

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Muffert, Johanna; Winkler, Erwin

Working Paper Using Machine Learning to Understand the Heterogeneous Earnings Effects of Exports

IZA Discussion Papers, No. 17667

Provided in Cooperation with: IZA – Institute of Labor Economics

Suggested Citation: Muffert, Johanna; Winkler, Erwin (2025) : Using Machine Learning to Understand the Heterogeneous Earnings Effects of Exports, IZA Discussion Papers, No. 17667, Institute of Labor Economics (IZA), Bonn

This Version is available at: https://hdl.handle.net/10419/314564

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 17667

Using Machine Learning to Understand the Heterogeneous Earnings Effects of Exports

Johanna Muffert Erwin Winkler

FEBRUARY 2025



Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 17667

Using Machine Learning to Understand the Heterogeneous Earnings Effects of Exports

Johanna Muffert University of Erlangen-Nuremberg

Erwin Winkler University of Erlangen-Nuremberg, IZA and LASER

FEBRUARY 2025

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9	Phone: +49-228-3894-0	
53113 Bonn, Germany	Email: publications@iza.org	www.iza.org

ABSTRACT

Using Machine Learning to Understand the Heterogeneous Earnings Effects of Exports^{*}

We study how the effects of exports on earnings vary across individual workers, depending on a wide range of worker, firm, and job characteristics. To this end, we combine a generalized random forest with an instrumental variable strategy. Analyzing Germany's exports to China and Eastern Europe, we document sharp disparities: workers in the bottom quartile (ranked by the size of the effect) experience little to no earnings gains due to exports, while those in the top quartile see considerable earnings increases. As expected, the workers who benefit the most on average are employed in larger firms and have higher skill levels. Importantly, however, we also find that workers with the largest earnings gains tend to be male, younger, and more specialized in their industry. These factors have received little attention in the previous literature. Finally, we provide evidence that the contribution to overall earnings inequality is smaller than expected.

JEL Classification: Keywords: C52, F14, J23, J24, J32 machine learning, earnings, inequality, exports, skills, labor market

Corresponding author:

Erwin Winkler University of Erlangen-Nuremberg Lange Gasse 20 90403 Nuremberg Germany E-mail: erwin.winkler@fau.de

^{*} We are grateful to Elliot Ash, Bernd Fitzenberger, Selina Gangl, Ingar Haaland, Martin Halla, Ines Helm, Christoph Moser, Markus Nagler, Oskar Nordström-Skans, Andreas Peichl, Regina Riphahn, Andreas Steinmayr, and participants at the CESifo Workshop on Big Data, EALE, EEA, and ETSG for helpful comments and suggestions. Winkler is grateful to the Joachim Herz Foundation for generous financial support.

1 Introduction

Who benefits and who loses from international trade? While the gains from trade are well-documented, their distribution across the population remains an important concern. For policymakers to be able to design effective policies, it is important to understand what characteristics allow workers to benefit most from trade and, analogously, what characteristics make workers vulnerable to fall behind. Workers may, for example, be differentially affected by trade depending on their skills and human capital, their industry, employer, job, demographics, and various interactions of these characteristics. In addition, the unequal effects of trade on workers can have broader implications for earnings inequality.

The recent literature has made substantial progress in documenting heterogeneous effects of trade on workers, often by focusing on one dimension of heterogeneity at a time, using sample splitting or interaction effects (e.g., Autor et al., 2014; Utar, 2018; Dauth et al., 2021; Keller and Utar, 2022; Kovak and Morrow, 2023, 2025). However, these approaches yield a coarse and potentially incomplete picture about the heterogeneity of effects across individuals. This is because the 'true' effect on an individual worker might simultaneously depend on many factors including, e.g., employer characteristics such as firm size (Kovak and Morrow, 2025), various types of skills (Utar, 2018; Dauth et al., 2021), and demographics like gender or age (Keller and Utar, 2022). Importantly, these factors might not simply add up, but interact in various and non-expected ways. In addition, the relative importance of different factors in explaining heterogeneous outcomes is still up to debate. For example, while workers with certain skills might benefit more from trade than others, it remains an open question what kind of skills become particularly valuable. Overall, a more comprehensive approach is needed to fully understand the heterogeneous earnings effects.

In this paper, we leverage advances in machine learning to better understand the heterogeneous earnings effects of trade shocks. Using a generalized random forest (Athey et al., 2019), we analyze the unequal earnings effects of Germany's exports to China and Eastern Europe since the 1990s. This method allows the effects of exports on earnings to differ flexibly across a large number of worker, firm, and job characteristics, as well as their interactions. We address three key questions. First, how large are the differences in earnings effects between those who benefit most and those who benefit least? Second, which characteristics explain who benefits most and who benefits

least? And third, to what extent do increased exports contribute to overall earnings inequality between workers? Germany's trade integration with China and Eastern Europe in the 1990s and 2000s provides an ideal setting to study these important questions. Following the fall of the Iron Curtain and China's transition into a market economy, Germany's exports to China and Eastern Europe increased sharply and unexpectedly fast - much faster than with the rest of the world.

We use detailed administrative data on worker careers and exploit variation in industry-level exports over time between manufacturing workers who are observationally identical in terms of their skills, demographics, as well as initial job and firm characteristics at baseline. Following an influential literature, we instrument Germany's industry-level trade flows with industry-level trade flows of a set of other high-income countries (e.g., Autor et al., 2013, 2014; Dauth et al., 2014, 2021).¹

Based on the existing theoretical and empirical literature, we identify 15 variables that we use to estimate heterogeneities in the effects of exports on workers' earnings. Using these variables, the algorithm looks for sample splits that maximize the heterogeneity in the estimated effect of exports on earnings across the resulting sub-samples. The algorithm builds a tree by splitting the data into increasingly smaller sub-samples to detect heterogeneity in the estimated effect. To minimize the risk of estimating spurious heterogeneities, the algorithm builds a large number of trees where each tree is based on a bootstrap sample of the original sample. As a result, we obtain fine-grained estimates of the effect of exports on earnings that vary at the level of individual workers as a function of worker, firm, and job characteristics. The method allows for potentially important non-linearities and interaction effects in a purely data-driven way and does not require us to make ex-ante choices on the exact functional form of the heterogeneity. Importantly, since the generalized random forest incorporates a large number of potential drivers of heterogeneity in a unified way, it enables us to perform a 'horse race' between different potential drivers of heterogeneous earnings effects.

Our first main result is that the positive earnings effects of exports are highly unevenly distributed across workers. We rank all workers according to the size of their individualized earnings effect, and it turns out that the bottom quartile of

¹We focus on the heterogeneous effects of exports, always controlling for the possibly heterogeneous effects of imports. We make this choice because our results as well as the results in the previous literature show that, in the context of Germany's trade with China and Eastern Europe, the effects of exports are of much larger economic significance than the effects of imports (Dauth et al., 2014, 2021).

workers experiences little to no earnings gains due to exports, while the top quartile experiences considerable positive earnings increases. Taking the estimates at face value, the export-induced earnings increase for a worker at the 90th percentile is about nine times larger than the increase for a worker at the 10th percentile and about twice as large as for a worker at the 50th percentile.

Our second main result is that the heterogeneity of earnings responses is multidimensional and cannot be summarized by one or two factors alone. As expected, we find that workers with the largest positive effects tend to be employed in larger firms and tend to have higher levels of (unobserved) skills, as measured by AKM worker effects estimated prior to our observation window (Abowd et al., 1999; Card et al., 2013). This result is consistent with models of firm and worker heterogeneity where the largest and most productive firms select into exporting and skilled workers benefit more from exporting through various channels.² However, we find that conventionally used firm- and skill-related variables and their interactions do not explain more than 40% of the total variation across workers in the estimated effects.

Importantly, we find that workers who benefit most are substantially more likely to be male, are younger, have lower tenure, and are employed in occupations that are highly specialized in the respective workers' industry. Overall, while these factors have received comparatively little attention in prior research, we show that they are at least as important as, e.g., measures of worker skills, in explaining who benefits most and who benefits least from exports. Note that all of these factors remain important predictors of large positive effects when conditioning on each other, suggesting that they partly reflect distinct underlying mechanisms.

In a final step, we provide evidence that the export shock had a very modest positive effect on overall inequality. In back-of-the-envelope calculations, we use our fine-grained estimates and compute the average estimated effect of exports on earnings within percentiles of the initial earnings distribution. It turns out that the earnings effect is most positive between the 50th and the 70th percentile of the initial earnings distribution. However, the average effect is robustly positive in all percentiles. We find that the Gini coefficient increases only slightly as a consequence. This is partly because the direction of the effects of tenure, age, and industry specialization

²See, for example, Egger and Kreickemeier (2009, 2012); Amiti and Davis (2011); Helpman et al. (2010); Sampson (2014); Burstein and Vogel (2017); Egger et al. (2020).

works towards a decrease in inequality.³ Put differently, a look at heterogeneity by skill-related variables alone might lead to an overestimation of the inequality effects of exports.

Our results provide a couple of novel insights which might be interesting from a policy perspective. First, while measures capturing unobserved worker skills are highly predictive of large positive earnings gains, formal education plays only a minor role. This means that exports might lead to an increase in skill demand, but not necessarily to an increase in the demand for workers with a university degree. Second, export shocks might not be gender-neutral. Importantly, we find that the gender differences cannot be explained by differential sorting into firms or occupations or gender differences in skills and human capital, suggesting that men and women differ in terms of their adjustment to export-induced demand shocks. We discuss a variety of mechanisms, including gender differences in job and geographic mobility, schedule flexibility, and bargaining behavior. Third, largely unequal earnings effects of exports do not necessarily translate into large increases in inequality between highand low-earnings workers. We find that most of the differences in the effects occur between workers with similar earnings rather than between high- and low-earnings workers.

This paper contributes to several strands of the literature. First, it contributes to the large literature on the labor market effects of trade shocks. Most closely related to our paper, a number of studies use longitudinal data to analyze the adjustment of workers to trade shocks and to trace out the consequences for workers' earnings (e.g., Autor et al., 2014; Utar, 2018; Dauth et al., 2021; Kovak and Morrow, 2025).⁴ Most of this literature focuses on import shocks, while a small but growing number of studies provide evidence on workers' adjustment to increased export opportunities (e.g., Huber and Winkler, 2019; Dauth et al., 2021; Kovak and Morrow, 2025). Closely related to this literature, a number of papers use regional variation to examine the labor market effects of trade shocks (e.g., Autor et al., 2013; Dauth et al., 2014; Balsvik et al., 2014; Hakobyan and McLaren, 2016; Dix-Carneiro and Kovak, 2019; Helm, 2020).⁵

³For example, young and low-tenure workers experience more positive effects (all else being equal), and have lower earnings to begin with than older and high-tenure workers.

⁴See also Hummels et al. (2014); Keller and Utar (2022); Arni et al. (2024); Menezes-Filho and Muendler (2011); Nilsson Hakkala and Huttunen (2018).

⁵See also the literature which uses a structural approach to study the adjustment of workers and labor market to trade (e.g., Dix-Carneiro, 2014; Dix-Carneiro and Kovak, 2017; Dix-Carneiro et al., forthcoming; Cosar et al., 2016; Traiberman, 2019).

To the best of our knowledge, this paper is the first one to use machine learning methods to study the unequal effects of trade shocks on workers' earnings. While the existing literature typically investigates one source of heterogeneity at a time in an isolated way using sample splitting or interaction effects, our approach offers a unified way to jointly study a large number of potential sources of heterogeneous earnings outcomes. This way, we can portray the full distribution of the effects across workers which results from the interplay of a variety of factors. Second, our approach provides a horse race between different potential drivers of heterogeneous earnings effects of exports. While the existing literature puts a strong focus on measures of worker skills or firm heterogeneity, our results point to an important role of factors which have received comparatively little attention, such as age, gender, or a worker's skill specialization. More generally, our results suggest that the heterogeneity of earnings responses is multidimensional.

We also contribute to the broader literature on trade and inequality. One strand of the literature uses matched employer-employee data to study the effects of trade on wages and wage inequality in the context of firm and worker heterogeneity (e.g., Schank et al., 2007; Verhoogen, 2008; Amiti and Davis, 2011; Frías et al., 2024).⁶ Much of this literature focuses on firm-level outcomes rather than worker-level outcomes. A recent strand of the literature uses a structural approach to quantify the contribution of trade to the skill premium and earnings inequality more generally (e.g., Burstein and Vogel, 2017; Cravino and Sotelo, 2019; Adao et al., 2022; Borusyak and Jaravel, 2024). Compared to the existing literature, our approach allows the effects of exports on earnings inequality to vary along a large number of dimensions.

We also contribute to a small but growing literature emphasizing gender differences in the effects of trade (e.g., Juhn et al., 2014; Saure and Zoabi, 2014; Bøler et al., 2018; Autor et al., 2019; Keller and Utar, 2022). Our method enables us to benchmark the gender inequality in the estimated impact of exports against other dimensions of inequality, for example by skill group. In addition, we quantify the extent to which these gender differences arise due to differences in characteristics between men and women and differential sorting across firms and jobs.

