

The Impact of Multinationals Along the Job Ladder*

Ragnhild Balsvik[†] Doireann Fitzgerald[‡] Stefanie Haller[§]

October 2023

Abstract

Multinational affiliates are more productive than domestic firms, so how does their presence in the labor market impact workers and domestic firms? We use matched employer-employee data for Norway to show that the labor market is characterized by a job ladder, with multinational affiliates located on the upper rungs. Motivated by these facts, we build a general equilibrium job ladder model with endogenous entry of multinational affiliates. We calibrate the model to match moments from the Norwegian data, including the size distribution of domestic and multinational establishments. We perform a counterfactual where multinationals face an infinite entry cost. This reduces payments to labor and increases total profits of domestic firms. But the impact is not uniform, because in the absence of multinational affiliates, competition for workers increases low down on the job ladder, while it decreases on the upper rungs. Meanwhile, unemployment is lower in the absence of multinationals, because the lower density of high productivity firms reduces the option value of continuing to search.

1 Introduction

There is considerable evidence that multinational affiliates are more productive than domestic firms.¹ Given that this is the case, how does the presence of multinational affiliates

*Expert research assistance was provided by Sang Min Lee and Rene Dias de Leon. We thank Manuel Amador, Nik Engbom, Rasmus Lentz, Jeremy Lise, Guido Menzio, Simon Mongey and Chris Moser for very helpful discussions, and seminar participants at various institutions and conferences 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, ragnhild.balsvik@hi.no

[‡]Federal Reserve Bank of Minneapolis, doireann.fitzgerald@gmail.com

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

¹See e.g. Antrás and Yeaple [2014], Criscuolo and Martin [2009], Greenaway and Kneller [2007], Haller [2012] and Tomiura [2007].

in a host country affect workers and domestic firms? The answer to this question matters because governments often use costly policy tools to provide incentives for multinational affiliates to locate in a particular jurisdiction. The tradeoffs in using these tools depend on how payments to labor and the profits of domestic firms are affected by multinational presence. While there are several additional channels through which multinational presence can affect the host economy, including output market competition, supply chain linkages, and productivity diffusion, we focus on their impact through their presence in the labor market.

How multinational presence in the labor market affects workers and domestic firms depends on the nature of labor market competition. We use matched employer-employee data for Norway to document that job-to-job transitions are frequent, and directed. Ranking establishments by their share of hires from employment (“poaching rank”), workers tend to move from establishments with a low rank to those with a high rank. That is, there is a job ladder in the poaching rank. Multinationals are located on the upper rungs of this ladder, in that they tend to hire more from employment than domestic establishments.

Based on these facts, we build a model that combines on-the-job search and wage bargaining as in Cahuc et al. [2006], endogenous vacancy posting by firms which are heterogeneous in productivity, and a firm free entry condition. We allow entry costs and the productivity distributions that potential entrants draw from to depend on multinational status. This can rationalize a productivity distribution for multinational affiliates which is shifted to the right relative to that for domestic firms, capturing selection into multinational status as in Helpman et al. [2004]. Conditional on a productivity draw, multinational affiliates and domestic firms in our model behave identically in the labor market, but multinational affiliates rebate profits to foreign owners, while the profits of domestic firms are owned by domestic residents.

In the model, firms pay workers a markdown on their marginal productivity. The markdown for a particular worker depends on that worker’s current outside option as well as the distribution of outside options in the labor market. Through on-the-job search, workers may improve their outside options within their current matches, or receive opportunities to make job-to-job transitions to firms that are better than their current matches. The distribution of outside options is an equilibrium object which depends on the entire firm productivity distribution, and therefore on the presence of multinational affiliates. By affecting outside options, multinational presence affects the wages and mobility of workers employed by domestic firms, and by extension, the incentives of domestic firms to post vacancies, and the profits they make.

We calibrate the model to match key moments of the Norwegian labor market, including

worker transitions, the nonemployment rate, the aggregate labor share, the distribution of average establishment-level wages, the size distribution of establishments, the share of establishments that are multinational, and the size distribution for multinational affiliates. Our calibration implies that multinational potential entrants pay a high cost to get a draw from a productivity distribution with a fat right tail, while domestic potential entrants 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 upper rungs of the job ladder, and rationalizes the higher wages that they pay compared to local establishments. We verify that the model does a good job of matching non-targeted moments, such as the poaching index distribution for domestic and multinational establishments.

In our counterfactual experiment, we set the entry cost for multinational affiliates to infinity, and use the domestic firm free entry condition to solve for the counterfactual mass of firms and active firm productivity distribution. Comparing baseline and counterfactual, we see that restricting multinational entry reduces average wages, as it reduces workers' productivity on average, while it increases the profits of domestic firms. Domestic income available for consumption (the sum of payments to labor and domestic firm profits, less steady state investment in creating new domestic firms) falls to 91% of its baseline level.

Notably, the impact of restricting multinational entry on wages is not uniform across the productivity distribution. Wages fall overall, but average wages actually rise at intermediate-productivity firms. This is because firms low down on the productivity distribution have a greater incentive to post vacancies in the absence of multinationals, because they find it easier to hire, and are less likely to lose workers to job-to-job transitions. This leads to a better set of outside options for workers employed at intermediate-productivity firms. In contrast, at high-productivity firms, outside options are worse in the absence of multinationals, leading to lower average wages. Meanwhile, unemployment falls in the absence of multinationals, because the lower density of high productivity firms reduces the option value of continuing to search. These impacts imply heterogeneous effects on firm profits across the productivity distribution.

