance and impact of exporting is substantially larger in a poor country like Vietnam. It would be interesting to collect more evidence on other relatively small economies in order to investigate this possibility. But another explanation is that these results are due to the methodology that is based on exploiting regional variation in trade exposure for identification.

This leads me to the following observation with which I will conclude the discussion of the second part of the book. While most empirical studies on the effect of trade on various aspects of inequality report statistically significant effects, few find effects that are large enough for one to conclude that trade had a major impact on inequality. Interestingly, the three papers included in this book that point to significant economic effects (Autor et al, Topalova, and McCaig) all rely on regional variation in trade exposure to identify the causal effect of trade. Of course, the effects estimated in these papers are only *relative* effects, but even so, their magnitudes are large relative to the aggregate trends. This raises the question of whether there is not something inherent in the "regional" approach that leads to the inference of large effects. Kovak's work suggests a possible explanation: All three papers exploit a measure of regional variation that includes workers employed in the non-traded sector; but if prices of non-traded goods adjust in equilibrium, this measure will differ from the theoretically appropriate measure of trade liberalization, and its use on the right hand side of regressions will impact the estimates of the effect of trade liberalization. This in turn raises the question of what the "appropriate" measure of liberalization is. This question cannot be answered outside a model. If one assumes Kovak's model (featuring perfectly competitive markets, flexible prices, sector-specific capital, and frictions to inter-regional movement of labor and capital), then the correct measure will be based on weights that include only workers employed in the traded sector. Alternatively, one can envision more realistic models that feature additional frictions in product and/or factor markets, for example price rigidities leading to incomplete price adjustment, that would lead to a different measure of trade liberalization, perhaps closer to the one employed by Topalova and the other papers discussed here. This would be an interesting and important question for future research, not only on intellectual grounds, but also because the particular measure of regional liberalization has been shown to matter for the estimated magnitudes of the effects.

Part III: Adjustment Costs and Labor Market Dynamics

One of the most important and robust insights of trade theory is that the distributional effects of trade depend on the time frame of the analysis. In the short run, factors of production are plausibly considered fixed and immobile across industries. Industry affiliation is all that matters for identifying winners and losers in this case: factors employed in industries that are adversely affected by trade lose and factors employed in the remaining industries win. In contrast, in the long run, factors are perfectly mobile across industries and this mobility implies that it is the identity of the production factor that matters for the distributional effects from trade. Knowing whether the world is best described by a "long run" or "short run" model is thus essential for identifying the winners and losers associated with trade reforms. At the same time, the literature emphasizes the importance of mobility and resource reallocation for the realization of the aggregate gains from trade: in the "long run", factors of production are supposed to move from the adversely affected industries/regions/firms towards the ones that have benefited from trade, and this reallocation is a crucial part of the mechanism through which the aggregate gains from trade are realized.

Empirical evidence suggests that in reality there are substantial barriers to mobility and reallocation. Indeed, the results in the papers by Autor et al, Goldberg and Pavcnik, Topalova and Kovak provide strong support for the view that the world is best described by the "short run" model and that this model provides an appropriate framework for analyzing the distributional effects of trade several years after a trade shock. Nevertheless, even with barriers to mobility, one would expect the economy to eventually adjust, albeit slowly, to the "long run" equilibrium. The interesting question is then: "how short is the short run?" In other words, "how long does it take for factors of production to reallocate in response to trade shocks and what are the impediments to such reallocation?" Are the effects documented in the previous papers only short-lived (in which case we could potentially downplay them as part of a normal adjustment process), or do they persist for many years with potentially severe consequences for the income distribution? Until recently, the literature had surprisingly little to say about these questions, as it was focused on steady state outcomes, which were interpreted either through the prism of the long run model (e.g., Porto), or through the prism of the short or medium-run models that underlie studies using industry or regional variation.