Finally, our paper is related to a growing literature that applies machine learning methods to study treatment effect heterogeneity in other settings. Recently, the generalized random forest has been applied to study the impact of summer jobs

⁶See also Munch and Skaksen (2008); Baumgarten (2013); Irrarazabal et al. (2013); Krishna et al. (2014); Helpman et al. (2017); Garin and Silverio (2023).

(Davis and Heller, 2017, 2020), the effects of job loss on crime (Britto et al., 2022), the earnings effects of job loss (Gulyas and Pytka, 2022; Athey et al., 2024), the effectiveness of training programs for unemployed individuals (Cockx et al., 2023), and to study education choices (Carlana et al., 2022). A small literature uses machine learning methods to study heterogeneities in the impact of trade agreements on trade flows (Breinlich et al., 2022; Baier and Regmi, 2023).

The remainder of this paper is structured as follows. In section 2, we introduce the data. In section 3, we then estimate the average effect of exports on earnings, using conventional IV methods. In section 4, we explain the use of machine learning to estimate the heterogeneous earnings effects of exports. Section 5 provides the results on heterogeneous earnings effects. Section 6 provides back-of-the-envelope calculations on the link between exports and inequality. Finally, section 7 concludes.

2 Data

We use two main data sets in our analysis. First, we use the 'Sample of Integrated Labour Market Biographies (SIAB)', provided by the Institute for Employment Research in Nuremberg (IAB) (version 7519 v1). The SIAB is a 2% random sample of the administrative social security records and covers about 80% of the German workforce, excluding self-employed persons, civil servants, and military personnel. The SIAB provides panel data on workers and allows us to follow career trajectories on a daily basis, with information on the employment status and daily wages. In addition, the data set provides information on individual characteristics, such as age, gender, nationality, education, and attributes of the establishment of a worker (e.g., industry, location, number of employees). Since wages in the SIAB are right-censored at the social security contribution ceiling, we use the imputation procedure of Card et al. (2013).⁷ We convert earnings into 2010 Euros using the consumer price index of the German Bundesbank. Following Dauth et al. (2021), we use two observation periods: 1990-2000 and 2000-2010. We focus on workers who are full-time employed in the manufacturing sector in the base year of each period (i.e., 1990 or 2000), who have at least two years of tenure, and who are aged between 22 and 55 years. We then follow these workers over ten years and do not impose any restrictions for the years following the base year. This means that workers might move between industries or sectors or become unemployed during the subsequent 10-year period. For each

⁷Around 13 % (15 %) of wage observations are censored in our base years 1990 (2000).

worker, we compute cumulative earnings during the 10-year period following the base year (1991-2000 and 2001-2010). Table A.1 in the Appendix shows descriptive statistics on the estimation sample.

Second, we combine the administrative labor market data with international trade data from the United Nations Commodity Trade Statistics Database (Comtrade). The Comtrade Database entails annual trade data for over 170 countries. Using the correspondence between the SITC rev.3 and the NACE codes provided by the UN Statistics Division, we can match the trade flows with the SIAB using the three-digit industry codes. Trade flows are also converted into 2010 Euro. To construct our main explanatory variables, we focus on Germany's exports (and imports) vis-à-vis China and a group of countries which we call 'Eastern Europe'.⁸ Panel (a) of Appendix Figure A.1 shows how Germany's trade with these countries has evolved over time. Before 1990, trade flows stagnated at a relatively low level. However, after 1990, during our period of analysis, trade flows increased rapidly and substantially. This rise in trade coincides with the sudden fall of the Iron Curtain and China's transformation into a market economy. Panel (b) of Appendix Figure A.1 shows that the share in total trade increased from below 5% before 1990 to 15-20% in 2010. This reflects the fact that Germany's trade with these countries increased substantially more than with the rest of the world during that period. Appendix Table A.2 shows the manufacturing industries with the largest increase in exports or imports, respectively. Note that we focus on the period 1990 to 2010 and do not include later years in order to obtain results that are comparable to the existing literature, which focused on this period of time (Dauth et al., 2021, 2014; Autor et al., 2013, 2014).

3 The average effect of exports on earnings

3.1 Empirical specification

In this section, we estimate the *average* effect of industry-level export (and import) exposure on cumulative earnings, using the empirical strategy along the lines of Autor et al. (2013, 2014) and Dauth et al. (2014, 2021). We estimate variants of the following specification:

⁸Following Dauth et al. (2014, 2021), we define 'Eastern Europe' as the set of the following countries: Bulgaria, Czech Republic, Hungary, Poland, Romania, Slovakia, Slovenia, the former USSR, and its successor states the Russian Federation, Belarus, Estonia, Latvia, Lithuania, Moldova, Ukraine, Azerbaijan, Georgia, Kazakhstan, Kyrgyzstan, Tajikistan, Turkmenistan, and Uzbekistan.

$$Earn_{ijt} = \beta^{Exp} Exp_{jt} + \beta^{Imp} Imp_{jt} + \Omega'_{ijt}\gamma + \epsilon_{ijt}$$
(1)

 $Earn_{ijt}$ denotes normalized earnings of worker *i* who was initially employed in 3-digit industry *j* in the base year (i.e., the first year) of period *t* (1990-2000 or 2000-2010). We compute $Earn_{ijt}$ as cumulative earnings over ten years (1991-2000 or 2001-2010), divided by base year earnings. We normalized by base year earnings instead of taking logs to allow workers with zero cumulative earnings after the base year to remain in the sample, following Autor et al. (2014).

 Exp_{jt} and Imp_{jt} are our measures of the change in export and import exposure, respectively, in industry *j* in which worker *i* was employed in the base year. To construct these variables, we compute the change in exports or imports, respectively, over ten years (1991-2000, 2001-2010), divided by the industry's base year wage sum to account for size differences across industries:

$$Exp_{jt} = \frac{\Delta Exports_{jt}}{WageSum_{jt}} \quad \text{and} \quad Imp_{jt} = \frac{\Delta Imports_{jt}}{WageSum_{jt}}$$
(2)

 Ω'_{ijt} contains an extensive set of worker, firm, and job/occupation controls, all measured in the base year. It includes education (university degree, vocational degree, less than vocational degree), gender, foreign nationality, age, experience (measured as total days in employment), firm tenure, industry tenure, occupation tenure, firm size (measured in number of total employees), pre-estimated AKM worker and firm effects, dummies for manager status and blue-collar workers, the routine share of the worker's 2-digit occupation, dummies for four broad industry groups (food products, consumer goods, industrial goods, capital goods), dummies for 141 labor market regions, and a dummy to differentiate between the two time windows, and the occupational employment share in the initial industry, respectively. We explain the construction of (and the intuition behind) the AKM effects, the routine share, and the occupational employment share in the initial industry in section 4.

The idea behind equation 1 is to compare the outcomes of workers who are very similar in terms of their socio-economic and demographic background, their previous labor market experience, as well as in terms of their initial plant and occupation, but are differently affected by the rise in trade with China and Eastern Europe because they happen to be employed in different 3-digit industries in the base year. Following Dauth et al. (2021), we cluster standard errors at the 3-digit industry x base year x labor market region level.

3.2 Instrument

A remaining threat to identification are industry-level demand and productivity shocks, which might be correlated with measured export and import exposure, and at the same time influence the workers' labor market outcomes. We implement the IV strategy proposed by Autor et al. (2013, 2014) and adapted to the German context by Dauth et al. (2014, 2021). More specifically, we instrument export and import exposure in Germany by export and import exposure in a set of third countries. To construct the instrument, we replace the change in German exports and imports in the numerator of equation 2 by the change in exports and imports in the following countries: Australia, Canada, Japan, Norway, New Zealand, Sweden, Singapore, and the United Kingdom. In addition, we lag the denominator by one further year, following Dauth et al. (2021). The idea of the IV strategy is that China and Eastern Europe experienced rapid productivity growth due to their transition to a market economy which went along with capital accumulation, migration to rural areas and improvement of the infrastructure. The productivity growth translated into a strong increase in import demand and in export capabilities. For China, this effect was amplified through its entry into the WTO at the beginning of the 2000s. This effect should not only be present for Germany, but also in other high-income countries. Then, instrumenting German industry-level trade with industry-level trade of these high-income countries should isolate the exogenous increase in export opportunities and import competition that is related to the productivity growth in China and Eastern Europe. For this strategy to be valid, export and import exposure of the instrument countries must not have a direct impact on German industries and industry-level supply and demand shocks in these countries should not be strongly correlated with those for German industries. The instrument group therefore does not contain any direct neighbors to Germany, no members of the European Monetary Union, and excludes the USA. It is noteworthy, however, that our heterogeneity estimates remain very similar when employing OLS rather than IV regressions.

3.3 Average effect of exports on earnings

Table 1 provides the results of the estimated average effect of increased exports on workers' cumulative earnings. The OLS estimate in column (1) points to a highly statistically significant positive link between export exposure and earnings. Column (2) shows that the point estimate becomes slightly larger when we apply our instrument.

The coefficient of 0.722 implies that a one standard deviation increase in export exposure (23.80, see Table A.1), translates into a cumulative earnings gain of about 17.2% of a workers' base year earnings over ten years (0.722 * 23.80 \approx 17.2). For a worker with mean base year earnings (44,762.66 EUR, see Table A.1), this implies a cumulative earnings gain of roughly 7,700 EUR over ten years (44,762.66 * 17.2/100 \approx 7,700) or on average 770 EUR per year. The OLS estimate of import exposure on earnings is substantially smaller in magnitude and statistically insignificant. The IV estimate is larger than the OLS estimate and statistically significant, consistent with the notion that unobserved demand shocks are biasing the coefficient towards zero. According to the estimated coefficient of -0.150, a one standard deviation increase in import exposure (40.92, see Table A.1) translates into a cumulative earnings loss of 6.1% of workers' base year earnings over ten years. The estimated negative effect of imports therefore is very modest, for example compared to estimates on the negative effects of automation of workers' cumulative earnings in the literature.⁹ Columns (3) and (4) provide a look at the first-stage estimates. The coefficients of the instruments are relevant, with F-statistics well beyond 100.

The estimated positive effects of exports on earnings are in line for example with models of rent sharing, where the additional rents from increased exports are partly shared with the workforce (Egger and Kreickemeier, 2009, 2012; Amiti and Davis, 2011). In Appendix Table A.3, we show that the positive earnings effects arise partly at the initial employer and partly at a different employer in the same 2-digit industry.¹⁰ This result is in line with a broad range of models with heterogeneous firms along the lines of Melitz (2003) which predict intra-industry reallocation of workers in response to trade. The negative estimates of import exposure on earnings are consistent with the notion that workers in import-competing industries experience increased rates of displacement and lose parts of their specific human capital upon moving between industries and sectors (e.g., Utar, 2018; Traiberman, 2019). Indeed, Appendix Table A.3 shows that import exposure is associated with a sharp decline in earnings at the initial employer. This is compensated partly (but not fully) by increased earnings at different employers, in particular in the service sector.

⁹For example, Bessen et al. (2025) find that automation events at firms trigger a cumulative earnings losses of around 9% of incumbent workers' base year earnings over only five years (rather than 10 years).

¹⁰We obtain these estimates by decomposing the dependent variable into different sources, e.g., counting only earnings that arise at the initial employer or counting only earnings that arise at a different employer in the same industry). This method is based on Autor et al. (2014) and Dauth et al. (2021).

	(1)	(2)	(3)	(4)
	OLS		IV	
		2nd stage	1st stage	1st stage
Dependent variable:	Earnings	Earnings	Exports	Imports
Exports	0.653***	0.722***		
	(0.063)	(0.127)		
Imports	-0.056	-0.150**		
	(0.037)	(0.062)		
Exports (instrument)			0.203***	0.159***
			(0.015)	(0.022)
Imports (instrument)			0.016***	0.227***
			(0.005)	(0.018)
Observations	159,213	159,213	159,213	159,213
R^2	0.111	0.111	0.417	0.504
F-Stat. of excl. instrument			189.457	150.618

Table 1: The average effect of exports on earnings

Note: In columns (1) and (2), the dependent variable is 100 x normalized cumulative earnings, i.e., cumulative earnings over 10 years divided by base year earnings. See equations 1 and 2 for the empirical specification. Standard errors, clustered by 3-digit industry x base year x local labor market, in parentheses. * p<0.10, ** p<0.05, *** p<0.01

Overall, the finding that the positive earnings effects of exports is substantially larger in magnitude than the negative earnings effects of imports is in line with previous evidence for Germany by Dauth et al. (2014, 2021). Due to the much larger economic significance of exports in this context, we choose to focus on the heterogeneity of export-induced earnings gains in the remainder of the paper, always controlling for the (possibly heterogeneous) earnings effects of imports.

4 Using machine learning to estimate the heterogeneous effects of exports

4.1 The generalized random forest

In order to analyze heterogeneities in the earnings effects of exports, we use recent advances in causal machine learning by Athey and Imbens (2016) and Athey et al. (2019). Generalized random forests (GRFs), as proposed by Athey et al. (2019), allow us to combine an instrumental variable regression with causal machine learning tailored for capturing meaningful differences in effects.