Our baseline model and calibration assumes homogeneous labor. We extend the model to allow for multiple worker skill types, and a production function with complementarity between worker skill and firm productivity. Complementarity generates sorting, which is a feature of the data. We calibrate the extended model with three skill types, and perform the same counterfactual of setting multinational entry costs to infinity. The impact is bigger than in the model with homogeneous labor: income for domestic consumption falls to 81%

of its level when multinationals are present. Given complementarity, the skill premium declines. In addition, there are heterogeneous impacts within skill groups, due to changes in the distribution of outside options, just as in the model with homogeneous labor. Overall inequality declines when multinational entry is restricted, with a substantial contribution from the decline in within-group inequality.

Our paper is closely related to Alfaro Ureña et al. [2021] and Setzler and Tintelnot [2021]. Both of these papers combine an instrumental variables strategy with differencing across labor markets to disentangle the impact of multinational presence on the labor market from the impact of the conditions or policies which led the multinationals to set up affiliates in the host region in the first place.² Although our methodological approach is quite different, our findings are complementary. Both papers 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 domestic firms. In addition, Setzler and Tintelnot [2021] find that these effects are bigger for high-paid workers than low-paid workers, consistent with the predictions of our counterfactuals.

Second, our work is related to several recent papers which apply general equilibrium job ladder models to a variety of questions: quantifying the contribution of sorting to wage dispersion (Bagger and Lentz [2018]), the relationship between productivity dispersion and the labor share (Gouin-Bonenfant [2022]), and the impact of a minimum wage on inequality (Engbom and Moser [2022]). 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 on the impact of outside options on wages, 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

²These papers also model multinational presence in a frictional labor market, but neither does so using a search model.

work is closest to Fajgelbaum [2020] and Ma et al. [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. The fifth section describes the calibration. In the sixth section we perform a counterfactual where we remove multinationals. The final section concludes.

2 Data description

We work with four different data sets administered by Statistics Norway. The most important data set for our analysis is the Population Register, which has annual files on the population aged between 16 and 74, with identifiers allowing us to follow individuals over time. These files include an establishment/workplace identifier for the main employer for people in the labor force, as well as the industry and municipality of the workplace, recorded in November of each year. From this source, we also 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.

Our second source of information is income tax files which include both establishment and firm identifiers, allowing us to allocate establishments to firms.³ 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 establishments as domestic- or foreign-owned based on the ownership of the associated firm identifier from the income tax files. Throughout the paper, we refer to domestic-owned establishments as “domestic,” and foreign-owned establishments 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 an establishment panel for the years 1996 to 2007 based on the establishment identifiers in the Population Register and income tax files. From this panel, we drop an establishment if it is not observed in both data sources for more than half of its years in the panel. We also drop establishments that have many years with missing location

³These files include information on job spells within a year, but the data are noisy, and we do not make use of them.

or industry affiliation. This affects 10% of the initial establishment-year observations. We further drop workplaces in the public sector, which account for 20% of the remaining sample.⁴ We also drop very small establishments 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 establishment-year observations.⁵

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 an establishment that is present in our final establishment sample.

Our dataset covers more than 1.2 million workers in nearly 117,000 establishments per year on average. It includes all sectors except the public sector, covering 12 NACE letter industries and 206 NACE 3-digit industries. These establishments are located in 160 local labor 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 establishment in two successive Novembers is a job stayer. A worker who is employed at one establishment in November of year t and at a different establishment 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 establishments 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 establishments, so as long as the misclassification is similar across establishments, especially along the dimension of ownership, this approach serves our purpose.

To measure wages, we make use of annual earnings variable from the Population Register. 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

⁴We also drop the very few workplaces that are classified as private households and extraterritorial organisations.

⁵The public sector accounts for around 40% of employment, while the remaining dropped observations account for about 2% of employment.

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 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 establishment.

Summary statistics on establishments and workers are presented in Table 1. We report statistics for all establishments and workers, for domestic establishments and their employees, and for multinational establishments and their employees, noting as above, that “multinational” refers to establishments that are foreign-owned. On average, about 6% of establishments 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, multinational establishments are bigger, pay more, and have better-educated workers than domestic establishments.

Table 1: Summary statistics on workers and establishments

	All		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	
	Establishment-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: ^aLog wage is the residual from a regression of log wage at the worker level on year dummies. ^bCognitive scores (1-9) are available from military records for men born between 1950 and 1993. ^cMedium-skilled workers are those with some high school, high school completed, or with a vocational degree. ^dHigh skilled workers have a BA or above.

In Figure 1, we report the industry composition of employment and establishments, for domestic and multinational establishments separately. Multinational affiliates are not concentrated in any one sector, but distributed across sectors in a pattern that is roughly similar to that for domestic establishments. The pattern for employment is similar.




Figure 1: Industry composition by ownership

Notes: Left panel shows share of employment by industry for workers employed by domestic establishments, and share of employment by industry for workers employed by multinational affiliate establishments. Right panel shows share of domestic establishments by industry, and share of multinational affiliate establishments by industry. See Appendix for corresponding table.

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 establishment 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. This points to a important role for job-to-job transitions in the functioning of the labor market.

Job-to-job transitions are not random, but have a strong directional component. Some establishments systematically 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: establishment size, average establishment-level wage, or an establishment’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 an establishment’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.

