The Impact of Multinationals Along the Job Ladder*

Ragnhild Balsvik† Doireann Fitzgerald‡ Stefanie Haller§

February 2023

Abstract

We ask what is the impact of the presence of multinational affiliates on local workers and firms through their interactions in the labor market. We use matched employer-employee data for Norway to show that the labor market is characterized by a job ladder, and that multinational affiliates are located on the upper rungs of this ladder. We then build a general equilibrium job ladder model with endogenous entry of multinational affiliates. We calibrate the model to match key moments from the Norwegian data, and find that the productivity distribution for multinational affiliates has a thicker right tail than that for domestic plants. We use the calibrated model to perform a counterfactual where multinationals face an infinite entry cost. We find that multinational presence increases output and the total wage bill in the economy. But wage inequality increases, due to an increase in the intensity of competition on the upper rungs of the job ladder, and a decrease in the intensity of competition lower down. Multinational presence also leads to higher unemployment, and reduces the profits of local firms.

1 Introduction

Multinational affiliates are on average more productive than local firms. Their presence in a host country therefore implies a thicker right tail of the overall firm productivity distribution.

*Expert research assistance was provided by Sang Min Lee and Rene Dias de Leon. We thank Manuel Amador, Rasmus Lentz, Jeremy Lise, Simon Mongey and Chris Moser for very helpful discussions, and seminar participants at various institutions for comments and suggestions. The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

†Institute of Marine Research

‡Federal Reserve Bank of Minneapolis and NBER, doireann.fitzgerald@gmail.com

§School of Economics, University College Dublin, stefanie.haller@ucd.ie

In this paper we ask, given their impact on the firm productivity distribution, how does the presence of multinational affiliates affect the host country through the labor market?

There are several ways in which the host country labor market can be affected by a thicker right tail of the firm productivity distribution. Access to more productive firms may raise average labor productivity, and depending on the nature of competition in the labor market, may also increase average wages. At the same time, more intense competition in the labor market could lead to exit of local firms, as well as depressing the profitability of those which survive.[1] In addition, the thicker right tail of the productivity distribution could contribute to a change in wage inequality, and could have differential effects on local firms depending on their productivity.

Assessing the impact of multinational presence is difficult, because the choice of a multinational firm to set up affiliates in a host country is precisely that, a choice. As a result, it is hard to disentangle the impact of multinational presence from the impact of the conditions or policies which led them to set up affiliates in the host country in the first place. Several recent papers have tried to deal with this problem by combining an instrumental variables strategy with differencing across labor markets. Using this approach, Alfaro Urena et al. [2021] and Setzler and Tintelnot [2021] find that not only do multinationals pay higher wages than local firms, their presence in a local labor market also increases the wages of workers employed by local firms.

However the approach of differencing across local labor markets cannot get at the general equilibrium impact of multinational presence in an integrated country-wide labor market. In addition, multinational presence can impact the economy through other channels, including forward and backward linkages and potential diffusion of technology. A reduced form strategy cannot separate out the impact through the labor market.

To assess the general equilibrium impact on the labor market specifically, without being subject to the endogeneity critique, we take a model-based approach. This requires us to build a model which captures both the nature of competition in the labor market, and where high-productivity multinational affiliates fit in.

To motivate our choice of model, we use matched employer-employee data for Norway to rank plants by their share of hires from employment ("poaching rank"). We show that job-to-job transitions are frequent, and workers tend to move from plants with a low share of hires from employment to plants with a high share of hires from employment. That is, there is a job ladder in the poaching rank. We find that multinationals are located on the

[1]Note that the profits of local firms are owned by domestic agents, but the profits of multinationals are owned by foreigners.

upper rungs of this ladder.

We then write down a general equilibrium job ladder model with multinational affiliates and local firms. The model combines endogenous vacancy posting by firms which are heterogeneous in productivity, with on-the-job search and wage bargaining as in Cahuc et al. [2006]. We allow for selection and entry into multinational status as in Helpman et al. [2004]. This feature of the model can endogenously generate a different distribution of productivity for multinational affiliates compared to local firms.

In this model, firms pay workers a markdown on their marginal productivity. The markdown depends on the worker's current outside option as well as the distribution of outside options in the labor market. The model generates wage dispersion across firms of different productivity, but also within firms. Meanwhile, workers can experience wage gains on the job as well as by making job-to-job transitions. The current version of our model assumes homogeneous labor, but we plan to add skill heterogeneity as well as complementarity between worker skill and firm productivity eventually.

We calibrate our model to match key features of the Norwegian labor market, including worker transitions, nonemployment, the labor share in output, the size distribution of plants, the distribution of average wages across plants, the share of plants that are multinational, and the relative size of multinational affiliates compared to locally-owned plants. Our calibration implies that potential entrant multinational affiliates pay a high cost to get a draw from a productivity distribution with a fat right tail, while potential entrant local plants pay a low cost to get a draw from a productivity distribution with a thinner right tail. This allows the model to match the position of multinational affiliates on the higher rungs of the job ladder, and rationalizes the higher wages that they pay compared to local plants.

To understand the general equilibrium impact of multinational presence through the labor market, we then use our model to perform a counterfactual exercise where we set the entry cost for multinational affiliates to infinity. We use the domestic free entry condition to solve for the counterfactual mass of firms and active firm productivity distribution holding fixed the domestic entry cost.

Comparing our baseline calibration with the counterfactual, we find that multinational presence increases payments to labor, and reduces domestic firm profits. On net, domestic income (the sum of payments to labor and domestic firm profits) is higher under multinational presence. At the same time, there are distributional consequences. Multinational presence reduces employment (by depressing entry of domestic firms) and leads to increased wage inequality. This increase in wage inequality arises through two channels. First, there is

a direct effect of having a fatter right tail of the productivity distribution of active firms. Second, the distribution of outside options improves for workers at high productivity firms, where competition for workers is more intense when multinationals are present, while it gets worse for workers at low productivity firms, where there is relatively less competition for workers. The impact on local firms is the exact inverse: high productivity local firms suffer from multinational presence, while low productivity local firms benefit from relatively less intense competition for workers.

Our paper is related to several literatures. First, and most obviously, it is related to Alfaro Urena et al. [2021] and Setzler and Tintelnot [2021]. As already noted, our approach is different from that taken by these papers. Nevertheless, our findings are very complementary. These two papers both find that in addition to the wage premium that multinationals pay their workers, there is also a positive impact on wages paid by local firms in labor markets with greater multinational presence. We discuss the relationship between our findings and theirs in greater detail when we describe the results of our counterfactual exercise.

Second, our work is related to several recent papers which use general equilibrium job ladder models to examine the contribution of sorting to wage dispersion (Bagger and Lentz [2018]), the relationship between productivity dispersion and the labor share (Gouin-Bonenfant et al. [2018]), and the impact of a minimum wage on inequality (Engbom and Moser [2021]). Our model is closest to that in Bagger and Lentz [2018], who also assume bargaining over wages as in Cahuc et al. [2006]. In contrast, the other two papers assume wage posting, building on Burdett and Mortensen [1998]. We choose the bargaining model of wage determination, because it allows us to speak more closely to the reduced form empirical evidence, as well as allowing for within-firm as well as cross-firm wage dispersion.

Third, our work is related to the empirical literature on job ladders, e.g. Haltiwanger et al. [2018] and the literature summarized in Moscarini and Postel-Vinay [2018]. Our findings on the job ladder in Norway are very similar to those for other countries.

Finally, our work is related to a literature on search and matching models of the distributional impact of international trade: Helpman et al. [2010], Coşar et al. [2016], and Helpman et al. [2017]. Relative to these papers, we investigate the impact of multinational presence rather than that of openness to international trade. We also differ in making use of a job ladder model where workers search on-the-job as well as from unemployment. In this, our work is closer to Fajgelbaum [2020].

The paper is organized as follows. In the second section, we describe our data. In the third section we use worker flows to characterize the job ladder, and show where multinationals are

on this ladder. In the fourth section, we describe our model. In the fifth section, we calibrate the model. In the sixth section we perform a counterfactual where we remove multinationals. The final section concludes.

2 Data description

We work mainly with three different data sets administered by Statistics Norway. For some robustness analysis we make use of a fourth data set.

Our first data source is the population register. This source has annual files on the population aged between 16 and 74, with identifiers that allow us to follow individuals over time. From this source, we obtain age, gender, years of education and highest level of education, total annual earnings and municipality of residence. For men born between 1950 and 1993 we observe in addition cognitive scores obtained from military records. These files include a plant/workplace identifier for the main employer for people in the labour force, as well as the industry and municipality of the workplace, recorded in November of each year.

Our second source of information is income tax files which include both plant and firm identifiers, allowing us to allocate plants to firms.[2] Our third data source is the register of foreign ownership interests in Norwegian firms (the SIFON register), which records foreign ownership shares at the firm level. We define a firm as foreign-owned if the total foreign ownership share is above 50% in the relevant year. We classify plants as domestic- or foreign-owned based on the ownership of the associated firm identifier from the income tax files. Note that throughout the paper, we refer to domestic-owned plants as "domestic," and foreign-owned plants as "multinational."

In order to identify domestic multinationals for robustness analysis, we use the Utenlandsoppgaven register, which covers Norwegian firms that have ownership shares in entities abroad. From this register we use information on whether a Norwegian firm has ownership interests in at least one entity abroad, with an ownership share of more than 50%.