The third part of the book contains a set of papers that attempt to make progress in addressing exactly this type of questions by explicitly focusing on the transition process associated with a trade reform. These papers are part of a growing literature that contains several more, still unpublished, contributions. This literature uses structural methods, but is motivated by reduced form evidence on the importance of mobility costs (see Goldberg and Pavcnik (2007) for an overview of this evidence). Methodologically, the papers develop structural models of labor and trade, which are either calibrated (e.g., Kambourov) or estimated (Artuc et al, Dix Carneiro), and then used in simulations in order to assess the distributional effects of trade and the impact of various hypothetical policies. As such, they do not really tell us "what happened" in an economy as a result of a trade liberalization, but rather what the assumed model implies about the effects of trade (very much in the same spirit as Porto). Hence, evaluation of their results and policy implications depend on the plausibility of the underlying model and key structural parameters.

Artuc et al develop a simple dynamic structural model of the labor market that assumes perfect competition in both product and factor markets. Workers are assumed to be infinitely-lived and homogeneous (both in terms of observed and unobserved to the econometrician characteristics). They face intersectoral mobility costs and iid idiosyncratic preference shocks for different sectors. These assumptions are restrictive, but result in a very tractable framework: they allow the authors to obtain a closed-form solution that involves an equation relating gross worker flows across sectors to intersectoral wage differentials. When they estimate this equation for the U.S., they find high intersectoral mobility costs (approximately six times the average annual wage) and a large variance of idiosyncratic preference shocks. They use these estimates to simulate the impact of a potential trade reform and find that such a reform benefits everyone, including the workers employed in the initially adversely affected sectors. This implication, which seems counterintuitive at first, is driven by the assumptions of the model: despite the large estimated mobility costs, workers are likely to move to a different sector at some point in their (infinite) lives, because the estimated variance of the iid preference shocks is so large. This reallocation happens independently of the trade reform. At the same time, the trade reform reduces prices and this benefits everyone.

Given the restrictive assumptions underlying the analysis, these implications are not to be taken literally. At best, they could be interpreted as capturing the very long run effects of trade – the infinitely lived agents of the model could be interpreted as agents who care as much about themselves as about their children and further descendants. The first generation may be adversely affected by a trade liberalization, but in the long run, their descendants will benefit from lower prices and will likely get jobs in different industries. Even with this interpretation in mind, the framework has several limitations, the most important of which is the assumption of homogeneous workers. The basic setup can be easily extended to accommodate heterogeneity in observable characteristics, but incorporating unobservable heterogeneity is substantially more challenging, and requires a different framework.

Such a framework is proposed by Dix-Carneiro, who considers a richer model featuring overlapping generations, observed and unobserved worker heterogeneity, and sector-specific experience that may not be perfectly transferrable across sectors. The cost of the richer structure is more complexity. The framework developed by Artuc et al is tractable and easy to incorporate in standard trade models. This is harder to achieve with Dix-Carneiro's model. However, the richer structure proves to be important, leading to different implications. Dix-Carneiro estimates the costs of mobility to be substantially lower than in Artuc et al. Intuitively, this is due to the role that worker heterogeneity, especially unobservable worker comparative advantage across sectors, plays in the inference of mobility costs. When workers are homogeneous (as in Artuc et al) and factor markets are perfectly competitive, the relevant counterfactual wage for a worker contemplating moving to a different sector is the average wage in this sector. But when workers are heterogeneous along dimensions that are not observed by the econometrician, the relevant counterfactual wage is known only to the worker. Self-selection based on unobservable heterogeneity implies that the average sectoral wage will then provide an overestimate of the actual counterfactual wage. Accordingly, the wage differentials based on average sectoral wages used to estimate mobility costs in the homogeneous worker framework tend to overstate the actual wage differentials that are relevant to workers' moving decisions, leading to an overstatement of mobility costs. The implications for the adjustment to a trade reform are also different. Dix-Carneiro considers a series of counterfactual experiments in order to assess the effects of trade liberalization on reallocation and worker welfare. Consistent with the finding of lower mobility costs, he finds considerable reallocation in response to a trade reform, though the speed of this reallocation crucially depends on the degree of capital mobility. The heterogeneous worker framework is ideal for assessing the effects of a trade shock on the wage distribution, as trade is shown to differentially affect workers with different characteristics. In contrast to Artuc et al, who find that trade liberalization benefits everyone, Dix-Carneiro's simulations suggest that workers initially employed in the adversely affected sector lose in terms of welfare. But perhaps the most interesting aspect of the simulations is that the model, despite its richness, cannot generate the lack of reallocation and low speed of adjustment implied by the results of Autor et al, or Goldberg and Pavcnik, Topalova, etc. This suggests the presence of additional frictions and/or departures from the assumption of perfect completion in product markets.