Intuitively, the GRF creates trees by splitting the sample into sub-samples based on the characteristics of the unit of observation in order to maximize differences in estimated effects between the resulting sub-samples. The resulting sub-samples at the bottom of the tree are called 'leaves' and the tree is grown deeper by recursively splitting into smaller and smaller sub-samples until a stopping criterion (e.g., a minimum number of observations within a leave) is reached. By construction, the tree allows for interaction effects and non-linearities in a data-driven way, without the need to ex-ante specify the functional form of the interaction of non-linearity. We employ the same control variables as in the previous section where we estimated the average effect of exports

Splitting the sample into many sub-samples may lead to a problem typically called 'overfitting'. This means that we might obtain estimates of heterogeneous effects which result from the peculiarities of the data set at hand and which may not hold in other samples. The GRF takes several steps to mitigate this risk. First, the GRF estimates a large number of trees, where each tree is based on a bootstrap sample of the original sample. The final estimates are an aggregation of the results of the single trees, making the final estimates less sensitive against the influence of single observations in the data. In addition, for a given tree, the GRF applies a principle called 'honesty'. This means that it splits the sample and uses different observations for constructing the tree and for estimating the heterogeneous effects across the leaves of the tree. We employ a GRF with 15,000 trees with a minimum leave size of 150 and otherwise standard settings of the package 'grf' in R. A more detailed description of the GRF and our chosen specification is provided in Appendix A.1.

Using a GRF, we obtain estimates of the effect of exports on earnings which vary across individual workers as a function of worker, firm, and job characteristics

and their interactions. We refer to the outcome of the GRF as the 'individualized effect of exports on earnings', denoted by $\hat{\beta}(x_i)$, where x_i denotes the characteristics of individual *i*.¹¹ Compared to the conventional approach of sample splitting or interaction effects based on one or two variables at a time, the GRF provides a much more fine-grained picture of the heterogeneous effects of exports on earnings. In our main analysis, we work with the individualized effects in order to study the characteristics which are associated with large/small earnings effects and to analyze the link between exports and earnings inequality.

4.2 Variables to study heterogeneous earnings effects

In what follows, we describe the set of worker, job and firm characteristics that we feed into the GRF in order to study the heterogeneous effects of exports. While the GRF is in principle able to accommodate a very large number of variables to study heterogeneous effects, we focus on 15 variables for which we have a prior (based on existing theories or empirical work) that these might play a role in explaining heterogeneous earnings effects. In what follows, we describe these priors. All of these variables are measured in the base year and are also used as control variables. Table 2 provides an overview of all variables.

4.2.1 Firm-level variables

We allow the earnings effects of industry-level exports to differ across workers initially employed in different firms. Models with heterogeneous firms along the lines of Melitz (2003) predict that the largest and most productive firms select into exporting. A range of models that combine firm heterogeneity with labor market frictions predict that workers employed in exporting firms benefit from increased revenues through different forms of rent sharing (e.g., Egger and Kreickemeier, 2009, 2012; Amiti and Davis, 2011).¹²

While we cannot measure or estimate productivity directly, we use two alternative proxies. Our preferred proxy for firm productivity is firm size, measured as the total

¹¹We omit the superscript 'Exp' and the time and industry subscripts for simplicity.

¹²One strand of the literature incorporates a fair-wage mechanism into a Melitz (2003)-type model of firm heterogeneity (e.g., Egger and Kreickemeier, 2009, 2012; Egger et al., 2020; Amiti and Davis, 2011). In these models, only the most productive firms within an industry select into exporting and thereby increase their profits. Increased profits in exporting firms go along with increased wages due to the fair-wage mechanism. A closely related strand of the literature combines the Melitz (2003)-framework with search frictions and bargaining of the surplus of the firm (Helpman et al., 2010, 2017).

number of employees (Kovak and Morrow, 2025). Firm size is strongly correlated with firm productivity (e.g., Syverson, 2011) and with export status (e.g., Bernard et al., 2012). Additionally, following recent work by Dauth et al. (2021), we also include the AKM firm wage premium into our analysis as a proxy for productivity.¹³ While the AKM firm wage premium is often used as a proxy for productivity, it might also reflect mechanisms unrelated to productivity. For example, it might capture strategic wage-posting behavior of firms (Burdett and Mortensen, 1998), compensating wage differentials for disamenities (Sorkin, 2018), or institutional factors such as collective bargaining agreements and the existence of works councils (Hirsch and Müller, 2020). While our firm variables are measured in the base year, some workers initially employed in small and unproductive firms might endogenously move towards exporting firms in response to trade and reap earnings gains. Our analysis allows for this type of mobility within the estimation window. Nevertheless, in a frictional labor market, the earnings effects might still depend crucially on the type of firm in which a worker is starting at baseline.

4.2.2 Human capital and skills

Existing theory suggests that increased international trade might lead to a rise in the return to skills or human capital. First, relative demand for skills might increase because trade favors a reallocation of employment towards the most productive firms which typically produce with higher skill-intensity (e.g., Burstein and Vogel, 2017). Second, more skilled workers might benefit relatively more from reallocation towards the most productive firms due to strong complementarities between worker skills and firm productivity (e.g., Sampson, 2014). Third, exports can make good worker-firm matches more valuable, leading to an increase in assortative matching between workers and firms and favoring more skilled workers (e.g., Davidson et al., 2014; Bombardini et al., 2019). Fourth, firms might endogenously become more skill-intensive due to

¹³AKM firm wage premia are obtained from the following regression: $y_{it} = \alpha_i + \psi_{J(it)} + z'_{it}\gamma + r_{it}$, where y_{it} denotes the log daily wage of worker *i* in year *t*. α_i reflects the worker effect. It captures all time-invariant observable and unobservable worker characteristics and the return to these characteristics. z'_{it} is a vector of year dummies and a cubic term in age fully interacted with education dummies. $\psi_{J(it)}$ is the AKM firm wage premium which reflects a proportional wage premium or wage discount that firm *J* pays to all its employees, regardless of their characteristics or skills. Estimated AKM effects for different time intervals are provided by the Institute for Employment Research (Lochner et al., 2023). In the base year 1990 (2000), we use the wage premia estimated based on the interval 1985-1992 (1993-1998).

trade-induced innovation and technical progress (e.g., Lilleeva and Trefler, 2010; Coelli et al., 2022).

While theory makes clear predictions on the rise in returns to skill due to trade, it is less clear what kind of skills or human capital are most relevant in this context. Our method is well-suited to run a horse race between alternative measures of skills and human capital. Following a vast empirical literature, we include a worker's formal education, differentiating between three levels of education: university degree, vocational degree, and low-skilled (i.e., less than a vocational degree). We also include the AKM worker fixed effect as a measure of workers' unobserved skills into the analysis, following (Dauth et al., 2021).

Increased exports might not only lead to a rise in the return to general skills or human capital (e..g., formal education), but also trigger an increase in the return to specific human capital and skills. This is because specific skills are costly to replace and a positive labor demand shock therefore generates rents for workers possessing specific skills as their bargaining position improves (e.g., Lazaer, 2009).

We include a measure of industry-specificity into the analysis, following recent work by Utar (2018) and Dauth et al. (2021). It is computed as the share of total employment of a 2-digit occupation which falls into a 3-digit industry. The idea of this measure is that the human capital of an occupation is more specific to a certain industry if its employment is very concentrated in that industry (Utar, 2018).

We also allow the effect of exports to vary depending on workers' initial firm, occupation, and industry tenure and experience (measured as total days in employment). Tenure is a potential measure of accumulated specific human capital. In addition, experience and tenure reflect a worker's labor market attachment. Existing research shows that the effects of industry-level trade shocks vary across workers with high versus low labor market attachment (e.g., Autor et al., 2014; Kovak and Morrow, 2025). More broadly, exporting might change the experience-wage profile of workers (Ma et al., forthcoming).

4.2.3 Job characteristics

The literature on knowledge-based hierarchies within a firm suggests that the wage effects of export-induced demand shocks might differ across the layers of hierarchy within a firm (e.g., Friedrich, 2022). We therefore include a dummy variable to differentiate between managers and non-managers. In addition, we include a dummy to differentiate between blue-collar and white-collar workers. First, this distinction

might also be relevant in terms of level of hierarchy. Second, the existing literature suggests that trade-induced increases in the skill premium might also show up as an increase in the relative wage for white-collar workers (e.g., Verhoogen, 2008). In addition, we include a variable capturing the routine-intensity of the initial 2-digit occupation of the worker.¹⁴ The routine content of a worker's job might be relevant in the context of trade-induced technological change (Lilleeva and Trefler, 2010), which might be routine-biased (Acemoglu and Autor, 2011).

4.2.4 Demographic characteristics

Finally, we allow the effects to vary across workers' age, gender, and nationality (German vs. foreign nationality). For example, differences in schedule flexibility between men and women might give rise to differential earnings effects of export-driven demand shocks (Bøler et al., 2018). Keller and Utar (2022) point to important interactions between gender and age in the context of trade shocks. Finally, workers with foreign nationality might be valuable for firms engaged in an international market (Parrotta et al., 2016).

¹⁴We compute this variable based on the BIBB/BAuA employment survey in 1992, in which workers indicate the tasks they perform on the job. The variable is computed as the share of routine tasks which are performed frequently among all tasks which are performed frequently. We aggregate the measure at the 2-digit industry-level. The following tasks are included: operating machines, manufacturing, storing, researching, programming, buying/selling, organizing, teaching, repairing, accommodating, caring, cleaning, protecting, writing, calculating, applying law, promoting, managing. The following tasks are coded as routine tasks: operating machines, manufacturing, storing, writing, and calculating. These data have recently been used to analyze the effects of AI on job tasks (Gathmann et al., 2024) and to study the wage premium to work-related stress (Nagler et al., 2024).

Variable	Explanation
Firm-level variables	
Firm size	Number of full-time employees in plant
AKM firm FE	Estimated firm wage premium
Human capital and skills	
Education	3 groups: university, vocational, <vocational< td=""></vocational<>
AKM worker FE	Estimated worker component of wage
Industry specificity	Share of occupational employment in industry
Firm tenure	Number of years in current plant
Industry tenure	Number of days in current 3-digit industry
Occupation tenure	Number of days in current 2-digit occupation
Experience	Total days in employment
Job characteristics	
Manager	Dummy, based on occupation classification
Blue collar	Dummy, based on occupation classification
Routine-intensity	Routine task share of 2-digit occupation
Demographic characteristics	
Age	Measured in years
Gender	Dummy
Foreign nationality	Dummy

Table 2: List of variables to estimate heterogeneous effects

Note: The table provides a list of all variables used to estimate heterogeneous effects with the generalized random forest. All of these variables are also used as control variables in the IV regressions and are measured in the worker's base year. We use AKM firm wage premia and worker effects estimated by Lochner et al. (2023). For the base year of 1990 (2000), we use AKM effects estimated using the years 1985-1992 (1993-1998).

5 The heterogeneous earnings effects of exports

5.1 Individualized effects

Figure 1 shows the distribution of the individualized effects of exports on earnings across all workers in the sample. Table 3 shows selected descriptive statistics on the distribution. The mean individualized effect of 0.69 is very close to the average effect estimated using a conventional IV estimation (0.72, see Table 1).

The results suggest that there is considerable heterogeneity in the effect of exports on earnings. A worker at the 90th percentile of the distribution has a coefficient of 1.25, whereas a worker at the 10th percentile has a coefficient of 0.14. To understand the magnitude of this difference, consider a one standard deviation increase of export exposure (23.80, see Table A.1). Taken at face value, for a worker at the 90th percentile of the effect, this implies a cumulative earnings gain of 29.75% of her base year earnings over 10 years (1.25 \times 23.80 \approx 29.75). For a worker at the 10th percentile of the effect, the corresponding cumulative earnings gain is 3.33% of her base year earnings $(0.14 * 23.80 \approx 3.33)$. For a worker with mean base year earnings (44.762.44 EUR, see Table A.1), these effects translate to earnings gains of 13,317 EUR and 1,491 EUR, respectively, over ten years. Appendix Figure A.2 shows that a wide distribution also arises when we directly take individual differences in base year earnings into account and compute the implied absolute earnings effects. The individualized effects and base year earnings are virtually uncorrelated, with a correlation coefficient of -0.06. In this section, we stick to the relative interpretation (i.e., relative to base year earnings), which is in the spirit of looking at changes in log earnings. However, when we analyze the contribution to earnings inequality in section 6, we also take differences in base year earnings across workers into account. Finally, Appendix Figure A.3 shows that we obtain a similar distribution across individuals when we run the GRF with an OLS instead of an IV regression.

Table 3: Descriptives of individualized effects

	Mean	SD	p10	p25	p50	p75	p90	N
Individualized effect $\widehat{\beta}(x_i)$	0.69	0.53	0.14	0.38	0.64	0.92	1.25	159,213

The table shows selected descriptive statistics of the individualized effect of exports on earnings $(\hat{\beta}(x_i))$.

Figure 1: Distribution of individualized effects of exports on earnings



Note: This Figure shows the distribution of the individualized effect of exports on earnings ($\hat{\beta}(x_i)$) across all individuals in the sample, based on the generalized random forest estimation. The dotted line reflects the sample mean of the individualized effect (0.68). Due to data disclosure regulations, the tails of the distribution are capped in order to ensure that every bin is based on at least 20 observations.