We start by constructing a plant panel for the years 1996 to 2007 based on the plant identifiers in the population register and income tax files. From this panel, we drop a plant if it is not observed in both data sources for more than half of its years in the panel. We also drop plants that have many years with missing location or industry affiliation. This affects 10% of the initial plant-year observations. We further drop workplaces in the public sector,

[2]These files include information on job spells within a year, but these data are noisy, and we do not make use of them.

which account for 20% of the remaining sample.[3] We also drop very small plants where all workers are recorded as self-employed or the total wage bill does not exceed 100,000 NOK in 2007 NOK. This affects a further 10% of plant-year observations.[4]

We rely on the population register record of main workplace in November to match employees to employers. Our employee sample consists of all individuals who are ever employed at a plant that is present in our final plant sample.

Our dataset covers more than 1.2 million workers in nearly 117,000 plants per year on average. It includes all sectors but the public sector, covering 12 NACE letter industries and 206 NACE 3-digit industries. These plants are located in 160 local labour markets which are defined on the basis of commuting zones by Statistics Norway.

To measure reallocation, we make use of the following definitions. A worker who is observed in the population register at the same plant in two successive Novembers is a job stayer. A worker who is employed at one plant in November of year t and at a different plant in November of year $t+1$ experiences a job-to-job transition. A worker who is employed in November of year t, and who is not employed in November of year $t+1$ experiences a transition from employment to non-employment. Similarly, a worker who is not employed in November of year t, but who is employed in November of year $t+1$ experiences a transition from non-employment to employment.

The annual frequency at which the data is observed implies some degree of misclassification. For example, a worker may pass through a period of non-employment (or indeed multiple spells of employment and non-employment) between one November and the next. But as long as he or she is employed at different plants in successive Novembers, we will count this as a job-to-job transition. Our main contribution is to make use of similarly defined transitions to make comparisons across plants, so as long as the misclassification is similar across plants, this approach serves our purpose.

To measure wages, we make use of annual earnings variable from the population register data. This is total pensionable earnings, including wages and benefits from all employers. It also includes payments for maternity leave, unemployment, and partial disability. We attribute all of these payments to the employer in November of the relevant calendar year.

As with our measures of transitions, there is measurement error in our wage variable, since it includes a variety of payments which may not be associated with employment at the

[3]We also drop the very few workplaces that are classified as private households and extraterritorial organisations.

[4]The public sector accounts for around 40% of employment, while the remaining dropped observations account for about 2% of employment.

employer listed in November of the relevant year. And as in the case of transitions, we rely on this error being similar for different kinds of plant.

Summary statistics on plants and workers are presented in Table 1. We report statistics for all plants and workers, for domestic plants and their employees, and for multinational plants and their employees, noting as above, that "multinational" refers to plants that are foreign-owned. On average, about 6% of plants are subsidiaries of multinationals. This increases from about 4.5% in 1998 to 6.6% in 2007. In line with what is found in the literature, on average multinational plants are bigger, pay better, and have better-educated workers than domestic plants.

Table 1: Summary statistics on workers and plants

	All plants		Domestic		MN	
	mean	sd	mean	sd	mean	sd
	Worker-years					
Log wage[a]	0.00	0.79	-0.03	0.79	0.15	0.75
Age	39.12	12.67	39.22	12.82	38.70	11.99
Yrs of education	12.65	2.02	12.60	2.01	12.88	2.03
Tenure	4.45	4.73	4.44	4.70	4.48	4.84
Ability[b]	5.28	1.80	5.24	1.80	5.46	1.79
Observations	12,001,918		9,815,230		2,186,688	
	Plant-years					
Log employment	1.45	1.13	1.39	1.09	2.35	1.32
Mean log wage	-0.14	0.63	-0.16	0.63	0.09	0.56
Share medium skilled[c]	0.53	0.33	0.53	0.34	0.53	0.26
Share high skilled[d]	0.19	0.30	0.19	0.30	0.23	0.27
Foreign owned	0.06					
Observations	1,166,918		1,091,231		75,687	

Notes: [a]Log wage is the residual from a regression of log wage at the worker level on year dummies. [b]Cognitive scores (1-9) are available from military records for men born between 1950 and 1993. [c]Medium-skilled workers are those with some high school, high school completed, or with a vocational degree. [d]High skilled workers have a BA or above.

3 Multinationals and the job ladder

3.1 Poaching index

The labor market in Norway, as in other countries, is characterized by a good deal of churn. At the plant level, the average annual separation rate is 24%, while new hires account for 26% of employment. Meanwhile, 72% of these new hires are hires from employment based on our definition.

Job-to-job reallocation is not random. Job-to-job transitions have a strong directional

component: some plants attract workers through these transitions, while others do not. The phenomenon where workers appear to share a ranking of employers, and move from less desirable to more desirable employers, is often referred to as a job ladder (see e.g. Burdett and Mortensen [1998], Postel-Vinay and Robin [2002] and Moscarini and Postel-Vinay [2018]). The empirical literature has used a several ways to measure the job ladder: plant size, average plant-level wage, or a plant's share of hires from employment, also known as the "poaching index."

We choose to work with the poaching index as a measure of the job ladder. This measure is consistent with a number of job ladder models. It is robust to the possibility that current wages may not fully capture a plant's attractiveness to workers, either because workers may be rewarded through a combination of their current wage and the option to move up the job ladder, as in Postel-Vinay and Robin [2002] and Cahuc et al. [2006], or because they may be rewarded through non-wage amenities. It is also robust to firm dynamics which are clearly present in the data, and which may make size a poor measure of the job ladder in practice. For example, Haltiwanger et al. [2018] find that young firms tend to be small, but that they also systematically poach workers from other businesses, whereas small firms that are old are less likely to poach.

The poaching index is constructed as follows. For each plant, we pool all hires in the sample period, and calculate the fraction of these hires that come from employment. This index is therefore fixed for a given plant over time. In making use of the poaching index, we restrict attention to plants making at least 10 hires over the sample period, one of which is from non-employment.[5] This follows Bagger and Lentz [2018]. To deal with the possibility of measurement error, for several of the exercises involving the job ladder, we work with deciles of the poaching index.

3.2 A job ladder in the poaching index

With the poaching index in hand, we circle back and show that workers making job-to-job transitions systematically move from plants with a low poaching index to plants with a similar or higher poaching index. This is illustrated in Figure 1, which uses a heatmap to illustrate the transition matrix for these workers across plants based on their poaching index decile. With the exception of transitions originating in the top decile, workers are always more likely to move sideways or up than to move down. And for transitions originating in

[5]These plants account for 82% of employment-years and 39% of plant-years in our baseline data. Summary statistics for this sample are similar to those presented in Table 1, see the Appendix.

the top decile, more than 60% are to plants in deciles 9 and 10. Overall, 66% of job-to-job transitions are to a plant in the same or higher decile of the poaching index relative to the origin plant (i.e. on or above the diagonal). This confirms that we have identified a job ladder. We also verify that the separation rate (including separations to non-employment as well as to employment) is mostly declining in a plant's position on the job ladder, again indicating that the poaching index captures the attractiveness of plants to workers (see Figure 2).

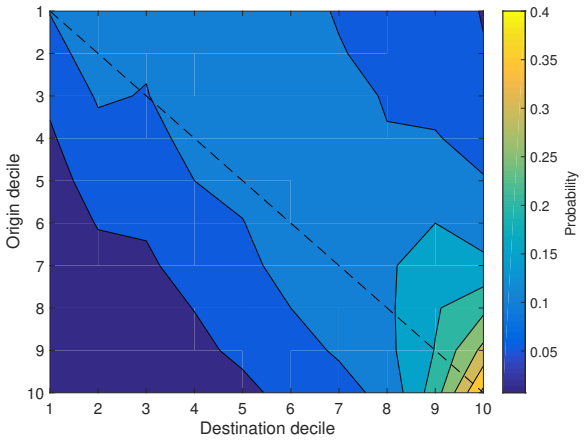


Figure 1: Transition heatmap

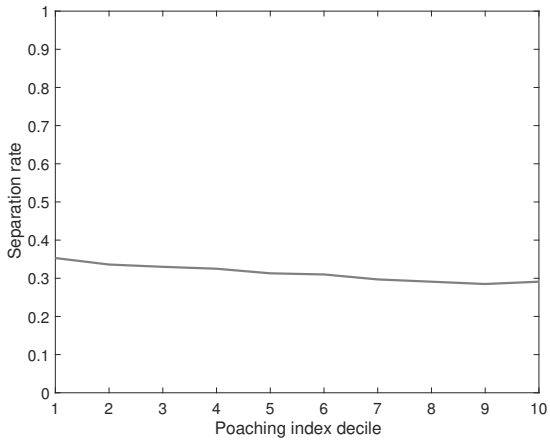


Figure 2: Separation rate

Notes: Figure 1 is a heat map of the probability distribution of job-to-job transitions originating in a plant of a given poaching index decile, and for which the poaching index decile of the destination plant is known, by poaching index decile of the destination plant. 13% of transitions originating in a plant for which the poaching index is defined are to a plant for which it is not defined. Figure 2 shows the separation rate by poaching index decile for plants for which the poaching index is defined.