In both Artuc et al and Dix-Carneiro's papers, "mobility costs" are a black box in the sense that they capture everything that may prevent a worker from moving to a different sector in response to a wage differential. Of course, the interpretation of the "black box" will depend on the model one used to infer mobility costs: In Artuc et al, the black box will contain anything other than wage differentials that may prevent workers from moving (including sector-specific experience that is imperfectly transferrable and unobserved comparative advantage). In Dix-Carneiro's richer model, mobility costs are net of the impact of sector-specific experience, compensating differentials, and unobserved comparative advantage. Still, from a policy point of view, we would like to know much more about the nature and sources of these reallocation costs in order to implement policies that will help reduce them.

The paper by Kambourov is a step in this direction. In contrast to the other two papers that estimate their respective empirical models, Kambourov relies on calibration. He develops and calibrates a two-sector small open economy model that features sector-specific human capital and import barriers. Frictions to mobility arise from labor market regulation, and specifically from the presence of firing costs. Simulating the model for Mexico suggests that labor market regulation prevents an economy from fully realizing potential gains from trade. The advantage of his approach is that it opens up the black box called "reallocation costs" and suggests a specific reason for constrained labor mobility: firing costs and labor market regulation. As such, it is more closely linked to policy. The natural policy implication of Kambourov's results is that labor deregulation is an important pre-requisite for an economy fully realizing the gains from trade liberalization.

The main focus of the three papers discussed above is on labor reallocation, which has important implications for inequality, and in particular for the identity of the winners and losers associated with a trade reform. As noted earlier, because they all employ simulations to examine the impact of trade on the labor market, they do not inform us on the actual effects of a trade liberalization episode. Nevertheless, they are helpful in highlighting the mechanisms through which adjustment (or lack thereof) to globalization takes place and in identifying issues that need further scrutiny. So far three such issues have emerged. First, it would be desirable to open the black box called "mobility costs" in order to understand which observable factors and policies inhibit reallocation. Second, the work to date has focused exclusively on *labor* reallocation, but Dix-Carneiro's simulations suggest that *capital* mobility is equally important for assessing the effects of trade on labor markets. Finally, even rich models such as Dix-Carneiro's have had a hard time generating the lack of mobility and reallocation observed in the data. This suggests that it is important to extend existing models to incorporate additional frictions in product markets (e.g., incomplete price adjustment in response to trade shocks). These are new and important areas of research and exciting work on these topics is already in progress.

Part IV: Theoretical Advances

Several new theoretical approaches have been developed in order to rationalize the recent empirical evidence on the relationship between trade and inequality and bridge the apparent discrepancy between the empirical findings and the predictions of the traditional theory. The fourth part of the book includes four influential papers, each representing a different theoretical framework for understanding the impact of trade on inequality.