We construct the poaching index following the approach of Bagger and Lentz [2018]. For each establishment, we pool all hires over the sample period, and calculate the fraction of these hires that come from employment. This index is therefore fixed for a given establishment over time. In making use of the poaching index, we restrict attention to establishments making at least 10 hires over the sample period, at least one of which is from non-employment.⁶ Because of 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 establishments with a low poaching index to establishments with a similar or higher poaching index. This is illustrated in Figure 2, which uses a heatmap to illustrate the transition matrix for these workers across establishments 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. For transitions originating in the top decile, more than 60% are to plants in deciles 9 and 10. Overall, 66% of job-to-job transitions are to an establishment in the same or higher decile of the poaching index relative to the origin establishment (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 broadly declining in an establishment’s position on the job ladder, again indicating that the poaching index captures the attractiveness of establishments to workers (see Figure 3).

⁶These establishments account for 82% of employment-years and 39% of establishment-years in our baseline data. Summary statistics for this sample are similar to those presented in Table 1 (reported in the Appendix).




Figure 2: Transition heatmap




Figure 3: Poaching index and separation rate

Notes: Figure 2 is a heat map of the probability distribution of job-to-job transitions originating in an establishment of a given poaching index decile, and for which the poaching index decile of the destination establishment is known, by poaching index decile of the destination establishment. 13% of transitions originating in an establishment for which the poaching index is defined are to an establishment for which it is not defined. Figure 3 shows the share of hires from employment (i.e. the poaching index) and the separation rate by poaching index decile for establishments for which the poaching index is defined. See Appendix for corresponding tables.


The rungs of the job ladder based on the poaching index are very similar to the rungs based on establishment mean wages, and are correlated with, though not identical to the rungs based on establishment size. This is illustrated in Figure 4. Panel (a) shows the mean poaching index by percentiles of establishment mean wage. This relationship is monotonically increasing. Panel (b) 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 establishment-level mean log wage.

Table 2: Spearman rank correlations for alternative establishment 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 establishment-years for which the poaching index is defined, $N = 395,551$. All correlations are significant at the 1% level.

In the Appendix, we document additional facts consistent with the job ladder view of the labor market. Worker age and tenure are increasing in the poaching index, as one would



(a) The poaching index and wages



(b) The poaching index and size

Figure 4: The poaching index and alternative measures of establishment rank

Notes: Panel (a) shows the relationship between the poaching index and establishment-level wages. To construct this panel, we average the mean log wage at the establishment level across all years the establishment appears in the sample. Establishments are then divided into percentiles based on this variable. The mean poaching index across all establishments in a given percentile for which the poaching index is defined is then plotted on the y-axis. Panel (b) shows the relationship between the poaching index and establishment size. To construct this panel, we average employment across all years the establishment appears in the sample. Establishments are then divided into percentiles based on this variable. The mean poaching index across all establishments in a given percentile for which the poaching index is defined is then plotted on the y-axis. There are no establishments in the lower percentiles of the right panel, as we require establishments 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.

expect if it takes time to climb the job ladder, and if workers high up on the ladder are less likely to meet with a new firm that is better than their current firm. We also show that mean wage gains conditional on making a job-to-job transition are greater than mean wage gains for job stayers.

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 establishments 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 establishments first-order stochastically dominates that for domestic establishments: multinationals sit disproportionately on the higher rungs of the job ladder.⁷ While they account for only 5% of establishments in the first decile, multinationals account for more than 18% of establishments in the 10th decile of the poaching index.

⁷This is not driven by a higher proportion of cross-establishment within-firm transfers for multinationals, as the picture is very similar at the firm level, see Appendix.




Figure 5: Poaching index distribution by ownership

Notes: Kernel density distribution of the poaching index by establishment ownership. The poaching index is constructed as described in the text.

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: The poaching index is constructed as described in the text. Summary statistics are calculated using establishment-years for which the poaching index is defined.

Given the existence of a job ladder, it should not come as a surprise that multinationals are on the upper rungs. As noted in the Introduction, there is a lot of evidence that multinationals have a productivity advantage over domestic establishments. 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.

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

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, we take elements from Helpman et al. [2004], which features selection into multinational status, and merge them with a general equilibrium version of Cahuc et al. [2006]. We first describe the baseline model, and then sketch two extensions.

Time is discrete. The agents in the model are firms (domestic and multinational) and workers. Workers are homogeneous: relaxing this is one of our extensions. They search for jobs when unemployed and on-the-job. There is random matching between searching workers and vacancies. Conditional on a match, firms and workers split match value according to the bargaining protocol in Cahuc et al. [2006], such that workers get their outside option, plus a fixed fraction of match surplus. Firms have heterogeneous productivity. They face a convex cost of posting vacancies, which delivers finite firm size. Firms die stochastically, and new firms enter to replace them. We look for a stationary equilibrium, with a constant measure of firms. The equilibrium measure 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 may pay a different cost to get a productivity draw, and their draw may come from a different distribution from that faced by domestic entrants. In addition, multinational profits are owned by foreigners, not domestic agents.

4.1 Assumptions: Workers

There is one worker type, with skill normalized to 1. There is a continuum of infinitely-lived workers 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 meeting with a firm, workers accept offers that make them better off, as we describe when we lay out the assumptions about bargaining over match value and the evolution of outside options. If they accept an offer, workers supply 1 unit of labor, and receive a wage, which depends on the details of bargaining over match value and outside options. Employed workers search for jobs with probability $s \leq 1$ each period. The match between a worker and a firm is broken with exogenous probability δ_m each period. Firms also die with probability δ_f . So from the worker perspective, matches die with exogenous probability $\delta = \delta_m + \delta_f - \delta_m\delta_f$ each period. Once a match breaks exogenously, the worker must 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.