The rungs of the job ladder based on the poaching index are very similar to the rungs based on plant mean wages, and are related to, but not identical to the rungs based on plant size. This is illustrated in Figures 3 and 4. Figure 3 (left panel) shows the mean poaching index by percentiles of plant wage. This relationship is increasing, except for at the very top. Figure 4 (right panel) shows the mean poaching index by percentiles of mean employment. This relationship is U-shaped. This is consistent with the findings of Haltiwanger et al. [2018] that small establishments combine young establishments which systematically poach workers from other businesses as well as older establishments which do not. Table 2 reports the Spearman rank correlations between the three potential measures of rungs of the job ladder. All are positively correlated with each other, but the strongest correlation is that between the poaching index and the plant-level mean log wage.

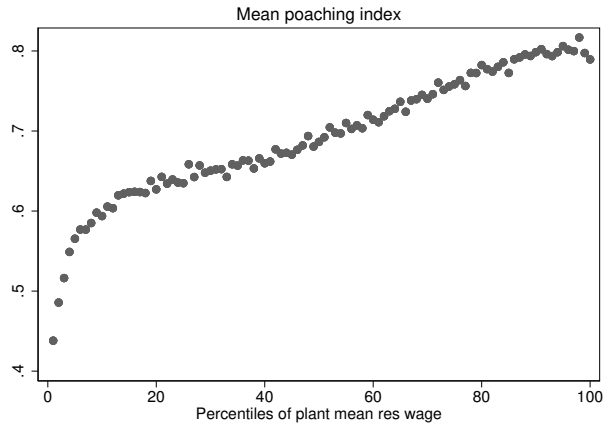


Figure 3: Wages & poaching index

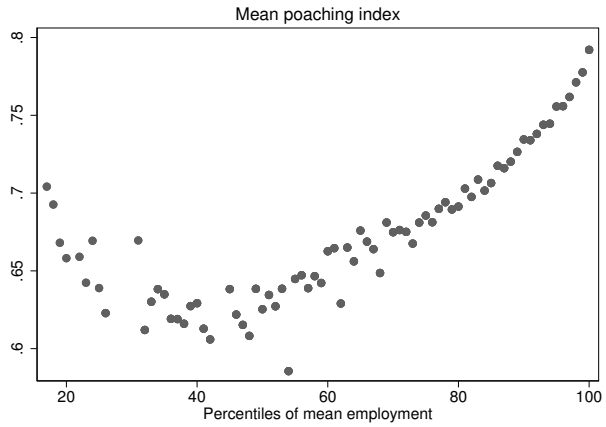


Figure 4: Employment & poaching index

Notes: To construct the left panel, we average the mean log wage at the plant level across all years the plant appears in the sample. Plants are then divided into percentiles based on this variable. The mean poaching index across all plants in a given percentile for which the poaching index is defined is then plotted on the y-axis. To construct the right panel, we average employment across all years the plant appears in the sample. Plants are then divided into percentiles based on this variable. The mean poaching index across all plants in a given percentile for which the poaching index is defined is then plotted on the y-axis. There are no plants in the lower percentiles of the right panel, as we require plants to make at least 10 hires over the sample period, one of which is from non-employment, in order for the poaching index to be defined.

Table 2: Spearman rank correlations for alternative plant rankings

Target	Log employment	Mean log wage	Poaching index
Log employment	1.000		
Mean log wage	0.229	1.000	
Poaching index	0.153	0.498	1.000

Notes: Correlations are calculated for plant-years for which the poaching index is defined, $N = 395,551$. All correlations are significant at the 1% level.

3.3 Multinational position on the job ladder

Having established that there is a job ladder in the poaching index, we now examine the position of multinational plants on this job ladder. Figure 5 shows the distribution of the poaching index by ownership, while Table 3 reports the associated summary statistics. The distribution for foreign-owned plants first-order stochastically dominates that for domestic plants: multinational plants sit disproportionately on the higher rungs of the job ladder.[6] While they account for only 5% of plants in the first decile, multinationals account for more than 18% of plants in the 10th decile of the poaching index.

Given the existence of a job ladder, it should not come as a surprise that multinationals are on the upper rungs. There is a lot of evidence that multinationals have a productivity

[6]This is not driven by a higher proportion of cross-plant within-firm transfers for multinationals, as the picture is very similar at the firm level, see Appendix.

advantage over domestic plants.[7] Meanwhile, in job ladder models with employer heterogeneity, position on the job ladder is increasing in productivity,[8]. In the Appendix, we show that domestic-owned multinationals, just like foreign-owned multinationals, are mainly on the upper rungs of the job ladder as we define it. Indeed, the distribution of the poaching index for Norwegian-owned multinationals is shifted to the right compared to the distribution for foreign-owned multinationals. This is consistent with what we know about the productivity advantage of multinationals.

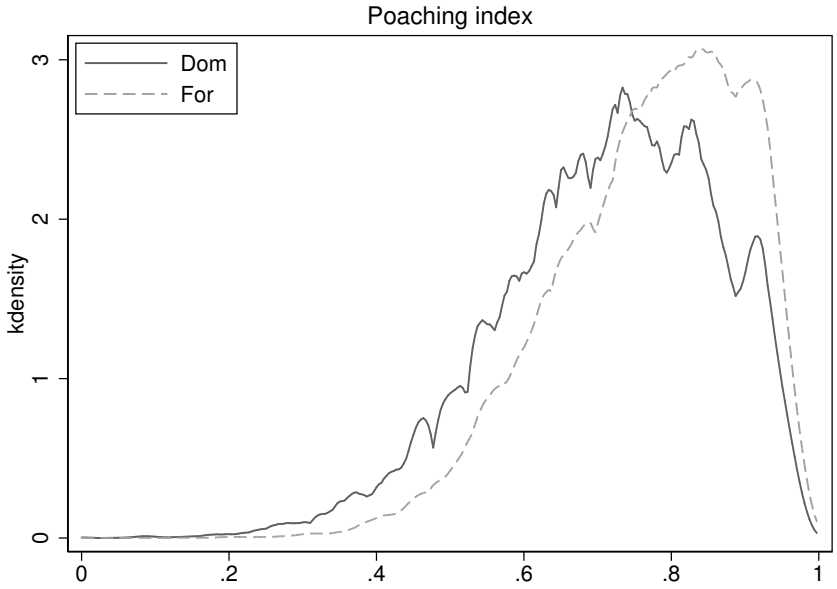


Figure 5: Poaching index distribution by ownership
Notes: Kernel density distribution of the poaching index by plant ownership.

Table 3: Summary statistics of poaching index distribution by ownership

	p10	p25	p50	p75	p90	mean	sd	N
Domestic	0.51	0.62	0.73	0.82	0.90	0.71	0.15	403,629
Foreign	0.59	0.68	0.79	0.87	0.92	0.77	0.13	50,977
Total	0.52	0.62	0.73	0.83	0.91	0.72	0.15	454,606

Notes: Summary statistics are calculated usiing plant-years for which the poaching index is defined.

3.4 Sorting of workers along the job ladder

We now investigate sorting of workers to plants along the job ladder. To measure the job ladder, we continue to use the poaching index. To measure worker type we use standard

[7]See e.g. Antrás and Yeaple [2014], Criscuolo and Martin [2009], Haller [2012], Tomiura [2007] See section 3.1 of Greenaway and Kneller [2007] for a summary.

[8]See e.g. Burdett and Mortensen [1998] and Moscarini and Postel-Vinay [2018].

observable measures of skill: years of education and ability. [9]Figure 6 (left panel) plots the mean number of years of education for workers at plants at different percentiles of the poaching index. Figure 7 (right panel) plots mean ability for male workers at plants at different percentiles of the poaching index. Average education and average ability of male workers are both increasing in the poaching index, suggesting positive assortative matching along the job ladder.

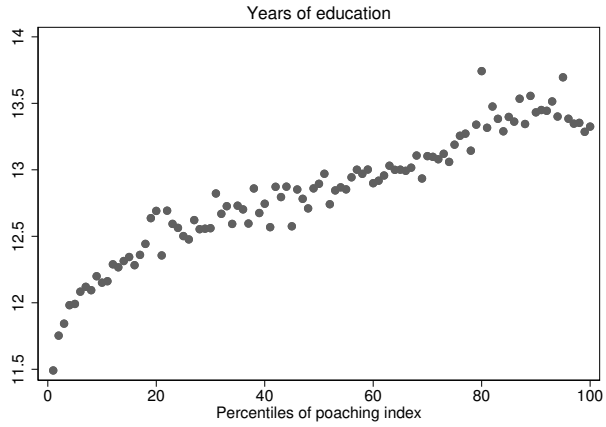


Figure 6: Education & poaching index

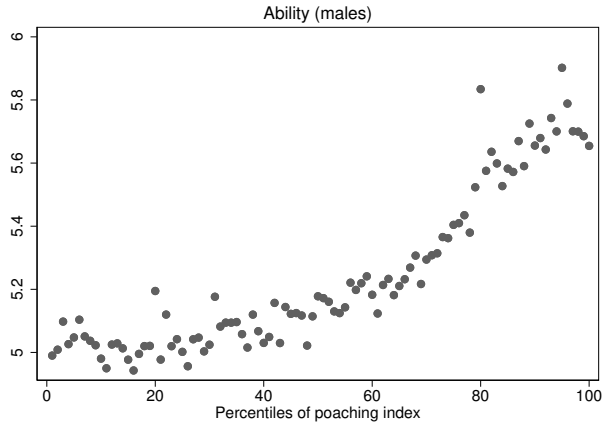


Figure 7: Ability & poaching index

Notes: To construct Figure 6 (left panel), we first calculate average years of education at the plant-year level. This is then averaged across all years that a plant appears in the sample. Plants are then divided into percentiles of the poaching index, and the average of the years of education variable for all plants in this bin is calculated. The mean poaching index across all plants in a given percentile for which the poaching index is defined is then plotted on the y-axis. Figure 7 (right panel) is constructed analogously, with average ability for male employees for which ability is available replacing average years of education.

