

NBER WORKING PAPER SERIES

DIVISION OF LABOR IN THE GLOBAL ECONOMY

Sascha O. Becker
Hartmut Egger
Michael Koch
Marc-Andreas Muendler

Working Paper 34549
<http://www.nber.org/papers/w34549>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2025

Hartmut Egger acknowledges financial support by the German Research Foundation for the project "Economic Change, the Global Division of Labor, and Prosperity" (No. 442940844). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Sascha O. Becker, Hartmut Egger, Michael Koch, and Marc-Andreas Muendler. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Division of Labor in the Global Economy

Sascha O. Becker, Hartmut Egger, Michael Koch, and Marc-Andreas Muendler

NBER Working Paper No. 34549

December 2025

JEL No. F12, F16, J30, L23

ABSTRACT

This paper links globalization, worker efficiency, and wage inequality within plants to internal labor market organization. Using German plant-worker data and information on the task content of occupations, we document that larger plants (i) use more occupations, (ii) assign fewer tasks per occupation, and (iii) exhibit greater wage dispersion. We develop a model where plants endogenously bundle tasks into occupations, improving worker-task matching at the cost of higher fixed span-of-control costs. Embedding this into a Melitz framework, we show that trade increases worker efficiency and wage inequality in exporting plants, whereas non-exporting plants experience the opposite effects. Structural estimation and simulations confirm the model's predictions and point to non-monotonic economy-wide effects.

Sascha O. Becker
University of Warwick
and Monash University
s.o.becker@warwick.ac.uk

Michael Koch
Aarhus University
Department of Economics
mkoch@econ.au.dk

Hartmut Egger
University of Bayreuth
hartmut.egger@uni-bayreuth.de

Marc-Andreas Muendler
University of California, San Diego
Department of Economics
and NBER
muendler@ucsd.edu

1 Introduction

More than two centuries ago, Adam Smith introduced the division of labor as a foundational principle for the productivity gains from specialization, emphasizing the assignment of narrower task ranges to individual workers as the key mechanism behind these gains.¹ Despite its early prominence, task-based modeling remained uncommon until the influential studies of Autor, Levy, and Murnane (2003) and Grossman and Rossi-Hansberg (2008), who considered task-based production to explain the distributional effects of technological change and offshoring. These and subsequent models building on them highlight how the routinization of tasks shapes occupational inequality. However, this narrow focus overlooks Smith’s operational view on how firms bundle cumulative tasks into occupations.

This paper revisits the cumulative nature of tasks and explores how task scope within occupations influences firm productivity and wage dispersion. Building on assignment models (see Sattinger 1993, for an overview), we assume that workers differ in their productivity across tasks. However, this heterogeneity is not fully observable to firms, leading to mismatches in task assignment and productivity losses. Since task bundling and worker assignment occur within firms, the resulting mismatch and wage dispersion are internal to the firm. This intra-firm dispersion explains a substantial share of wage inequality not accounted for by observable characteristics (see Abowd et al. 2001, Lemieux 2006).

We extract information on cumulative tasks in Adam Smith’s operational sense from German labor force surveys of the working population. These surveys allow us to construct time-consistent measures of workplace operations and multitasking, including the number of tasks workers perform *within* their occupations. We combine the task information by occupation, industry, location and plant size with German linked plant–worker data. This enables us to document three striking facts: First, larger plants adopt more occupations. Second, workers at larger plants perform a narrower range of tasks within the same occupation. Third, wages are more dispersed *within* occupations at larger plants.

These three facts suggest two key conclusions. First, heterogeneity in worker efficiency to perform diverse workplace tasks is inevitably linked to wage inequality. Second, observed worker efficiency and wage dispersion are shaped by plant size. Our hypothesis is that workers differ in their ability to carry out the tasks of an occupation, such that match quality determines worker efficiency within occupations. Ability mismatches gen-

¹To put it in the words of Smith (1776, Book I, Chapter I): “[M]aking a pin is . . . divided into about eighteen distinct operations. . . . [T]en persons . . . could make among them upwards of forty-eight thousand pins in a day. . . . But if they had all wrought separately and independently . . . they certainly could not each of them have made twenty, perhaps not one pin in a day.”

erate wage inequality. Reducing such mismatches is costly and more attractive for plants with higher revenue potential.

To formalize our hypothesis, we propose a model of endogenous occupation choice and task assignment by the employer. Employers can organize the full range of tasks required for production into either fewer or more occupations. Fewer occupations imply that each encompasses a broader range of tasks, requiring workers to multitask more. Conversely, a greater number of occupations narrows the task range per occupation, reflecting a finer division of labor. We postulate that each worker has a core ability that makes them most efficient at one particular task, with efficiency declining monotonically as tasks deviate from this core ability. As a result, narrower task ranges reduce mismatch, since all tasks are closer to a worker’s core ability. This enhances worker efficiency and benefits the plant. However, increasing the number of occupations entails a span-of-control fixed cost that rises with occupational complexity. Our model thus captures a trade-off between lower variable costs from better task matching and higher fixed costs from finer labor division—echoing the foundational insight of Becker and Murphy (1992) on the tension between multitasking and specialization.

In a Melitz (2003)-type model with exogenous plant heterogeneity rooted in elemental productivity differences, more productive plants can recover higher span-of-control fixed costs through greater operating profits. As a result, they optimally adopt a larger number of occupations. In particular, productive plants that self-select into exporting choose more occupations with narrower task ranges than non-exporters. Consistent with our stylized facts, the model also predicts that plants with a finer internal division of labor exhibit greater wage inequality within occupations. This arises because the quality of the plant–worker match affects how sensitively worker performance responds to task mismatch.

Because several key parameters in our model are not directly observable in the data, we organize our formal analysis to facilitate their structural estimation within a coherent theoretical framework. This enables us to quantify the mechanisms we propose in a numerical simulation of our model. Specifically, we model unobserved plant-level heterogeneity using stochastic parameters drawn from a multivariate log-normal distribution (see Helpman et al. 2017, Fernandes et al. 2023). These latent characteristics include elemental productivity as well as the fixed costs associated with both production and foreign market entry. In contrast, we treat the sensitivity of worker performance to task mismatch as a common deterministic parameter. This modeling choice reflects a structural constraint: the joint estimation of the sensitivity of worker performance alongside plant-specific stochastic terms is not feasible within the current framework.

Although the presence of multiple stochastic parameters complicates the formal analysis, the model still yields analytical predictions consistent with our empirical findings. For example, in an open economy, exporters adopt more occupations with narrower task ranges than non-exporters. This reduces mismatch, increases worker efficiency, and, if individual worker performance responds to mismatch with an elasticity smaller than one, it leads to greater wage dispersion. Moreover, the general equilibrium structure of the model allows us to examine the broader effects of trade liberalization. While trade increases aggregate welfare, its impact on average worker efficiency and economy-wide wage dispersion is non-monotonic (see Helpman et al. 2017, for non-monotonic effects of trade on wage inequality between firms).

To estimate the parameters of the multivariate log-normal distribution alongside other model parameters, we use a maximum likelihood (ML) estimator tailored to our setting. Our approach extends previous implementations of structural heterogeneous-firm estimation in two important ways. First, we depart from the common assumption of a fixed mass of potential entrants and a Chaney (2008)-type framework in which firms jointly earn positive profits. Instead, we adopt the original two-stage entry mechanism of Melitz (2003), with an unbounded pool of potential entrants. This allows us to explore how the extensive margin from adjustments in the mass of firms drawing elemental productivity affects parameter identification. Second, we address a censoring problem that does not conform to standard Tobit models such as Carson and Sun (2007), since censoring in our case arises from unobserved stochastic plant characteristics rather than from an observed outcome variable. We show that our model is point identified and we derive an ML estimator that is broadly applicable to a wide range of Melitz (2003)-type models.

Since our structural estimation does not discipline the sensitivity of worker performance to task mismatch, we cannot separately identify worker efficiency and wage dispersion. To address this, we apply a second-step method of moments estimator, using computed and observed realizations of economy-wide worker efficiency and wage dispersion. This allows us to quantify the impact of trade liberalization on both outcomes. We find that the simulated model closely matches observed realizations of worker efficiency and wage dispersion in out-of-sample periods. Moreover, we show that trade liberalization due to a uniform reduction in fixed exporting costs can significantly raise economy-wide worker efficiency, particularly at intermediate levels of fixed costs, where the efficiency gains are sizable relative to overall welfare effects. In contrast, the impact on wage dispersion is modest.

The evolving task composition within occupations has been shown to be closely linked

with recent labor market developments, including the polarization of employment (Autor, Katz, and Kearney 2006, Dustmann, Ludsteck, and Schönberg 2009, Goos, Manning, and Salomons 2009) and the offshorability of jobs (Levy and Murnane 2004, Blinder 2006). The assignment of tasks in an open economy, and the implications for welfare and wage inequality, have been studied from a theoretical perspective in industry-level models, including the Heckscher-Ohlin (Grossman and Rossi-Hansberg 2008, 2010) and the Ricardian framework (Rodríguez-Clare 2010, Acemoglu and Autor 2011). Our model complements the industry-level perspective with a plant-level view and emphasizes productivity gains from trade due to average improvements in the match quality between worker abilities and job requirements in exporting plants.

The improvements in match quality highlighted in this paper differ from the efficiency and distributional effects studied in models of heterogeneous worker–firm matching. These include frameworks with market thickness effects (Amiti and Pissarides 2005, Davidson, Matusz, and Shevchenko 2008), Roy-type assignment mechanisms (Ohnsorge and Trefler 2007, Costinot and Vogel 2010), positive assortative matching (Sampson 2014, Grossman, Helpman, and Kircher 2017), and monopsonistic labor market structures (Jha and Rodríguez-Lopez 2021, Egger et al. 2022). Empirical support for these mechanisms comes from Bombardini, Orefice, and Tito (2019), who use French employer–employee data to show that exporters achieve better match quality. Our model emphasizes a distinct source of efficiency gains: plants improve match quality by narrowly assigning tasks to workers whose core abilities best fit those tasks. This interpretation aligns with the idea that human capital is occupation-specific, as supported by empirical evidence from Kambourov and Manovskii (2009).

Focusing on the plant-internal assignment of heterogeneous workers to heterogeneous tasks links our analysis to a literature pointing to the role of human resource management practices to explain variation in plant and firm productivity within and across countries (Bloom and Van Reenen 2011). Yet, aspects of the internal labor market and residual wage inequality are difficult to observe directly. Recent studies of the firm’s internal labor market have therefore turned to the importance of observable hierarchies (Caliendo and Rossi-Hansberg 2012, Caliendo, Monte, and Rossi-Hansberg 2015) and their response to firm-level trade. Our model complements the hierarchical approach to a firm’s internal organization with a perspective on the horizontal differentiation of worker abilities and their tasks within hierarchical layers. In fact, we find that most employer-level residual wage inequality in the German data occurs within occupations, suggesting that an important component of the differences in both worker efficiency and residual wages exists

within hierarchies.

In our model, relatively more productive plants enhance their elemental productivity through a stricter division of labor, which raises worker efficiency and amplifies plant size differences beyond inherent productivity dispersion. While the selection of more productive plants into exporting remains a core mechanism (see, e.g., Clerides, Lach, and Tybout 1998), our model introduces a feedback loop: exporting increases internal specialization, improving match quality and worker efficiency, akin to a learning-by-exporting effect (see Crespi, Criscuolo, and Haskel 2008). This labor market feedback resembles screening in Helpman, Itskhoki, and Redding (2010), innovation-driven efficiency in Aw, Roberts, and Xu (2011), and team-based specialization in Chaney and Ossa (2013), who build on Becker and Murphy (1992). In all these cases, firms face a trade-off between higher fixed costs and lower variable costs, with the choice to incur higher fixed costs being more appealing to exporters due to their access to larger markets.

Most closely related to our mechanism, Chaney and Ossa (2013) show that a market size increase induces task reassignment toward more specialized teams, thereby reinforcing the returns to exporting. Our model complements this mechanism by introducing employer and worker heterogeneity, linking firm-level specialization to differences in worker efficiency and wage inequality within occupations. In this respect, our model is also related to the team production framework of Jarosch, Oberfield, and Rossi-Hansberg (2021), who explore learning from coworkers as a determinant of workers' future wage growth.

The remainder of this paper is organized as follows. In Section 2, we present data and descriptive evidence to motivate our analysis. In Section 3, we model production with task assignments to occupations. We derive the equilibrium for a closed economy in Section 4 and for two symmetric open economies in Section 5. In Section 6 we structurally estimate key model parameters and use them to simulate the impact of trade liberalization on worker efficiency and wage inequality in Germany. Section 7 concludes.

2 Data and Evidence

The two main sources for our micro-level evidence on plant-level task assignments are *(i)* linked plant–worker data and *(ii)* labor force surveys of the working population. In this section, we elicit three descriptive facts from these two datasets to motivate a theory that can explain the within-plant division of labor and the resulting wage dispersion within occupations.

2.1 Linked plant–worker data

To obtain detailed information on workers and their employers, we use data from the German Federal Employment Office’s Institute for Employment Research (IAB): the linked plant–worker data LIAB. The LIAB data combine administrative records on workers from the German social security system with the IAB establishment panel, which provides plant information from surveys on an annual basis since 1993. Since information on plants in East Germany is only available since 1996, we focus on observation years after 1996 to cover the German economy as a whole. At the plant level we use information on revenues, exporting and employment as well as region and industry categories.² At the worker level, LIAB offers a comprehensive set of characteristics. We use demographic, tenure and education indicators, occupation characteristics, and data on workers’ monthly wages. Wage information in the social security records is right-censored, so we replace censored by imputed wages, following the procedure proposed by Card, Heining, and Kline (2013). Since we do not observe work time, we restrict the sample to full-time workers and use daily wages as the most granular measure of earnings for our analysis. Larger plants are over-represented in the establishment panel, so we use weighting factors provided by IAB to make the plant-level data representative of the German economy.


We can use the LIAB data to establish:

Fact 1. *Both the count of occupations at a plant and the residual wage dispersion within plant-occupations increase with plant employment.*

To relate the number of distinct occupations (from a comprehensive list of 357 occupations according to the German nomenclature KldB-88) to plant size, we project the observed count of occupations of a plant on sector, region, time, occupation and worker characteristics. Panel A of Figure 1 plots the thus cleaned count of occupations (on the horizontal axis) against plant employment by size category (on the vertical axis). We focus on the four observation years 1999, 2006, 2012, and 2018, which are also covered by our second dataset, and we normalize the occupation count (on the horizontal axis) by subtracting the count at the smallest plants with 1 to 4 workers. The figure shows that the occupation count increases monotonically with plant size. Around the average occupation count per plant-size category, the figure draws thick, medium, and thin lines that represent the 99, 95, and 90 percent confidence intervals, but those are largely invisible given only minor dispersions of the occupation counts within size categories. Of course, the reported

²We construct 39 longitudinally consistent industries for all data sources, based on an aggregation of 2-digit industries from the German nomenclature WZ 2003 (see Becker and Muendler 2015).

Figure 1: Occupation Counts and Internal Wage Dispersion by Plant Size



Source: LIAB 1999, 2006, 2012 and 2018.

Notes: Panel A: Prediction of occupation count by plant employment category, controlling for sector, region, time, occupation and worker characteristics. Panel B: Prediction of CV of residual daily wages by plant employment category, controlling for sector, region, occupation and worker characteristics. Results in both panels are differences to smallest plant-size category with 1 to 4 workers. Thick, medium, and thin lines represent the 99, 95, and 90 percent confidence intervals.

positive link between plant size and occupation count is not necessarily the result of profit maximization. A similar pattern would arise when plants of different size just randomly assign workers to the available occupations.