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 productivity p is p . Like workers, firms discount the future at rate β , but in addition, firms die with probability δ_f each period. Firms are born with no workers. Matches between workers and surviving firms are broken with probability δ_m each period. Each period a firm pays cost $c(v)$ in units of output to post vacancies $v \in \mathbb{R}$, with $c(0) = 0$, $c'(v) > 0$ and $c''(v) > 0$. A firm of productivity 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 assumptions about bargaining over match value. We describe this problem in more detail after laying out the assumptions about bargaining over match value and the evolution of outside options. Due to the linearity of the production function, the vacancy posting decision does not depend on the firm's current stock of workers.⁹

Let $B(v(p), p) = B(p)$ denote the value to a firm with productivity p of posting its optimal measure of vacancies. The value of an entrant (i.e. a firm with no employees) of productivity p in a stationary equilibrium is equal to the present value of the flow $B(p)$, appropriately discounted.

Let \tilde{m}_i be the measure of potential entrants of type $i \in \{D, F\}$, i.e., domestic or multinational. We assume that each potential entrant of type i pays a cost $C_i(\tilde{m}_i)$ in units of output to get a productivity draw from a distribution with cdf $\tilde{\Gamma}_i(p)$ and pdf $\tilde{\gamma}_i(p)$, defined on $[b, \bar{p}]$. Entry costs are convex: $C'_i(\cdot) > 0$, $C''_i(\cdot) > 0$.¹⁰ The free entry condition for a potential entrant of type i is:

$$C_i(\tilde{m}_i) = \int_b^{\bar{p}} \frac{B(p)}{1 - (1 - \delta_f)\beta} \tilde{\gamma}_i(p) dp \quad (1)$$

As we 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 potential entrant receiving a productivity draw below \underline{p} exits immediately. This implies that \tilde{m}_i , the measure of potential

⁹Linearity of the marginal product of labor is key to tractability of the model.

¹⁰We assume convexity and sufficient assumptions on $C_D(0)$ and $C_F(0)$ to guarantee an equilibrium with positive measure of both domestic and multinational firms.

entrants of type i , and m_i , the measure of actual entrants of type i , are related as follows:

$$m_i = \frac{\tilde{m}_i}{1 - \tilde{\Gamma}_i(\underline{p})} \quad (2)$$

In the stationary equilibrium, the mass of entrants of type i (i.e. potential entrants with $p \geq \underline{p}$) is equal to the mass of exiting firms, $\delta_f M_i = m_i$. The measures of active domestic and multinational firms, M_D and M_F , are then pinned down by the domestic and multinational free entry conditions together with the cutoff \underline{p} . The 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:

$$\gamma(p) = \frac{M_D}{M_D + M_F} \left(\frac{\tilde{\gamma}_D(p)}{1 - \tilde{\Gamma}_D(\underline{p})} \right) + \frac{M_F}{M_D + M_F} \left(\frac{\tilde{\gamma}_F(p)}{1 - \tilde{\Gamma}_F(\underline{p})} \right) \quad (3)$$

4.3 Assumptions: Matching

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

$$V = M \int_{\underline{p}}^{\bar{p}} v(p) \gamma(p) dp. \quad (4)$$

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

$$S = u + s(1 - \delta)(1 - u). \quad (5)$$

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

$$\lambda = \frac{\mu(S, V)}{S}, \quad (6)$$

while the probability that an employed worker meets a vacancy is λs . The probability that a vacancy matches with a worker is:

$$\chi = \frac{\mu(S, V)}{V}. \quad (7)$$

4.4 Assumptions: Bargaining between workers and firms

Bargaining between workers and firms, and the associated implementation through wages, is as described in Cahuc et al. [2006]. When a worker and a firm match, they split the value of the match, i.e. the appropriately discounted flow of p , where discounting takes account of future separations. The worker is offered a value equal to their outside option, plus a fraction ϕ of match surplus, i.e. the value of the match less the value of their outside option. We assume this split is implemented by workers receiving a constant wage w until their outside option changes.¹¹

Let U denote the value of unemployment. Let $w_0(p)$ denote the wage of a worker at firm of productivity p whose outside option is unemployment. Let $W(w, p)$ denote the value to a worker of receiving wage w at firm of productivity p . Let $w(q, p)$ denote the wage of a worker at firm of productivity p whose outside option is a firm of productivity q . An unemployed worker meeting a firm of productivity p accepts an offer $w_0(p)$ such that:

$$W(w_0(p), p) = U + \phi(W(p, p) - U). \quad (8)$$

Suppose a worker at firm of productivity p whose outside option is q meets a firm of productivity 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:

$$W(w(p', p), p) = W(p', p') + \phi(W(p, p) - W(p', p')). \quad (9)$$

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

$$W(w(p, p'), p') = W(p, p) + \phi(W(p', p') - W(p, p)). \quad (10)$$

4.5 Equilibrium definition

An equilibrium consists of a mass of domestic firms M_D , a mass of multinational affiliates M_F , a lower bound \underline{p} for productivity below which no firm operates, and a vacancy posting function $v(p)$ describing the measure of vacancies posted by firms of type p such that:

¹¹This is an assumption, not a result. Transitions are pinned down by the assumption on the surplus split, but the same transitions could be supported by alternative wage policies consistent with the same surplus sharing rule.

1. Worker optimality: Given match surpluses and meeting rates, worker mobility decisions maximize utility.
2. Firm optimality: Given match surpluses and meeting rates, firms choose vacancy posting $v(p)$ to maximize $B(v, p)$.
3. Labor market clearing: Worker mobility decisions and optimal vacancy posting deliver a stationary distribution of workers across firms of different productivity.
4. Free entry: Given $B(v(p), p)$, the mass of domestic and multinational potential entrants is such that the respective free entry conditions are satisfied.