4 Model

Motivated by the existence of a job ladder in the Norwegian data, we now turn to our model. As described in the Introduction, the model takes elements from Helpman et al. [2004], which features selection in multinational status, and merges them with a general equilibrium version of Cahuc et al. [2006]. We first describe the simplest possible version of the model (assumptions and results) to build intuition, and then sketch some extensions.

The agents in our models are firms and workers. In the baseline version of the model, workers are homogeneous (we add worker heterogeneity in an extension). They search for jobs both when unemployed, and on-the-job. There is random matching between searching

[9]Plant and worker fixed effects from log wage regressions are often used as measures of plant and worker types. (see Abowd et al. [1999]). These measures rely on monotonicity of wages in plant and worker types, and on workers not selecting into plants based on the idiosyncratic component of wages, assumptions which may be violated in job ladder models. See, e.g. Postel-Vinay and Robin [2002] and Bagger and Lentz [2018].

workers and vacancies. Conditional on a match, wages are set according to the wage bargaining protocol in Cahuc et al. [2006], such that workers get their outside option, plus a fraction of the match surplus. Firms have heterogeneous productivity. They face a convex cost of posting vacancies, which delivers finite firm size. The mass of firms is determined by a free entry condition, which equates the cost of getting a draw of productivity with the expected value of the draw. Multinational affiliates are distinguished from domestic firms by the fact that they pay a higher cost to get a productivity draw, and their draw comes from a different distribution from that faced by local entrants. In addition, multinational profits are not owned by domestic agents. The model is set up in discrete time, and we look for a stationary equilibrium.

4.1 Assumptions: Workers

There is only one worker type with skill normalized to 1. There is a continuum of infinitely-lived workers of this type on $[0, 1]$. Workers have linear utility, and discount the future at rate β. Unemployed workers receive a flow of utility b every period of unemployment. Conditional on being matched with a firm, workers supply 1 unit of labor, and receive wage w, which we will presently characterize. The match between a worker and a firm is broken with probability δ each period. The worker must then pass through one period of unemployment before searching for a job. After this first period, unemployed workers search for jobs with probability 1 each period. Employed workers search for jobs with probability $s \leq 1$ each period. Conditional on meeting with a firm, workers accept offers that make them better off, as we describe when we lay out the assumptions about wage setting.

4.2 Assumptions: Firms

All firms produce the same final good. A firm is a draw of productivity p. Output per worker employed at firm of type p is p.[10] Each firm pays cost $c(v)$ in units of output to post vacancies $v \in \mathbb{R}$, with $c(0) = 0$, $c'(v) > 0$ and $c''(v) > 0$. Like workers, firms discount the future at rate β. A firm of type p chooses the optimal measure of vacancies to post, $v(p)$, in order to maximize $B(v, p)$, the value of these vacancies, given the wage setting protocol. We describe this problem in more detail after laying out the assumptions about wage setting. Note that the vacancy posting decision does not depend on the firm's current stock of workers, or the distribution of wages across its current stock of workers.

[10] The linearity of the production function is key to tractability of the problem.

Let $B(v(p), p) = B(p)$ denote the value to firm with productivity p of posting its optimal amount of vacancies. The value of an entrant (i.e. a firm with no employees) with productivity p in a stationary equilibrium is therefore equal to the present value of $B(p)$. We assume that domestic potential entrants can pay a cost C^D (in units of output) to get a productivity draw from a distribution with cdf $\tilde{\Gamma}^D(p)$ and pdf $\tilde{\gamma}^D(p)$, defined on $[b, \bar{p}]$. Meanwhile multinational potential entrants can pay a cost C^M to get a productivity draw from a distribution with cdf $\tilde{\Gamma}^F(p)$ and pdf $\tilde{\gamma}^F(p)$, also defined on $[b, \bar{p}]$. The free entry condition for a potential entrant of type i is:

$$C^i = \int_b^{\bar{p}} \frac{B(p)}{1-\beta} \tilde{\gamma}^i(p)\, dp$$

As we will presently derive, there is a cutoff level of productivity $\underline{p} > b$ below which no firm will be able to attract workers, so $B(p) = 0$ for $p \leq \underline{p}$. Any entrant receiving a productivity draw below this level exits immediately. The measure of active domestic firms, M^D and the measure of active multinational firms, M^F, are pinned down by the domestic and multinational free entry conditions together with this cutoff. The resulting total measure of active firms is $M = M^D + M^F$. Let $\gamma(p)$ denote the pdf of all active firms, while $\Gamma(p)$ is the corresponding cdf. Then $\gamma(p)$ is given by the mixture distribution:

$$\gamma(p) = \frac{M^D}{M^D + M^F} \left(\frac{\tilde{\gamma}^D(p)}{1 - \tilde{\Gamma}^D(\underline{p})} \right) + \frac{M^F}{M^D + M^F} \left(\frac{\tilde{\gamma}^F(p)}{1 - \tilde{\Gamma}^F(\underline{p})} \right)$$

We proceed under the assumption that we are in the stationary equilibrium of the model. There is no firm exit, and therefore no firm entry. The free entry conditions hold, but all entry costs are sunk, and do not enter into the resource constraint of the economy. Firm output net of the vacancy posting cost is divided between payments of wages to labor and payments of profits to owners of the firm.

4.3 Assumptions: Matching

The total measure of vacancies in this economy is V, where

$$V = M \int_{\underline{p}}^{\bar{p}} v(p)\, \gamma(p)\, dp$$

The total measure of searching workers is S, where, remembering that newly unemployed workers cannot search,

$$S = u + s(1-\delta)(1-u)$$

Searching workers and vacancies match randomly, with a constant returns to scale matching function given by $\mu(S,V)$. We denote the probability that an unemployed worker meets a vacancy by

$$\lambda = \frac{\mu(S,V)}{S}$$

while the probability that an employed worker meets a vacancy is given by λs. We denote the probability that a vacancy matches with a worker by

$$\chi = \frac{\mu(S,V)}{V}$$

4.4 Assumptions: Wage setting protocol

The wage setting protocol is as described in Cahuc et al. [2006]. When a worker and a firm match, the worker is offered a value equal to their outside option, plus a fraction ϕ of the match surplus (workers and firms value surplus in the same way). This is implemented by workers receiving a constant wage w until their outside option changes.

Let U denote the value of unemployment. Let $W(w,p)$ denote the value to a worker of receiving wage w at firm of type p. Let $w(q,p)$ denote the wage of a worker at firm of type p whose outside option is a firm of type q.

An unemployed worker meeting a firm of type p accepts an offer $w_0(p)$ such that:

$$W(w_0(p),p) = U + \phi(W(p,p) - U)$$

Suppose a worker at firm of type p whose outside option is q meets a firm of type p'. There are three possibilities:

1. If $p' \leq q \leq p$, nothing happens.

2. If $q < p' \leq p$, the worker stays at the current firm, but receives a new wage $w(p',p)$ such that:
$$W(w(p',p),p) = W(p',p') + \phi(W(p,p) - W(p',p'))$$

3. If $p < p'$, the worker moves to the firm of type p', and receives wage $w(p, p')$ such that:

$$W(w(p, p'), p') = W(p, p) + \phi(W(p', p') - W(p, p))$$

4.5 Results: Transitions

Matched workers and firms separate with probability δ each period. Unemployed searchers (i.e. those who have been unemployed for more than 1 period) meet a firm with probability λ, and accept all offers. A worker employed at a firm of type p meets a new firm with probability λs each period. If the new firm has productivity $p' > p$, the worker moves to the new firm. Otherwise it remains employed at the original firm.

4.6 Results: Wages

Let

$$f(x) = \frac{v(x)\gamma(x)}{\int_{\underline{p}}^{\bar{p}} v(y)\gamma(y)\,dy}$$

where $v(x)$ is the measure of vacancies posted by a firm of type x. This is the pdf of the job offer distribution. The corresponding cdf is $F(x)$.

In the Appendix, we derive the following expression for the wage of worker at firm of type p who has an outside option q:

$$w(q, p) = \phi p + (1-\phi)q - \int_q^p \frac{(1-\phi)^2 \beta(1-\delta)\lambda s(1-F(x))}{1-\beta(1-\delta)(1-\phi\lambda s)(1-F(x))}dx$$

Workers are paid a markdown on their marginal product. This markdown depends on their current outside option, and on the option value of negotiating a wage increase within their current firm. Note that wages need not be monotonic in p, and workers could potentially experience a wage reduction on moving from their current firm to a firm with higher productivity if the option value of wage increases in the new firm is high enough.