The first paper by Costinot and Vogel is motivated by the observation that the traditional factor proportions theory yields stark predictions regarding the effect of trade on the wage distribution only in its 2x2 version. More general versions featuring multiple goods and factors have the unfortunate property that the predictions depend critically on the number of goods and factors. From an empirical point of view, it is not ideal for results to depend on the particular way goods and factors are defined. Against this background, one might posit that the entire empirical literature on trade and inequality has had a loose relationship to the traditional Heckscher-Ohlin theory, as this theory in its multi-dimensional version, strictly speaking, has nothing definitive to say about the impact of trade on inequality. Even papers like the one of Porto that has a close connection to theory interpret the empirical findings using the *intuition* conveyed by the 2x2 model rather than rigorous theoretical results. This is where Costinot and Vogel come in. One of the main contributions of their work is that they provide strong theoretical results for the multi-dimensional version of the factor proportions theory.

The clever idea of the paper is to transform the problem of analyzing a competitive equilibrium to a matching problem, in which workers are matched to tasks. This allows the authors to derive predictions for a world with many factors and goods, which can be taken to the data. Interestingly, if one focuses on a setting in which trade is driven only by differences in factor endowments across countries, then the predictions of this more general framework are similar to those derived in the stylized 2x2 model: trade will have opposite effects on inequality in developed and developing countries. This result however disappears if one considers a second source of comparative advantage, technological differences across countries. In this case, it is shown that the effect of globalization on inequality depends on the correlation between factor endowments and technological differences. There is abundant evidence that skill-abundant countries tend to employ skill-biased technologies. This correlation together with the paper's theoretical results can rationalize a number of empirical findings. Moreover, the framework can be employed to analyze additional phenomena such as global technological change or offshoring.

While Costinot and Vogel's framework is based on traditional notions of comparative advantage, the second paper by Helpman et al provides a more drastic departure from the standard theory, by proposing a framework that draws on recent trade theories emphasizing firm heterogeneity and frictions. The model introduces search and matching frictions into a Melitz model of international trade along with ex-post match-specific heterogeneity in workers' abilities. Larger, more productive firms pay higher wages and exporting increases the wages paid by firms of a given productivity. The framework delivers a series of interesting predictions that match empirical regularities. Most importantly, it suggests a mechanism that generates changes in *within-group* and *within-industry* inequality – dimensions that have been shown to be important in the data. Another result with significant implications for inequality is that the impact of a gradual trade liberalization on inequality is shown to be non-monotonic: inequality first increases and then decreases as trade barriers go down. Finally, the framework can also be used to analyze the effects of trade on unemployment. The theoretical results suggest that the effect is ambiguous as trade can lead to either an increase or decrease in unemployment depending on whether or not trade increases the tightness of the labor market.

The last two papers in this part of the book focus on another line of explanation for the rising inequality that emphasizes the changing structure of production and globalization of production activities. Antras et al develop a framework in which heterogeneous agents are organized in hierarchical teams. Less-skilled workers perform routine production tasks, while high-skilled workers perform knowledge-intensive tasks such as managing. Globalization is modeled as a process that blurs national borders in labor markets so that agents in different countries can join together in teams. Countries differ in their distributions of worker ability. Globalization allows high-skilled workers in advanced economies to leverage their knowledge at lower cost by working with less skilled workers in developing countries on routine tasks while high-ability workers in developing economies benefit from working in international teams. The model captures the main features of offshoring and explores its implications for inequality. It is shown that globalization increases inequality among workers in developing countries, but not necessarily in developed countries where the effect is ambiguous.

Costinot et al in the last paper featured in the book also focus on offshoring with the difference that instead of relying on the distinctions between routine and knowledge-intensive tasks and managers versus non-managers, they emphasize the fragmentation of the production process into sequential stages. Production of the final good requires a series of stages to be performed sequentially. At each stage, the output of intermediate goods depends on the skill of workers employed at the previous stage. These input-output linkages are essential. As in Antras et al, countries in developed and developing countries differ in the distribution of worker ability. Globalization is modelled as a process of integration of intermediate and final product markets that gives rise to "global supply chains". The paper shows that this process leads workers in developing countries to move into earlier stages of production. The result is a non-monotonic effect of globalization on inequality in these countries, with inequality decreasing at the bottom of the wage distribution, but increasing at the top.