However, pointing to a so-far unexplored link between plant size and the dispersion of daily wages within plant-occupations, the second finding in Fact 1 speaks against a random process explaining the observed pattern between plant size and occupation counts. To assess the dispersion in daily wages, we first remove observed demographic, education and tenure information together with time, industry and region effects from log daily wages in a Mincer regression, and compute residual log daily wages. Similar to other studies we find that half of the wage dispersion remains unexplained by observed worker characteristics. To control for other potential explanations of wage dispersion mentioned in the literature, we also eliminate wage differences *between* industries, plants, and occupations. This refined measure of plant-occupation internal wage dispersion still amounts to almost one-fourth of the overall dispersion of raw wages in our data and, as illustrated by a detailed wage decomposition in Appendix A.1, is large compared to the wage dispersion between plants and between hierarchical layers within plants highlighted by previous research (see Helpman, Itskhoki, and Redding 2010, Caliendo and Rossi-Hansberg 2012).

To relate wage dispersion to plant size, we then project the coefficient of variation,

CV , of the (exponentiated) residual daily wages within a plant-occupation on sector, region, occupation and worker characteristics. Panel B of Figure 1 plots this cleaned CV of daily wages within a plant-occupation in logs, after subtracting the coefficient of daily wage variation at plants with up to four workers, (on the horizontal axis) against plant employment by size category (on the vertical axis). Naturally, we cannot compute the coefficient of variation of wages for plants with less than two employees. Therefore, we focus on plants with at least two (full-time) workers in the construction of the figure. There is a clearly positive relationship: workers within the same occupation are subject to more wage dispersion within their occupation at larger plants.


To provide an explanation for the systematic relationship between a plant's choice of occupation counts and the resulting plant-occupation internal wage dispersion, we need to look inside occupations and get a better understanding of what workers actually do in their jobs. German labor force survey data provide this information.

2.2 Labor force survey data

For a meaningful analysis of the properties of occupations, we take information on the organization of the workplace from four German labor force surveys of working population conducted over the years 1999 through 2018 by Germany's Federal Institute for Vocational Education and Training (BIBB) in collaboration with Germany's Federal Institute for Occupational Safety and Health (BAuA). Each wave selects a random sample of around one promille of the German labor force with more than 20 hours of work per week. The BIBB-BAuA data report detailed information on workplace properties, worker characteristics, the industry, occupation and earnings, as well as rudimentary information on the employer, such as the size of a worker's plant in seven categories. Most importantly, we observe workers' responses to survey questions that regard the tasks they perform in their occupation. Following the time consistent definitions in Becker and Muendler (2015), we employ the 1999, 2006, 2012 and 2018 survey data and make use of the questions that point to *what* operations (tasks) a worker carries out on the job. A worker may report these cumulative operations as performed or not.³

³We can discern 15 time-consistent workplace operations: 1. Manufacture, Produce Goods; 2. Repair, Maintain; 3. Entertain, Accommodate, Prepare Foods; 4. Transport, Store, Dispatch; 5. Measure, Inspect, Control Quality; 6. Gather Information, Develop, Research, Construct; 7. Purchase, Procure, Sell; 8. Program a Computer; 9. Apply Legal Knowledge; 10. Consult and Inform; 11. Train, Teach, Instruct, Educate; 12. Nurse, Look After, Cure; 13. Advertise, Promote, Conduct Marketing and PR; 14. Organize, Plan, Prepare Others' Work; 15. Control Machinery and Technical Processes. We report frequencies by individual task in the Supplemental Appendix and show that the number of tasks performed by workers has increased considerably between 1999 and 2006 from an average of 5.28 to an average of 7.38, while

Figure 2: Plant Size, Number of Tasks per Occupation, and Residual Wage inequality



Source: BIBB-BAuA and LIAB with imputed task numbers 1999, 2006, 2012 and 2018.

Notes: Panel A: Prediction of number of tasks within plant-occupation by plant employment category, controlling for sector, region, time, occupation and worker characteristics. Results are displayed relative to the smallest plant-size category (1 to 4 workers). Panel B: Prediction of coefficient of variation of daily wage residual (exponentiated Mincer residual) CV within plant-occupation by task number, controlling for sector, region, time, occupation and worker characteristics. Results are displayed relative to the smallest task-number category (0 to 3 tasks). Thick, medium, and thin lines represent the 99, 95, and 90 percent confidence intervals.

The BIBB-BAuA data allows us to establish:

Fact 2. *The task number within an occupation at a plant decreases with plant size.*

To produce Fact 2, we compute the number of tasks that workers in their respective occupations report in the BIBB-BAuA data. We then project the reported number of tasks per occupation on the same sector, region, time, occupation and worker characteristics as before. Panel A of Figure 2 plots the thus cleaned number of tasks per occupation (on the horizontal axis) against the plant's employment by size category (on the vertical axis). We normalize the number of tasks by subtracting the number of tasks at the smallest plants with 1 to 4 workers.

The figure shows that the number of tasks decreases with plant size. In other words, larger plants assign fewer tasks to their workers (who fill more occupations by Fact 1). In magnitude, the reduction in the number of tasks from small plants, with 1 to 4 workers, to large plants, with 100 or more workers, is 0.4 tasks per worker out of 15 possible tasks. Above the threshold of 100 to 499 workers, plants assign roughly similar task ranges to remaining at a relatively constant level afterwards.

their workers. The pattern shown in Panel A of Figure 2 is well in line with Adam Smith’s tenet that workers engage in less multitasking at more pin-factory like larger plants.

2.3 Data combination

To conduct an employer-level analysis of task assignment, we need to combine the BIBB-BAuA labor force survey information with the LIAB linked plant–worker records through imputation. A large set of worker characteristics and plant attributes overlaps between the BIBB-BAuA survey and the LIAB records. We opt for regression-based imputation to preserve the within-occupation and time variation of task-related information from BIBB-BAuA in the LIAB data.

We first run a linear (OLS) model on the BIBB-BAuA data, regressing the number of tasks (the sum over the 15 task indicators) on a set of worker, occupation and plant attributes that are jointly observed in the BIBB-BAuA and in the LIAB data.⁴ Using the estimated coefficients, we perform an out-of-sample linear prediction in the LIAB data using all common variables. Finally, by computing the mean over all individuals within a plant, we obtain a measure of the average number of tasks per occupation within a plant.⁵

Furthermore, we can also make an out-of-sample prediction regarding the probability that a worker performs a specific task in the LIAB data. For this purpose, we run 15 probit regressions (one for each task) with the same set of explanatory variables as in the regression for the number of tasks outlined above. With these out-of-sample predictions, we then construct a measure for the overall number of distinct tasks performed at a plant in LIAB. Due to the chosen estimation approach, the total number of distinct tasks must be smaller than 15 and it is larger than zero if our mapping was successful for at least one worker at the plant. We then divide the average number of tasks by the full count of distinct tasks observed at the plant to obtain a normalized measure of the number of tasks, a real number on the unit interval.⁶

The imputation of BIBB-BAuA task information into LIAB allows us to establish:

⁴The covariates used in the regression are log daily wage, job experience, squared job experience together with indicators for (*i.*) gender, (*ii.*) 7 schooling and vocational training indicators, (*iii.*) 16 regions, (*iv.*) 34 sectors, (*v.*) 7 plant-size categories, and (*vi.*) 335 occupations. We estimate the number of tasks separately for the four available survey years of 1999, 2006, 2012, and 2018 and compute year-specific predictions.

⁵The imputed average number of tasks per occupation at the plant level varies between 0.60 and 10.85, with a mean of 6.33 and a standard deviation of 1.54. In the BIBB-BAuA data the observed average number of tasks varies between 0 and 15, with a mean of 6.90 and standard deviation 2.93.

⁶The imputed number of distinct tasks varies between a minimum of 0.91 and a maximum of 14.77, with a mean of 8.63 and standard deviation 2.08, whereas the imputed normalized number of tasks varies between 0.08 and almost 1.00 with a mean of 0.75 and a standard deviation of 0.14 (see the descriptive statistics in the Supplemental Appendix).

Fact 3. *Residual wage dispersion within a plant-occupation decreases with the number of tasks per plant-occupation.*

To produce Fact 3, we project the coefficient of variation CV of the (exponentiated) residual daily wages within a plant-occupation on sector, region, time, occupation and worker characteristics, as before. Panel B of Figure 2 plots this cleaned CV of daily wages within a plant-occupation after subtracting the coefficient of daily wage variation in the range of less than or equal to three imputed tasks (on the horizontal axis) against the number of tasks (on the vertical axis), for plants with at least two workers. There is a clear negative relationship: wage dispersion decreases as the number of tasks per plant-occupation increases. Workers within the same occupation are subject to more wage dispersion within their occupation at the same employer if they are assigned narrower task ranges. The theoretical model introduced in the next section is devised to relate the more pronounced within plant-occupation wage dispersion to the plant's internal division of labor.⁷

3 Production with Task Assignment

3.1 Consumption

We consider an economy with a population of L risk-neutral individuals. The representative consumer has preferences over a continuum of differentiated consumption goods $c(\omega)$ and maximizes utility

$$U = \left[\int_{\omega \in \Omega} c(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right]^{\frac{\sigma}{\sigma-1}},$$

subject to the economy-wide budget constraint $\int_{\omega \in \Omega} p(\omega)c(\omega) d\omega = Y$, where Ω is the set of available varieties, $p(\omega)$ is the price of variety ω , Y is aggregate income, and $\sigma > 1$ is the elasticity of substitution between varieties. The resulting economy-wide demand for variety ω of the consumption good is:

$$c(\omega) = \left(\frac{p(\omega)}{P} \right)^{-\sigma} \frac{Y}{P}, \quad (1)$$

⁷In the Supplemental Appendix, we further examine Facts 1 to 3 using econometric methods that address omitted variable bias by controlling for granular fixed effects. To mitigate endogeneity in the relationship between plant size, task scope, and residual wage dispersion, we instrument plant size with time-varying industry-level third-country exports to China, using information from CEPII and the OECD.

where $P \equiv \left[\int_{\omega \in \Omega} p(\omega)^{1-\sigma} d\omega \right]^{1/(1-\sigma)}$ is the CES price index. A producer of variety ω faces total demand $c(\omega)$ for its product.

3.2 Production

We characterize a plant ω by its variable and fixed input requirements, which are satisfied by domestic labor. Following Melitz (2003), we assume that variable labor input depends on the plant-specific elemental productivity $\tilde{\varphi}(\omega) > 0$. The fixed labor input is also plant-specific, employed at a common wage rate w , and comprises three elements: (i) the deterministic common fixed costs of lottery participation, $wf_e > 0$; (ii) the stochastic fixed costs of production, $w\tilde{\zeta}(\omega)f(\omega) > 0$; and (iii) the stochastic fixed costs for foreign market access in the open economy, $w\tilde{\zeta}(\omega)\tilde{\delta}(\omega)f_x > 0$. Both $\tilde{\zeta}(\omega)$ and $\tilde{\delta}(\omega)$ are drawn in a lottery along with $\tilde{\varphi}(\omega)$ and are therefore plant-specific. The fixed costs of production include a non-stochastic, endogenous element $f(\omega)$, which results from a plant's decision on the extent of labor division. This decision determines the variable labor input requirement for production beyond the realization $\tilde{\varphi}(\omega)$.

Labor division results from a plant's assignment of tasks to occupations. There exists an exogenous range of tasks that are uniformly distributed around an activity circle with length 1. The plant's technology requires it to cover a segment of this circle of length $\tilde{\beta}(\omega) < 1$. We can think of $\tilde{\beta}(\omega)$ as being stochastic and jointly drawn with the other technology parameters in a single lottery. By choosing more occupations, with their total count given by $n(\omega) + 1$, a plant lowers the number of tasks a worker must execute to $b(\omega) < \tilde{\beta}(\omega)$, leading to greater worker specialization. We associate occupations with the bundling of adjacent tasks and consider them to provide a symmetric division of the task segment on which a plant operates, imposing the following functional relationship:

$$\frac{b(\omega)}{\tilde{\beta}(\omega)} = \frac{1}{\nu n(\omega) + 1}, \quad (2)$$

where $\nu \in (0, 1]$ gives the exogenous degree of overlap in the set of tasks assigned to occupations, a common parameter beyond a plant's control. In the limiting case of $\nu = 1$, the sets of tasks executed in the various occupations are disjoint, whereas in the polar case of $\nu = 0$, they coincide with each occupation executing the whole segment of tasks $\tilde{\beta}(\omega)$.

To link a task to worker efficiency, we assume that its execution requires specific abilities based on its position on the activity circle. Worker abilities are horizontally differentiated and uniformly distributed along this circle, with each worker's location indicating their core ability. Following Becker and Murphy (1992), workers allocate equal time to all

tasks in their occupation. Efficiency declines with the distance between a worker's core ability at location i and the tasks in the interval $[0, b(\omega)]$ covered by an occupation, giving the average distance between workers and their tasks an interpretation of mismatch. Assuming workers are not systematically misallocated and have their core ability within their occupation's task range, mismatch $m[i, b(\omega)]$ for worker i can be captured by:

$$m[i, b(\omega)] = \frac{1}{b(\omega)} \left[\int_0^i (i - t) dt + \int_i^{b(\omega)} (t - i) dt \right] = \frac{b(\omega)^2 - 2i[b(\omega) - i]}{2b(\omega)}, \quad (3)$$

where t indexes task location. Mismatch is lowest when the worker is located in the middle of the task interval and highest at its boundaries.

There needs to be an inverse relationship between mismatch $m[i, b(\omega)]$ and a worker i 's efficiency $\lambda(i, \omega)$, which we define as

$$\lambda(i, \omega) \equiv \eta + \frac{\tilde{\beta}(\omega)}{m[i, b(\omega)]} = \eta + \frac{2\tilde{\beta}(\omega)b(\omega)}{b(\omega)^2 - 2i[b(\omega) - i]}, \quad (4)$$

with η as the exogenous sensitivity of worker performance to task mismatch. For worker efficiency to be well-defined, we impose that the sensitivity of performance satisfies $\eta > -2$, so that workers from the interval $[0, b(\omega)]$ have positive efficiency in the execution of all tasks covered by the occupation, irrespective of interval length $b(\omega) \leq 1$. Ceteris paribus, there exists an inherent efficiency disadvantage of workers in plants that stochastically draw a higher level of $\tilde{\beta}(\omega)$, because, given the number of occupations, plants covering a larger task segment exhibit less division of labor. Multiplying the reciprocal of the mismatch function by $\tilde{\beta}(\omega)$ neutralizes this scaling effect.

Plant ω can choose to hire a measure $\ell_j(i, \omega)$ of workers with core ability i into occupation j . Average worker efficiency in occupation j is then

$$\lambda_j(\omega) = \frac{1}{\ell_j(\omega)} \int_0^{b(\omega)} \lambda(i, \omega) \ell_j(i, \omega) di, \quad \text{where} \quad \ell_j(\omega) \equiv \int_0^{b(\omega)} \ell_j(i, \omega) di \quad (5)$$

denotes the total amount of labor hired for occupation j at a plant with task range $b(\omega)$ per occupation. Occupation-level output is then $q_j(\omega) = \lambda_j(\omega) \ell_j(\omega)$.

The plant combines outputs $q_j(\omega)$ of all distinct occupations $j = 1, \dots, n(\omega) + 1$ using a Cobb-Douglas production function

$$q(\omega) = \tilde{\varphi}(\omega) [n(\omega) + 1] \exp \left[\frac{1}{n(\omega) + 1} \sum_{j=1}^{n(\omega)+1} \ln q_j(\omega) \right], \quad (6)$$

where $q(\omega)$ is the quantity produced by plant ω . In the case of symmetric levels of output $q_j(\omega)$ across all occupations, the production function in Eq. (6) simplifies to $q(\omega) = \tilde{\varphi}(\omega)[n(\omega) + 1]q_j(\omega)$. Thus, keeping total plant-wide employment and the distribution of workers over the task interval constant, productivity does not change in our model simply because a plant adds new occupations $n(\omega)$. Only if workers specialize on a smaller range of tasks when new occupations are added does worker efficiency, and therefore plant productivity, increase with the addition of these occupations.