To what extent are the machine learning-based estimates $\hat{\beta}(x_i)$ predictive of 'actual' differences in the effect of exports on earnings? A potential concern is that the distribution shown in Figure 1 reflects just noise around the mean effect of exports on earnings. Testing the accuracy of the generalized random forest is difficult since we do not observe a 'ground truth' that we can use our estimates to test against. To provide evidence on the reliability of the GRF-based estimates, we group our sample into four quartiles, based on the individualized effect $\hat{\beta}(x_i)$. We then estimate conventional IV regressions of the effect of exports on earnings, separately for each quartile. If the machine learning-based method 'works', we would expect the estimates of the conventional IV regressions to follow the ranking that we obtain using the generalized random forest (i.e., highest estimate in the 4th quartile, lowest estimate in the 1st quartile).

Figure 2 shows that there is a close link between the ranking of effects implied by the generalized random forest and the corresponding conventional IV estimates. First, reassuringly, the quartile which is predicted by the GRF to have the largest (smallest) effect also has the largest (smallest) effect in a conventional IV estimation. The confidence intervals show that the differences between the first and the fourth quartile are statistically significant. In addition, the results show that the point estimates of the conventional IV estimation are very close to the mean coefficients based on the GRF. The GRF estimate is somewhat more conservative than the conventional IV estimate in the top quartile. A potential explanation is that the GRF, compared to the conventional IV estimation, is more robust against the influence of individual outlier observations. More specifically, the bootstrap approach and the honest approach of the GRF make sure that the estimate only captures heterogeneity in the estimates which arises consistently across different bootstrapped sub-samples. Overall, we conclude from Figure 2 that the GRF identifies actual differences in the effect of exports on earnings.

Finally, Figure 2 stresses an important quantitative point on the extent of heterogeneity in the effects of exports. According to the estimates, workers in the bottom quartile see little to no nominal earnings gains due to increased exports, whereas workers in the top quartile experience considerable positive effects.¹⁵ The point estimate of the conventional IV for the top quartile is around three times larger than the point estimate for the second quartile. A similar (but somewhat more extreme) conclusion on the existence of workers who see no nominal earnings gains is reached when looking at individualized t-values. The GRF yields an estimated variance of the individualized effect for each worker, and Appendix Figure A.4 shows the distribution of individualized t-statistics. It turns out that the point estimate is not statistically different from zero at the 5% (10%) significance level for 47.8% (38.7%) of workers in the sample. However, of course, we cannot rule of the existence of type-2-errors for individual observations. In what follows, we aim to understand which characteristics make workers likely to enter the top or bottom of the distribution of estimated effects, respectively.

¹⁵We stress the term 'nominal' because we cannot rule out trade-induced decreases in prices which translate into a real earnings increase also for workers in the bottom group.



Figure 2: Conventional IV, by quartile of machine learning-based estimate

Note: This Figure shows estimates of conventional IV-regressions of the effect of exports on earnings, separately by quartiles of the individualized effects ($\hat{\beta}(x_i)$, see Figure 1). The green circle reflects the mean individualized effect of the GRF within the respective quartile.

5.2 Who benefits most and who benefits least from exports?

In a first step, we simply compare the characteristics of workers who, according to the GRF, benefit most from exports to the characteristics of workers who benefit least. To this end, we compute differences in characteristics between the top and the bottom quartile or decile of the effect size distribution. To facilitate comparison across variables, we standardize all continuous variables to a mean of zero and a standard deviation of one.¹⁶

Figure 3 illustrates that workers in the top and bottom of the distribution differ in a variety of worker, job, and firm characteristics. Panel (a) shows that workers who benefit most from increased exports on average are employed in larger firms, are employed in occupations which are more specialized on the workers' initial industry, are younger, have less labor market experience, and have lower levels of firm, industry,

¹⁶See Appendix Table A.5 for the corresponding values of all variables.

and occupation tenure. The magnitude of the difference in characteristics is substantial. For example, firm size differs by 0.9 standard deviations between the best and the worst decile. Age differs by 0.7 standard deviations between the best and the worst decile.



Figure 3: Differences in characteristics between top and bottom groups

Note: This Table shows differences in characteristics between workers in the top and bottom quartile (decile) of the distribution of the individualized effects of exports on earnings ($\hat{\beta}(x_i)$, see Figure 1). To facilitate comparison between variables, the continuous variables in Panel (a) are standardized to have a mean of 0 and a standard deviation of 1.

Panel (b) of Figure 3 reveals a striking gender difference in the estimated effect of exports on earnings. The share of females in the top decile is 40 percentage points lower than in the bottom decile, suggesting that females are considerably underrepresented among the 'winners' of increased exports. This result suggests that exports might increase gender inequality among manufacturing workers. We discuss the potential underlying mechanism in more detail in the following sections. In addition, interestingly, the Figure shows that workers at the top of the distribution are more likely to be blue-collar workers. In terms of formal education, the Figure shows a slightly non-linear pattern. Low-skilled workers (i.e., those with less than a vocational degree) are underrepresented among the most positively affected group. When looking at the difference between the top and bottom decile, workers are less likely to have a university degree and more likely to have a vocational degree.

5.3 Accounting for the potential correlation between characteristics

The previous section provides information on which worker, job, and firm characteristics are highly predictive of large or small export-induced earnings gains. However, since these characteristics might be correlated, an unconditional comparison of characteristics of the top and bottom groups is not fully informative about the underlying mechanism. For example, higher earnings gains for workers in larger firms might be the result of 'pure firm effects' where workers in large and highly productive firms benefit from increased export opportunities due to rent sharing (e.g., Amiti and Davis, 2011; Egger and Kreickemeier, 2012). Alternatively, higher earnings gains in larger firms might be driven by a composition effect where firm size is correlated with worker or job characteristics that are associated with higher gains (such as measures of worker skills). Analogously, it is a relevant question to what extent the gender differences in export-induced earnings gains reflect differential sorting of men and women into jobs or firms which make them more or less likely to see large earnings gains. Note that these types of questions typically also arise in the case of conventional sample splits or interaction effects as soon as the interaction or splitting variable of interest is correlated with other potentially important variables. Appendix Table A.7 shows the correlation between all variables that we use to estimate heterogeneous effects.

To account for the potential correlation between variables, we report estimates of regressions of the individualized effects on worker, job, and firm characteristics in Figures 4 and 5. The Figures report bivariate regressions as well as regressions in which we control for the other respective variables.¹⁷ By conditioning on the other variables, we ask how the effect of exports on earnings changes if we change one variable, holding all other variables constant. We rank all variables by the size of the conditional coefficient.

¹⁷We make two exceptions. First, in the case of firm size, we do not control for the AKM firm wage premium (and vice versa), since these two variables are both used as a proxy for firm productivity. Second, in the case of the tenure variables, we do not control for the respective other tenure variables, due to the very high positive correlation and given that we consider all tenure variables as measures of a similar underlying mechanism.



Figure 4: Individualized effects and worker/job/firm characteristics (metric variables)

Note: This Figure shows results of regressions of the individualized effect $\hat{\beta}(x_i)$ on worker/job/firm characteristics. The blue circle shows estimates based on a bivariate regression. The red diamond shows estimates where we control for all other variables, with three exceptions. In case of the tenure variables, we do not control for the other tenure variables. In case of firm size and AKM firm wage premium, we do not control for the respective other variable. The same is true for education and AKM worker effect. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.

Figure 4 confirms the importance of firms in understanding the heterogeneous earnings effects of exports. Interestingly, the coefficient of firm size hardly changes after conditioning on the other variables, suggesting that workforce or job composition do not play a big role in explaining why workers in larger firms see larger earnings gains. The conditional coefficient of firm size is the largest among all continuous variables in Figure 4. The coefficient of industry specialization decreases only slightly after conditioning on other worker, job and firm characteristics. This reinforces the notion that workers in more specialized occupations benefit more from exports. The coefficient of the AKM worker effect becomes even larger after conditioning for worker, job, and firm variables. A potential interpretation is that unobserved skills rather than observed skills (e.g., education) become more valuable due to exports (e.g., Helpman et al., 2010). Interestingly, higher AKM firm effects are associated with lower earnings gains, both unconditionally and conditionally. Finally, and not surprisingly, the coefficients for age, experience, and the tenure variables decrease after conditioning, reflecting the correlation between these variables.



Figure 5: Individualized effect and worker/job/firm characteristics (dummy variables)

Note: This Figure shows results of regressions of the individualized effect $\hat{\beta}(x_i)$ on worker/job/firm characteristics. The blue circle shows estimates based on a bivariate regression. The red diamond shows estimates where we control for all other variables, with three exceptions. In case of the tenure variables, we do not control for the other tenure variables. In case of firm size and AKM firm wage premium, we do not control for the respective other variable. The same is true for education and AKM worker effect. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.

Figure 5 confirms the existence of an important gender difference in the earnings effects of exports. On average, the individualized coefficient between men and women differs by 0.27. The gender differences might reflect differences in skills or differential sorting into jobs and firms. When we control for all other worker, job, and firm characteristics, the gender difference shrinks from 0.27 to 0.20. This implies that sorting into different jobs and firms as well as differences in other characteristics cannot fully explain why men benefit more from exports than women. This finding also holds when we control for firm variables more non-parametrically and when we control for interactions between firm and job as well as firm and worker characteristics in Appendix Table A.6. In other words, the evidence suggests that men reap higher

earnings gains from exports than women even when starting out in the same or similar firm and job and conditional on having the same socio-demographic characteristics. We discuss potential reasons for this in section 5.6.

The estimates in Figure 5 also suggest that the earnings effects are non-linear in education as workers with medium levels of education (i.e., those with a vocational degree) exhibit more positive earnings effects than workers with a university degree and low-skilled workers. In line with these result, Figure 5 also shows that blue-collar workers (and non-managers) tend to have higher positive effects. However, the conditional regression for blue-collar status suggests that the positive effects for blue-collar workers are mostly not due to blue-collar status per se but rather due to a correlation with other characteristics.

In Appendix Figures A.5, A.6, A.7, and A.8 we provide estimates of linear probability models where the dependent variable is an indicator for being in the top/bottom quartile of individualized effects. The pattern is qualitatively very similar. For example, a one standard deviation increase in firm size is associated with a 10 percentage points increase of reaching the top quartile (see Figure A.7). In addition, for comparison to our machine learning-based method, Appendix Figures A.9 and A.10 provide conventional IV estimates with interaction effects with one variable at a time.

5.4 Non-linearity in firm size

In light of selection of the most productive firms into exporting, the heterogeneity by firm size might be highly non-linear. More specifically, one might expect the positive link between firm size and earnings to be driven mostly by workers in the largest firms. The data-driven approach of the GRF flexibly allows for non-linearities and is therefore well-suited to identify the shape of the heterogeneity by firm size.

Figure 6 shows that the heterogeneity by firm size is indeed highly non-linear. The Figure shows the results of a regression of the individualized effect on dummies for firm size decile, controlling for all other worker, job, and firm characteristics. It turns out that the heterogeneity by firm size is relatively flat below the 9th decile. In line with selection of the largest firms into exporting, positive earnings effects of exporting are substantially larger in the 9th and 10th decile. According to the estimates, the individualized effect among workers in the 10th decile of firm size is on average around 0.7 higher than the individualized effect of workers in the 5th decile (the

reference group). This difference corresponds to a switch from the median to the 90th percentile of the overall distribution of individualized effect (see Table 3).



Figure 6: Link between individualized effects and firm size deciles.

Note: The Figure shows estimates of a regression of individualized effects $\hat{\beta}(x_i)$ on firm size deciles, controlling for all other variables excpect for the AKM firm wage premium. Firm size decile 5 is the reference group. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.

5.5 What do we miss when focusing on conventional firm- and skill-related variables?

The previous sections point to an important role of firm heterogeneity and worker skills in understanding who benefits most. However, the results also show that other dimensions, including demographics, matter for the understanding of heterogeneous earnings effects. A natural question is to what extent do we 'miss' part of the unequal earnings effects of exports when focusing on conventionally used firm- and skill-related variables alone.

In Table 4, we provide evidence that a focus on conventional skill- and firm-related variables misses a substantial share of the total heterogeneity detected by the GRF. The table shows the R^2 and adjusted R^2 of a regression of the individualized effects on firm size, AKM worker effects, AKM firm wage premia, and formal education. Even in column (5) where we allow the effects of firm size to be non-linear and where

we account for all possible interactions between the respective variables, we cannot explain more than 40% of the total variation of export-induced earnings effects.