In the Appendix, we also derive the following expression for the wage of a worker at firm of type p who is hired from unemployment:

$$w_0(p) = w(\underline{p}, p) = \phi p + (1-\phi)\underline{p} - \frac{(1-\phi)^2}{\beta\lambda\phi(1-(1-\delta)s)}(\underline{p} - b)$$

where $\underline{p}$ is implicitly defined as the level of productivity such that the unemployed are indifferent between taking an offer from this firm and remaining unemployed. It is given by

(again see the Appendix for derivation):

$$\underline{p} = b + \beta\lambda\phi\left(1-(1-\delta)\,s\right)\int_{\underline{p}}^{\bar{p}} \frac{1-F\left(x\right)}{1-\beta\left(1-\delta\right)\left(1-\phi\lambda s\right)\left(1-F\left(x\right)\right)}dx$$

Meanwhile, the value to a worker at firm of type p of receiving wage w is given by:

$$W\left(w,p\right) = \frac{w+\beta\delta U}{1-\beta\left(1-\delta\right)} + \frac{\beta\left(1-\delta\right)\lambda s}{1-\beta\left(1-\delta\right)}\left[\begin{array}{c}(1-\phi)\int_{q(w,p)}^{p}\frac{(1-F(x))}{1-\beta(1-\delta)(1-\phi\lambda s)(1-F(x))}dx+\\ \\ +\phi\int_{p}^{\bar{p}}\frac{(1-F(x))}{1-\beta(1-\delta)(1-\phi\lambda s)(1-F(x))}dx\end{array}\right]$$

where $q\left(w,p\right)$ is defined by $w\left(q,p\right)$. The value of unemployment, U, is given by:

$$U = \frac{b}{1-\beta} + \frac{\beta\lambda\phi}{1-\beta}\int_{\underline{p}}^{\bar{p}} \frac{1-F\left(x\right)}{1-\beta\left(1-\delta\right)\left(1-\phi\lambda s\right)\left(1-F\left(x\right)\right)}dx$$

4.7 Results: Employment and wage distributions

Let $L\left(p\right)$ be the probability that an employed worker works at a firm with productivity $\leq p$. Let $l\left(p\right)$ be the associated pdf. Note that this is a distribution across *workers*, not across *firms*. Note also that since there are $(1-u)$ employed workers, $(1-u)\,L\left(p\right)$ is the *number* of workers working at firms with productivity $\leq p$. In steady state, the outflow of workers from firms of type p must equal the inflow of workers into firms of type p. In the Appendix we show that this implies that:

$$l\left(p\right) = \frac{1+\left(\frac{1-\delta}{\delta}\right)s\lambda}{\left(1+\left(\frac{1-\delta}{\delta}\right)s\lambda\left(1-F\left(p\right)\right)\right)^{2}}f\left(p\right)$$

Meanwhile, let $G\left(q|p\right)$ be the cdf of the steady state distribution of outside options for workers at firms of type p, and let $g\left(q|p\right)$ be the associated pdf. In the Appendix, we show that:

$$G\left(w|p\right) = \left(\frac{1+\left(\frac{1-\delta}{\delta}\right)s\lambda\left(1-F\left(p\right)\right)}{1+\left(\frac{1-\delta}{\delta}\right)s\lambda\left(1-F\left(q\right)\right)}\right)^{2}$$

4.8 Results: Vacancies and profits

Define $J\left(q,p\right)$ to be the value to a firm of productivity p of employing a worker with outside option q. This is equal to the firm's outside option (zero) plus its share of match surplus,

i.e. $(1-\phi)(W(p,p)-W(q,q))$. In the Appendix, we show that $J(q,p)$ is given by:

$$J(q,p)=\frac{(1-\phi)}{1-\beta(1-\delta)}\left[(p-q)-\int_q^p\frac{\beta(1-\delta)\lambda s\phi(1-F(x))}{1-\beta(1-\delta)(1-\phi\lambda s)(1-F(x))}dx\right]$$

The value to a firm with productivity p of posting v vacancies is then:

$$B(p,v)=\max_v\left\{\chi v\left[\frac{u}{S}J\left(\underline{p},p\right)+\frac{(1-u)(1-\delta)s}{S}\left(\int_{\underline{p}}^p J(x,p)\,l(x)\,dx\right)\right]-c(v)\right\}$$

The first order condition for the optimal vacancy posting decision is therefore:

$$\chi\left[\frac{u}{S}J\left(\underline{p},p\right)+\frac{(1-u)(1-\delta)s}{S}\left(\int_{\underline{p}}^p J(x,p)\,l(x)\,dx\right)\right]=c'(v)$$

This first order condition implicitly defines $v(p)$, the optimal measure of vacancies posted by firm of type p.

When all firms are at their steady state size, the number of workers employed at a firm of type p is given by the total number of workers employed at firms of type p, i.e. $(1-u)\,l(p)$, divided by the number of firms of type p, i.e. $\gamma(p)\,M$. This implies that:

$$e(p)=\frac{(1-u)\,l(p)}{M\gamma(p)}$$

Firm flow profits net of vacancy posting costs are then given by:

$$\pi(p)=\left(p-\underbrace{\int_{\underline{p}}^p w(x,p)\,g(x,p)\,dx}_{\text{average wage at firm } p}\right)e(p)-c(v(p))$$

4.9 Results: Ranking firms

In the absence of data on productivity, it is useful to have a way of ranking firms. Since $w(q,p)$ need not be monotonic in p, there is no guarantee that the average wage at firm of type p is monotonic in p. Although *steady state* employment $e(p)$ is increasing in p, we do not know whether firms in the data have reached their steady state size. But conveniently, a firm's share of hires from employment is monotonically increasing in p. Denote this share

by $poach\,(p)$. It is given by:

$$poach\,(p) = \frac{(1-u)\,(1-\delta)\,s\int_{\underline{p}}^{p} l\,(x)\,dx}{u+(1-u)\,(1-\delta)\,s\int_{\underline{p}}^{p} l\,(x)\,dx}$$

The intuition for monotonicity of the poaching index is that because of random matching, the share of matches who are unemployed is the same for all firms, but conditional on a match, firms with higher p will induce a higher share of employed searchers to make the job-to-job transition.

4.10 Extension: Capital in the production function

Suppose that the production function in firm of type $\hat{p}$ is $y = \hat{p}k^{\kappa}l^{1-\kappa}$. Under the assumption that all firms face the same rental price of capital (exogenous, set on world markets), and there are no frictions in the rental market for capital, the marginal product of capital is equalized across all workers:

$$MP_k\,(\hat{p}) = \kappa\hat{p}k^{\kappa-1}l^{1-\kappa} = R$$

This implies that the optimal amount of capital hired by firm of type $\hat{p}$ is given by:

$$k\,(\hat{p}) = \left(\frac{\kappa\hat{p}}{R}\right)^{\frac{1}{1-\kappa}} l\,(\hat{p})$$

Meanwhile, the marginal product of labor in firm of type $\hat{p}$ is given by:

$$MP_l\,(\hat{p}) = (1-\kappa)\,\hat{p}\,(k\,(\hat{p})\,/l\,(\hat{p}))^{\kappa} = (1-\kappa)\,\hat{p}^{\frac{1}{1-\kappa}}\left(\frac{\kappa}{R}\right)^{\frac{\kappa}{1-\kappa}}$$

So making use of the optimal amount of capital, the marginal product of labor at firm of type $\hat{p}$ is:

$$MP_l\,(\hat{p}) = (1-\kappa)\left(\frac{\kappa}{R}\right)^{\frac{\kappa}{1-\kappa}}\hat{p}^{\frac{1}{1-\kappa}} = p$$

Payments to capital from firm of type $\hat{p}$ as a share of total output are given by:

$$\frac{Rk\,(\hat{p})}{\hat{p}k\,(\hat{p})^{\kappa}\,l\,(\hat{p})^{1-\kappa}} = \frac{R}{\hat{p}}\left(\frac{k\,(\hat{p})}{l\,(\hat{p})}\right)^{1-\kappa} = \kappa$$

This implies that our model can be reinterpreted as one where there is a standard Cobb-Douglas production function in capital and labor, capital gets share κ of output, and the

remaining $(1-\kappa)$ share is divided between labor and firm profits. Meanwhile, marginal productivity p is the marginal productivity of equipped labor, and is a function of true underlying TFP $\hat{p}$, the rental price of capital R, and the capital share κ.

4.11 Extension: Labor heterogeneity & sorting

The model can be extended in a straightforward way to allow for a finite number of labor skill types, $h = 1, \ldots, H$. There is a fixed supply a_h of each labor type with $\sum_{h=1}^{H} a_h = 1$. Labor types are observable, both to workers and firms. Workers search in the market for their labor type, while firms choose the measure of vacancies to post vacancies in each of these markets. They face a convex cost of vacancy posting in each market, and the matching function operates at the market level.

Reservation utility, the rate of exogenous separations, and the search intensity of currently matched workers may differ across worker types, i.e. $\{b_h, \delta_h, s_h\}$. The worker share in match surplus may also differ across worker types, ϕ_h. In addition, vacancy posting costs and the matching function could differ across skill markets.

Worker marginal productivity at firm of type p may depend on skill type. Marginal productivity of type 1 is normalized to p. Marginal productivity of type h is given by $\eta_h p^{\nu_h}$, with $1 < \eta_2 < \ldots < \eta_H$, and $1 \leq \nu_2 \leq \ldots \leq \nu_H$. If $\nu_h > 1$ for some h, this will induce sorting.

Since firm vacancy posting decisions are independent across skill markets, most of the results of the model are unchanged. However for the poaching index to be guaranteed monotonic in firm type p, we require all firms to post vacancies (and hire) in all skill markets.

5 Calibration

In this section we describe our calibration strategy and results. We calibrate a version of the model with capital in the production function, but without labor heterogeneity. Eventually we plan to calibrate a version of the model with multiple skill types to match a richer set of data moments.