3.3 Wage dispersion and the sensitivity of performance

To accommodate the link between the normalized number of tasks executed by workers and plant-occupation internal wage dispersion displayed in Panel B of Figure 2, we need to impose a mechanism that connects the individual wage earned by a worker with core ability i , $w(i, \omega)$, to the worker's efficiency in performing the various tasks in the occupation, $\lambda(i, \omega)$. A particularly tractable link is obtained by setting

$$w(i, \omega) = w \frac{\lambda(i, \omega)}{\lambda_j(\omega)}, \quad (7)$$

where $\lambda(\omega)$ is the employment-weighted average of worker efficiency across the $n(\omega) + 1$ distinct occupations of plant ω , such that $\lambda(\omega) \sum_{j=1}^{n(\omega)+1} \ell_j(\omega) = \sum_{j=1}^{n(\omega)+1} \lambda_j(\omega) \ell_j(\omega)$, whereas w represents the wage paid to workers employed for the provision of the fixed input of production with an efficiency equal to one.

A possible mechanism linking wages to productivity is the efficiency-wage theory, which relies on information asymmetries that make effort non-contractible. In our model, this mechanism must accommodate heterogeneity in both workers and tasks, involving two types of asymmetry. First, before employment, workers cannot fully observe the task content of occupations, and firms cannot assess workers' exact abilities—only whether their core ability lies within the occupation's task range, based on a binary signal. Second, workers choose between full effort (equal to one) or none, and effective efficiency is given by $e(i)\lambda(i, \omega)$. Since effort is unobservable to third parties, contracts can only be written on performance $e(i)\lambda(i, \omega)$. If effort entails a small disutility, firms then optimally offer the wage schedule in eq. (7), paying a uniform rate per efficiency unit, $w/\lambda_j(\omega)$, to all workers in occupation j . This makes them indifferent among applicants whose core abilities fall within the task range.⁸

⁸Despite information asymmetries, the model does not generate unemployment, as workers are fully compensated for productivity differences. A model extension in the Supplemental Appendix introduces

To illustrate how mismatch between workers' core abilities and task requirements affects wages, consider a uniform distribution of workers over the task interval $b(\omega)$. In this case, average worker efficiency is identical across occupations and given by

$$\lambda(\omega) = \frac{1}{b(\omega)} \int_0^{b(\omega)} \lambda(i, \omega) di = \eta + \pi \frac{\tilde{\beta}(\omega)}{b(\omega)}, \quad (8)$$

according to eq. (5). Individual wages are then

$$w(i, \omega) = \frac{w}{\eta + \pi \tilde{\beta}(\omega)/b(\omega)} \left[\eta + \frac{2b(\omega)\tilde{\beta}(\omega)}{b(\omega)^2 - 2i[b(\omega) - i]} \right], \quad (9)$$

according to eq. (7). A positive η attenuates the sensitivity of worker efficiency to mismatch, resulting in an elasticity below unity. Thus, smaller task ranges, by producing lower mismatch, lead to greater dispersion in worker efficiency and wages when $\eta > 0$, and less dispersion when $\eta < 0$.

To further elaborate on the role of the technology parameter η in shaping wage variability within plant-occupations, we use the coefficient of variation $cv_j(\omega)$ as a concise measure of wage dispersion at the occupation level. Assuming a uniform distribution of workers over the task interval $b(\omega)$, we compute

$$cv_j(\omega) = \sqrt{4 - \pi(\pi - 2)} \frac{\tilde{\beta}(\omega)/b(\omega)}{\eta + \pi \tilde{\beta}(\omega)/b(\omega)} \equiv cv(\omega). \quad (10)$$

Eq. (10) confirms that a finer division of the task segment, i.e., more occupations with narrower task ranges, leads to higher wage dispersion when $\eta > 0$. The negative relationship between the normalized task range and wage variation, shown in Panel B of Figure 2, thus suggests a positive value of η . Supporting evidence in the Supplemental Appendix shows that greater sensitivity of worker performance increases the financial risk for employers from small mistakes.

4 Division of Labor in the Closed Economy

As outlined in Section 3, plants are characterized by four stochastic technology parameters: $\tilde{\varphi}(\omega)$, $\tilde{\zeta}(\omega)$, $\tilde{\delta}(\omega)$, $\tilde{\beta}(\omega)$. For our analysis, we focus on $\tilde{\varphi}(\omega)$, $\tilde{\zeta}(\omega)$, and $\tilde{\delta}(\omega)$, which directly influence labor division but are unobservable in the data and require structural es-

search frictions and wage bargaining à la Stole and Zwiebel (1996), showing that our main results hold even with equilibrium unemployment.

timization. In contrast, $\tilde{\beta}(\omega)$ matters only through its ratio with the endogenous task range $b(\omega)$, which is determined by the plant's occupation count $n(\omega)$ via eq. (2).

4.1 Profit maximization in the closed economy

Plants make entry and production decisions in three stages. In stage one, a plant ω pays fixed costs wf_e to enter the technology lottery, receiving a draw of $\{\tilde{\varphi}(\omega), \tilde{\zeta}(\omega), \tilde{\delta}(\omega)\}$, which are immediately sunk. In stage two, conditional on the technology draw, the plant chooses the number of occupations $n(\omega)$ and pays a fixed operating cost of $w\tilde{\zeta}(\omega)f(\omega)$, where $f(\omega) = f_0 + \{\eta + \pi([\nu n(\omega) + 1])\}^\gamma$, with $\gamma > 0$ capturing the convex cost of managing more occupations and narrower task ranges. In stage three, plants hire workers $\ell_j(i, \omega)$ for occupations $j = 1, \dots, n(\omega) + 1$, produce output $q(\omega)$, and sell to consumers. We solve the plant's problem by backward induction.

Stage 3: Profit-maximizing employment choice

Given the wage schedule from eq. (7), the plant chooses employment $\ell_j(i, \omega)$ to maximize operating profits:

$$\psi(\omega) = p(\omega)q(\omega) - w \sum_{j=1}^{n(\omega)+1} \frac{q_j(\omega)}{\lambda_j(\omega)} - w\tilde{\zeta}(\omega) \{\eta + \pi[\nu n(\omega) + 1]\}^\gamma - w\tilde{\zeta}(\omega)f_0, \quad (11)$$

subject to consumer demand in eq. (1), market clearing condition $c(\omega) = q(\omega)$, the worker efficiency in eq. (5), occupation output $q_j(\omega) = \lambda_j(\omega)\ell_j(\omega)$, the production function in eq. (6), and common non-negativity constraints. Profit maximization yields a first-order condition that can be transformed into $[n(\omega) + 1]w\ell_j(\omega) = r(\omega)[(\sigma - 1)/\sigma]$, where $r(\omega) = p(\omega)q(\omega)$ is plant revenue. This implies equal employment across occupations: $\ell_j(\omega) = \ell(\omega)/[n(\omega) + 1]$ (as anticipated in section 3.3), with $\ell(\omega)$ as total employment of plant ω .

Importantly, while the first-order condition determines occupation-level employment, it does not affect the distribution of worker abilities within occupations. Due to the wage schedule and information asymmetry, both plants and workers are indifferent among matches within the relevant task range. As a result, hiring resembles a random draw, and worker abilities are uniformly distributed across each occupation's task interval. Thus, $\ell_j(i, \omega)$ is constant across i , and plant output simplifies to $q(\omega) = \tilde{\varphi}(\omega)[\eta + \pi(\nu n(\omega) + 1)]^\gamma$.

1)] $\ell(\omega)$. The profit-maximizing price follows as a constant markup over marginal cost:

$$p(\omega) = \frac{\sigma}{\sigma - 1} \frac{w}{\tilde{\varphi}(\omega) \{\eta + \pi[\nu n(\omega) + 1]\}}. \quad (12)$$

Stage 2: Profit-maximizing choice of occupation counts

Plants anticipate stage-three profits as a function of their chosen occupation count. Substituting eq. (12) into eq. (11) and applying eq. (1), profits become:

$$\begin{aligned} \psi(\omega) = \frac{1}{\sigma} \frac{Y}{P^{1-\sigma}} \left[\frac{\sigma}{\sigma - 1} \frac{w}{\tilde{\varphi}(\omega) \{\eta + \pi[\nu n(\omega) + 1]\}} \right]^{1-\sigma} \\ - w\tilde{\zeta}(\omega) \{\eta + \pi[\nu n(\omega) + 1]\}^\gamma - w\tilde{\zeta}(\omega)f_0. \end{aligned} \quad (13)$$

Increasing the occupation count reduces marginal costs but raises overhead costs.

Treating $n(\omega)$ as continuous for purposes of exposition, the first-order condition for the profit-maximization problem at stage two is given by:

$$r(\omega) \frac{\sigma - 1}{\sigma} = \gamma w \tilde{\zeta}(\omega) \{\eta + \pi[\nu n(\omega) + 1]\}^\gamma. \quad (14)$$

We assume $\gamma > \sigma - 1$ to ensure an interior maximum. Additionally, we impose conditions such that every plant benefits from choosing $n(\omega) > 0$, returning to this formally below.

Eqs. (1) and (12), together with market clearing $c(\omega) = q(\omega)$, establish a link between relative revenues and occupation counts across plants. Eq. (14) provides a second relationship. For any two plants ω_1 and ω_2 , we obtain:

$$\frac{r(\omega_1)}{r(\omega_2)} = \left(\frac{\tilde{\varphi}(\omega_1) \{\eta + \pi[\nu n(\omega_1) + 1]\}}{\tilde{\varphi}(\omega_2) \{\eta + \pi[\nu n(\omega_2) + 1]\}} \right)^{\sigma-1}, \quad \frac{r(\omega_1)}{r(\omega_2)} = \left(\frac{\eta + \pi[\nu n(\omega_1) + 1]}{\eta + \pi[\nu n(\omega_2) + 1]} \right)^\gamma \frac{\tilde{\zeta}(\omega_1)}{\tilde{\zeta}(\omega_2)}.$$

Solving these jointly yields:

$$\frac{r(\omega_1)}{r(\omega_2)} = \left(\frac{\tilde{\theta}(\omega_1)}{\tilde{\theta}(\omega_2)} \right)^\xi \frac{\tilde{\zeta}(\omega_1)}{\tilde{\zeta}(\omega_2)}, \quad \frac{\eta + \pi[\nu n(\omega_1) + 1]}{\eta + \pi[\nu n(\omega_2) + 1]} = \left(\frac{\tilde{\theta}(\omega_1)}{\tilde{\theta}(\omega_2)} \right)^{\xi/\gamma}, \quad (15)$$

where $\tilde{\theta}(\omega) \equiv \tilde{\varphi}(\omega)/\tilde{\zeta}(\omega)^{1/(\sigma-1)}$ is an auxiliary stochastic variable that captures the interaction of technology parameters, and $\xi \equiv \gamma(\sigma - 1)/(\gamma - \sigma + 1) > 0$ is the elasticity of revenue with respect to $\tilde{\theta}(\omega)$. We refer to $\tilde{\theta}(\omega)$ as *augmented productivity*, as it plays a role similar to elemental productivity in standard heterogeneous firm models.

Plants produce only if their draws of $\tilde{\theta}(\omega)$ and $\tilde{\zeta}(\omega)$ yield non-negative profits. Con-

ditional on $\tilde{\zeta}(\omega)$, a threshold $\hat{\theta}_a^*$ separates active from inactive plants. From eqs. (13) and (14), this zero-profit cutoff is given by:

$$\frac{r(\omega)}{\tilde{\zeta}(\omega)} = \frac{\sigma \xi f_0}{\sigma - 1} w \equiv \bar{r} \left(\hat{\theta}_a^* | \tilde{\zeta}(\omega) \right). \quad (16)$$

Eq. (15) shows that plant outcomes depend only on $\tilde{\theta}(\omega)$ and $\tilde{\zeta}(\omega)$, allowing us to drop ω and index plants by these parameters from now on. Since the marginal firm earns zero profits, its occupation count is minimal. From eqs. (14) and (16), this minimum is determined by:

$$\nu n \left(\hat{\theta}_a^* \right) + 1 = \frac{1}{\pi} \left[(\xi f_0 / \gamma)^{1/\gamma} - \eta \right], \quad (17)$$

implying strictly positive occupation counts for all plants if $\xi f_0 / \gamma > (\eta + \pi)^\gamma$.

Stage 1: Profit-maximizing entry decision

Entry into the technology lottery depends on expected surplus under profit-maximizing behavior by all market participants. Unlike standard models of heterogeneous producers, our framework involves simultaneous draws of three technology parameters. We assume these follow a trivariate normal distribution with density:

$$g_{u,v,z} \equiv \frac{1}{\sqrt{(2\pi)^3 \det(\Sigma)}} \exp \left[-\frac{1}{2} \tilde{\mathbf{x}}^T \Sigma^{-1} \tilde{\mathbf{x}} \right], \quad (18)$$

where $u = \xi \ln \tilde{\theta}$, $v = \ln \tilde{\zeta}$, and $z = \ln \tilde{\delta}$ are auxiliary variables, and $\tilde{\mathbf{x}} = (u, v, z)^T$. The variance-covariance matrix Σ is formed over standard deviations $\sigma_u, \sigma_v, \sigma_z$ and correlation coefficients $\rho_{uv}, \rho_{uz}, \rho_{vz}$, with determinant $\det(\Sigma) = \sigma_u^2 \sigma_v^2 \sigma_z^2 (1 - \rho_{uv}^2 - \rho_{uz}^2 - \rho_{vz}^2 + 2\rho_{uv}\rho_{uz}\rho_{vz})$.

Since $\tilde{\delta}$ is irrelevant in autarky, we focus on the marginal distribution of (u, v) for the moment. The free-entry condition equates expected returns to the cost of lottery participation:

$$\begin{aligned} & \int_{\ln \hat{\theta}_a^*}^{\infty} \int_{-\infty}^{\infty} \exp[v] \left[\frac{\sigma - 1}{\sigma \xi} \bar{r} \left(\exp[u] | \exp[v] \right) - w f_0 \right] \\ & \times \frac{1}{2\pi \sigma_u \sigma_v \sqrt{1 - \rho_{uv}^2}} \exp \left\{ -\frac{1}{2(1 - \rho_{uv}^2)} \left[\left(\frac{u}{\sigma_u} \right)^2 - 2 \frac{\rho_{uv} u v}{\sigma_u \sigma_v} + \left(\frac{v}{\sigma_v} \right)^2 \right] \right\} dv du = w f_e. \end{aligned}$$

Using eqs. (15) and (16), and defining the auxiliary function

$$F(x, \iota) \equiv \exp \left[\frac{(\iota \sigma_u)^2 + 2\iota \rho_{uv} \sigma_u \sigma_v + \sigma_v^2}{2} \right] \Phi \left(-\frac{\ln x - \rho_{uv} \sigma_v \sigma_u + \iota \sigma_u^2}{\sigma_u} \right),$$

with $\Phi(\cdot)$ as the standard normal cumulative distribution function, the free-entry condition simplifies to:

$$\frac{1}{\hat{\theta}_a^*} F(\hat{\theta}_a^*, 1) - F(\hat{\theta}_a^*, 0) = \frac{f_e}{f_0}. \quad (19)$$

As shown in Appendix B.1, eq. (19) has a unique interior solution $\hat{\theta}_a^* > 0$, with $\partial \hat{\theta}_a^* / \partial f_0 > 0$ and $\partial \hat{\theta}_a^* / \partial f_e < 0$.