4.6 Results: Worker transitions

Matched workers separate to unemployment with probability $\delta = \delta_m + \delta_f - \delta_m \delta_f$ 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 productivity 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.

4.7 Results: Wages

The pdf of the job offer distribution faced by all searching workers is:

$$f(x) = \frac{v(x) \gamma(x)}{\int_{\underline{p}}^{\bar{p}} v(y) \gamma(y) dy}, \quad (11)$$

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

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

$$w(q, p) = \phi p + (1 - \phi) q - \int_q^p \frac{(1 - \phi)^2 \beta (1 - \delta) \lambda s (1 - F(x))}{1 - \beta (1 - \delta) (1 - \phi \lambda s (1 - F(x)))} dx. \quad (12)$$

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 (i.e. climbing the job ladder) 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 productivity p who is hired from unemployment:

$$w_0(p) = w(\underline{p}, p) = \phi p + (1 - \phi) \underline{p} - \int_{\underline{p}}^p \frac{(1 - \phi)^2 \beta (1 - \delta) \lambda_s (1 - F(x))}{1 - \beta (1 - \delta) (1 - \phi \lambda_s (1 - F(x)))} dx, \quad (13)$$

where \underline{p} is implicitly defined as the level of productivity such that the unemployed are indifferent between taking an offer from a firm with this productivity and remaining unemployed, i.e. (see the Appendix for derivation):

$$\underline{p} = b + \int_{\underline{p}}^{\bar{p}} \frac{\beta \phi \lambda (1 - (1 - \delta) s) (1 - F(x))}{1 - \beta (1 - \delta) (1 - \phi \lambda_s (1 - F(p)))} dx. \quad (14)$$

In the Appendix we show that the value to a worker of receiving wage w at firm of productivity p is:

$$W(w, p) = \frac{1}{1 - \beta (1 - \delta)} \left(w + \beta \delta U + \left(\int_{q(w, p)}^p \frac{(1 - \phi) \beta (1 - \delta) \lambda_s (1 - F(x))}{1 - \beta (1 - \delta) (1 - \phi \lambda_s (1 - F(x)))} dx + \int_p^{\bar{p}} \frac{\phi \beta (1 - \delta) \lambda_s (1 - F(x))}{1 - \beta (1 - \delta) (1 - \phi \lambda_s (1 - F(x)))} dx \right) \right), \quad (15)$$

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

$$U = \frac{1}{1 - \beta} \left(b + \int_{\underline{p}}^{\bar{p}} \frac{\beta \phi \lambda (1 - F(x))}{1 - \beta (1 - \delta) (1 - \phi \lambda_s (1 - F(p)))} dx \right). \quad (16)$$

4.8 Results: Employment and wage distributions

Let $L(p)$ be the probability that an employed worker works at a firm with productivity $\leq p$. Let $l(p)$ 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, $(1 - u) L(p)$ is the *measure* 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:

$$l(p) = \left(\frac{\delta + (1 - \delta) \lambda_s L(p)}{\delta + (1 - \delta) \lambda_s (1 - F(p))} \right) f(p). \quad (17)$$

Meanwhile, let $G(q|p)$ be the cdf of the distribution of outside options for workers at firms of type p , and let $g(q|p)$ be the associated pdf. In steady state, the inflow of workers employed by firms of productivity p with outside option less than or equal to q must be equal to the outflow. In the Appendix, we show that this implies:

$$G(q|p) = \left(\frac{\delta + (1 - \delta) \lambda_s L(q)}{\delta + (1 - \delta) \lambda_s (1 - F(q))} \right) \frac{f(p)}{l(p)}. \quad (18)$$

4.9 Results: Optimal vacancy posting

Define $J(q, p)$ 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:

$$J(q, p) = \frac{(1 - \phi)}{1 - \beta(1 - \delta)} \left((p - q) - \int_q^p \frac{\phi \beta (1 - \delta) \lambda_s (1 - F(x))}{1 - \beta(1 - \delta) (1 - \phi \lambda_s (1 - F(x)))} dx \right). \quad (19)$$

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

$$B(v, p) = \max_v \left\{ \chi v \left[\frac{u}{S} J(\underline{p}, p) + \frac{(1 - u)(1 - \delta)s}{S} \left(\int_{\underline{p}}^p J(x, p) l(x) dx \right) \right] - c(v) \right\}. \quad (20)$$

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

$$\chi \left[\frac{u}{S} J(\underline{p}, p) + \frac{(1 - u)(1 - \delta)s}{S} \left(\int_{\underline{p}}^p J(x, p) l(x) dx \right) \right] = c'(v). \quad (21)$$

This first order condition implicitly defines $v(p)$, the optimal measure of vacancies posted by firm of type p . Note that $v(p)$ depends on the offer distribution $f(p)$, and therefore on the vacancies posted by all the firms in the economy. But it does not depend on the firm's current employment, or the distribution of outside options of employees of the firm.

Substituting $v(p)$ into $B(v, p)$, we obtain $B(p) = B(v(p), p)$, the value to a firm of type p of posting the optimal measure of vacancies in each period. This is what appears in the free entry condition, equation (1).

4.10 Results: Firm size distribution

In steady state, there is a size distribution of firms of productivity p , with firms that have just been born being smaller than older firms. Firms of productivity p which have survived

to age a have size

$$e(p, a) = h(p) \left(\frac{1 - z(p)^a}{1 - z(p)} \right), \quad (22)$$

where $h(p)$ is per period hires by firms of productivity p :

$$h(p) = v(p) \chi \left(\frac{u + (1 - u)(1 - \delta) s L(p)}{S} \right) \quad (23)$$

and

$$z(p) = (1 - \delta_m)(1 - \lambda s(1 - F(p))) < 1. \quad (24)$$

The fraction of firms of age a is given by $(1 - \delta_f)^{a-1} \delta_f$ (the same for all p).