5.1 Functional forms

Following the literature, we assume that the matching function is Cobb-Douglas:

$$\mu(S, V) = AS^{\theta}V^{1-\theta}$$

Following e.g. Bagger and Lentz [2018] we assume that the cost of vacancy posting takes the form:

$$c(v) = \frac{v^{1+\frac{1}{\alpha}}}{1+\frac{1}{\alpha}}$$

Note that the intercept in the matching function and the intercept in the cost of vacancy posting are not separately identified (see the first order condition for vacancies), so we normalize the intercept in the cost of vacancy posting to 1. Note also that A must be restricted such that the matching probabilities $\{\lambda, \chi\}$ both lie on the unit interval.

We assume that domestic and multinational firms each take draws from a truncated Pareto distribution, with scale parameter equal to b, and shape parameter equal to σ^D and σ^M respectively.

5.2 Solution algorithm

We set the following parameters exogenously. The length of a period in our calibration is a quarter, so we set $\beta = 0.95^{1/4}$. We set the capital share $\kappa = 1/3$. We normalize the flow value of unemployment, $b = 1$. Following the literature, we set the exponent in the matching function $\theta = 0.5$. We set $\delta = 0.038$ to match the quarterly rate of employment to nonemployment transitions in Norway over the period 2011-2019 (from Eurostat).

It turns out to be more convenient to treat ω and the measure of active firms M as parameters, and use them to recover the entry costs C^D and C^M. So given a vector of values for $\left\{s, \phi, A, \alpha, \sigma^D, \sigma^F, M, \omega\right\}$, we discretize the domestic and multinational Pareto distributions for productivity. We then solve the model by first guessing an aggregate measure of vacancies, V. Given V,we use the matching function together with the steady state relationship between λ, δ and u to recover a guess of the unemployment rate. This gives us guesses of the worker and firm matching probabilities λ and χ. We then use policy function iteration on the first order condition for the optimal choice of vacancies to recover the full vector of vacancies for each firm type, updating our guesses of $\{V, u, \lambda, \chi\}$ in each iteration. Finally, we iteratively arrive at the cutoff value for productivity $\underline{p} > b$ below which no firm can attract any workers.

Outside of this, we use a particle swarm algorithm to search for the values of the parameter vector $\left\{s, \phi, A, \alpha, \sigma^D, \sigma^F, M, \omega\right\}$ which minimize the sum of squared differences between target moments in the model and in the data.

5.3 Targets and calibrated parameters

We choose these 8 parameters to match 8 target moments. The targets we pick include labor market transition rates, the nonemployment rate, the labor share, and moments of the plant size and wage distributions, including differences between domestic and multinational plants. Although all parameters affect all moments, roughly speaking there is a one-to-one mapping between certain parameters and moments. Table 4 lists the target moments, where we obtain them from, their values in the data, the fitted values in the model, the corresponding parameter, and its fitted value.

Table 4: Calibration targets and parameter estimates

Target	Data	Model	Parameter	Estimate
Outside data (source)				
EE quarterly transition rate (Eurostat)	0.03	0.03	s	0.52
Labor share (Statistics Norway)	0.60	0.60	ϕ	0.85
Nonemployment rate 25-54 (Statistics Norway)	0.15	0.15	A	0.46
Our data				
Std dev log plant-level emp	1.13	1.13	α	0.27
Std dev log plant-level avg wage	0.63	0.63	σ^D	1.67
Average plant size	10.29	10.29	M	0.08
Diff in avg emp betw domestic & MN plants	19.90	19.91	σ^F	0.71
Share of plants that are domestic	0.94	0.94	ω	0.01

The shape parameters of the domestic and multinational productivity distributions are consistent with the hypothesis that the productivity distribution for multinational affiliates has a fatter right tail than that for domestic plants. Given the calibrated parameters, we can recover the entry costs C^D and C^F by making use of the free entry conditions. We find that $C^F/C^D = 53$, so multinational affiliates pay a higher cost to obtain a draw from a productivity distribution with a higher mean.

It is worth noting that while the shape parameters we obtain are quite extreme, these are for the distribution of the marginal productivity p of *equipped* labor, which is more dispersed than that of underlying TFP $\hat{p}$, which is proportional to $p^{1-\kappa}$.

5.4 Nontargeted moments

Having obtained parameter estimates, we simulate our quarterly model for 1,000,000 workers over a period of 10 years. We then use the simulated data to calculate labor transitions, the poaching index, and wages as in the data. That is, we use a single quarterly cross-

section within the year to link workers to firms. We use these annual cross-sections to code transitions, and we attribute all earnings within a calendar year to the firm the worker is matched with in that cross-section.

In Figure 8, we reproduce the distributions of the poaching index for domestic and multinational firms in the data, and in Figure 9 we plot the same distributions for plants in the model.

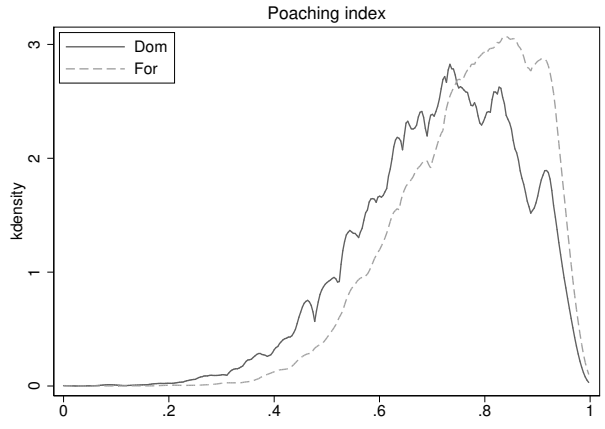


Figure 8: Poaching index dist: Data

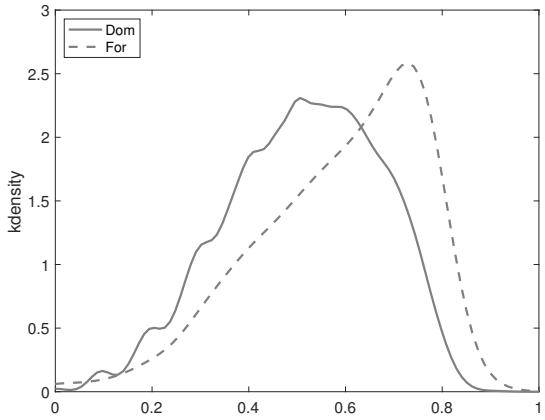


Figure 9: Poaching index dist: Model

Qualitatively, the pictures look very similar, though the mean poaching index in the model is lower than that in the data, both for plants whose productivity is drawn from the domestic distribution, and for plants whose productivity is drawn from the multinational distribution. In a future version of the paper, we plan to address this by adding an additional parameter which allows the search intensity of the nonemployed to be below 1, and choose the value of this parameter to match quarterly rates of nonemployment to employment transitions.

In Figure 10, we plot the mean poaching index for percentiles of the mean residual log wage at the plant level. In Figure 11, we plot the corresponding figure based on the simulated data.

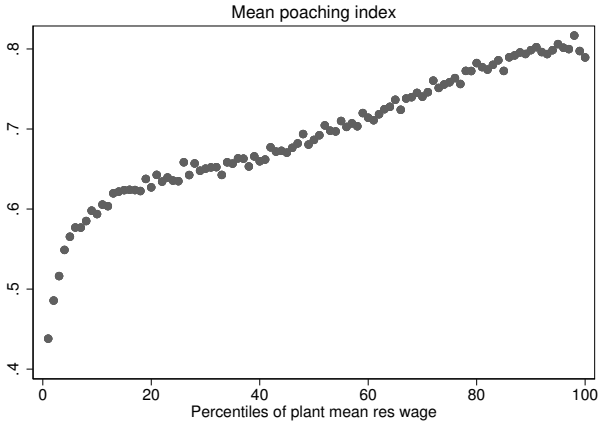


Figure 10: Poaching index & wages: Data

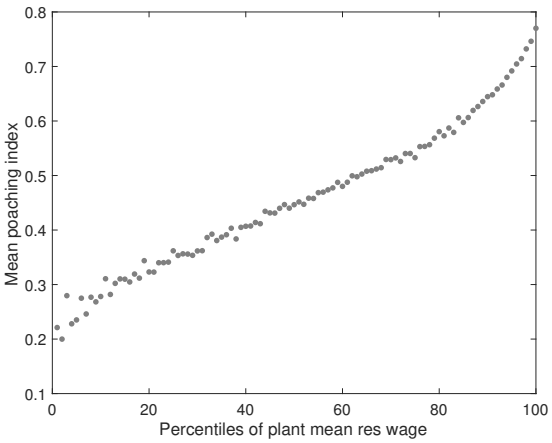


Figure 11: Poaching index & wages: Model

In both cases, the relationship is positive, though it is steeper in the simulated data than in the actual data, particularly at the top of the poaching index distribution.

6 Counterfactual: no multinationals

We now examine how the presence of high-productivity multinational affiliates in the Norwegian labor market affects workers and domestic firms. We do this by implementing the counterfactual exercise of setting the cost of entry for multinationals to infinity in our calibrated model, while holding fixed the entry cost for domestic firms. We then solve for the counterfactual measure of firms M' such that the domestic firm free entry condition is satisfied when $\tilde{\Gamma}(p) = \tilde{\Gamma}^D(p)$. This also yields a counterfactual $\underline{p}'$, and the counterfactual productivity distribution for active firms, which has pdf given by:

$$\gamma'(p) = \frac{\tilde{\gamma}^D(p)}{1 - \tilde{\Gamma}^D\left(\underline{p}'\right)}$$

We can now compare outcomes across the stationary equilibria of the model with and without multinational entry. In each, we treat the entry cost as sunk, and so entry costs (and who pays them) do not enter into the resource constraint of the economy.