4.2 The autarky equilibrium

With the solution to the three-stage profit maximization problem, we can now characterize the autarky equilibrium. The free-entry condition implies zero aggregate profits, so total revenues equal total labor income. As in other models with entry à la Melitz (2003), the real wage therefore serves as a utilitarian welfare measure. Combining plant-level revenues from eq. (13), optimal occupation counts from eq. (14), and the zero-cutoff profit from eq. (16), we derive⁹

$$\left(\frac{w}{P} \right)_a = \left(\frac{\sigma - 1}{\sigma} \right)^{\frac{\sigma}{\sigma-1}} \left(\frac{L}{\gamma} \right)^{\frac{1}{\sigma-1}} \left(\frac{\gamma \hat{\theta}_a^*}{\xi f_0} \right)^{\frac{1}{\xi}}. \quad (20)$$

From eq. (20), $(w/P)_a$ increases with a better plant composition, reflected in a higher $\hat{\theta}_a^*$.

To compute an economy-wide measure of worker efficiency, we combine eq. (8) with eqs. (2), (15) and (17), yielding $\lambda(\hat{\theta}) = (\hat{\theta}/\hat{\theta}_a^*)^{1/\gamma} (\xi f_0/\gamma)^{1/\gamma}$. Averaging over plants, weighted by their employment, gives

$$\Lambda_a = \left(\frac{\xi f_0}{\gamma \hat{\theta}_a^*} \right)^{\frac{1}{\gamma}} \frac{F(\hat{\theta}_a^*, 1 + 1/\gamma)}{F(\hat{\theta}_a^*, 1)}. \quad (21)$$

This measure declines with increasing $\hat{\theta}_a^*$, regardless of the sign of η , reflecting a trade-off: higher $\hat{\theta}_a^*$ improves plant composition (higher average elemental productivity) but

⁹From eq. (13), we can write plant-level revenues in the closed economy as a function of $\tilde{\varphi}(\omega)$ and $n(\omega)$. Combining the resulting expression with the first-order condition in eq. (14), we can solve for $r(\omega)/\tilde{\zeta}(\omega)$ as a function of $\hat{\theta}(\omega) = \tilde{\theta}(\omega)^\xi$. Evaluated for a plant with $\hat{\theta}(\omega) = \hat{\theta}_a^*$, using eq. (16) and substituting $Y = wL$, then establishes the real wage in eq. (20).

reduces labor division, as the marginal producer's occupation count remains fixed while their productivity rises.

For wage dispersion, combining eq. (10) with eqs. (2), (15) and (17) establishes $cv(\hat{\theta}) = \sqrt{4 - \pi(\pi - 2)}[1 - \eta/\lambda(\hat{\theta})]/\pi$. Averaging over plants, weighted by their employment, gives

$$CV_a = \frac{\sqrt{4 - \pi(\pi - 2)}}{\pi} \left[1 - \eta \left(\frac{\gamma \hat{\theta}_a^*}{\xi f_0} \right)^{\frac{1}{\gamma}} \frac{F(\hat{\theta}_a^*, 1 - 1/\gamma)}{F(\hat{\theta}_a^*, 1)} \right]. \quad (22)$$

This economy-wide measure of plant-occupation internal wage dispersion may increase or decrease with $\hat{\theta}_a^*$, depending on the sign of η . If $\eta > 0$, stronger labor division raises wage dispersion, so CV_a decreases (along with Λ_a) if $\hat{\theta}_a^*$ increases. If $\eta < 0$, labor division reduces wage dispersion, and CV_a increases with rising $\hat{\theta}_a^*$. If $\eta = 0$, wage dispersion is uniform across plants, making CV_a independent of $\hat{\theta}_a^*$.

The relationships characterized in this section must be interpreted with care, as all variables, including the productivity threshold $\hat{\theta}_a^*$ are endogenous in general equilibrium. To clarify these mechanisms, the Supplemental Appendix explores the effects of exogenous changes in f_0 and f_e , showing that increases in either parameter lower welfare, while raising worker efficiency to moderate the welfare loss, and, if $\eta > 0$, increasing wage dispersion.

5 Division of Labor in the Global Economy

To derive global equilibrium relationships, we focus on trade in differentiated consumption goods between two symmetric countries with equal wages, w . Consumption and production follow the framework outlined earlier. Trade entails two types of costs: variable iceberg transport costs, meaning that $\tau > 1$ units must be shipped for one unit to arrive abroad, and fixed exporting costs, $w\tilde{\zeta}(\omega)\tilde{\delta}(\omega)f_x > 0$, associated with establishing a foreign distribution network. These costs interact with heterogeneity in augmented productivity $\hat{\theta}(\omega)$ and lead more productive plants to self-select into exporting, assuming that trade costs are sufficiently high. However, fixed exporting costs are also plant-specific in our model, resulting in overlapping productivity distributions between exporters and non-exporters. Moreover, the export decision influences a plant's optimal occupation count, and thereby affects worker efficiency. To account for this, we distinguish variables for exporters (superscript e) and non-exporters (superscript d), and use subscript T to denote total market activity, encompassing both domestic and foreign operations.

5.1 The plant's problem in the open economy

Holding economy-wide variables constant, access to exporting does not affect profit maximization of non-exporters. For exporters, however, foreign revenue influences specialization in the internal labor market. Let $r^e(\omega)$ denote an exporter's domestic revenue. In symmetric countries, foreign revenue equals $\tau^{1-\sigma}r^e(\omega)$. The exporter's optimal occupation count $n^e(\omega)$ therefore satisfies:

$$(1 + \tau^{1-\sigma}) r^e(\omega) \frac{\sigma - 1}{\sigma} = \gamma w \tilde{\zeta}(\omega) \{ \eta + \pi [\nu n^e(\omega) + 1] \}^\gamma. \quad (23)$$

Eq. (23) is uniform across exporters, so the ratios in eq. (15) apply for plants with the same export status. However, comparing two otherwise identical plants, one exporting and one not, yields¹⁰

$$\frac{r^e(\omega)}{r^d(\omega)} = (1 + \tau^{1-\sigma})^{\frac{\xi}{\gamma}}, \quad \frac{\eta + \pi [\nu n^e(\omega) + 1]}{\eta + \pi [\nu n^d(\omega) + 1]} = (1 + \tau^{1-\sigma})^{\frac{\xi}{\gamma(\sigma-1)}}. \quad (24)$$

Exporting generates additional revenues and therefore induces a plant to adopt more occupations: $n^e(\omega) > n^d(\omega)$. The resulting finer division of labor makes exporters more efficient and lowers their unit production costs, leading to lower prices by eq. (12). The added efficiency raises an exporter's sales in both the domestic and the foreign market, establishing $r^e(\tilde{\varphi}) > r^d(\tilde{\varphi})$ in eq. (24). In summary, there is a positive feedback effect of exporting on domestic revenues, and this effect raises a plant's incentives to export beyond the benchmark Melitz (2003) model. While exporters raise their productivity by adopting more occupations, the associated increase in efficiency units of labor does not fully accommodate the added labor demand from higher overall sales, so that an exporting plant expands employment. This follows from the constant markup pricing rule generating a direct link between revenues and employment in our model.

Making use of the profit-maximizing choices for employment and occupation counts, we can express the total profits of plant ω under exporting as

$$\psi_T^e(\omega) = t \frac{\sigma - 1}{\sigma \xi} r^d(\omega) - w \tilde{\zeta}(\omega) \tilde{f}_0 - w \tilde{\zeta}(\omega) \tilde{\delta}(\omega) f_x, \quad (25)$$

with $t \equiv (1 + \tau^{1-\sigma})^{\frac{\xi}{\sigma-1}} > 2$ as a composite trade cost parameter. Total profits under non-exporting are given by $\psi_T^d(\omega) = [(\sigma - 1)/(\sigma \xi)] r^d(\omega) - w \tilde{\zeta}(\omega) f_0$. The decision to

¹⁰The plant adopts a single degree of specialization in its internal labor market regardless of the destinations of its products, so $n^e(\omega)$ and $n^d(\omega)$ do not carry a subscript T .

export depends on a plant's draw of augmented productivity $\hat{\theta}(\omega)$ and the stochastic fixed cost parameter $\tilde{\delta}(\omega)$. Some plants produce only due to favorable exporting fixed costs and, combining equations (16) and (25), the zero-cutoff profit condition for this group of exporters is:

$$\frac{r^d(\omega)}{\tilde{\zeta}(\omega)} = \frac{\sigma\xi}{\sigma-1} \frac{w(f_0 + \tilde{\delta}(\omega)f_x)}{t} \equiv \bar{r}^d \left(\hat{\theta}^*(\tilde{\delta}) | \tilde{\zeta} \right), \quad (26)$$

with $\hat{\theta}_0^*$ as the lower bound of $\hat{\theta}^*(\tilde{\delta})$ for $\tilde{\delta} = 0$, characterizing the lowest augmented productivity level among exporters. Equations (15) and (26) then imply $\hat{\theta}^*(\tilde{\delta}) = (1 + \tilde{\delta}f_x/f_0)\hat{\theta}_0^*$.

Similar to the closed economy, there exists also a productivity threshold for non-exporters, denoted by $\hat{\theta}_1^*$ and determined by the zero-cutoff profit condition in eq. (16), establishing $[\sigma\xi/(\sigma-1)]wf_0 = \bar{r}^d(\hat{\theta}_1^*|\tilde{\zeta})$. Since exporting provides access to foreign consumers and generates a positive feedback effect on domestic sales, we have $\hat{\theta}_0^* < \hat{\theta}_1^*$, according to eqs. (15) and (24). Moreover, $\hat{\theta}_1^*$ serves as an upper bound for $\hat{\theta}^*(\tilde{\delta})$, defined by the condition that a plant earns zero profits under both exporting and non-exporting. Setting $\hat{\theta}_1^* = \hat{\theta}^*(\tilde{\delta})$ determines a critical value of $\tilde{\delta}(\omega)$, given by: $\delta_x \equiv (t-1)f_0/f_x$.

In summary, for $\tilde{\delta}$ -draws in interval $[0, \delta_x)$, we only observe exporters in our model. In contrast, for $\tilde{\delta} > \delta_x$ exporters and non-exporters coexist. Using eqs. (14) to (16), (23) and (24), we express a plant ω 's added profit from exporting, $\Delta\psi_T(\omega) \equiv \psi_T^e(\omega) - \psi_T^d(\omega)$, as follows:

$$\Delta\psi_T(\omega) = (t-1) \frac{\hat{\theta}(\omega)}{\hat{\theta}_1^*} w\tilde{\zeta}(\omega)f_0 - w\tilde{\zeta}(\omega)\tilde{\delta}(\omega)f_x. \quad (27)$$

Firms with $\tilde{\delta}(\omega) \geq \delta_x$ opt to export when $\Delta\psi_T(\omega) \geq 0$, which is more likely for plants with higher draws of augmented productivity $\hat{\theta}(\omega) = \tilde{\theta}(\omega)^\xi$, according to eq. (15). Setting eq. (27) equal to zero gives for any realization of $\tilde{\delta}(\omega) \geq \delta_x$ a lower threshold of $\hat{\theta}(\omega)$ that must be surpassed by plants to make exporting attractive to them. This lower productivity threshold is given by $\hat{\theta}_x^*(\tilde{\delta}) = \hat{\theta}_1^*\tilde{\delta}(\omega)/\delta_x$ and thus larger than $\hat{\theta}_1^*$ if $\tilde{\delta}(\omega) > \delta_x$.

As in the closed economy we determine the free entry condition by setting equal the ex ante expected profits of plants to the common cost of participating in the technology lottery, wf_e . However, computing these ex ante expected profits is a tedious task in our model, because we have to distinguish three possible realizations of technology draws. The first one is given by $\tilde{\delta}(\omega) \leq \delta_x$ and $\hat{\theta}(\omega) > \hat{\theta}^*(\tilde{\delta})$ and leads the plant to export. The second one is given by $\tilde{\delta}(\omega) > \delta_x$ and $\hat{\theta}(\omega) > \hat{\theta}_x^*(\tilde{\delta})$ and leads the plant to export as well. The third one is given by $\tilde{\delta}(\omega) > \delta_x$ and $\hat{\theta}(\omega) \in [\hat{\theta}_1^*, \hat{\theta}_x^*(\tilde{\delta}))$ and leads the plant to remain a non-exporter.

In the Supplemental Appendix, we derive the free entry condition for the model with plant-specific $\tilde{\delta}(\omega)$ and show that it can be expressed in the following way:

$$\begin{aligned} \frac{f_e}{f_0} = & \int_{\ln \delta_x}^{\infty} \left\{ \frac{1}{\hat{\theta}_1^*} \tilde{F}(z, \hat{\theta}_1^*, 1) - \tilde{F}(z, \hat{\theta}_1^*, 0) \right. \\ & \left. + \frac{f_x \exp[z]}{f_0} \left[\frac{\tilde{F}(z, \hat{\theta}_x^*(\exp[z]), 1)}{\hat{\theta}_x^*(\exp[z])} - \tilde{F}(z, \hat{\theta}_x^*(\exp[z]), 0) \right] \right\} dz \quad (28) \\ & + \int_{-\infty}^{\ln \delta_x} \left(1 + \frac{f_x \exp[z]}{f_0} \right) \left[\frac{\tilde{F}(z, \hat{\theta}^*(\exp[z]), 1)}{\hat{\theta}^*(\exp[z])} - \tilde{F}(z, \hat{\theta}^*(\exp[z]), 0) \right] dz, \end{aligned}$$

where

$$\begin{aligned} \tilde{F}(z, x, \iota) \equiv & \exp \left[\frac{(\iota \sigma_u)^2 + 2\iota \rho_{uv} \sigma_u \sigma_v + \sigma_v^2}{2} \right] \frac{1}{\sigma_z} \phi \left(\frac{z - \rho_{vz} \sigma_v \sigma_z - \iota \rho_{uz} \sigma_u \sigma_z}{\sigma_z} \right) \\ & \times \Phi \left(- \frac{\ln x - z \rho_{uz} \sigma_u / \sigma_z - \iota \sigma_u^2 (1 - \rho_{uz}^2) - \sigma_u \sigma_v (\rho_{uv} - \rho_{uz} \rho_{vz})}{\sigma_u \sqrt{1 - \rho_{uz}^2}} \right) \end{aligned}$$

is an auxiliary function with $\phi(\cdot)$ as the probability density function of the standard normal distribution. We can show that the right-hand side of Eq. (28) is monotonically decreasing in $\hat{\theta}_1^*$, falling from infinity if $\hat{\theta}_1^* \rightarrow 0$ to zero if $\hat{\theta}_1^* \rightarrow \infty$. This ensures a unique interior open-economy equilibrium. Higher fixed costs of exporting lower the probability of firms successfully entering the productivity lottery: $d\hat{\theta}_1^*/df_x > 0$. Moreover, $\lim_{f_x \rightarrow \infty} \hat{\theta}_1^* = \hat{\theta}_a^*$ and $\lim_{f_x \rightarrow 0} \hat{\theta}_1^* = t\hat{\theta}_a^*$ define the bounds within which this productivity threshold can vary.