4.11 Results: Ranking firms

In the absence of data on productivity, it is useful to have an alternative 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 productivity p is monotonic in p . Although *long run* employment is monotonically increasing in p , *actual* employment also depends on firm age, which we do not observe. But conveniently, a firm's share of hires from employment, which we do observe, is monotonically increasing in p . Denote this share by *poach* (p). It is given by:

$$poach(p) = \frac{(1 - u)(1 - \delta) s \int_{\underline{p}}^p l(x) dx}{u + (1 - u)(1 - \delta) s \int_{\underline{p}}^p l(x) dx}. \quad (25)$$

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, and therefore, the share of *hires* from employment will be higher for high p firms.

4.12 Extension: Capital in the production function

Suppose that the production function in firm with TFP \hat{p} is $y = \hat{p} k^\kappa l^{1-\kappa}$. If all firms face the same rental price of capital R (exogenous, set on world markets), and there are no frictions in the rental market for capital, 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. Marginal

productivity of labor 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 . See Appendix for details.

4.13 Extension: Labor heterogeneity and sorting

The model can be extended to allow for a finite number of labor skill types, $h = 1, \dots, H$. Assume there is a fixed supply z_h of each skill type, with $\sum_{h=1}^H z_h = 1$. Skill types are observable, both to workers and firms. Workers search in the market for their skill type, while firms choose the measure of vacancies to post in each skill market. Firms face a convex cost of vacancy posting in each market, $c_h(v_h)$, and the matching function, $\mu_h(S_h, V_h)$ operates at the level of the skill market.

Reservation utility, the rate of exogenous separations, and on-the-job search intensities may differ across skill types, i.e. $\{b_h, \delta_h, s_h\}$. The worker share in match surplus, ϕ_h , may also differ across skill types.

Worker marginal productivity at firm of type p may depend on the worker's skill type. Marginal productivity of type $h = 1$ is normalized to p . Marginal productivity of type h is given by $\eta_h p^{\nu_h}$, with $1 \leq \eta_2 \leq \dots \leq \eta_H$, and $1 \leq \nu_2 \leq \dots \leq \nu_H$. If $\nu_h > 1$ for some h (i.e. complementarity between firm and worker type), this will induce sorting, i.e. the share of employment by skill type will differ across firms with different productivity, with higher productivity firms having higher employment shares for high skill types than low productivity firms do.¹²

Since firm vacancy posting decisions are independent across skill markets, most of the results of the model are unchanged. However for the unconditional poaching index (i.e. the poaching index calculated unconditional on skill) to be guaranteed monotonic in firm type p , all active firms must post vacancies in all skill markets.

5 Calibration

We now describe our calibration strategy, parameter values, and fit, for the model with homogeneous labor and capital in the production function. In Section 6 we outline a calibration of the model with labor heterogeneity.

¹²Allowing for unobserved within-skill-group worker heterogeneity as in Postel-Vinay and Robin [2002] and Cahuc et al. [2006] would improve the model's ability to match worker-level wage moments. But in the absence of complementarity between worker skill and firm productivity, it would not have any interesting interaction with multinational presence, so we abstract from it.

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] and Engbom and Moser [2022] we assume that the cost of vacancy posting takes the form:

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

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 the intercept in the matching function, A , must be such that matching probabilities $\{\lambda, \chi\}$ both lie on the unit interval.

We assume that domestic entrants take productivity draws from a truncated Pareto distribution with scale parameter b , and shape parameter σ_D , while multinational firms take draws from a truncated Lognormal distribution, with parameters $\{\mu_F, \sigma_F\}$. The truncation point \bar{p} is such that fraction 0.995 of the mass of the distribution with the thickest right tail lies to the left of \bar{p} .

We do not need to take a stand on the functional forms for domestic and foreign entry costs in order to calibrate the model. For the counterfactual, we do need to make assumptions about domestic entry costs. We describe these in Section 6.

5.2 Solution algorithm

We set some 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 $\kappa = 0.25$.¹³ 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). We set $\delta_f = 0.01$ based on an annual exit rate of 4% for Norwegian manufacturing establishments reported in Balsvik and Haller [2010].

¹³The literature on production function estimation leaves a lot of uncertainty about this parameter, as it is very sensitive to the estimation approach used. See, e.g. Collard-Wexler and De Loecker [2016]. We examine the robustness of our results to different choices for this parameter.

It turns out to be convenient to treat the share of multinationals in potential entrants (call this ω), and M , the measure of active firms, as parameters. The parameters to be calibrated are then $\{s, \phi, A, \alpha, M, \omega, \sigma_D, \mu_F, \sigma_F\}$. Given a vector of values for these parameters, we discretize the domestic and multinational productivity distributions. We solve the model by first recovering the $\{\lambda, \chi, S, V\}$ consistent with the target unemployment rate. Given V , we guess a vector $\{v_i\}$ of vacancies for each productivity type i . We then use policy function iteration on the first order condition for the optimal choice of vacancies to recover the vector of vacancies such that for each firm type, the vacancy posting decision is optimal given the vacancy posting decisions of all other firms. We update our guesses of $\{\lambda, \chi, S, V\}$ in each iteration. Finally, we iteratively arrive at the cutoff value for productivity, $\underline{p} > b$, below which no firm can attract any workers. Given $\{v_i\}$, \underline{p} , M and ω , we can solve for all of our target moments. Using the free entry conditions, we can also recover entry costs at the equilibrium values of \tilde{m}_D and \tilde{m}_F , i.e. $C_D = C_D(\tilde{m}_D)$ and $C_F = C_F(\tilde{m}_F)$.