A potential concern is that the remaining 60% of the variation reflect noise generated by the GRF rather than factors like demographics or occupational characteristics. However, in Figure 7, we show that the remaining 60% of the variation in GRF-based estimates do have explanatory power of 'actual' differences in export-induced earnings effects. In this graph, we show estimates of conventional IV regressions, based on quartiles of the GRF-based individualized effects, analogously to Figure 2. However, in contrast to Figure 2, we first residualize the individualized effects from the skill- and firm-related variables and their interactions in column (5) of Table 4. Therefore, by construction, the residual captures variation which is uncorrelated to the variables mentioned in column (5) of Table 4. As expected, the difference in estimates across quartiles becomes smaller, but in particular the difference between the first and the fourth quartile is still sizable and highly statistically significant.¹⁸ In other words, there are systematic differences in export-induced earnings effects which can be explained by variables other than conventionally used firm- and skill-related variables, including for example demographic variables such as gender and age, or occupational characteristics.

	Dep. var.: Individualized effect					
	(1)	(2)	(3)	(4)	(5)	
R^2	0.06	0.13	0.20	0.32	0.39	
Adj. R ²	0.06	0.13	0.20	0.32	0.39	
Firm size (metric)	Х	Х				
Firm size (10 dummies)			Х	Х	Х	
AKM worker FE (metric)	Х	Х	Х			
AKM worker FE (10 dummies)				Х	Х	
Education		Х	Х	Х	Х	
AKM firm FE (metric)		Х	Х	Х	Х	
Interaction effects					Х	

Table 4: Explanatory power of conventional firm- and skill-related variables

Note: The Table shows the R^2 and adjusted R^2 of a regression of the individualized effects $\hat{\beta}(x_i)$ on firm- and skill-related variables. Metric variables are standardized to have a mean of zero and standard deviation of one. The last column includes all possible interactions between the variables.

¹⁸In Figure 2, the difference between the estimate in the fourth and first quartile is roughly 1.8. In Figure 7, it is roughly 1.0.

Figure 7: Conventional IV by quartile of GRF-based estimate, residualized from firmand skill-related variables



Note: This Figure shows estimates of conventional IV-regressions of the effect of exports on earnings, separately by quartiles of the individualized effects ($\hat{\beta}(x_i)$, residualized from the skill- and firm-related variables and their interactions in column (5) of Table 4.

5.6 Interpretation and discussion of heterogeneities

The previous results suggest that the heterogeneity in the effects of exports on earnings is multidimensional and cannot be summarized by one or two variables alone. In what follows, we put the results into perspective and discuss potential underlying mechanisms.

5.6.1 Firm size and AKM worker effect

In line with the vast theoretical and empirical literature in trade, we find that workers initially employed in larger firms and workers with higher skill levels (reflected by higher AKM worker effects) reap larger gains from increased industry-level exports. For example, a prominent class of models predicts that the largest and most productive firms select into exporting and workers employed by these firms benefit in terms of higher wages due to rent sharing (e.g., Egger and Kreickemeier, 2009, 2012; Amiti and Davis, 2011). Importantly, the link between positive earnings effects and firm

size holds also after accounting for a potential correlation between firm size and worker and job characteristics which might be linked to higher earnings gains. Our results are therefore related to a large literature in labor economics which points to the existence of sizable firm wage premia (conditional on worker composition) (e.g., Card et al., 2013; Song et al., 2019) and with recent evidence finding that exports lead to an increase in these wage premia (Frías et al., 2024). The positive between AKM worker effects and exports is in line with a large literature in economics pointing to an increase in the skill premium due to trade (e.g., Burstein and Vogel, 2017; Dauth et al., 2021).¹⁹

Interestingly, the AKM worker effect, which captures unobserved skills in addition to observed skills, turns out to be more relevant than formal education for understanding who benefits most from exports. In other words, while export expansion goes along with an increase in the demand for skills, it might not necessarily go along with a strong increase in the demand for university-educated workers. Instead, the rise in skill demand might be more subtle, reflecting an increased need for workers with high levels of unobserved skills. This finding may be of interest from the perspective of policymakers.²⁰

The previous section suggests that these variables explain at most 40% of the total heterogeneity of earnings effects due to exports. In what follows, we discuss a variety of factors which turn out to be important for understanding the unequal earnings effects of exports, but have received much less attention in the existing literature.

5.6.2 Gender

On the more surprising end of the results, we find that men see larger earnings gains from exports than women. What explains this gender difference? First, it might be driven by differential sorting across firms and jobs or differences in skills. For example, women might be less likely to be employed by large firms. However, the conditional coefficients in Figure 5 and Table A.6 suggest that only a small share of

¹⁹Higher AKM firm effects are associated with lower earnings gains, in contrast to the notion that they are a proxy for firm productivity. A potential explanation is that AKM firm effects also reflect other factors unrelated to productivity, such as strategic wage-posting behavior of firms (Burdett and Mortensen, 1998), compensating wage differentials for disamenities (Sorkin, 2018), or institutional factors such as collective bargaining agreements and the existence of works councils (Hirsch and Müller, 2020).

²⁰While analyzing the exact nature of these unobserved skills is beyond the scope of this paper, the literature in labor economics for example point to increased returns for social skills (Deming, 2017).

the gender difference seems to be driven by differential sorting across jobs and firms or by differences in other characteristics such as skills.

Second, the literature points to important gender differences in job mobility and job search behavior (Le Barbanchon et al., 2021; Cortes et al., 2023). Men might be more able or might be more willing to take advantage of higher earnings possibilities by switching to exporting firms. In Appendix Table A.4, we estimate conventional IV regressions, splitting the sample between men and women. We decompose overall earnings into earnings at the initial employer and earnings at a different employer. It turns out that the gender difference in export-induced earnings effects arises exclusively at the initial employer, suggesting that differences in job mobility are not a main driver either.

Gender differences in earnings effects at the initial employer might be driven by differences in bargaining behavior between men and women (Caldwell et al., 2024). Men could therefore be able to reap a larger share of the rents accruing from increased export opportunities. Alternatively, gender differences in earnings outcomes might arise due to a higher schedule flexibility of men compared to women. Increased exports, especially to different time zones, might require more temporal flexibility among employees, translating into gender differences in earnings outcomes (Bøler et al., 2018).

5.6.3 Age, experience, tenure

The results point to a substantial heterogeneity of the effects in terms of age, experience, and firm/occupation/industry tenure. Similar to gender, these variables have received little attention in the existing research on the heterogeneous earnings effects of exports, but turn out to be quantitatively important. We consider all of these heterogeneities as capturing the same or a very similar underlying mechanism.

First, low age, experience, and tenure are proxies for low labor market attachment of workers (Kovak and Morrow, 2025). These workers tend to be more vulnerable to negative shocks, e.g., from import competition (Autor et al., 2014). The presence of tenure-based layoff rule in Germany makes low-tenure workers more likely to be laid off due to negative shocks. Therefore, these workers might benefit more from a positive shock due to exports, especially in light of declining manufacturing employment and associated costs for laid-off workers (Helm et al., 2024). Second, a potential explanation is that young workers accumulate more human capital at exporting firms and industries than at non-exporting firms and industries (Ma et al., forthcoming). Third, in light of steeper wage growth at younger ages, the possibilities to negotiate higher wages and participate in the rents from exports might be higher for younger workers.

5.6.4 Industry specialization

Our results also show that workers employed in occupations which are more specialized on their industry see larger earnings gains due to exports. Our measure can be interpreted as a measure of industry-specific human capital and has been employed prominently by Utar (2018) who shows that workers in more specialized occupations incur larger earnings losses from import competition. The positive interaction between industry specialization and exports is in line with findings by Dauth et al. (2021). One way to interpret our results is that specific human capital could be more important or at least as important as general human capital in explaining the heterogeneous effect of exports. A potential mechanism behind the strong positive interaction with exports is that workers who are more specialized on their industry become particularly valuable in light of a positive demand shock (Lazaer, 2009). This is true especially in a world with search frictions where replacing skilled workers is costly and time-consuming.

6 Exports and inequality

In a last step, we use our granular estimates to provide back-of-the-envelope calculations on the link between exports and overall earnings inequality. Based on the heterogeneous effects found in the previous sections, it is not clear to what extent the export shock contributes to overall inequality. A look at the correlations of the variables of interest with base year earnings in Appendix Table A.7 suggests that some factors could work towards an increase in inequality (e.g., firm size, AKM worker effect, gender), while others could work towards a decrease of inequality (e.g., age, experience, tenure, industry specilization). To shed light on the link between the export shock, we perform different exercises, building on the individualized effects as well as the actual export exposure on workers.

For each worker in our sample, we compute three alternative versions of an estimated effect of exports on cumulative earnings:

$$\widehat{\beta}(x_i) \times \overline{Exp_j}$$
 (3)

$$\widehat{\beta}(x_i) \times \overline{Exp_j} \tag{4}$$

$$\widehat{\beta}(x_i) \times Exp_j$$
 (5)

where *i* denotes a worker and *j* denotes the industry in which worker *i* is initially employed.²¹ $\overline{\hat{\beta}(x_i)}$ denotes the sample average of individualized effects of exports (0.68, see Table 3) and $\overline{Exp_j}$ denotes the sample mean of increased export exposure (26.97, see Table A.1).

Expression 3 employs the sample mean of both, the individualized effects and the export exposure measures. The corresponding predicted effect on cumulative earnings is therefore, by construction, the same across all individuals and can be considered as an estimate of the average effect on cumulative earnings. Expression 4 allows the beta coefficient to be person-specific and holds the increase in export exposure constant across workers. Finally, expression 5 also allows the increase in export exposure to vary across individuals, based on their initial industry.

In Figure 8, we plot the average estimated effect within percentiles of the base year earnings distribution separately for all three measures. The green crosses show that the average estimated effect amounts to roughly 19% of workers' base year earnings over ten years. When we allow the effects of a one-unit increase in export exposure to be individual-specific (blue circles), it turns out that workers between the 60th and the 80th percentile are most positively affected on average. The estimated effect increases monotonically between the 20th and the 70th percentile and decreases above the 80th percentile. Finally, when we also allow the export exposure to vary across individual, based on their industry affiliation (red diamonds), the pattern is qualitatively similar but becomes more pronounced. This is because workers with base year earnings above the median tend to be employed in industries with a larger increase in export exposure, on average.

Overall, Figure 8 suggests that the inequality-increasing effect of exports is much less pronounced than expected based on the existing literature. Looking at the average estimate by percentile, it turns out that no percentile exhibits a mean estimated effect below 10%, suggesting that even workers at the bottom of the initial earnings

²¹We pool both time windows and drop the subscript t for simplicity.

distribution on average gain from exports. This result reflects the fact that some factors which predict large export-induced earnings gains are in fact negatively correlated with base year earnings (e.g., age, tenure, experience, industry specialization). Nevertheless, workers in the upper part of the earnings distribution benefit more on average. The bootstrap estimates in Panel (A) of Table 5 suggest that the differences, for example between the 75th and the 25th percentile, are also statistically significant.



Figure 8: Earnings effects by initial earnings percentile

Note: This Figure plots the mean earnings effect, by percentile of initial earnings, for different scenarios. The green crosses show the effects, holding the coefficients constant at the sample mean (0.68) and holding export exposure constant at the sample mean (26.97). The blue circles reflect the effect using individualized betas, holding export exposure constant at the sample mean. The red diamonds show the effects using individualized betas and using actual export exposure, based on the individual's initial industry. The lines reflect a local polynomial smoothing.

Based on the differences across percentiles, it is not clear what the export shock implies for measures of aggregate inequality like the Gini coefficient. In a next step, we provide a back-of-the-envelope calculation on the consequences for the Gini coefficient. To this end, we have to translate our estimated earnings effect, which is measured as percent of base year earnings, into an absolute value. We do this by multiplying our estimate by base year earnings:

$$\frac{\hat{\beta}(x_i) \times Exp_j}{100} \times BaseEarnings_i \times \frac{1}{10}$$
(6)

Note that we divide by ten in order to obtain the average annual earnings change over ten years. We compare the Gini of base year earnings to the Gini of base year earnings plus the annual earnings change from equation 6.

	(A) Earnii	ngs effects	by initial earnings percentile
	estimate	(SE)	[95%-CI]
p10	13.16***	(0.64)	[11.90 ; 14.42]
p25	13.54***	(0.50)	[12.56 ; 14.53]
p50	20.65***	(0.57)	[19.52 ; 21.78]
p75	26.67***	(0.67)	[25.36 ; 27.98]
p90	14.68***	(0.45)	[13.79 ; 15.56]
		(B) (Gini coefficient
	estimate	(SE)	[95%-CI]
Base year	0.2785***	(0.0010)	[0.2765 ; 0.2806]
Base year + exports	0.2794***	(0.0010)	[0.2774 ; 0.2814]
Difference	0.0008***	(0.0000)	[0.0008 ; 0.0009]

Table 5: Exports and earnings inequality

Note: The table shows the estimated earnings effects (individual coefficients and individual exposure) for selected percentiles of initial earnings (Panel (A)) and a back-of-the-envelope calculation for the effects on the Gini coefficient (Panel (B)). Standard errors and confidence intervals are computed using 200 bootstrap replications. * p<0.10, ** p<0.05, *** p<0.01

In Panel (B) of Table 5 we provide evidence that the export shock goes along with a very modest increase in the Gini coefficient of earnings. The Gini of base year earnings is equal to 0.2785. After adding the estimated cumulative earnings change, the Gini amounts to 0.2794. While the difference is statistically significant based on the bootstrap standard errors, it is quantitatively small and, depending on the prior, smaller than expected. This finding is of interest for policymakers who worry that unequal earnings effects of exports translate into a sharp increase in inequality between workers at the top and workers at the bottom of the earnings distribution. Our results suggest that most of the differences in the effects occur between workers with similar earnings rather than between high- and low-earnings workers.