Since the elaborate model with plant-specific realizations of $\tilde{\delta}(\omega)$ has a complicated structure, we additionally consider a more parsimonious setting in which $\tilde{\delta}(\omega) = 1$ holds for all plants, implying exporting fixed costs of $w\tilde{\zeta}(\omega)f_x$. Since in this scenario our model shares important properties with Melitz (2003), we refer to it as *canonical*. Using auxiliary function $F(x, \iota)$ from section 4, we express the free entry condition in the canonical model as follows

$$\frac{1}{\hat{\theta}_1^*} F(\hat{\theta}_1^*, 1) - F(\hat{\theta}_1^*, 0) + \frac{f_x}{f_0} \left[\frac{1}{\hat{\theta}_x^*} F(\hat{\theta}_x^*, 1) - F(\hat{\theta}_x^*, 0) \right] = \frac{f_e}{f_0}, \quad (29)$$

with $\hat{\theta}_1^*$ ($\hat{\theta}_x^*$) as the augmented productivity thresholds that must be surpassed to make production (exporting) attractive for the plant and with $\hat{\theta}_x^* = \hat{\theta}_1^* f_x / [f_0(t-1)]$ if $f_x > f_0(t-1)$ and $\hat{\theta}_x^* = \hat{\theta}_1^*$, otherwise. Eq. (29) is derived in Appendix B.3, where we demonstrate

that the free-entry condition has a unique solution in $\hat{\theta}_1^* > 0$, with $d\hat{\theta}_1^*/df_x > 0$ if $f_x < f_0(t-1)$ and $d\hat{\theta}_1^*/df_x < 0$ if $f_x > f_0(t-1)$. Moreover, we find $\lim_{f_x \rightarrow \infty} \hat{\theta}_1^* = \hat{\theta}_a^*$ and $\lim_{f_x \rightarrow 0} \hat{\theta}_1^* = \hat{\theta}_a^*$. A parameter domain with $f_x > f_0(t-1)$ refers to an outcome with sharp selection of more profitable plants into exporting, which, conditional on the draw of $\tilde{\zeta}(\omega)$, are plants with higher realizations of augmented productivity $\hat{\theta}(\omega)$. The finding that $\hat{\theta}_1^*$ is non-monotonic in f_x highlights a key distinction between the canonical model with identical $\tilde{\delta}(\omega) = 1$ and the more elaborate framework with a stochastic, plant-specific fixed cost parameter $\tilde{\delta}(\omega)$.

With these insights into plant decision-making, we can next determine the general equilibrium in the open economy. Given its simpler structure, we will first examine the canonical model, analyze the open economy equilibrium and the impact of trade liberalization for the case of a common $\tilde{\delta}(\omega) = 1$, and discuss how the results have to be modified when considering plant-specific realizations of $\tilde{\delta}(\omega)$ afterwards.

5.2 The open economy equilibrium in a canonical model

The real wage in the open economy with $\tilde{\delta}(\omega) = 1$ is given by

$$\frac{w}{P} = \begin{cases} \left(\frac{\sigma-1}{\sigma} \right)^{\frac{\sigma}{\sigma-1}} \left(\frac{L}{\gamma} \right)^{\frac{1}{\gamma}} \left(\frac{\gamma \hat{\theta}_1^*}{\xi f_0} \right)^{\frac{1}{\xi}} & \text{if } f_x > f_0(t-1) \\ \left(\frac{\sigma-1}{\sigma} \right)^{\frac{\sigma}{\sigma-1}} \left(\frac{L}{\gamma} \right)^{\frac{1}{\gamma}} \left(\frac{\gamma t \hat{\theta}_1^*}{\xi f_0 + f_x} \right)^{\frac{1}{\xi}} & \text{otherwise} \end{cases}. \quad (30)$$

The first line follows directly from eq. (20), as the marginal producer remains a non-exporter when $f_x > f_0(t-1)$. The second line is derived analogously, recognizing that the marginal producer becomes an exporter with total revenues of $tr^d(\omega)$ and incurs fixed costs of $w\tilde{\zeta}(\omega)(f_0 + f_x)$ to operate in both domestic and foreign markets. Since $d\hat{\theta}_1^*/df_x < 0$ when $f_x > f_0(t-1)$, and $0 < d\hat{\theta}_1^*/d(f_0 + f_x) \times (f_0 + f_x)/\hat{\theta}_1 < 1$ when $f_x < f_0(t-1)$, trade unambiguously raises welfare.

While gains from trade are anticipated from prior findings that the resource allocation is undistorted in a one-sector model with iso-elastic demand and monopolistic competition (see, e.g., Dhingra and Morrow 2019), we now examine how the division-of-labor mechanism in our model contributes to these gains. We begin by analyzing plant-level effects under the scenario where only the most productive firms export, i.e., $f_x > f_0(t-1)$. For a non-exporting plant with augmented productivity $\tilde{\theta}(\omega)$, trade reduces the occupation count and thus worker efficiency: $\lambda^d(\hat{\theta})/\lambda^a(\hat{\theta}) = (\hat{\theta}_a^*/\hat{\theta}_1^*)^{1/\gamma}$. This is intuitive, as the marginal producer (earning zero profits) retains its occupation count after the country opens up to trade. However, since $\hat{\theta}_1^* > \hat{\theta}_a^*$, the marginal producer in the open economy

has a higher elemental productivity level and was infra-marginal under autarky, implying a lower occupation count for this plant. Given that relative worker efficiency among non-exporters is a positive function of their productivity ratio, this reduction extends to all non-exporters.

For a plant that chooses to export under the parameter condition $f_x > f_0(t-1)$, access to trade expands market size, creating an initial incentive to increase the occupation count. This effect counteracts the general tendency to reduce occupation counts shared with non-exporters. Combining eqs. (15) and (24), we obtain $\lambda^e(\hat{\theta})/\lambda^a(\hat{\theta}) = \left(t\hat{\theta}_a^*/\hat{\theta}_1^*\right)^{1/\gamma}$ and show in Appendix B.4 that $\lambda^e(\hat{\theta}) > \lambda^a(\hat{\theta})$ when $f_x > f_0(t-1)$. This implies that exporters increase their occupation count and achieve higher worker efficiency in the open economy compared to autarky. Although this result is derived under the assumption that only the most productive plants export, due to $f_x > f_0(t-1)$, it generalizes naturally to the case $f_x \leq f_0(t-1)$, where all firms export.

In summary, we identify an asymmetric response between exporters and non-exporters in their plant-level labor market adjustments following trade liberalization. The following proposition formalizes these findings.

Proposition 1. *In the open economy, relative to autarky, exporting plants increase the number of occupations and narrow the task range per occupation within the internal labor market. This adjustment reduces mismatch and enhances worker efficiency. In contrast, non-exporting plants reduce the number of occupations and widen the task range per occupation, leading to greater mismatch and lower worker efficiency.*

Proof. Analysis in the text and formal derivation details in Appendix B.4. □

In the open economy, exporters increasingly resemble Adam Smith's pin factory, while non-exporters diverge from that model. This asymmetric adjustment in internal labor markets has implications for wage dispersion across plant-occupations. Building on the closed economy derivations, we express the coefficient of variation of wages as a function of a plant's occupation count: $cv(\hat{\theta}) = \sqrt{4 - \pi(\pi - 2)}[1 - \eta/\lambda(\hat{\theta})]/\pi$. A lower occupation count in non-exporting plants not only reduces worker efficiency but also affects internal wage dispersion, which decreases if $\eta > 0$ and increases if $\eta < 0$. Conversely, exporters raise their occupation count, which enhances worker efficiency and increases or decreases wage dispersion depending on the sign of η . The following proposition summarizes the effects of trade on plant-level wage dispersion.

Proposition 2. *In the open economy, relative to autarky, exporting plants increase plant-occupation internal wage dispersion if and only if $\eta > 0$. In this case, non-exporting plants reduce internal wage inequality.*

Proof. Analysis in the text and formal derivation details in Appendix B.4. $\square$

Due to asymmetric plant-level responses, trade openness generates opposing effects on economy-wide labor efficiency Λ and plant-occupation internal wage dispersion CV . Using the auxiliary function $F(x, \iota)$ from Section 4, we compute

$$\Lambda = \begin{cases} \left(\frac{\xi f_0}{\gamma \hat{\theta}_1^*} \right)^{\frac{1}{\gamma}} \frac{F(\hat{\theta}_1^*, 1+1/\gamma) + (t^{1+1/\gamma} - 1)F(\hat{\theta}_x^*, 1+1/\gamma)}{F(\hat{\theta}_1^*, 1) + (t-1)F(\hat{\theta}_x^*, 1)} & \text{if } f_x > f_0(t-1) \\ \left(\frac{\xi(f_0 + f_x)}{\gamma \hat{\theta}_1^*} \right)^{\frac{1}{\gamma}} \frac{F(\hat{\theta}_1^*, 1+1/\gamma)}{F(\hat{\theta}_1^*, 1)} & \text{otherwise} \end{cases}. \quad (31)$$

Derivation details are provided in Appendix B.5. There, we show that Λ is non-monotonic in f_x . Specifically, $\lim_{f_x \rightarrow 0} \Lambda = \lim_{f_x \rightarrow \infty} \Lambda = \Lambda_a$, and average worker efficiency reaches a maximum at some $f_x \geq f_0(t-1)$. The effect of increasing f_x is most transparent when f_x is low enough that all plants choose to export. In this case, the marginal producer, earning zero profits, is also an exporter and, like all other plants, chooses a higher occupation count in the open economy than under autarky. As a result, average worker efficiency rises. Changes in f_x affect Λ only through shifts in the composition of active plants when $f_x < f_0(t-1)$. In this case, a higher export fixed cost raises the productivity threshold $\hat{\theta}_1^*$, which on its own reduces Λ . However, this negative effect is offset by the higher occupation count chosen by the marginal plant in response to increased fixed costs. Since Λ increases with f_x in this domain, it must eventually decline for higher values of $f_x > f_0(t-1)$ to ensure that Λ converges back to its autarky level as $f_x \rightarrow \infty$.

For $\eta > 0$, the discussion above extends to the effect of export fixed costs on economy-wide wage dispersion, which in the open economy is given by:

$$CV = \begin{cases} \frac{\sqrt{4-\pi(\pi-2)}}{\pi} \left[1 - \eta \left(\frac{\gamma \hat{\theta}_1^*}{\xi f_0} \right)^{\frac{1}{\gamma}} \frac{F(\hat{\theta}_1^*, 1-1/\gamma) + (t^{1-1/\gamma} - 1)F(\hat{\theta}_x^*, 1-1/\gamma)}{F(\hat{\theta}_1^*, 1) - (t-1)F(\hat{\theta}_x^*, 1)} \right] & \text{if } f_x > f_0(t-1) \\ \frac{\sqrt{4-\pi(\pi-2)}}{\pi} \left[1 - \eta \left(\frac{\gamma \hat{\theta}_1^*}{\xi(f_0 + f_x)} \right)^{\frac{1}{\gamma}} \frac{F(\hat{\theta}_1^*, 1-1/\gamma)}{F(\hat{\theta}_1^*, 1)} \right] & \text{otherwise} \end{cases}. \quad (32)$$

Similar to Λ , we find that CV is non-monotonic in f_x . Specifically, $\lim_{f_x \rightarrow 0} CV = \lim_{f_x \rightarrow \infty} CV = CV_a$. For $\eta > 0$, economy-wide wage dispersion reaches a maximum at some $f_x \geq f_0(t-1)$, whereas for $\eta < 0$, it reaches a minimum at some $f_x \geq f_0(t-1)$.

The following proposition summarizes the effects of increasing export fixed costs on economy-wide worker efficiency and wage dispersion.

Proposition 3. *Lower levels of export fixed costs f_x , while monotonically increasing total welfare, exert a non-monotonic effect on average economy-wide worker efficiency, which reaches a maximum at some $f_x \geq f_0(t - 1)$. Similarly, higher export fixed costs have a non-monotonic effect on economy-wide wage dispersion, which reaches a maximum (minimum) at some $f_x \geq f_0(t - 1)$ if $\eta > (<) 0$.*

Proof. Analysis in the text and formal derivation details in Appendix B.5. □

The analysis in this section highlights that firm-level responses to trade shocks can be asymmetric and, as such, may but need not contribute to aggregate welfare and a less egalitarian wage distribution. Importantly, the gains from trade, in the form of a stronger internal division of labor and reduced worker mismatch, as well as the associated effects on plant-occupation wage variation, are most pronounced at intermediate stages of globalization, when exporting is widespread but the costs of foreign market access remain substantial. In contrast, relative to autarky, plant-level adjustments in task assignment have only limited aggregate effects when export costs are either very high or very low.

While the canonical assumption of a common $\tilde{\delta}(\omega) = 1$ across all producers is analytically convenient, it limits the model's ability to capture key features of exporting behavior observed in our data. Consistent with evidence from other countries (see, e.g., Armenter and Koren 2015), we find that high-productivity non-exporters coexist with low-productivity exporters. For example, in our combined dataset, over three percent of plants from the lowest revenue quartile are exporters, while nearly two-thirds of plants in the highest revenue quartile are non-exporters. This pattern contradicts the selection mechanism implied by the canonical model. Moreover, although fewer than one-fifth of plants export, the variance of log revenues among exporters is 2.37, compared to 1.70 among non-exporters. As shown in the Supplemental Appendix, this observation is inconsistent with the canonical framework, suggesting that it may be too restrictive to match key empirical moments.

5.3 An open economy with plant-specific realizations of $\tilde{\delta}(\omega)$

Allowing for stochastic, plant-specific realizations of $\tilde{\delta}(\omega)$ makes the model better suited for capturing the rich exporting patterns in our data. However, it also increases mathematical complexity and reduces analytical tractability. For this reason, we present the formal analysis of the extended model in the Supplemental Appendix and focus here on

discussing the key insights informally. An important difference between the canonical model with $\tilde{\delta}(\omega) = 1$ for all plants and the extended model with stochastic, plant-specific $\tilde{\delta}(\omega)$ lies in the entry behavior of firms. In the extended model, some plants with low augmented productivity may enter production solely because a favorable draw of $\tilde{\delta}(\omega)$ enables them to export. Consequently, exporters do not necessarily exhibit a stronger internal division of labor than non-exporters, as some exporters are relatively small compared to highly productive non-exporters with unfavorable draws of $\tilde{\delta}(\omega)$. In other words, under non-sharp selection, there is overlap between exporters and non-exporters in both the distributions of augmented productivity and worker efficiency.

Despite their asymmetry in firm entry, the two model variants still exhibit similar qualitative effects of trade on aggregate outcomes. One such similarity concerns the welfare implications of trade. In both models, gains from trade arise because the real wage and thus utilitarian welfare are determined by the revenues of the marginal non-exporter earning zero profits from domestic sales and therefore remain to be given by eq. (30). A second similarity concerns the role of the common export fixed cost parameter f_x in shaping economy-wide worker efficiency and wage dispersion. As formally shown in the Supplemental Appendix, these aggregates can be expressed in the extended model as:

$$\begin{aligned}\Lambda &= \left(\frac{\xi f_0}{\gamma \hat{\theta}_1^*} \right)^{\frac{1}{\gamma}} \frac{G\left(\hat{\theta}_1^*, \frac{f_x}{f_0}, 1 + \frac{1}{\gamma}\right)}{G\left(\hat{\theta}_1^*, \frac{f_x}{f_0}, 1\right)}, \\ CV &= \frac{\sqrt{4 - \pi(\pi - 2)}}{\pi} \left[1 - \eta \left(\frac{\gamma \hat{\theta}_1^*}{\xi w f_0} \right)^{\frac{1}{\gamma}} \frac{G\left(\hat{\theta}_1^*, \frac{f_x}{f_0}, 1 - \frac{1}{\gamma}\right)}{G\left(\hat{\theta}_1^*, \frac{f_x}{f_0}, 1\right)} \right],\end{aligned}\tag{33}$$

with $G(\hat{\theta}_1^*, f_x/f_0, \iota) \equiv \int_{\ln \delta_x}^{\infty} [\tilde{F}(z, \hat{\theta}_1^*, \iota) + (t^\iota - 1)\tilde{F}(z, \hat{\theta}_x^*(\exp[z]), \iota)] dz + \int_{-\infty}^{\ln \delta_x} t^\iota \tilde{F}(z, \hat{\theta}^* \exp[z], \iota) dz$. The effects of f_x on Λ and CV are non-monotonic, with both variables converging to their autarky levels as $f_x \rightarrow 0$ or $f_x \rightarrow \infty$. Furthermore, increasing f_x from a very low level raises both economy-wide worker efficiency and wage dispersion when $\eta > 0$. This confirms the key insight from the canonical model: worker efficiency and, provided $\eta > 0$, also wage dispersion attain an interior maximum at some $f_x > 0$.