Outside of this loop, we use a particle swarm algorithm to search for the values of the parameter vector $\{s, \phi, A, \alpha, M, \omega, \sigma_D, \mu_F, \sigma_F\}$ 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 the 9 parameters $\{s, \phi, A, \alpha, M, \omega, \sigma_D, \mu_F, \sigma_F\}$ to match 9 target moments. The targets we pick include labor market transitions, the nonemployment rate, labor share, moments of the establishment size and wage distributions, and moments of multinational presence in the economy.

The precise set of target moments is motivated by several considerations. We focus on nonemployment rather than unemployment, because we do not model the labor force participation margin, and in any case many hires are made from workers formally designated as nonparticipants. Though there are alternative ways of calibrating worker bargaining power ϕ , for our question, we think it is important to match the labor share. We do not model worker-level heterogeneity in unobserved skills or match-specific amenities, so we target moments of establishment-level wages, rather than worker-level wages. We anticipate that measurement error for establishment-level wages is particularly likely in the lower tail, where there are a lot of hires from nonemployment, so we do not try to match the full range of establishment-level wages.

Table 4 lists the target moments, the source for each moment, values in the data, fitted values in the model, the (roughly) corresponding parameters, and their fitted values. The

parameters we recover for the domestic and multinational productivity distributions are consistent with the hypothesis that the productivity distribution for multinational affiliates has a higher mean, and thicker right tail than that for domestic establishments. They also imply that $C_F(\tilde{m}_F)/C_D(\tilde{m}_D) = 407$, so multinational affiliates pay a higher cost to obtain a productivity draw.

Table 4: Calibration targets and parameter estimates

Target	Data	Model	Parameter	Value
Outside data (source)				
EE quarterly transition rate (Eurostat)	0.03	0.03	s	0.56
Labor share (Statistics Norway)	0.60	0.60	ϕ	0.58
Nonemployment rate 25-54 (Statistics Norway)	0.155	0.157	A	0.32
Our data				
P99 log establishment employment	4.73	4.90	α	0.52
Average establishment size	10.29	10.26	M	0.08
Share of active establishments that are domestic	0.94	0.94	ω	0.0008
P99-P25 establishment avg log wage	1.52	1.52	σ^D	2.57
Average establishment size, MN	28.89	29.01	μ^F	0.48
P99 log establishment employment, MN	5.78	5.56	σ^F	1.64

5.4 Nontargeted moments

Having obtained the calibrated vector of parameters, we simulate our quarterly model for a panel of 1,200,000 workers over a period of 15 years (i.e. 60 quarters). We start with a cross-section where workers are allocated to firms based on the ergodic distribution, and then simulate transitions. We use the last 10 years of the simulated data to calculate the poaching index, wages and firm size as in the data. As in the data, we use a single quarterly cross-section within the year to link workers to firms. We use these cross-sections to code transitions on an annual basis, and we attribute all earnings within a calendar year (including the per-period flow b in any quarters of unemployment) to the firm the worker is matched with in that quarterly cross-section.

Panel (a) of Figure 6, plots the establishment wage distribution in the data, and the corresponding distribution in the simulated data. Panel (b) plots the corresponding picture for the establishment size distribution. The relevant targets are also plotted.¹⁴ Despite the parsimony of the target vector of moments, the model can match the shape of establishment wage and size distributions.

¹⁴Note that the calibration targets moments of the ergodic rather than simulated distributions.




Figure 6: Establishment size and wage distributions: data & model

Notes: Panel (a) plots the distribution of the establishment mean log wage in the data (normalized to 0 at the 25th percentile) and in the simulated data generated from the calibrated model. The circles indicate the targets of the calibration. Panel (b) plots the distribution of log establishment size in the data, and in the model. The circle indicates the target of the calibration. The calibration also targets average firm size, the dotted line on the right panel.

The only moments related to transitions that we target are the quarterly rates of employment to nonemployment and job-to-job transitions, taken from Eurostat. We do not make use of any moments related to or conditioned on the poaching index in calibrating the model. In Figure 7 we reproduce the mean poaching index and separation rates by deciles of the poaching index in the data and in the model. We construct the poaching index in the simulated data in the same way as in the actual data, requiring that an establishment have 10+ hires over the 10-year sample period, with at least one hire from nonemployment. Based on this criterion, the average rates of both job-to-job transitions and job separations are somewhat lower in the simulated data than in the actual data. But otherwise the model pictures look very similar to the data.

In Figure 8, we reproduce the distributions of the poaching index for domestic and multi-national establishments in the data, and the same distributions for firms in the simulated data. Qualitatively, the pictures look very similar. In the Appendix, we report performance of the model in matching the joint distribution of the poaching index, establishment-level average log wages, and establishment size. The model does a good job of replicating the patterns in the data, giving us the confidence to proceed with our counterfactual analysis.




Figure 8: Poaching index distribution: domestic & multinational

Notes: Panel (a) plots the distribution of the poaching index for domestic-owned establishments and multinational affiliates from the data, reproducing Figure 5. Panel (b) plots the distribution of the poaching index for domestic and foreign firms from the simulated data generated from the calibrated model.




Figure 7: Poaching index and separations

Notes: Panel (a) plots the mean poaching index and mean separation rate by poaching index decile from the data, reproducing Figure 3. Panel (b) plots the mean poaching index and mean separation rate by poaching index decile from the simulated data generated from the calibrated model.