Setting the cost of multinational entry to infinity results in a leftward shift in the aggregate productivity distribution and a 12% increase in the mass of firms. The shift in the productivity distribution is illustrated in Figure 12.

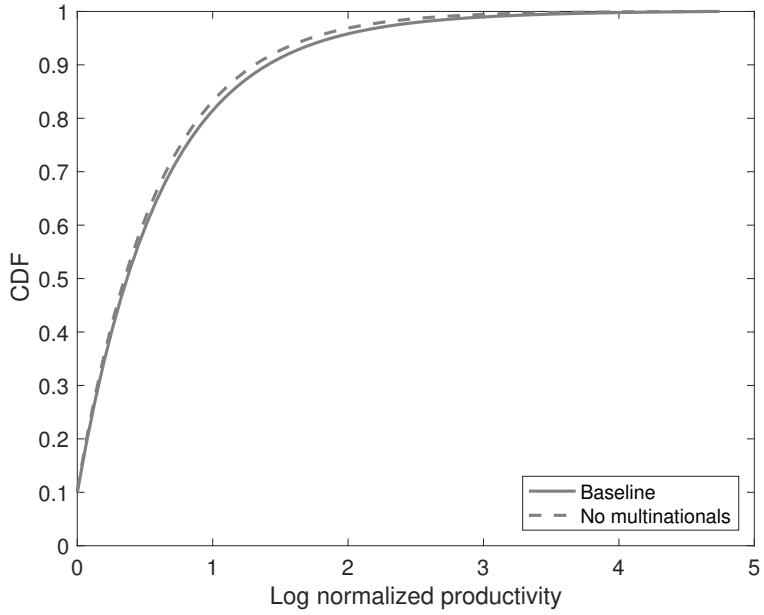


Figure 12: Shift in cdf of aggregate productivity under counterfactual of no multinationals

6.1 Impact on components of output

Table 5 reports the impact of restricting multinational entry on output, and on how output is distributed across workers, owners of domestic firms, owners of multinational firms, and owners of capital. Restricting multinational entry reduces output to 73% of its level in the baseline. This is intuitive, as average productivity falls in the economy. Wages fall by a little less than output, as the labor share increases marginally. Profits at domestic firms rise relative to the baseline economy with multinationals, while obviously, there are no profits at multinational firms in the counterfactual. Note that the reduction in the wage bill more than offsets the increase in domestic firm profit, so income of domestic agents falls when multinational entry is restricted (this ignores ownership of capital). On aggregate, multinational presence benefits domestic workers, and hurts domestic firms.

Table 5: Impact of restricting multinational entry on components of output

	Level		Share of output	
	Baseline	No MN	Baseline	No MN
Payments to labor	1	0.74	0.60	0.60
Domestic firm profit	1	1.18	0.03	0.05
Foreign firm profit	1	0.00	0.03	0.00
Payments to capital	1	0.73	0.01	0.01
Hiring cost	1	0.70	0.33	0.33
Output	1	0.73		

6.2 Impact on workers

Table 6 summarizes the impact of restricting multinational entry on the labor market. Employment rise when multinational entry is restricted. This may seem a little counterintuitive, but the intuition is straightforward within the context of the model. Remember that the value to a worker of unemployment is given by:

$$U = \frac{b}{1-\beta} + \frac{\beta\lambda\phi}{1-\beta}\int_{\underline{p}}^{\bar{p}} \frac{1-F(x)}{1-\beta(1-\delta)(1-\phi\lambda s)(1-F(x))}dx$$

Multinational presence shifts the productivity distribution in such a way that it increases the option value of being hired out of unemployment. This acts like an increase in the utility flow in unemployment, b. As workers are more picky, entry and vacancy posting respond in equilibrium in such a way that employment falls.

Because employment expands when multinational entry is restricted, the average wage (0.73) falls by more than the wage bill (0.74) relative to the case where multinationals are present.

Finally, multinational presence leads to greater wage inequality, as the wage Gini coefficient is 0.61 with multinationals relative to 0.55 when multinational entry is restricted. This is illustrated in Figure 13 which plots the share of wage income against the share of workers ranked by income in the baseline and in the counterfactual. To understand the intuition for this shift in inequality, it helps to consider the expression for the wage for a worker at firm of type p with outside option $q \leq p$:

$$w(q,p) = \phi p + (1-\phi)q - \int_q^p \frac{(1-\phi)^2\beta(1-\delta)\lambda s(1-F(x))}{1-\beta(1-\delta)(1-\phi\lambda s)(1-F(x))}dx$$

The wage has three components. The first is linear in the productivity of the worker's

current firm, p. The second is linear the worker's outside option, q. The final integral term is a discount due to the option value of moving up the job ladder within the current firm which depends on both p and q as well as the distribution of outside options and the matching probability λ. It turns out that this latter term reduces wages by less than 1% at the combination of p and q for which the discount is biggest. So though this discount declines when multinational entry is restricted compared to the baseline, its contribution to the change in wage inequality is small.

Instead, the two main contributors to the decrease in wage in equality when multinational entry is restricted are (a) the thinner tail of the distribution of p, and (b), shifts in the distribution of workers across outside options conditional on p. These shifts are illustrated in Figure 14, where we plot for a low value of p the distribution of workers across outside options in the baseline and when multinational entry is restricted, and in Figure 15, where we plot for a high value of p the distribution of workers across outside options in these two cases. The outside offer distribution shifts in opposite directions in these two cases. For workers at low productivity plants, restricting multinational entry improves the distribution of outside options. For workers high at high productivity plants, restricting multinational entry makes the distribution of outside options worse. Together, these shifts reinforce the reduction in inequality from the thinning of the tail of the firm productivity distribution.

Table 6: Labor market impact of restricting multinational entry

	Level	
	Baseline	No MN
Employment	0.850	0.855
Output per worker	1	0.73
Average worker-level wage	1	0.73
Wage Gini coefficient	0.61	0.55