6 Structural Estimation and Quantification of the Model

We revisit the theoretical framework outlined in Sections 3 and 5 for the more elaborate model variant, in which each plant is characterized by a triple of stochastic determinants: $\{\tilde{\theta}(\omega), \tilde{\zeta}(\omega), \tilde{\delta}(\omega)\}$. The plant's profit-maximizing decisions yield a system of estimable

equations that link these stochastic variables to realizations observed in our combined dataset, which merges information from the BIBB-BAuA labor force surveys with the linked plant–worker data in LIAB. One possible representation of the model is a three-equation system involving the log of plant revenues, $y(\omega) \equiv \ln r(\omega)$, the coefficient of variation of daily wage residuals within plant-occupations combined with the normalized number of tasks per occupation, $o(\omega) \equiv \ln[CV(\omega)b(\omega)/z(\omega)]$, and an export indicator, $I(\omega) \equiv \mathbf{1}_x(\omega)$ as well as the jointly normally distributed disturbances $u(\omega) = \xi \ln \tilde{\theta}(\omega)$, $v(\omega) = \ln \tilde{\zeta}(\omega)$, $z(\omega) = \ln \tilde{\delta}(\omega)$, where $u(\omega)$ is truncated from below at $\ln \hat{\theta}_0^*$. To simplify notation, we drop the plant identifier and represent the equation system as

$$y = \begin{cases} \mu_Y^e + u + v & \text{if } I = 1 \\ \mu_Y^d + u + v & \text{if } I = 0 \end{cases}, \quad (34a)$$

$$o = \begin{cases} \mu_O^e - (1/\gamma)u & \text{if } I = 1 \\ \mu_O^d - (1/\gamma)u & \text{if } I = 0 \end{cases}, \quad (34b)$$

$$I = \begin{cases} 1 & \text{if } \mu_X + u - z \geq 0 \\ 0 & \text{if } \mu_X + u - z < 0 \end{cases}, \quad (34c)$$

$$y, o, I = \text{missing} \quad \text{if } u < \ln \hat{\theta}_0^*, \quad (34d)$$

Eq. (34a) follows from eqs. (15), (16), and (24) and the definitions $\mu_Y^d \equiv \ln[w\sigma\xi f_0/(\sigma - 1)] - \ln \hat{\theta}_1^*$ and $\mu_Y^e \equiv \mu_Y^d + \ln t$. Eq. (34b) follows from eqs. (2), (10), and (14) with $\mu_O^d \equiv (1/2) \ln[4 - \pi(\pi - 2)] + (1/\gamma) \ln \gamma + (1/\gamma) \ln[(\sigma - 1)/\sigma] - (1/\gamma)\mu_Y^d$ and $\mu_O^e \equiv \mu_O^d - (1/\gamma)(\mu_Y^e - \mu_Y^d)$. Eq. (34c) follows from eq. (27) with $\mu_X \equiv \ln(t - 1) + \ln f_0 - \ln \hat{\theta}_1^* - \ln f_x$. Eq. (34d) gives a truncation condition, indicating that plants with productivity draws lower than $\hat{\theta}_0^*$ are inactive.¹¹

We can estimate equation system (34) using a maximum-likelihood (ML) estimator. Allowing for truncation by augmented productivity (censoring) requires that we use the observed maximum of composite variable $\ln[cv(\omega)b(\omega)/\tilde{\beta}(\omega)]$ in the ML estimation to recover the truncation point for augmented productivity $\hat{\theta}_0^*$. Given the parameter estimates for first and second moments related to composite variable $\ln[cv(\omega)b(\omega)/\tilde{\beta}(\omega)]$ and productivity $\xi \ln \tilde{\theta}(\omega)$, we can infer from the observed maximum of $\ln[cv(\omega)b(\omega)/\tilde{\beta}(\omega)]$ the internally consistent cutoff of augmented productivity $\ln \hat{\theta}_0^*$. We thereby make use of

¹¹Equation system (34) omits time indices because we estimate the parameters using a cross-section of data from 2006. Additionally, domestic and foreign market sizes do not appear due to the assumption of symmetric countries in Section 5. An alternative specification with asymmetric countries would imply that model parameters reflect relative market sizes in addition to trade costs τ .

the important insight that for unconstrained $\tilde{\delta}(\omega)$ the least productive firm is an exporter, which is the case for 2006, the year we use to estimate the model parameters.¹²

6.1 Implementation of estimation model

With the ML estimation of equation system (34) we aim to determine 14 parameters of our theoretical model. These include the means $\mu_Y^e, \mu_Y^d, \mu_O^e, \mu_O^d, \mu_x$, the second moments of the three stochastic variables $\sigma_u, \sigma_v, \sigma_z, \rho_{uv}, \rho_{uz}, \rho_{vz}$, the two fundamental model parameters σ, γ , and the truncation point $\hat{\theta}_0^*$. As pointed out by Maddala (1986), one of the variance parameters remains undetermined by our estimator. We use this insight and set the variance of the combined stochastic term $u - z$ equal to one. This conditions σ_z on other model parameters by $\sigma_u \sqrt{1 - 2\rho_{uz}\sigma_z/\sigma_u + \sigma_z^2/\sigma_u^2} = 1$.

We derive the likelihood functions for our estimator in the Supplemental Appendix and report, in Table 1, the result of the ML estimation for the fourteen underlying model parameters, with standard errors computed using the Delta method. All of the estimated parameters are highly significant and of reasonable size. For instance, the reported value for demand elasticity σ is with a value of 5.67 in the range of parameter estimates reported by previous research (see Broda and Weinstein 2006). Combining our parameter estimates, we can determine a theory-consistent iceberg trade cost parameter of $\tau = 1.24$, which is at the lower bound of estimates reported by Novy (2013). The estimated truncation cutoff is low and implies that only a small fraction of plants fails to enter.

The parameter estimates reported in Table 1 provide a crucial input for quantifying our model. However, they are not sufficient to fully exploit the general equilibrium structure, which accounts for firm entry and exit in both domestic and foreign markets. To address this, we take an additional step and use the estimates from Table 1 to compute theory-consistent values for the (deterministic) operational and export fixed costs: $\ln(wf_0) = \mu_Y^d + \xi \ln \theta_1^* - \ln[\sigma\xi/(\sigma - 1)]$ and $\ln(wf_x) = \ln\{\exp[\mu_Y^e - \mu_Y^d] - 1\} + \mu_Y^d - \mu_x - \ln[\sigma\xi/(\sigma - 1)]$. Together with the free entry condition in eq. (28), we then derive the fixed costs of entering the technology lottery, wf_e , using the Gauss-Kronrod quadrature formula as a numerical integration technique. Our structural estimation procedure does not separately identify worker efficiency and wage dispersion, as it does not discipline the mismatch sensitivity parameter η . To gain further insight into the role of trade for worker efficiency and wage dispersion, we next determine a theory-consistent value for η using a method-of-moments estimator. Specifically, we equate the model-implied values of

¹²We also use observed export shares for 2006 to discipline the trade cost elasticity $\sigma - 1$ by our data (see the Supplemental Appendix for further details).

Table 1: Maximum Likelihood Parameter Estimates

		Model parameters				
γ	σ	$\ln \hat{\theta}_0^*$				
6.924*** (0.080)	5.666*** (0.271)	-14.802*** (0.112)				
		Distributional parameters				
<i>First moments:</i>		μ_Y^e	μ_Y^d	μ_O^e	μ_O^d	μ_X
14.453*** (0.103)	13.502*** (0.031)	-2.221*** (0.029)	-2.084*** (0.021)	-1.007*** (0.032)		
<i>Higher moments:</i>		σ_u	σ_v	σ_z	ρ_{uv}	ρ_{uz}
5.924*** (0.172)	6.025*** (0.165)	6.105*** (0.167)	-0.975*** (0.001)	0.987*** (0.001)	-0.967*** (0.003)	
log Pseudo-Likelihood		-2,140,964.979				
Observations		7,259				

Source: LIAB and BIBB-BAuA, 2006. Plants with at least two full-time workers, excluding marginal workers with a daily wage below a minimum threshold (see Stüber, Dauth, and Eppelsheimer 2023) and excluding workers with imputed daily wages above EUR 2,000 in 1998 prices. Plant observations are weighted by sampling frequencies.

Notes: To account for parameter constraints imposed by our model, we estimate μ_W^d and μ_X in levels, $\mu_W^e - \mu_W^d$, γ , σ_u , and σ_v in logs, and correlation coefficients as log-transformed variables $\log[(1+x)/(1-x)]$. The log Pseudo-Likelihood refers to the estimation of these transformed parameters. Fundamental parameters are inferred using the functional relationships outlined in the text. Standard errors in parentheses are computed with the Delta method. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

economy-wide Λ and CV from eq. (33) with their empirical counterparts observed in the 2006 data. These two equations jointly identify the previously undetermined mismatch sensitivity parameter η and the composite parameter $(\xi f_0/\gamma)^{1/\gamma}$.¹³

In the next subsection, we use our parameter estimates for a counterfactual analysis, in which we change the fixed export cost parameter, f_x , and analyze the impact of this change on economy-wide worker efficiency and wage dispersion. In this context, we are interested in answering three specific questions: (i) To what extent is our quantitative model able to explain observed patterns of worker efficiency and wage dispersion in observation years different from 2006? (ii) How strong is the impact of varying fixed costs of exporting on worker efficiency and wage dispersion? (iii) To what extent do the quantitative effects of trade on worker efficiency and wage dispersion differ between

¹³We can set the parameter composite independently of the estimates in Table 1, as the latter pin down overhead fixed costs in 2006 prices, while Λ and CV are invariant to the specific choice of base year.

the elaborate model with stochastic, plant-specific realizations of $\tilde{\delta}(\omega)$ and the canonical model with a common $\tilde{\delta}(\omega) = 1$ for all plants?

6.2 The impact of f_x on worker efficiency and wage dispersion

For the quantification of our model, we first assess its ability to replicate observed patterns of worker efficiency and wage dispersion in sample years not used for the structural parameter estimation described above. For this purpose, we employ the observed exporter shares in these years and equate them to the theoretical solution for the exporter share, χ , which, as formally derived in the Supplemental Appendix, can be expressed as

$$\chi = \frac{\int_{\ln \delta_x}^{\infty} \hat{F}\left(z, \hat{\theta}_x^*(\exp[z])\right) dz + \int_{-\infty}^{\ln \delta_x} \hat{F}\left(z, \hat{\theta}^*(\exp[z])\right) dz}{\int_{\ln \delta_x}^{\infty} \hat{F}\left(z, \hat{\theta}_1^*\right) dz + \int_{-\infty}^{\ln \delta_x} \hat{F}\left(z, \hat{\theta}^*(\exp[z])\right) dz}, \quad (35)$$

where $\hat{F}(z, x) \equiv \frac{1}{\sigma_z} \phi\left(\frac{z}{\sigma_z}\right) \Phi\left(-\frac{\ln x - z \rho_{uz} \sigma_u / \sigma_z}{\sigma_u \sqrt{1 - \rho_{uz}^2}}\right)$. This procedure allows us to solve for theory-consistent, time-varying values of the fixed export cost parameter f_x , which we then use alongside the other time-invariant parameter estimates to predict worker efficiency and wage dispersion for the four observation years in which we can link the BIBB-BAuA and LIAB datasets. In Table 2, we report predicted values relative to their observed counterparts. Since worker efficiency is not directly observable in our data, we report the transformed measure $\Lambda/\pi - \eta/\pi$ instead of Λ . We also include results for the year 2006, which was used to calibrate our parameter estimates. As expected, the model fits the data perfectly in this case, confirming the validity of our numerical approach. For the other three sample years, the calibrated model performs reasonably well in capturing both worker efficiency and wage dispersion. However, in 1999, we underestimate worker efficiency by more than 20 percent. This discrepancy arises because, as reported in footnote 3, there was a substantial shift toward multitasking between 1999 and 2006 that was not matched by a corresponding change in the exporter share.

In a next step, we use our calibrated model to quantify the effects of changes in the export cost parameter f_x on worker efficiency and wage dispersion. To do so, we vary the log of fixed export costs over the interval $(-30, 30)$ in small increments, compute the corresponding values for worker efficiency Λ , wage dispersion CV , and the exporter share χ using eqs. (33) and (35), and present the results of this counterfactual analysis in Figure 3 (solid black lines). Figure 3 illustrates the key insight from Section 5: the effects of trade on labor division and wage dispersion are non-monotonic and most pronounced

Table 2: Model fit over four observation years

	1999	2006	2012	2018
Transformed worker efficiency $\Lambda/\pi - \eta/\pi$	0.794	1.000	1.016	1.049
Wage dispersion CV	1.002	1.000	0.996	0.928

Source: LIAB and BIBB-BAuA, 1999, 2006, 2012 and 2018. Plants with at least 2 full-time workers, excluding marginal workers with a daily wage below a minimum threshold (see Stüber, Dauth, and Epelsheimer 2023) and excluding workers with imputed daily wages above EUR 2,000 in 1998 prices. Plant observations are weighted by sampling frequencies.

Notes: Calibration of sensitivity to mismatch parameter η and composite parameter $(\xi f_0/\gamma)^{1/\gamma}$ match average coefficient of variation of residual wages and average worker efficiency observed for 2006.


at intermediate levels of χ . Panel A shows that our model predicts a substantial impact of trade on worker efficiency, reaching more than six percent at intermediate exporter shares. This effect is sizable relative to the overall gains from trade, which peak at less than seven percent when $\chi = 1$. In contrast, Panel B reveals that the effects of trade on wage dispersion are modest.

In Figure 3, we also illustrate the effects of varying export fixed costs in the canonical model, which assumes a common $\tilde{\delta}(\omega) = 1$ for all plants (dashed gray lines).¹⁴ Comparing the canonical model to the more elaborate version with plant-specific $\tilde{\delta}(\omega)$, we find that the former significantly overstates the effects of trade on worker efficiency at low levels of χ (Panel A). This occurs because the sharp selection mechanism, where only highly productive plants export, eliminates a counteracting effect in the extended model, if low-productivity plants export due to a favorable draw of $\tilde{\delta}(\omega)$. Intuitively, this difference is most pronounced at low exporter shares, where the assumption of sharp selection is particularly restrictive. Similarly, Panel B shows that the canonical model also exaggerates the impact of trade on wage dispersion at low levels of χ . However, even in the canonical model, the overall effect of changing f_x on CV remains moderate.¹⁵

¹⁴For the counterfactual analysis in the canonical model, we rely on the parameter estimates reported in Table 1, except for the augmented productivity threshold $\hat{\theta}_1^*$, which we adjust to align the model-predicted exporter share with the observed share in the combined dataset for 2006.

¹⁵To assess the role of endogenous adjustments in plant-internal labor division, we also compare the welfare effects from our model to those from an alternative framework that excludes feedback effects of trade through the internal reassignment of workers to tasks by setting $\nu = 0$. Evaluated at estimated parameter values, we find that ignoring this adjustment margin overstates the gains from trade considerably. For the observed exporter share in 2006, our model predicts moderate welfare gains of 0.19 percent relative to autarky. In contrast, the alternative model lacking the feedback effect on internal labor division implies a more sizable welfare gain of 2.76 percent. This difference arises because a reduced division of labor in non-exporting plants dampens aggregate welfare gains in our more comprehensive framework.