7 Conclusion

How do labor markets and workers adjust to international trade and what are the distributional consequences? This question is of large importance in a highly globalized world. From the perspective of policymakers, it is particularly important to understand what characteristics make workers likely to reap large gains from international trade and what characteristics make workers vulnerable to incur (relative) earnings losses. The goal of this paper is to use novel machine learning methods to better understand the full heterogeneity of the effects of trade along a variety of dimensions and to study the extent to which positive export shocks influence existing earnings inequalities. To do so, we focus on a major episode of globalization from the perspective of the German economy, namely the substantial increase in trade with China and Eastern European countries starting in the 1990s.

The finding that the earnings effects of exports are vastly heterogeneous is not surprising in light of existing research. However, a main takeaway from our analysis is that the heterogeneity of the earnings effects is multidimensional and therefore cannot be summarized by one or two factors alone. The important role of firm size and workers' unobserved skills is largely consistent with a large body of research on the distributional effects of trade. Yet, our results point to further important dimensions of heterogeneity which have received much less attention in prior research, including demographic variables like age and gender. In addition, our results point to an important role of a workers' skill specialization. Workers employed in occupations that are more concentrated in the respective industry reap larger gains from exports. In contrast, workers with a university degree do not necessarily fare better than those without a university degree in the context of export shocks. From a policy perspective, this result is interesting as it implies that a worker's skill specialization might be more important than her degree per se in determining the extent to which she benefits from exports.

Our results also suggest that exports are not a major driver of inequality across initial earnings percentiles. In fact, most of the heterogeneity of the earnings effects occur within groups with similar initial earnings rather than between earnings groups. This result is valuable from the perspective of policymakers who might worry about substantially widening gaps in earnings across the earnings distribution.

References

- ABOWD, J., F. KRAMARZ, AND D. MARGOLIS (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–333.
- 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, Elsevier, vol. 4B, 1043–1171.
- ADAO, R., P. CARRILLO, A. COSTINOT, D. DONALDSON, AND D. POMERANZ (2022): "Imports, Exports, and Earnings Inequality: Measures of Exposure and Estimates of Incidence," *Quarterly Journal of Economics*, 137, 1553–1614.
- AMITI, M. AND D. R. DAVIS (2011): "Trade, Firms, and Wages: Theory and Evidence," *The Review of Economic Studies*, 79, 1–36.
- ARNI, P., P. EGGER, K. ERHARDT, M. GUBLER, AND P. SAURE (2024): "Heterogeneous Impacts of Trade Shocks on Workers," *IZA Discussion Paper No. 16895*.
- ATHEY, S. AND G. IMBENS (2016): "Recursive partitioning for heterogeneous causal effects," *Proceedings of the National Academy of Sciences*, 113, 7353–7360.
- ATHEY, S., L. K. SIMON, O. N. SKANS, J. VIKSTROM, AND Y. YAKYMOVYCH (2024): "The Heterogeneous Earnings Impact of Job Loss Across Workers, Establishments, and Markets," *mimeo*.
- ATHEY, S., J. TIBSHIRANI, AND S. WAGER (2019): "Generalized Random Forests," *The Annals of Statistics*, 47, 1148–1178.
- AUTOR, D. H., D. DORN, AND G. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–2168.
- AUTOR, D. H., D. DORN, G. HANSON, AND J. SONG (2014): "Trade Adjustment: Worker Level Evidence," *Quarterly Journal of Economics*, 129, 1799–1860.

^{——— (2019): &}quot;When Work Disappears: Manufacturing Decline and the Falling Marriage-Market Value of Young Men," *American Economic Reviw: Insights*, 1, 161–178.

- BAIER, S. AND N. R. REGMI (2023): "Using Machine Learning to Capture Heterogeneity in Trade Agreements," *Open Economies Review*, 34, 863–894.
- BALSVIK, R., S. JENSEN, AND K. G. SALVANES (2014): "Made in China, sold in Norway: Local labor market effects of an import shock," *Journal of Public Economics*, 127, 137–144.
- BAUMGARTEN, D. (2013): "Exporters and the rise in wage inequality: Evidence from German linked employer–employee data," *Journal of International Economics*, 90, 201–217.
- BERNARD, A., B. JENSEN, S. R. REDDING, AND P. K. SCHOTT (2012): "The Empirics of Firm Heterogeneity and International Trade," *Annual Review of Economics*, 4, 283–313.
- BESSEN, J., M. GOOS, A. SALOMONS, AND W. VAN DEN BERGE (2025): "What happends to workers at firms that automate?" *The Review of Economics and Statistics*, 107, 125–141.
- BØLER, E. A., B. JAVORCIK, AND K. H. ULLVEIT-MOE (2018): "Working across time zones: Exporters and the gender wage gap," *Journal of International Economics*, 111, 122–133.
- Вомваядіні, М., G. Orefice, and M. Tito (2019): "Does Exporting Improve Matching? Evidence from French Employer-Employee Data," *Journal of International Economics*, 229–241.
- BORUSYAK, K. AND X. JARAVEL (2024): "Are trade wars class wars? The importance of trade-induced horizontal inequality," *Journal of International Economics*, 150, 103935.
- BREIMAN, L. (2001): "Random forests," Machine learning, 45, 5–32.
- BREINLICH, H., V. CORRADI, N. ROCHA, MICHELE RUTA, J. S. SILVA, AND T. ZYLKIN (2022): "Machine Learning in International Trade Research - Evaluating the Impact of Trade Agreements," *mimeo*.
- BRITTO, D. G., P. PINOTTI, AND B. SAMPAIO (2022): "The effect of job loss and unemployment insurance on crime in Brazil," *Econometrica*, 90, 1393–1423.
- BURDETT, K. AND D. T. MORTENSEN (1998): "Wage Differentials, Employer Size and Unemployment," *International Economic Review*, 39, 257–273.
- BURSTEIN, A. AND J. VOGEL (2017): "International Trade, Technology, and the Skill Premium," *Journal of Political Economy*, 125, 1356–1412.

- CALDWELL, S., I. HAEGELE, AND J. HEINING (2024): "Bargaining and Inequality in the Labor Market," *mimeo*.
- CARD, D., P. KLINE, AND J. HEINING (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 128, 967–1015.
- CARLANA, M., E. L. FERRARA, AND P. PINOTTI (2022): "Goals and Gaps: Educational Carreers of Immigrant Children," *Econometrica*, 90, 1–29.
- COCKX, B., MICHAEL LECHNER, AND J. BOLLENS (2023): "Priority to unemployed immigrants? A causal machine learning evaluation of training in Belgium," *Labour Economics*, 80, 102306.
- COELLI, F., A. MOXNES, AND K. H. ULLTVEIT-MOE (2022): "Better, Faster, Stronger: Global Innovation and Trade Liberalization," *The Review of Economics and Statistics*, 54, 205–216.
- CORTES, P., J. PAN, L. PILOSSOPH, E. REUBEN, AND B. ZAFAR (2023): "Gender Differences in Job Search and the Earnings Gap: Evidence from the Field and Lab," *Quarterly Journal of Economics*, 128, 2069–2126.
- COSAR, K. A., N. GUNER, AND J. TYBOUT (2016): "Firm Dynamics, Job Turnover, and Wage Distributions in an Open Economy," *American Economic Review*, 106, 625–663.
- CRAVINO, J. AND S. SOTELO (2019): "Trade-Induced Structural Change and the Skill Premium," *American Economic Journal: Macroeconomics*, 11, 289–326.
- 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, 1–33.

——— (2021): "Adjusting to Globalization in Germany," *Journal of Labor Economics*, 39, 263–302.

- DAVIDSON, C., F. HEYMAN, S. MATUSZ, F. SJÖHOLM, AND S. C. ZHU (2014): "Globalization and imperfect labor market sorting," *Journal of International Economics*, 94, 177–194.
- DAVIS, J. M. V. AND S. B. HELLER (2017): "Using Causal Forests to Predict Treatment Heterogeneity: An Application to Summer Jobs," *American Economic Review: Papers and Proceedings*, 107, 546–550.

- (2020): "Rethinking the Benefits of Youth Employment Programs: The Heterogeneous Effects of Summer Jobs," *The Review of Economics and Statistics*, 102, 664–677.
- DEMING, D. (2017): "The Growing Importance of Social Skills in the Labor Market," *The Quarterly Journal of Economics*, 132, 1593–1640.
- DIX-CARNEIRO, R. (2014): "Trade Liberalization and Labor Market Dynamics," *Econometrica*, 82, 825–885.
- DIX-CARNEIRO, R. AND B. K. KOVAK (2017): "Trade Liberalization and Regional Dynamics," *American Economic Review*, 107, 2908–2946.

——— (2019): "Margins of labor market adjustment to trade," Journal of International Economics, 117, 125–142.

- DIX-CARNEIRO, R., J. PESSOA, R. REYES-HEROLES, AND S. TRAIBERMAN (forthcoming): "Globalization, Trade Imbalances, and Labor Market Adjustment," *Quarterly Journal of Economics*.
- EGGER, H., P. EGGER, U. KREICKEMEIER, AND C. MOSER (2020): "The Exporter Wage Premium When Firms and Workers Are Heterogenous," *European Economic Review*, 130, 1–27.
- EGGER, H. AND U. KREICKEMEIER (2009): "Firm Heterogeneity and the Labor Market Effects of Trade Liberalization," *International Economic Review*, 50, 187–216.

——— (2012): "Fairness, trade, and inequality," *Journal of International Economics*, 86, 184–196.

- FRÍAS, J. A., D. S. KAPLAN, E. VERHOOGEN, AND D. ALFARO-SERRANO (2024): "Exports and Wage Premiums: Evidence from Mexican Employer-Employee Data," *The Review* of Economics and Statistics, 106, 305–321.
- FRIEDRICH, B. (2022): "Trade shocks, firm hierarchies, and wage inequality," *The Review* of Economics and Statistics, 104, 652–667.
- GARIN, A. AND F. SILVERIO (2023): "How Responsive Are Wages to Firm-Specific Changes in Labour Demand? Evidence from Idiosyncratic Export Demand Shocks," *Review of Economic Studies*, 00, 1–40.

- GATHMANN, C., F. GRIMM, AND E. WINKLER (2024): "AI, Task Changes in Jobs, and Worker Reallocation," CESifo Discussion Paper No. 11585.
- GULYAS, A. AND K. PYTKA (2022): "Understanding the Sources of Earnings Losses After Job Displacement: A Machine-Learning Approach," *mimeo*.
- HAKOBYAN, S. AND J. MCLAREN (2016): "Looking for Local Labor Market Effects of NAFTA," *The Review of Econmics and Statistics*, 98, 728–741.
- HELM, I. (2020): "National Industry Trade Shocks, Local Labor Markets, and Agglomeration Spillovers," *Review of Economic Studies*, 87, 1399–1431.
- HELM, I., A. KUEGLER, AND UTASCHOENBERG (2024): "Displacement Effects in Manufacturing and Structural Change," *mimeo*.
- HELPMAN, E., O. ITSKHOKI, M.-A. MUENDLER, AND S. REDDING (2017): "Trade and Inequality: From Theory to Estimation," *The Review of Economic Studies*, 84, 357–405.
- HELPMAN, E., O. ITSKHOKI, AND S. REDDING (2010): "Inequality and Unemployment in a Global Economy," *Econometrica*, 78, 1239–1283.
- HIRSCH, B. AND S. MÜLLER (2020): "Firm wage premia, industrial relations, and rent sharing in Germany," *ILR Review*, 73, 1119–1146.
- HUBER, K. AND E. WINKLER (2019): "All you need is love? Trade shocks, inequality, and risk sharing between partners," *European Economic Review*, 111, 305–335.
- HUMMELS, D., R. JOERGENSEN, J. MUNCH, AND C. XIANG (2014): "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, 104, 1597–1629.
- IRRARAZABAL, A., A. MOXNES, AND K. H. ULLVEIT-MOE (2013): "Heterogeneous Firms or Heterogeneous Workers? Implications for Exporter Premiums and the Gains from Trade," *The Review of Economics and Statistics*, 95, 839–849.
- JAMES, G., D. WITTEN, T. HASTIE, R. TIBSHIRANI, G. JAMES, D. WITTEN, T. HASTIE, AND R. TIBSHIRANI (2021): "Statistical learning," An introduction to statistical learning: with applications in R, 15–57.