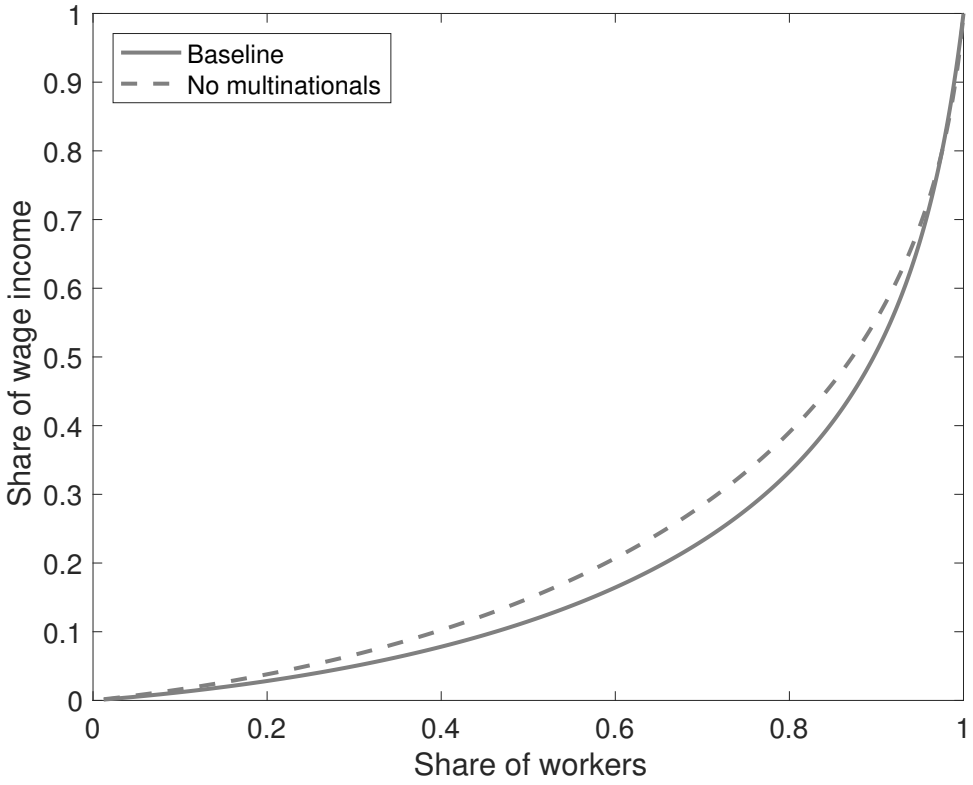


Figure 13: Impact of removing multinationals on wage inequality

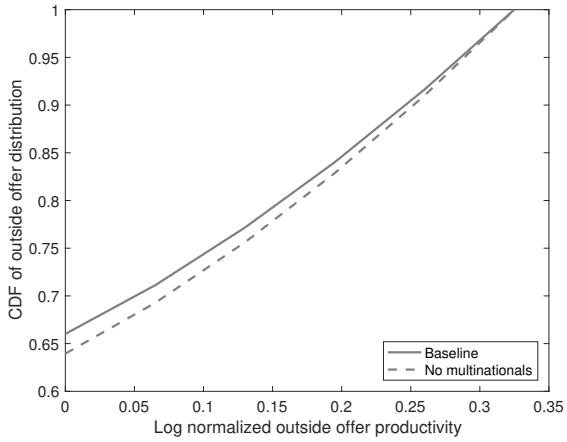


Figure 14: Low p firm

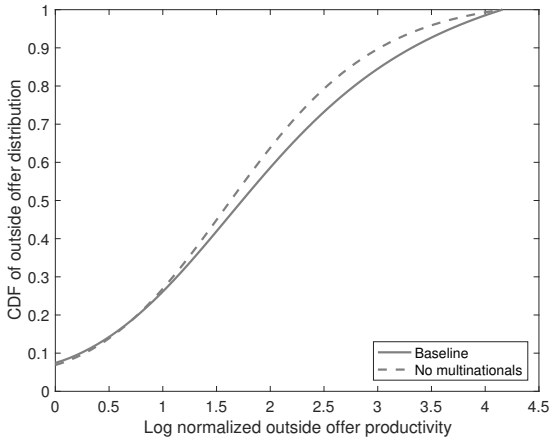


Figure 15: High p firm

6.3 Impact on local firms

Finally, we examine the impact on local firms of removing the competition in the labor market provided by multinational affiliates. Table 7 reports the impact on the number of local firms, and their average size. Both the number of local firms and the average size of local firms is higher when multinational entry is restricted compared with the baseline.

However the impact is not uniform across local firms of all types. Both very low productivity local firms, and very high productivity local firms are actually smaller in the absence of multinationals than when multinationals are present. This is illustrated in Figure 16, which shows how firm size conditional on productivity differs in the baseline case from the case when multinational entry is restricted.

Meanwhile, although overall profits of domestic firms are higher when multinational entry is restricted than in the baseline case, the impact on profits varies across the productivity distribution. Low productivity local firms are actually more profitable when multinationals are present than when they are not, while the opposite is true for high productivity local firms. This is illustrated in Figure 17, which shows how profit conditional on productivity differs in the baseline case from the case when multinational entry is restricted. This is because labor market competition is actually less intense at the low end of the job ladder when multinationals are present, while being more intense at the top of the job ladder, as Figures 14 and 15 illustrate.

Table 7: Impact on local firms of restricting multinational entry

	Level	
	Baseline	No MN
Average firm size	10.29	9.28
Average domestic firm size	9.08	9.28
Mass of firms	1	1.12
Mass of domestic firms	1	1.19

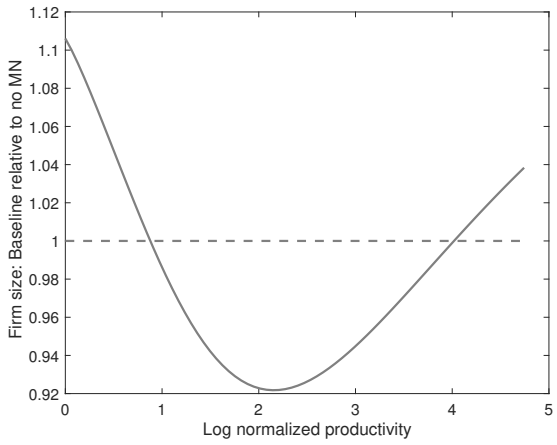


Figure 16: MN presence and firm size

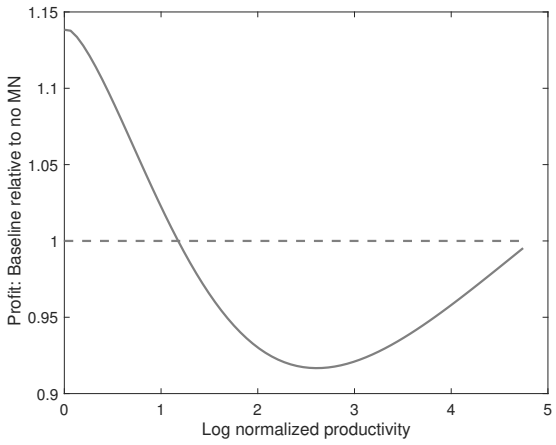


Figure 17: MN presence and firm profit

6.4 Robustness

To be completed.

6.5 Relation to the empirical literature

The results of our counterfactual exercise are consistent with the findings of the empirical literature on the impact of multinationals through the labor market. Our model is consistent with the existence of a multinational wage premium, which has been extensively documented using data for many countries.

The findings of our counterfactual that multinational presence has an impact on the wages workers employed by local firms through their outside options is also supported by the findings of the empirical literature. Both Alfaro Urena et al. [2021] and Setzler and Tintelnot [2021] find a positive impact of (instrumented) multinational presence in a local labor market on the wages of those employed at domestic firms. Perhaps even more relevant to our results, Setzler and Tintelnot [2021] find that the impact is bigger for high-paid workers. This is consistent with what we illustrate in Figures 14 and 15. Because multinationals increase competition for workers at the top of the job ladder relative to the bottom, their presence has a bigger impact on the wages of high-paid workers than that of low-paid workers.

7 Conclusions and future work

Governments, both at the level of countries, and at the level of states and cities, frequently provide incentives for multinational firms to set up affiliates in their jurisdictions. We show that on aggregate, multinational presence can benefit local economies through increasing the total wage bill as well as labor productivity, and that this may outweigh any negative effects on local firms due to increased competition in the labor market. At the same time, multinational presence may increase inequality by increasing wages for workers at high-wage high-productivity firms relative to workers at low-wage low-productivity firms. It may also increase unemployment through exit of locally owned firms.

These impacts do not necessarily justify government intervention, as subsidies to multinationals must be funded by (potentially distortionary) taxation. However we have developed a framework that allows us to provide a positive analysis of the impact of taxes on / subsidies to multinationals in the presence of an imperfectly competitive labor market.

References

John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999.

Alonso Alfaro Urena, Isabela Manelici, and Jose Vasquez. The effects of multinationals on workers: evidence from costa rican microdata. Technical report, Princeton University, Department of Economics, Center for Economic Policy, 2021.

Pol Antrás and Stephen R. Yeaple. Multinational Firms and the Structure of International Trade. In G. Gopinath, E. Helpman, and K. Rogoff, editors, *Handbook of International Economics*, volume 4 of *Handbook of International Economics*, pages 55–130. Elsevier, 2014.

Jesper Bagger and Rasmus Lentz. An Empirical Model of Wage Dispersion with Sorting. *The Review of Economic Studies*, 86(1):153–190, 05 2018. ISSN 0034-6527.

Kenneth Burdett and Dale T. Mortensen. Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2):257–273, 1998. ISSN 00206598, 14682354.

Pierre Cahuc, Fabien Postel-Vinay, and Jean-Marc Robin. Wage bargaining with on-the-job search: Theory and evidence. *Econometrica*, 74(2):323–364, 2006.

A Kerem Coşar, Nezih Guner, and James Tybout. Firm dynamics, job turnover, and wage distributions in an open economy. *American Economic Review*, 106(3):625–63, 2016.

Chiara Criscuolo and Ralf Martin. Multinationals and US productivity leadership: Evidence from Great Britain. *Review of Economics and Statistics*, 91(2):263–281, 2009.

Niklas Engbom and Christian Moser. Earnings inequality and the minimum wage: Evidence from brazil. Technical report, National Bureau of Economic Research, 2021.

Pablo D Fajgelbaum. Labour market frictions, firm growth, and international trade. *The Review of Economic Studies*, 87(3):1213–1260, 2020.

Emilien Gouin-Bonenfant et al. Productivity dispersion, between-firm competition, and the labor share. In *Society for Economic Dynamics. Meeting Papers*, volume 1171, 2018.

David Greenaway and Richard Kneller. Firm heterogeneity, exporting and foreign direct investment. *Economic Journal*, 117(517):F134–F161, 2007.

Stefanie A. Haller. Intra-firm trade, exporting, importing, and firm performance. *Canadian Journal of Economics/Revue canadienne d'économique*, 45(4):1397–1430, 2012. ISSN 1540-5982.

John C. Haltiwanger, Henry R. Hyatt, Lisa B. Kahn, and Erika McEntarfer. Cyclical Job Ladders by Firm Size and Firm Wage. *American Economic Journal: Macroeconomics*, 10 (2):52–85, April 2018.

Elhanan Helpman, Marc J Melitz, and Stephen R Yeaple. Export versus fdi with heterogeneous firms. *American economic review*, 94(1):300–316, 2004.

Elhanan Helpman, Oleg Itskhoki, and Stephen Redding. Inequality and unemployment in a global economy. *Econometrica*, 78(4):1239–1283, 2010.

Elhanan Helpman, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen J Redding. Trade and inequality: From theory to estimation. *The Review of Economic Studies*, 84(1):357–405, 2017.

Giuseppe Moscarini and Fabien Postel-Vinay. The cyclical job ladder. *Annual Review of Economics*, 10(1):165–188, 2018.

Fabien Postel-Vinay and Jean-Marc Robin. Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, 70(6):2295–2350, 2002.

Bradley Setzler and Felix Tintelnot. The effects of foreign multinationals on workers and firms in the united states. *The Quarterly Journal of Economics*, 136(3):1943–1991, 2021.

Eiichi Tomiura. Foreign outsourcing, exporting, and FDI: A productivity comparison at the firm level. *Journal of International Economics*, 72(1):113–127, 2007.