Figure 3: Counterfactual analysis



Sources: LIAB and BIBB-BAuA 2006.

Notes: Counterfactual analysis of changing fixed exporting costs on worker efficiency and the coefficient of variation of residual wages.

7 Concluding Remarks

We document empirically that workers in larger plants perform fewer tasks and that the resulting specialization of workers on fewer tasks is inherently linked to higher wage dispersion within plant-occupations. Based on these observations, we build a model of the internal labor market, where the employer chooses the division of labor by assigning task ranges to occupations, workers of different ability match to occupations and the match quality determines the wage dispersion within plant-occupations. We embed this rationale into a heterogeneous-firm model of trade to relate global product-market conditions to the employer's optimal choice of the internal division of labor. A plant that commands a larger market share can achieve greater worker efficiency by incurring higher fixed costs to narrow the range of tasks performed per occupation and simultaneously raising the count of occupations to which it assigns tasks. In equilibrium, a priori more productive plants and exporters adopt a stricter division of labor and thus increase their productivity.

We use German plant-worker data, combined with detailed German survey information on time-varying tasks performed by workers within their occupations, to structurally estimate key model parameters and to quantify the impact of trade liberalization on economy-wide worker efficiency and wage dispersion. Our results indicate a non-monotonic effect on both of these variables. We find the effect of trade liberalization on worker efficiency to be potentially important, while the impact on wage dispersion is modest.

Our framework isolates the within-plant and within-occupation changes that globalization induces. Beyond identifying an important feedback effect of trade liberalization

on worker productivity through finer divisions of labor and lower task mismatch in a setting with multitasking, our model speaks to a dominant, so far largely unexplored, part of residual wage inequality that materializes within plant-occupations. This within-plant perspective complements the between-plant perspective that is dominant in the literature and emphasizes the reallocation of labor between employers as a key mechanism explaining positive effects of trade on aggregate productivity with possibly detrimental distributional effects in the presence of labour market distortions. A joint consideration of within- and between-plant allocation mechanisms is left for future research.

References

- Abowd, John M., Francis Kramarz, David N. Margolis, and Kenneth R. Troske. 2001. "The Relative Importance of Employer and Employee Effects on Compensation: A Comparison of France and the United States." *Journal of the Japanese and International Economies* 15 (4):419–436.
- Acemoglu, Daron and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of labor economics*, vol. 4B, edited by David Card and Orley Ashenfelter, chap. 12. Amsterdam, New York and Oxford: Elsevier Science, North-Holland, 1043–1171.
- Amiti, Mary and Christopher A. Pissarides. 2005. "Trade and Industrial Location with Heterogeneous Labor." *Journal of International Economics* 67 (2):392–412.
- Armenter, Roc and Miklos Koren. 2015. "Economies of Scale and the Size of Exporters." *Journal of the European Economic Association* 13 (3):482–511.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review: Papers and Proceedings* 96 (2):189–194.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118 (4):1279–1333.
- Aw, Bee Yan, Mark J. Roberts, and Daniel Yi Xu. 2011. "R&D Investment, Exporting, and Productivity Dynamics." *American Economic Review* 101 (4):1312–1344.

- Becker, Gary S. and Kevin M. Murphy. 1992. "The Division of Labor, Coordination Costs, and Knowledge." *Quarterly Journal of Economics* 107 (4):1137–1160.
- Becker, Sascha O. and Marc-Andreas Muendler. 2015. "Trade and Tasks: An Exploration Over Three Decades in Germany." *Economic Policy* 30 (84):589–641.
- Blinder, Alan S. 2006. "Offshoring: The Next Industrial Revolution?" *Foreign Affairs* 85 (2):113–28.
- Bloom, Nicholas and John Van Reenen. 2011. "Human Resource Management and Productivity." In *Handbook of Labor Economics, Handbooks in Economics*, vol. 4B, edited by Orley Ashenfelter and David Card, chap. 19. Amsterdam: Elsevier, 1697–1767.
- Bombardini, Matilde, Gianluca Orefice, and Maria D. Tito. 2019. "Does Exporting Improve Matching? Evidence from French Employer-Employee Data." *Journal of International Economics* 117:229–241.
- Broda, Christian and David E. Weinstein. 2006. "Globalization and the Gains from Variety." *Quarterly Journal of Economics* 121 (2):541–85.
- Caliendo, Lorenzo, Ferdinando Monte, and Esteban Rossi-Hansberg. 2015. "The Anatomy of French Production Hierarchies." *Journal of Political Economy* 123 (4):809–52.
- Caliendo, Lorenzo and Esteban Rossi-Hansberg. 2012. "The Impact of Trade on Organization and Productivity." *Quarterly Journal of Economics* 127 (3):1393–1467.
- Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128 (3):967–1015.
- Carson, Richard T. and Yixiao Sun. 2007. "The Tobit Model with a Non-zero Threshold." *Econometrics Journal* 10 (3):488–502.
- Chaney, Thomas. 2008. "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *American Economic Review* 98 (4):1707–21.
- Chaney, Thomas and Ralph Ossa. 2013. "Market Size, Division of Labor, and Firm Productivity." *Journal of International Economics* 90 (1):177–80.

- Clerides, Sofronis K., Saul Lach, and James R. Tybout. 1998. "Is Learning by Exporting Important? Micro-dynamic Evidence from Colombia, Mexico, and Morocco." *Quarterly Journal of Economics* 113 (3):903–47.
- Costinot, Arnaud and Jonathan Vogel. 2010. "Matching and Inequality in the World Economy." *Journal of Political Economy* 118 (4):747–86.
- Crespi, Gustavo, Chiara Criscuolo, and Jonathan Haskel. 2008. "Productivity, Exporting and the Learning-by-Exporting Hypothesis: Direct Evidence from U.K. Firms." *Canadian Journal of Economics* 41 (2):619–638.
- Davidson, Carl, Steven J. Matusz, and Andrei Shevchenko. 2008. "Globalization and Firm Level Adjustment with Imperfect Labor Markets." *Journal of International Economics* 75 (2):295–309.
- Dhingra, Swati and John Morrow. 2019. "Monopolistic Competition and Optimum Product Diversity under Firm Heterogeneity." *Journal of Political Economy* 127 (1):196–232.
- Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg. 2009. "Revisiting the German Wage Structure." *Quarterly Journal of Economics* 124 (2):843–881.
- Egger, Hartmut, Udo Kreickemeier, Christoph Moser, and Jens Wrona. 2022. "Exporting and Offshoring with Monopsonistic Competition." *Economic Journal* 132 (644):1449–1488.
- Fernandes, Ana M., Peter J. Klenow, Sergii Meleshchuk, Martha Denisse Pierola, and Andrés Rodríguez-Clare. 2023. "The Intensive Margin in Trade: How Big and How Important?" *American Economic Journal: Macroeconomics* 15 (3):320–54.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2009. "Job Polarization in Europe." *American Economic Review* 99 (2):58–63.
- Grossman, Gene M., Elhanan Helpman, and Philipp Kircher. 2017. "Matching, Sorting, and the Distributional Effects of International Trade." *Journal of Political Economy* 125 (1):224–264.
- Grossman, Gene M. and Esteban Rossi-Hansberg. 2008. "Trading Tasks: A Simple Theory of Offshoring." *American Economic Review* 98 (5):1978–97.

- . 2010. “External Economies and International Trade Redux.” *Quarterly Journal of Economics* 125 (2):829–58.
- Helpman, Elhanan, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen J. Redding. 2017. “Trade and Inequality: From Theory to Estimation.” *Review of Economic Studies* 84 (1):357–405.
- Helpman, Elhanan, Oleg Itskhoki, and Stephen J. Redding. 2010. “Inequality and Unemployment in a Global Economy.” *Econometrica* 78 (4):1239–1283.
- Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg. 2021. “Learning from Coworkers.” *Econometrica* 89 (2):647–767.
- Jha, Pryaranjan and Antonio Rodriguez-Lopez. 2021. “Monopsonistic Labor Markets and International Trade.” *European Economic Review* 140:103939.
- Kambourov, Gueorgui and Iourii Manovskii. 2009. “Occupational Specificity of Human Capital.” *International Economic Review* 50 (1):63–115.
- Lemieux, Thomas. 2006. “Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?” *American Economic Review* 96 (3):461–498.
- Levy, Frank and Richard J. Murnane. 2004. *The New Division of Labor*. Princeton: Princeton University Press.
- Maddala, G. S. 1986. “Disequilibrium, Self-selection, and Switching Models.” In *Handbook of Econometrics, Handbooks in Economics*, vol. 3, edited by Zvi Griliches and Michael D. Intriligator, chap. 28. Amsterdam, Oxford and Tokyo: North-Holland, 1633–88.
- Melitz, Marc J. 2003. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” *Econometrica* 71 (6):1695–1725.
- Novy, Dennis. 2013. “Gravity Redux: Measuring International Trade Costs with Panel Data.” *Economic Inquiry* 51 (1):101–21.
- Ohnsorge, Franziska and Daniel Treffer. 2007. “Sorting It Out: International Trade with Heterogeneous Workers.” *Journal of Political Economy* 115 (5):868–92.
- Rodríguez-Clare, Andrés. 2010. “Offshoring in a Ricardian World.” *American Economic Journal: Macroeconomics* 2 (2):227–58.

- Sampson, Thomas. 2014. “Selection into Trade and Wage Inequality.” *American Economic Journal: Microeconomics* 6 (3):157–202.
- Sattinger, Michael. 1993. “Assignment Models of the Distribution of Earnings.” *Journal of Economic Literature* 31 (2):831–880.
- Smith, Adam. 1776. *An inquiry into the nature and causes of the wealth of nations*. London: Printed for W. Strahan and T. Cadell. Reprinted 1966 by A. M. Kelley, New York.
- Stole, Lars A. and Jeffrey Zwiebel. 1996. “Intra-firm Bargaining under Non-binding Contracts.” *Review of Economic Studies* 63 (3):375–410.
- Stüber, Heiko, Wolfgang Dauth, and Johann Eppelsheimer. 2023. “A Guide to Preparing the Sample of Integrated Labour Market Biographies (SIAB, version 7519 v1) for Scientific Analysis.” *Journal for Labour Market Research* 57 (7):1–11.

Appendix

A Empirical Appendix

A.1 Decomposition of wage variation in linked plant-worker data

Using our linked plant–worker data, we can shed light on key factors explaining wage dispersion in the German labor market. We focus on the year 2006 and eliminate observed demographic, education and tenure information together with region effects from log daily wages in a Mincer regression to compute residual log daily wages. About one-half of the wage dispersion remains unexplained by this procedure. We next decompose the variance of log residual daily wages and summarize the main insights from this analysis in Column 1 of Table A1.

In a first step, we eliminate residual wage differences between industries, which reduces the wage variance by 14 percent. Controlling for wage differences between occupations further reduces wage variance by 16 percent, while eliminating variation between plants still leaves more than 60 percent of the residual wage variance unexplained. Controlling for wage differences between hierarchies within plants reduces the unexplained part of residual wage dispersion only to a small extent, while additionally controlling for differences of wages between plant-occupations leaves 47 percent of the residual wage variation unexplained. Hence, one quarter of the total observed wage variance is within

Table A1: Employer-Employee Data: Residual Log Wage Inequality 2006

Within component (%)	Plants			Firms			
	DE	SE	BR	SE	BR	DK	FR
industry	86	97	95	96	95	95	92
industry-occupation	70	91	79	90	79	81	60
employer	63	75	51	77	58	74	58
employer-layer	59	69	43	72	50	61	35
employer-occupation	47	60	32	63	36	48	26

Source: LIAB Germany (DE), Sweden (SE), RAIS Brazil (BR), Denmark (DK), France (FR) 2006.

Notes: Wage variance is decomposed into a within and a between component for groups g using

$$(1/L) \sum_{i=1}^L (\ln w_i - \overline{\ln w})^2 = (1/L) \sum_{g \in \mathbb{G}} \sum_{i=1}^{L_g} (\ln w_i - \overline{\ln w_g})^2 + \sum_{g \in \mathbb{G}} (L_g/L) (\overline{\ln w_g} - \overline{\ln w})^2.$$

The reported figures show shares of within components for varying worker groups in percent. Residual log daily wage from standard Mincer regression, conditioning on demographics, education and tenure as well as region effects (excluding industry effects). Industry aggregates: 38 from NACE, 90 SNI, 60 CNAE, 82 Branchekode, 82 Activit  Principale de l'Entreprise; occupations: 357 KldB-88, 186 SISCO, 348 CBO, 205 DISCO, 427 PSE-ESE; regions: 16 German federal states, 21 Swedish regions, 26 Brazilian states, 11 Danish, 35 French regions; education groups: 6 Germany, 6 Sweden, 4 Brazil, 5 Denmark, none in France; hierarchies: 4 layers based on Caliendo, Monte, and Rossi-Hansberg (2015).

plant-occupations. Table A1 shows that the wage dispersion inside employer-occupations also explains a large part of overall residual wage dispersion in other countries.¹⁶

B Theoretical Appendix

B.1 Derivation and discussion of productivity threshold $\hat{\theta}_a^*$

We first note that

$$\begin{aligned} & \frac{1}{2\pi\sigma_u\sigma_v\sqrt{1-\rho_{uv}^2}} \exp \left\{ -\frac{1}{2(1-\rho_{uv}^2)} \left[\left(\frac{u}{\sigma_u} \right)^2 - 2\frac{\rho_{uv}uv}{\sigma_u\sigma_v} + \left(\frac{v}{\sigma_v} \right)^2 \right] \right\} dv du \\ &= \frac{1}{\sqrt{2\pi}\sigma_u} \exp \left[-\frac{1}{2} \left(\frac{u}{\sigma_u} \right)^2 \right] \frac{1}{\sqrt{2\pi}\tilde{\sigma}_v} \exp \left[-\frac{1}{2} \left(\frac{v-\mu_v}{\tilde{\sigma}_v} \right)^2 \right] dv du, \end{aligned}$$

with $\mu_v \equiv u\rho_{uv}\sigma_v/\sigma_u$, $\tilde{\sigma}_v \equiv \sigma_v\sqrt{1-\rho_{uv}^2}$. Making use of eqs. (15) and (16) and acknowledging $u = \xi \ln \tilde{\theta}(\omega)$, $v = \ln \tilde{\zeta}(\omega)$, we find that $[(\sigma-1)/(\sigma\xi)]\bar{r}(\exp[u]|\exp[v]) =$

¹⁶We would like to thank Anders  kerman and L a Marshal for conducting the wage decomposition on Swedish and French data.

$wf_0 \exp[u]/\hat{\theta}_a^*$. This allows us to express the free entry condition under autarky as

$$f_0 \int_{\ln \hat{\theta}_a^*}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_u} \left[\frac{\exp[u]}{\hat{\theta}_a^*} - 1 \right] \exp \left[-\frac{1}{2} \left(\frac{u}{\sigma_u} \right)^2 \right] \\ \times \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\tilde{\sigma}_v} \exp[v] \exp \left[-\frac{1}{2} \left(\frac{v - \mu_v}{\tilde{\sigma}_v} \right)^2 \right] dv du = f_e,$$

which, solving the integrals, establishes eq. (19).

Moreover, differentiating the left-hand side of eq. (19) with respect to productivity threshold $\hat{\theta}_a^*$ provides $dLHS_{eq. (19)}/d\hat{\theta}_a^* = -(1/\hat{\theta}_a^*)^2 F(\hat{\theta}_a^*, 1)wf_0 < 0$, with $F(x, \iota)$ given in the main text. Noting further that $\lim_{\hat{\theta}_a^* \rightarrow 0} LHS_{eq. (19)} = \infty$ and $\lim_{\hat{\theta}_a^* \rightarrow \infty} LHS_{eq. (19)} = 0$, it follows from the intermediate value theorem that $LHS_{eq. (19)} = f_e$ has a solution in $\hat{\theta}_a^*$ on interval $(0, \infty)$. It follows from the monotonicity of $LHS_{eq. (19)}$ that the solution is unique. Finally, applying the implicit function theorem establishes $d\hat{\theta}_a^*/df_0 > 0$ and $d\hat{\theta}_a^*/df_e < 0$, which completes the proof.