Overall, the new directions explored in the recent theoretical work on trade and inequality lead to two observations. First, in contrast to the standard Heckscher-Ohlin theory that delivered stark predictions about the impact of trade on inequality in developed and developing countries, the new theories yield more nuanced results with effects that are often ambiguous, non-monotonic in trade barrier reductions, and dependent on several factors such as initial conditions, the extent of liberalization, frictions in other markets (most importantly labor markets), technological differences across countries, the structure of production, and so forth. Precisely because these results are more nuanced, they do a much better job than the standard Heckscher-Ohlin model in matching the richness and complexity of the empirical evidence. At the same time, from a policy perspective, they suggest that unqualified statements about the effects of globalization on inequality are unwarranted. Each case is different, and an informed perspective on this topic requires a careful study of the institutional setting, the production structure, the functioning of the markets in each country and the degree and nature of liberalization.

Second, these studies raise the question whether the mechanisms highlighted in each case can generate quantitatively the changes in inequality observed in the last three decades. *A-priori*, it seems that the scope of trade to impact inequality may be greater in developing countries, especially poorer ones, given that trade in many such countries constitutes a significant part of their economic activity. Against this background, it is perhaps no surprise that some of the largest effects of opening up to trade have been documented in a poor country such as Vietnam that is heavily dependent on trade. But the evidence on this issue is still scant. We need many more empirical, but theoretically informed, studies in order to assess the relevance of each mechanism, not only in qualitative, but also quantitative terms, in explaining the rise in within-country inequality.

References for Introduction

Acemoglu, Daron, and David H. Autor (2011), "Skills, Tasks, and Technologies: Implications for Employment and Earnings." In Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, vol. 4, 1043-171. Amsterdam: Elsevier-North Holland.

Autor, David (2014): "Skills, Education and the Rise of Earnings Inequality Among the 'Other 99 Percent'," *Science*, 23 May 2014: Vol. 344 no. 6186, pp. 843-851.

Autor, David H., and David Dorn (2013), "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market," *American Economic Review*, 2013, 103(5), pp. 1553–1597.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2006), "The Polarization of the U.S. Labor Market," *American Economic Review* 96, no. 2 (May), pp. 189-94.

Bartik, Timothy J. (1991), "Who Benefits from State and Local Economic Development Policies?", W.E. Upjohn Institute for Employment Research, 1991.

Blinder, Alan S. and Alan B. Krueger (2013), "Alternative Measures of Offshorability: A Survey Approach", *Journal of Labor Economics*, 31(2): S97-128.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux (2011), "Occupational Tasks and Changes in the Wage Structure," University of British Columbia Working Paper.

Goldberg, Pinelopi and Nina Pavcnik (2007): "Distributional Effects of Globalization in Developing Countries," *Journal of Economic Literature*, March 2007, 45(1), pp. 39-82.

Goos, Maarten, and Alan Manning (2007), "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics* 89, no. 1 (February), pp. 118-33.

Goos, Maarten, Manning, Alan and Salomons, Anna (2014), "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, August 2014, pp. 2509-2526.

Kaplan, Steven and Joshua Rauh (2013): "It's the Market: The Broad-Based Rise in the Return to Top Talent," *Journal of Economic Perspectives,* Volume 27, Number 3, Summer 2013, pp. 35–56.

Lucas, Robert, Jr. (2004): "The Industrial Revolution: Past and Future," *The Region*, The Federal Reserve Bank of Minneapolis, Banking and Policy Issues Magazine, May 2004, 2003 Annual Report Issue.

Schott, Peter (2003), "One Size Fits All? Heckscher-Ohlin Specialization in Global Production," *American Economic Review*, 93(2), June 2003, pp. 686-708.

Schott, Peter (2004), "Across-Product versus Within-Product Specialization in International Trade," *Quarterly Journal of Economics* 119(2), May 2004, pp. 647-678.