- JUHN, C., G. UJHELYI, AND C. VILLEGAS-SANCHEZ (2014): "Men, women, and machines: How trade impacts gender inequality," *Journal of Development Economics*, 106, 179–193.
- KELLER, W. AND H. UTAR (2022): "Globalization, Gender, and the Family," *The Review* of *Economic Studies*, 89, 3381–3409.
- KOVAK, B. AND J. MORROW (2023): "Distributional Impacts of the Canada-US Free Trade Agreement," *AEA Papers and Proceedings*, 113, 585–589.
- ——— (2025): "The Long-Run Labour Market Effects of the Canada-U.S. Free Trade Agreement," *Review of Economic Studies, forthcoming*.
- KRISHNA, P., J. P. POOLE, AND M. Z. SENSES (2014): "Wage Effects of Trade Reform with Endogenous Worker Mobility," *Journal of International Economics*, 93, 239–252.
- LAZAER, E. (2009): "Firm-Specific Human Capital: A Skill-Weights Approach," *Journal* of Political Economy, 117, 914–940.
- LE BARBANCHON, T., R. RATHELOT, AND A. ROULET (2021): "Gender Differences in Job Search: Trading off Commute against Wage," *Quarterly Journal of Economics*, 136, 381–426.
- LILLEEVA, A. AND D. TREFLER (2010): "Improved Access to Foreign Markets Raises Plant-Level Productivity...For Some Plants," *Quarterly Journal of Economics*, 151, 1051–1099.
- LOCHNER, B., S. SETH, AND S. WOLTER (2023): "AKM effects for German labour market data 1985-2021," FDZ Methodenreport 01/2023.
- MA, X., M.-A. MUENDLER, AND A. NAKAB (forthcoming): "Exporting, Wage Profiles, and Human Capital: Evidence from Brazil," *Review of Economics and Statistics*.
- MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.
- MENEZES-FILHO, N. A. AND M.-A. MUENDLER (2011): "Labor Reallocation in Response to Trade Reform," *mimeo*.
- MUNCH, J. R. AND J. R. SKAKSEN (2008): "Human capital and wages in exporting firms," *Journal of International Economics*, 75, 363–372.

- NAGLER, M., J. RINCKE, AND E. WINKLER (2024): "High-Pressure, High-Paying Jobs?" *The Review of Economics and Statistics*, 1–45.
- NILSSON HAKKALA, K. AND K. HUTTUNEN (2018): "Worker-level Consequences of Firm-level Import Shocks: Evidence from Finland," *mimeo*.
- PARROTTA, P., D. POZZOLI, AND D. SALA (2016): "Ethnic diversity and firms'export behavior," *European Economic Review*, 89, 248–263.
- SAMPSON, T. (2014): "Selection into Trade and Wage Inequality," *American Economic Journal: Microeconomics*, 6, 157–202.
- SAURE, P. AND H. ZOABI (2014): "International trade, the gender wage gap and female labor force participation," *Journal of Development Economics*, 111, 17–33.
- SCHANK, T., C. SCHNABEL, AND J. WAGNER (2007): "Do exporters really pay higher wages? First evidence from German linked employer–employee data," *Journal of International Economics*, 72, 52–74.
- Song, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2019): "Firming Up Inequality," *Quarterly Journal of Economics*, 134, 1–50.
- SORKIN, I. (2018): "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, 133, 1331–1393.
- SYVERSON, C. (2011): "What Determines Productivity?" *Journal of Economic Literature*, 49, 326–365.
- TRAIBERMAN, S. (2019): "Occupations and Import Competition: Evidence from Denmark," *American Economic Review*, 109, 4260–4301.
- UTAR, H. (2018): "Workers Beneath the Floodgates: Low-wage Import Competition and Workers' Adjustment," *Review of Economics and Statistics*, 100, 621–647.
- VERHOOGEN, E. (2008): "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *Quarterly Journal of Economics*, 489–530.
- WAGER, S. AND S. ATHEY (2018): "Estimation and inference of heterogeneous treatment effects using random forests," *Journal of the American Statistical Association*, 113, 1228–1242.

A Appendix



Figure A.1: The evolution of trade with China and Eastern Europe over time

Note: Panel (a) shows annual exports to and imports from China and Eastern Europe, measured in billions of 2010-euros. Panel (b) shows the share of exports to and imports from China and Eastern Europe in total German exports and imports. The vertical solid line indicates the start of our period of main analysis. Following Dauth et al. (2014, 2021) define 'Eastern Europe' as the following set of countries: Bulgaria, Czech Republic, Hungary, Poland, Romania, Slovakia, Slovenia, the former USSR, and its successor states the Russian Federation, Belarus, Estonia, Latvia, Lithuania, Moldova, Ukraine, Azerbaijan, Georgia, Kazakhstan, Kyrgyzstan, Tajikistan, Turkmenistan, and Uzbekistan.





Cumulative earnings effects over 10 years

Note: This Figure shows the distribution of implied earnings effects over ten years of a one standard deviation increase in export exposure (23.80). The values are obtained by the following formula: $\frac{\hat{\beta}(x_i) \times 23.80}{100} \times BaseEarnings_i$. The dotted line reflects the sample mean. Due to data disclosure regulations, the tails of the distribution are capped in order to ensure that every bin is based on at least 20 observations.



Figure A.3: Distribution of individualized effects: OLS Forest

Note: This Figure shows the distribution of the individualized effect of exports on earnings ($\hat{\beta}(x_i)$) across all individuals in the sample, based on an OLS Forest. The dotted line reflects the sample mean of the individualized effect. Due to data disclosure regulations, the tails of the distribution are capped in order to ensure that every bin is based on at least 20 observations.



Figure A.4: Distribution of individualized t-statistics

Note: This Figure shows the distribution of the individualized t-statistics which are computed by dividing the individualized effect ($\hat{\beta}(x_i)$) by the individualized standard error. The dashed line reflects the threshold of 1,96. Due to data disclosure regulations, the tails of the distribution are capped in order to ensure that every bin is based on at least 20 observations.



Figure A.5: Probability of belonging to the bottom quartile and worker/job/firm characteristics (metric variables)

Note: This Figure shows results of regressions of an indicator of belonging to the top quartile based on the ranked individualized effect $\hat{\beta}(x_i)$ on worker/job/firm characteristics. The blue circle shows estimates based on a bivariate regression. The red diamond shows estimates where we control for all other variables, with three exceptions. In case of the tenure variables, we do not control for the other tenure variables. In case of firm size and AKM firm wage premium, we do not control for the respective other variable. The same is true for education and AKM worker effect. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.

Figure A.6: Probability of belonging to the bottom quartile and worker/job/firm characteristics (dummy variables)



Note: This Figure shows results of regressions of an indicator of belonging to the top quartile based on the ranked individualized effect $\hat{\beta}(x_i)$ on worker/job/firm characteristics. The blue circle shows estimates based on a bivariate regression. The red diamond shows estimates where we control for all other variables, with three exceptions. In case of the tenure variables, we do not control for the other tenure variables. In case of firm size and AKM firm wage premium, we do not control for the respective other variable. The same is true for education and AKM worker effect. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.



Figure A.7: Probability of belonging to the top quartile and worker/job/firm characteristics (metric variables)

Note: This Figure shows results of regressions of an indicator of belonging to the bottom quartile based on the ranked individualized effect $\hat{\beta}(x_i)$ on worker/job/firm characteristics. The blue circle shows estimates based on a bivariate regression. The red diamond shows estimates where we control for all other variables, with three exceptions. In case of the tenure variables, we do not control for the other tenure variables. In case of firm size and AKM firm wage premium, we do not control for the respective other variable. The same is true for education and AKM worker effect. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.





Note: This Figure shows results of regressions of an indicator of belonging to the bottom quartile based on the ranked individualized effect $\hat{\beta}(x_i)$ on worker/job/firm characteristics. The blue circle shows estimates based on a bivariate regression. The red diamond shows estimates where we control for all other variables, with three exceptions. In case of the tenure variables, we do not control for the other tenure variables. In case of firm size and AKM firm wage premium, we do not control for the respective other variable. The same is true for education and AKM worker effect. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.

Figure A.9: Interaction effects of worker/job/firm characteristics and export exposure with conventional IV regressions (metric variables)



Note: This Figure shows the interaction effects of the export exposure with standardized worker/job/firm characteristics, estimated with conventional IV regressions. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.

Figure A.10: Interaction effects of worker/job/firm characteristics and export exposure with conventional IV regressions (dummy variables)



Note: This Figure shows the interaction effects of the export exposure with worker/job/firm characteristics, estimated with conventional IV regressions. The lines reflect 95% confidence intervals. Standard errors are clustered at the 3-digit industry x base year x labor market region level.

	Mean	SD	p25	p75	Ν
Cumulative earnings (10 years)	388,520.63	261,682.12	227,150.89	467,824.69	159,213
Base year earnings	44,762.66	33,045.27	30,289.71	47,742.27	159,213
100 x Normalized earnings	888.53	386.08	679.41	1081.80	159,213
University degree	0.09	0.28	0.00	0.00	159,213
Vocational degree	0.74	0.44	0	1.00	159,213
Low-skilled	0.18	0.38	0.00	0.00	159,213
Female	0.22	0.41	0.00	0.00	159,213
Age	39.42	8.55	32.00	47.00	159,213
Foreign nationality	0.11	0.31	0.00	0.00	159,213
Experience (years)	13.19	5.94	9.00	16.00	159,213
Firm tenure	9.87	6.01	5.00	15.00	159,213
Industry tenure	10.65	6.10	5.00	15.00	159,213
Occupation tenure	10.18	6.06	5.00	15.00	159,213
Firm size	3,675.35	9,426.59	97.00	1,678.00	159,213
Routine share	0.56	0.20	0.35	0.74	159,213
Blue-collar job	0.71	0.45	0.00	1.00	159,213
Manager	0.02	0.14	0.00	0.00	159,213
Industry specialization	0.74	0.99	0.09	0.89	159,213
AKM worker FE	4.33	0.30	4.16	4.48	159,213
AKM firm FE	0.25	0.16	0.18	0.35	159,213
Export exposure	26.97	23.80	10.88	38.20	159,213
Import exposure	24.80	40.92	5.30	32.72	159,213

Table A.1: Sample descriptives

Note: The table shows descriptive statistics on the estimation sample. Cumulative earnings are computed as the sum of total earnings within the 10-year period (1991-2000, 2001-2010). Base year earnings are earnings in the first year of the interval (1990, 2000). Normalized earnings are computed as cumulative earnings divided by base year earnings.

Dep.	Dep. Var.: 100 x Normalized cumulative earnings							
	All employers	S	ame sector	-	Other sector			
Same employer		Yes	No	No	No			
Same 2-digit industry		Yes	Yes	No	No			
	(1)	(2)	(3)	(4)	(5)			
Exports	0.722***	0.523***	0.411***	0.041	-0.253***			
	(0.127)	(0.257)	(0.164)	(0.077)	(0.108)			
Imports	-0.150***	-0.532***	-0.128***	0.074***	0.436***			
	(0.062)	(0.133)	(0.068)	(0.032)	(0.071)			
Obs.	159,213	159,213	159,213	159,213	159,213			
R^2	0.111	0.053	0.027	0.036	0.050			

Table A.3: Worker Adjustment

Note: The table shows the estimated link between industry-level import and export competition and normalized cumulative earnings. See equation 1 and section 3 for the empirical specification, the list of control variables, and an explanation of the IV strategy. * p<0.10, ** p<0.05, *** p<0.01

3-digit industry	Change 1990-2010 (Bil.)
(a) Change in exports	
Manufacture of motor vehicles	17.92
Manufacture of parts and accessoires of motor vehicles	12.85
Manuf. of machinery for production & use of mechanical power	8.42
Other special purpose machinery	7.72
Manufacture of electricity distribution and control apparatus	3.66
(b) Change in imports	
Manufacture of office machinery and computers	13.56
Manufacture of motor vehicles	8.69
Manufacture of parts and accessoires of motor vehicles	8.60
Manufacture of velves and tubes	8.22
Manufacturing of radio, television, sound and video	6.50

Table A.2: Change in exports and imports, by industry

Notes: The table shows the top 3-digit industries, in terms of change in exports or imports between 1990 and 2010, respectively, measured in billions of Euros.

Dep. V	ar.: 100 x Norm	alized cum	ulative ea	rnings	
	All employers	Sa	ame sector		Other sector
Same employer		Yes	No	No	No
Same 2-digit industry		Yes	Yes	No	No
	(1)	(2)	(3)	(4)	(5)
Males	0.765***	0.690***	0.370***	0.031	-0.327***
	(0.124)	(0.262)	(0.167)	(0.089)	(0.115)
Females	0.413	0.008	0.467^{*}	0.019	-0.081
	(0.259)	(0.394)	(0.240)	(0.114)	(0.172)

Table A.4: Adjustment to Exports: Differences by Gender

Note: The table shows the estimated link between industry-level export exposure and normalized cumulative earnings, separately by gender. See equation 1 and section 3 for the empirical specification, the list of control variables, and an explanation of the IV strategy. * p<0.10, ** p<0.05, *** p<0.01

Quartile:	(1)	(2)	(3)	(4)		
	B:]	Dummy	variabl	es		
University degree	0.10	0.07	0.09	0.09		
Low-skilled	0.28	0.16	0.12	0.14		
Female	0.38	0.20	0.16	0.15		
Foreign nationality	0.13	0.11	0.09	0.11		
Blue-collar job	0.62	0.69	0.72	0.80		
Manager	0.04	0.02	0.02	0.01		
	B: Metric variables					
Experience (years)	0.39	0.14	-0.19	-0.35		
Firm tenure	0.36	0.06	-0.24	-0.18		
Industry tenure	0.35	0.07	-0.23	-0.19		
Occupation tenure	0.36	0.11	-0.19	-0.28		
Firm size	-0.16	-0.15	-0.09	0.39		
Age	0.52	0.04	-0.27	-0.30		
Routine share	0.03	-0.04	-0.12	0.13		
Industry specialization	-0.11	-0.10	-0.06	0.27		
AKM worker FE	-0.13	0.03	0.07	0.04		
AKM firm FE	0.09	0.05	-0.01	-0.13		

Table A.5: Descriptives for quartiles of $\hat{\beta}(x_i)$

The table shows mean descriptive statistics for quartiles of the ranked $\hat{\beta}_i$ predicted by the GRF. Metric variables are normalized.