B.2 Derivation and discussion of eqs. (21) and (22)

For a derivation of eq. (21), we can first note that $\lambda(\exp[u]) = (\exp[u]/\hat{\theta}_a^*)^{1/\gamma} (\xi f_0/\gamma)^{1/\gamma}$ follows from the main text of section 4.2 and the definition of $u = \xi \ln \tilde{\theta}(\omega)$. Moreover, acknowledging that ξf_0 is the employment of the least productive producer according to the constant markup pricing rule in eq. (12) and the zero-cutoff profit condition in eq. (16), while the employment of a firm with augmented productivity $\tilde{\theta}$ relative to the least productive plant follows as $\exp[u]/(\xi f_0)$ from eq. (15), we can compute the employment weighted average of endogenous worker efficiency according to

$$\Lambda_a = \frac{N_a \xi f_0 \hat{\theta}_a^*}{L_v^a} \left(\frac{\xi f_0}{\gamma \hat{\theta}_a^*} \right)^{\frac{1}{\gamma}} \int_{\ln \hat{\theta}_a^*}^{\infty} \exp \left[\left(1 + \frac{1}{\gamma} \right) u \right] \frac{1}{\sqrt{2\pi}\sigma_u} \exp \left[-\frac{1}{2} \left(\frac{u}{\sigma_u} \right)^2 \right] \\ \times \int_{-\infty}^{\infty} \exp[v] \frac{1}{\sqrt{2\pi}\tilde{\sigma}_v} \exp \left[-\frac{1}{2} \left(\frac{v - \mu_v}{\tilde{\sigma}_v} \right)^2 \right] dv du,$$

with L_v^a as economy-wide employment of labor in its use as a variable production input under autarky. We compute $L_v^a = N_a \xi f_0 \hat{\theta}_a^* F(\hat{\theta}_a^*, 1)$. Moreover, acknowledging

$$\int_{\ln \hat{\theta}_a^*}^{\infty} \exp \left[\left(1 + \frac{1}{\gamma} \right) u \right] \frac{1}{\sqrt{2\pi}\sigma_u} \exp \left[-\frac{1}{2} \left(\frac{u}{\sigma_u} \right)^2 \right]$$

$$\times \int_{-\infty}^{\infty} \exp[v] \frac{1}{\sqrt{2\pi}\tilde{\sigma}_v} \exp\left[-\frac{1}{2}\left(\frac{v-\mu_v}{\tilde{\sigma}_v}\right)^2\right] dv du = F(\hat{\theta}_a^*, 1 + 1/\gamma)$$

establishes eq. (21).

Second, we provide derivation details for eq. (21). We can make use of $cv(\exp[u]) = \sqrt{4 - \pi(\pi - 2)}[1 - \eta(\hat{\theta}/\hat{\theta}_a^*)^{-1/\gamma}(\xi f_0/\gamma)^{-1/\gamma}]/\pi$ from the main text of section 4.2 and the definition of u , whose employment-weighted average can be computed according to

$$CV_a = \frac{N_a \xi f_0 \hat{\theta}_a^*}{L_v^a} \int_{\ln \hat{\theta}_a^*}^{\infty} cv(\exp[u]) \exp[u] \frac{1}{\sqrt{2\pi}\sigma_u} \exp\left[-\frac{1}{2}\left(\frac{u}{\sigma_u}\right)^2\right] \\ \times \int_{-\infty}^{\infty} \exp[v] \frac{1}{\sqrt{2\pi}\tilde{\sigma}_v} \exp\left[-\frac{1}{2}\left(\frac{v-\mu_v}{\tilde{\sigma}_v}\right)^2\right] dv du.$$

Making use of

$$\int_{\ln \hat{\theta}_a^*}^{\infty} \exp\left[\left(1 - \frac{1}{\gamma}\right)u\right] \frac{1}{\sqrt{2\pi}\sigma_u} \exp\left[-\frac{1}{2}\left(\frac{u}{\sigma_u}\right)^2\right] \\ \int_{-\infty}^{\infty} \exp[v] \frac{1}{\sqrt{2\pi}\tilde{\sigma}_v} \exp\left[-\frac{1}{2}\left(\frac{v-\mu_v}{\tilde{\sigma}_v}\right)^2\right] dv du = F(\hat{\theta}_a^*, 1 - 1/\gamma),$$

and the solution for L_v^a from above, we obtain eq. (22). This completes the proof.

B.3 Derivation details for eq. (29)

Considering a parameter domain with $f_x > f_0(t - 1)$ and using the revenues of the marginal firm with augmented productivity $\hat{\theta}_1^*$, we compute for economy-wide revenues

$$R = N \frac{\sigma \xi}{\sigma - 1} \frac{w f_0}{\hat{\theta}_1^*} \int_{\ln \hat{\theta}_1^*}^{\ln \hat{\theta}_x^*} \frac{1}{\sqrt{2\pi}\sigma_u} \exp[u] \exp\left[-\frac{1}{2}\left(\frac{u}{\sigma_u}\right)^2\right] \\ \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\tilde{\sigma}_v} \exp[v] \exp\left[-\frac{1}{2}\left(\frac{v-\mu_v}{\tilde{\sigma}_v}\right)^2\right] dv du \\ + N t \frac{\sigma \xi}{\sigma - 1} \frac{w f_0}{\hat{\theta}_1^*} \int_{\ln \hat{\theta}_x^*}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_u} \exp[u] \exp\left[-\frac{1}{2}\left(\frac{u}{\sigma_u}\right)^2\right] \\ \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\tilde{\sigma}_v} \exp[v] \exp\left[-\frac{1}{2}\left(\frac{v-\mu_v}{\tilde{\sigma}_v}\right)^2\right] dv du,$$

where the first line represents total revenues generated by non-exporters, while the second line represents total revenues of exporters. Solving the integrals, gives $R = N[w\sigma\xi/(\sigma -$

$$1)] \left[(f_0/\hat{\theta}_1^*)F(\hat{\theta}_1^*) - (f_x)\hat{\theta}_x^*)F(\hat{\theta}_x^*, 1) \right].$$

We next compute aggregate fixed costs of operation, $w\hat{\zeta}(\omega)f_0$, and exporting, $w\hat{\zeta}(\omega)f_x$, as $wf_0F(\hat{\theta}_1^*, 0) + wf_xF(\hat{\theta}_x^*, 0)$. Noting that profits of non-exporters and exporters are given by $\psi_T^d(\omega) = [(\sigma - 1)/(\sigma_x)]r^d(\omega) - w\hat{\zeta}(\omega)f_0$ and $\psi_T^e(\omega) = [(\sigma - 1)/(\sigma_x)]tr^d(\omega) - w\hat{\zeta}(\omega)f_0 - w\hat{\zeta}(\omega)f_x$, respectively, we can express the average expected profits of plants prior to their participation in the lottery as $\bar{\psi}_T = wf_0[(1/\hat{\theta}_1^*)F(\hat{\theta}_1^*, 1) - F(\hat{\theta}_1^*, 1)] + wf_x[(1/\hat{\theta}_x^*)F(\hat{\theta}_x^*, 1) - F(\hat{\theta}_x^*, 1)]$. Following similar derivation steps, we compute for a parameter domain with $f_x \leq f_0(t - 1)$ an expected profit equal to $\bar{\psi}_T = w(f_0 + f_x)[(1/\hat{\theta}_1^*)F(\hat{\theta}_1^*, 1) - F(\hat{\theta}_1^*, 0)]$. The free entry condition in eq. (29) equates $\bar{\psi}_T$ with the entry costs of the lottery wf_e . Existence and uniqueness of $\hat{\theta}_1^*$ then follows from noting $\lim_{\hat{\theta}_1^* \rightarrow 0} \bar{\psi}_T = \infty$, $\lim_{\hat{\theta}_1^* \rightarrow 0} \bar{\psi}_T = \infty$, and $d\bar{\psi}_T/d\hat{\theta}_1^* < 0$. Noting further that $\partial\bar{\psi}_T/\partial f_x < (>) 0$ if $f_x > (<) f_0(t - 1)$, the sign of $\hat{\theta}_1^*/df_x$ follows from the implicit function theorem. Finally, the free entry condition in the open economy coincides with the free entry condition in the closed economy in the two limiting cases of $f_x \rightarrow 0$ and $f_x \rightarrow \infty$. This completes the proof.

B.4 Firm-level effects of openness in the canonical model

We want to show that $\lambda^e(\hat{\theta})/\lambda^a(\hat{\theta}) = \left(t\hat{\theta}_a^*/\hat{\theta}_1^*\right)^{1/\gamma} > 1$ holds for all possible $f_x > f_0(t - 1)$. For this purpose, we note from Appendix B.3 that $t\hat{\theta}_a^*/\hat{\theta}_1^*$ has a minimum at $f_x = f_0(t - 1)$. Therefore, $\lambda^e(\hat{\theta})/\lambda^a(\hat{\theta})$ is larger than one for any $f_x > f_0(t - 1)$ if $\lambda^e(\hat{\theta})/\lambda^a(\hat{\theta}) \geq 1$ holds for $f_x = f_0(t - 1)$. Evaluating the free-entry condition in eq. (29) at $f_x = f_0(t - 1)$ gives $f_0G(\hat{\theta}_1^*, t) = f_e$, with $G(\hat{\theta}_1^*, t) \equiv t \left[(1/\hat{\theta}_1^*)F(\hat{\theta}_1^*, 1) - F(\hat{\theta}_1^*, 0) \right]$. We have $\partial G(\cdot)/\partial t > 0$ and $\partial G(\cdot)/\partial \hat{\theta}_1^* < 0$, while applying the implicit function theorem establishes $d\hat{\theta}_1^*/dt \times t/\hat{\theta}_1^* = 1 - F(\hat{\theta}_1^*, 0)/[(1/\hat{\theta}_1^*)F(\hat{\theta}_1^*, 1)] \in (0, 1)$. Since $f_0G(\hat{\theta}_1^*, t) = f_e$ has a solution at $\hat{\theta}_1^* = t\hat{\theta}_a^*$ if $t = 1$, it follows that $\hat{\theta}_1^* < t\hat{\theta}_a^*$ must hold for all $t > 1$, which is sufficient for $\lambda^e(\hat{\theta})/\lambda^a(\hat{\theta}) > 1$. This completes the proof.

B.5 Derivation and discussion of eqs. (31) and (32)

We first sketch how to derive eq. (21) and begin with a parameter domain of $f_x > f_0(t - 1)$. In this case, we have $\lambda^d(\exp[u]) = \left(\frac{\exp[u]}{\hat{\theta}_1^*}\right)^{1/\gamma} \left(\frac{\xi f_0}{\gamma}\right)^{1/\gamma}$ and $\lambda^e(\exp[u]) = \left(\frac{t \exp[u]}{\hat{\theta}_1^*}\right)^{1/\gamma} \left(\frac{\xi f_0}{\gamma}\right)^{1/\gamma}$, where $u = \xi \ln \tilde{\theta}(\omega)$ has been used. Moreover, acknowledging that ξf_0 is the employment of the least productive producer according to the constant markup pricing rule in eq. (12) and the zero-cutoff profit condition in eq. (16), the employment of

a non-exporter with augmented productivity $\hat{\theta}(\omega) = \exp[u]$ relative to the least productive plant follows as $\exp[u]/(\xi f_0)$ according to eq. (15), whereas the employment of an exporter with augmented productivity $\hat{\theta}(\omega) = \exp[u]$ relative to the least productive plant follows as $t \exp[u]/(\xi f_0)$ according to eqs. (15) and (24). Employing the auxiliary function $F(x, \iota)$ for $\iota = 1 + 1/\gamma$, and following the derivation steps from the closed economy, we can then compute the employment weighted average of worker efficiency in the first line of eq. (31). For the alternative case of $f_x \leq f_0(t-1)$, we can note that all firms export, implying that $\lambda^e(\exp[u]) = (\exp[u]/\hat{\theta}_1^*)^{1/\gamma}[\xi(f_0 + f_e)/\gamma]^{1/\gamma}$. Moreover, employment of an exporter with augmented productivity $\hat{\theta}(\omega) = \exp[u]$ relative to the least productive plant (an exporter itself) is given by $\exp[u]/[\xi(f_0 + f_x)]$ according to eq. (15). Following the derivation steps of the closed economy and making use of auxiliary function $F(x, \iota)$ with $\iota = 1 + 1/\gamma$, we can solve for the second line in eq. (31).

To determine the solution for the average economy-wide wage dispersion in eq. (32), we can note that under a parameter domain with $f_x > f_0(t-1)$, the marginal plant making zero profits is a non-exporter. Moreover, noting that the plant-occupation internal wage dispersion of non-exporters and exporters can be expressed as $cv^d(\exp[u]) = \sqrt{4 - \pi(\pi - 2)}[1 - \eta/\lambda^d(\cdot)]/\pi$ and $cv^e(\exp[u]) = \sqrt{4 - \pi(\pi - 2)}[1 - \eta/\lambda^e(\cdot)]/\pi$, respectively, we can compute the first line of eq. (32) following the steps outlined above. Similarly, noting that for the alternative parameter domain with $f_x \leq f_0(t-1)$ the marginal plant is a non-exporter, we can express the occupation-plant internal wage dispersion for all producers as $cv^e(\exp[u])$. Following the derivation steps outlined above, we can then compute the second line in eq. (32).

To determine the impact of openness on Λ , we first note that $\lim_{f_x \rightarrow \infty} \hat{\theta}_x^* = \infty$ and thus $\lim_{f_x \rightarrow \infty} F(\hat{\theta}_x^*, \iota) = 0$ for $\iota = \{1, 1 + 1/\gamma\}$. Acknowledging $\lim_{f_x \rightarrow \infty} \hat{\theta}_1^* = \hat{\theta}_a^*$, it follows from eq. (21) and the first line of eq. (31) that $\lim_{f_x \rightarrow \infty} \Lambda = \Lambda^a$. Moreover, noting that $\lim_{f_x \rightarrow 0} \hat{\theta}_1^* = \hat{\theta}_a^*$, it follows from eq. (21) and the second line of eq. (31) that $\lim_{f_x \rightarrow 0} \Lambda = \Lambda^a$. Moreover, for $f_x < f_0(t-1)$ we compute

$$\frac{\hat{\theta}_1^*}{\Lambda} \frac{\partial \Lambda}{\partial \theta_1^*} = -\frac{1}{\sigma_u} H \left(\frac{\ln \hat{\theta}_1^* - \rho_{uv} \sigma_u \sigma_v - (1 + 1/\gamma) \sigma_u^2}{\sigma_u} \right) + \frac{1}{\sigma_u} H \left(\frac{\ln \hat{\theta}_1^* - \rho_{uv} \sigma_u \sigma_v - \sigma_u^2}{\sigma_u} \right) > 0$$

according to the second line in eq. (31). Then, noting that $0 < d\hat{\theta}_1^*/d(f_x + f_0) \times (f_x + f_0)/\hat{\theta}_1^* < 1$, it follows that $d\Lambda/d(f_x + f_0) > 0$ if $f_x < f_0(t-1)$. This is sufficient for Λ to having an interior maximum at some $f_x \geq f_0(t-1)$. Following a similar line of reasoning, we can show that CV has an interior maximum (minimum) at some $f_x \geq f_0(t-1)$ if $\eta > (<) 0$. This completes the proof.