	Dep. var.: Individualized effect						
	(1)	(2)	(3)	(4)	(5)		
Female	-0.27***	-0.25***	-0.28***	-0.20***	-0.19		
	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)		
R^2	0.04	0.12	0.33	0.38	0.57		
Skills & Demographics		Х	Х	Х	Х		
Firms			Х	Х	Х		
Occupations				Х	Х		
Firms x Skills & Demographics					Х		
Firms x Occupations					Х		

Table A.6: Gender differences: the role of differences in characteristics and sorting

The Table shows estimates of a regression of the individualized effect on a gender dummy and control variables. Skill and demographic controls include education, AKM worker effects, age, firm/occupation/industry tenure, experience, and nationality. Firm controls include dummies for deciles of firm size and AKM firm wage premium. * p<0.10, ** p<0.05, *** p<0.01

	Female	Foreign	Age	Low- skilled	Uni- versity degree	Firm size	Expe- rience (years)	Firm tenure	Indu- stry tenure	AKM worker FE	AKM firm FE	Indu- stry spece- lisation	Rou- tine share	Occu- pation tenure	Manager	Blue- collar job	Base year earnings
Female Foreign Age Low-skilled University degree Firm size Experience (years) Firm tenure Industry tenure AKM worker FE AKM firm FE Industry specialization Routine share Occupation tenure Manager Blue-collar job	$\begin{array}{c} 1.00\\ 0.01\\ 0.01\\ 0.03\\ 0.18\\ 0.09\\ -0.09\\ -0.09\\ -0.04\\ -0.04\\ -0.04\\ -0.04\\ -0.04\\ -0.04\\ -0.04\\ -0.03\\ -0.04\\ -0.03\\ -$	1.00 0.03 0.03 0.03 0.02 0.01 0.01 0.01 0.01 0.01 0.01 0.02 0.01 0.02 0.01 0.02 0.03 0.02 0.01 0.02 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.02 0.01 0.03 0.03 0.01 0.03 0.01 0.03 0.03 0.01 0.03 0.03 0.01 0.03 0.01 0.03 0.01 0.03 0.01 0.03 0.01 0.03 0.03 0.03 0.01 0.03	$\begin{array}{c} 1.00\\ 0.02\\ 0.02\\ 0.02\\ 0.04\\ 0.07\\ 0.04\\ 0.07\\ 0.05\\ 0.04\\ 0.07\\ 0.07\\ 0.02\\ 0.07\\ 0.07\\ 0.07\\ 0.07\\ 0.07\\ 0.07\\ 0.07\\ 0.05\\ 0.07\\ 0.05\\ 0.07\\ 0.05\\$	$\begin{array}{c} 1.00\\ -0.14\\ -0.04\\ 0.02\\ 0.07\\ -0.02\\ 0.07\\ 0.02\\ 0.03\\ 0.02\\ 0.03\\ 0.04\\ 0.03\\ 0.04\\ 0.03\\ 0.04\\ 0.03\\ 0.04\\ 0.$	$\begin{array}{c} 1.00\\ 0.06\\ -0.09\\ -0.11\\ -0.11\\ -0.15\\ -0.38\\ -0.38\\ -0.38\\ 0.21\\ 0.21\\ 0.21\\ \end{array}$	$\begin{array}{c} 1.00\\ 0.15\\ 0.15\\ 0.12\\ 0.12\\ 0.11\\ 0.12\\ 0.03\\ 0.02\\ 0.02\\ 0.02\\ 0.02\\ 0.02\\ 0.02\\ 0.02\\ 0.02\\ 0.03\\ 0.02\\ 0.02\\ 0.03\\ 0.02\\ 0.02\\ 0.03\\ 0.02\\ 0.03\\ 0.02\\ 0.03\\ 0.02\\ 0.03\\$	1.00 0.64 0.72 0.17 0.17 0.17 0.22 -0.05 0.03 0.03	1.00 0.89 0.15 0.18 -0.03 0.61 -0.01	1.00 0.17 0.18 0.08 0.00 0.00 0.00	1.00 0.19 -0.47 0.16 0.23	-0.01 -0.07 -0.08 -0.08 -0.01	1.00 0.49 0.04 0.09	1.00 -0.03 -0.21	1.00 -0.07 0.03	1.00	1.00	
Base year earnings Note: The Table s heterogeneous effect; All variables are mea	-0.22 shows th s. In add sured in	-0.07 ne correl lition, it i the base	0.11 ation l nclude year.	-0.17 betweer s the co	0.5 1 all va rrelation	0.15 riables s with	0.13 used t base yea	0.09 o estim ar earnii	0.11 ate ıgs.	0.61	0.31	-0.18	-0.37	0.08	0.2	-0.27	1.00

Table A.7:

A.1 Details on the Generalized Random Forest

A.1.1 Regression trees

We shortly introduce the concept of a regression tree, adapted to the purpose of finding heterogeneity in treatment effects by Athey and Imbens (2016). A tree starts with the entire sample of observations and recursively splits the sample into smaller sub-samples. The sub-samples are selected in order to maximize the heterogeneity of the treatment effects between the resulting groups.

Figure A.11 shows an example of a simple regression tree. Here, the tree first splits on the variable x_1 . Each observation with values $x_1 \le t_1$ will be part of the sub-sample in the left branch of the tree, and each observation with $x_1 > t_1$ will be part of the sub-sample in the right branch of the tree. The sub-samples after the last split of this tree are called leaves or final regions (R_1 to R_5 in Figure A.11). It is common practice to specify a stopping criterion, e.g., a minimum number of observations in a leaf, to avoid spurious estimates and over-fitting. A tree is greedy in that with each new split it chooses the variable and value that maximizes the heterogeneity of the effect between the groups without considering later splits (so-called top-down greedy approach). The final estimates of the treatment effect are estimated within the leaves and are estimates of observations with similar characteristics.





Note: This Figure shows a simple example of a regression tree, provided by James et al. (2021). The variables X_1 and X_2 are used to estimate heterogeneous effects. R_1 - R_5 are called leaves or final regions.

A.1.2 Building a (generalized) forest

A single regression tree is prone to provide unreliable or spurious estimates of the true heterogeneity in the estimated impact of the variable of interest. A more reliable estimate can be obtained by using so-called random forests. Random forests, first proposed by Breiman (2001), build a large number of trees, where each tree uses a bootstrap sample of the original sample to reduce the influence of single observations. In addition, each split within a tree only uses a randomly chosen subset of all available variables. The final estimate is a combination of the results of all trees. A key concept in the context of causal trees and forests is called honesty (Athey and Imbens, 2016). Honesty refers to the practice of using two different samples for building a tree and for estimating the effect of interest within the final leaves of the tree. In the case of a causal random forest, the bootstrap sample is used to construct the tree and the observations which have not been drawn into the bootstrap sample are used to estimate the effects within the leaves.

The generalized random forest (GRF) procedure proposed by Athey et al. (2019) enables us to combine a causal random forest with an instrumental variable regression (instrumental forest). Let us denote the effect of exports on earnings as β . In addition, considering a set of variables *x* used to study heterogeneities in the effect of exports on earnings. With the GRF, we can estimate a local version of our instrumental variable regression within the leaves of the trees and learn how the export-induced earnings effect β is distributed over the values (*x*)

$$\beta(x) = \frac{Cov[Y_i, Z_i | X_i = x]}{Cov[W_i, Z_i | X_i = x]}.$$
(7)

An important new feature in the GRF algorithm is the inclusion of similarity weights, showing the importance of a single observation for the estimation of the local treatment effect. Analogously to traditional local estimation, $\beta(x)$ is estimated by weighting each observation by its importance (e.g. in kernel regression as the distance to a local point x). In the GRF procedure, weights $\alpha_i(x)$ are found in a data-driven way. Given a forest where the set of grown trees is indexed by b = 1, ..., B, the weights are obtained as the average of the importance of a single worker *i* for the estimation of the local effect $\beta(x)$. In a single tree *b*, the estimation is given as follows:

$$\alpha_{bi}(x) = \frac{1(\{X_i \in L_b(x)\})}{|L_b(x)|},\tag{8}$$

which equals 1 for a single worker in a final region of a tree divided by $L_b(x)$, the number of workers falling into that region. Coming to the forest, the weight for a single worker *i* is given by:

$$\alpha_i(x) = \frac{1}{B} \sum_{b=1}^B \alpha_{bi}(x).$$
(9)

The final weights represent the average of the importance of a worker for estimating the local effect $\beta(x)$ over all trees the worker was part of. Having these (similarity) weights, we are able to estimate a weighted local conditional average treatment effect by:

$$(\hat{\beta}(x), \hat{\nu}(x)) \in \arg\min_{\beta, \nu} \left\{ \left\| \sum_{n=1}^{i=1} \alpha_i(x) \psi_{\beta, \nu}(O_i) \right\|_2 \right\},$$
(10)

where $\nu(O_i)$ is a nuisance parameter and O_i contains the information required for the estimation of $\hat{\beta}(x)$. In the instrumental forest O_i includes the outcome Y_i , the endogeneous variable of interest W_i (in our case export exposure) and the instrument Z_i . Estimating (10) with the forest-based weights as in (9), we receive the local effect of an increased export exposure at every worker characteristic $X_i = x$. An important feature of the GRF is that we obtain a fine-grained estimate at the individual level, the individualized effect of exports on earnings $\beta(x_i)$ for a worker *i*.

Similar to the procedure proposed by Athey and Imbens (2016) and Wager and Athey (2018), the trees still split observations to maximize the difference between the treatment effect of the estimated data partitions. In a nutshell, at a new split, the maximum of the difference is found but estimating the effects for the potential sub-groups (referred to as the children C_j) based on every partitioning variable and their realizations $X_i = x$.²² The splitting criterion is given by:

$$\hat{\Delta}(C_1, C_2) = \sum_{j=1}^2 \frac{1}{|\{i : X_i \in C_j\}|} \left(\sum_{\{i : X_i \in C_j\}} \rho_i\right)^2.$$
(11)

Even if only allowing a random subset of partitioning variables this procedure would be very computationally expensive. Instead of maximizing the difference of the resulting treatment effects directly, Athey et al. (2019) use a gradient-based approximation of the true effect. These pseudo effects ρ_i are based on the moment

²²This is a simplification of the splitting criterion described in Athey et al. (2019) which also includes an error term stabilizing the construction of trees as well as keeping them from favoring splits that would complicate the estimation of the model.

conditions of the sample before conducting the split (called parent *P*). Taking the first split in a tree as an example and referring to an application of an instrumental forest, a tree would conduct possible splits and build potential child nodes (C_1 , C_2). But instead of solving the instrumental regression in the child nodes directly, the pseudo outcomes are estimated by relying on the treatment effect $\hat{\beta}_P$ and moment conditions of the parent node as the best possible approximation of the true effect, with ρ_i :

$$\rho_i = (Z_i - \bar{Z}_P)((Y_i - \bar{Y}_P) - (W_i - \bar{W}_P)\hat{\beta}_P),$$
(12)

where Y_i is the outcome, W_i is the endogeneous variable of interest and Z_i is the instrument of a worker *i* who is part of the child node C_j . \bar{Y}_P, \bar{W}_P and \bar{Z}_P denote the moments the parent node *P*. After having found the split where $\hat{\Delta}(C_1, C_2)$ is maximized, the true effect β as in (10) can be estimated in the child nodes with the forest weights as in (8) and (9).

In the GRF, we do not use the original outcomes but conditionally centered versions by leave-one-out estimations, with $\tilde{Y}_i = Y_i - \hat{y}^{(-i)}(X_i)$, $\tilde{W}_i = W_i - \hat{w}^{(-i)}(X_i)$ and $\tilde{Z}_i = Z_i - \hat{z}^{(-i)}(X_i)$, to regress out the marginal effect of X_i . This step is done by running three separate linear regressions to obtain $\hat{y}^{(-i)}(X_i)$, $\hat{w}^{(-i)}(X_i)$ and $\hat{z}^{(-i)}(X_i)$ and estimating the centered outcomes \tilde{Y}_i , \tilde{W}_i , \tilde{Z}_i . The orthogonalization of Y_i , W_i and Z_i before the estimation of the GRF has the advantage of allowing us to consider variables, such as regional indicators, without having to include them in the instrument forest and making the moment conditions more robust to confounding effects.