6 Counterfactual: No multinational affiliates

We now examine the impact of multinational presence on workers and domestic firms by setting the entry cost for multinationals in our calibrated model to infinity, and solving for the new stationary equilibrium. This requires us to take a stand on the entry cost function for domestic firms, $C_D(\tilde{m}_D)$. In our baseline counterfactual we assume that the domestic entry

cost is insensitive to the measure of potential entrants paying the cost: $C_D(\tilde{m}_D) = C_D, \forall \tilde{m}_D$. We see this as a best-case-scenario counterfactual. In robustness analysis we present results for the case where the measure of potential entrants paying the entry cost is fixed at its level in the baseline calibration.¹⁵

Beyond our assumption on the behavior of entry costs, the results of the counterfactual also rely on (a) the thickness of the right tail of the multinational productivity distribution relative to that of the domestic productivity distribution and (b) the latent left tail of the domestic productivity distribution. For (a), our calibration relies on tail moments for identification. For (b) we rely on functional form restrictions, i.e. the assumption that domestic potential entrants draw productivity from a Pareto distribution.

Given an assumption about $C_D(\tilde{m}_D)$, we solve for the counterfactual measure of firms M^{CF} such that the domestic firm free entry condition holds. This yields a counterfactual \underline{p}^{CF} , and the counterfactual productivity distribution for active firms, with pdf:

$$\gamma^{CF}(p) = \frac{\tilde{\gamma}_D(p)}{1 - \tilde{\Gamma}_D(\underline{p}^{CF})}.$$

Having solved for the counterfactual stationary equilibrium, we compare outcomes with and without multinational affiliates.

6.1 Shift in the productivity distribution

Setting the cost of multinational entry to infinity results in a leftward shift in the active firm productivity distribution, illustrated in Figure 9. The productivity level below which no firm finds it worthwhile to post vacancies (\underline{p}) is lower in the counterfactual than in the baseline. Intuitively, local firms which would not find it profitable to operate when they have to compete for workers with productive multinational affiliates can survive when multinationals are not present in the economy.

¹⁵The assumption of a fixed measure of potential entrants is similar to that made by Cahuc et al. [2022].




Figure 9: CDF of aggregate productivity under baseline and counterfactuals

Notes: Figure plots the CDF of active firm productivity in the baseline, where productivity is a mixture over domestic and multinational productivity, and in the counterfactual, where productivity depends only on the distribution from which domestic entrants draw, along with the endogenous cutoff below which no firm can attract workers. The dotted vertical line indicates \underline{p} , the level of productivity below which no firm can attract any workers, in the baseline.

6.2 Aggregate impact

Table 5 compares output and its components in the baseline calibration with their levels in the counterfactual with no multinational affiliates. Output in the counterfactual with no multinational affiliates is 89% of its level in the baseline calibration.¹⁶ The fact that output is lower in the absence of multinational affiliates is intuitive, as average labor productivity falls. For context, note that in the baseline calibration, as in the data, 94% of establishments are domestic-owned, and they account for 82% of employment.

¹⁶Using equations (15) and (16) to calculate worker values, and the appropriate distributions to aggregate across workers, worker welfare in the counterfactual is 89% of its level in the baseline.

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

	Level		Share of output	
	Base	No MN	Base	No MN
Output	1	0.89	1	1
Payments to labor	1	0.89	0.599	0.602
Domestic firm profit	1	1.20	0.07	0.09
Foreign firm profit	1	0	0.03	0.00
Payments to capital	1	0.89	0.25	0.25
Hiring cost	1	0.88	0.06	0.06
Labor + dom profit	1	0.93	0.67	0.69
Labor + dom profit - dom entry cost	1	0.91	0.63	0.65

Notes: The counterfactual results (No MN) in this table refer to the case where domestic firm entry is assumed elastic. Left panel reports various different aggregates relative to their levels in the baseline economy. Right panel reports each of these aggregates as shares of total output.

The wage bill falls by roughly the same amount as output, while aggregate profits at domestic firms rise to 120% of their level in the baseline. This is both because there are more domestic firms, and because conditional on productivity, profits are higher at most firms. By construction, there are no profits at multinational firms in the counterfactual economy. The reduction in the wage bill more than offsets the increase in domestic firm profits, so ignoring ownership of capital (on which we have not taken a stand), income of domestic agents, i.e. wage bill plus domestic firm profit, is 93% of its baseline level when multinational entry is restricted. Resources available to domestic agents for consumption, i.e. the wage bill plus domestic firm profits less expenditures on domestic firm entry, is 91% of its baseline level.

6.3 Distributional impact

Table 6 reports the impact on workers and domestic firms in more detail. Restricting multinational entry *reduces* unemployment. This has a straightforward explanation in the context of our search model. The value of unemployment, U , depends on the flow utility of unemployment, b , and the option value of continuing to search, see equation (16). Unemployment is lower in the absence of multinationals because the option value of continuing to search is lower, making workers less picky about the jobs they accept. This is consistent with the reduction in \underline{p} illustrated in Figure 9. As a matter of accounting, the reduction in payments to labor as a result of restricting multinational entry is a combination of the positive but modest impact on employment, and the strong negative impact on the average worker-level wage.

Table 6: Impact on workers and local firms of restricting multinational entry

	Base	No MN
Nonemployment rate	0.157	0.144
Average worker-level wage	1	0.88
Wage Gini coefficient	0.32	0.32
Measure of firms	1	1.43
Measure of domestic firms	1	1.52
Average firm size	10.26	7.74
Average domestic firm size	9.00	7.74

Notes: The counterfactual results (No MN) in this table refer to the case where domestic firm entry is assumed elastic.

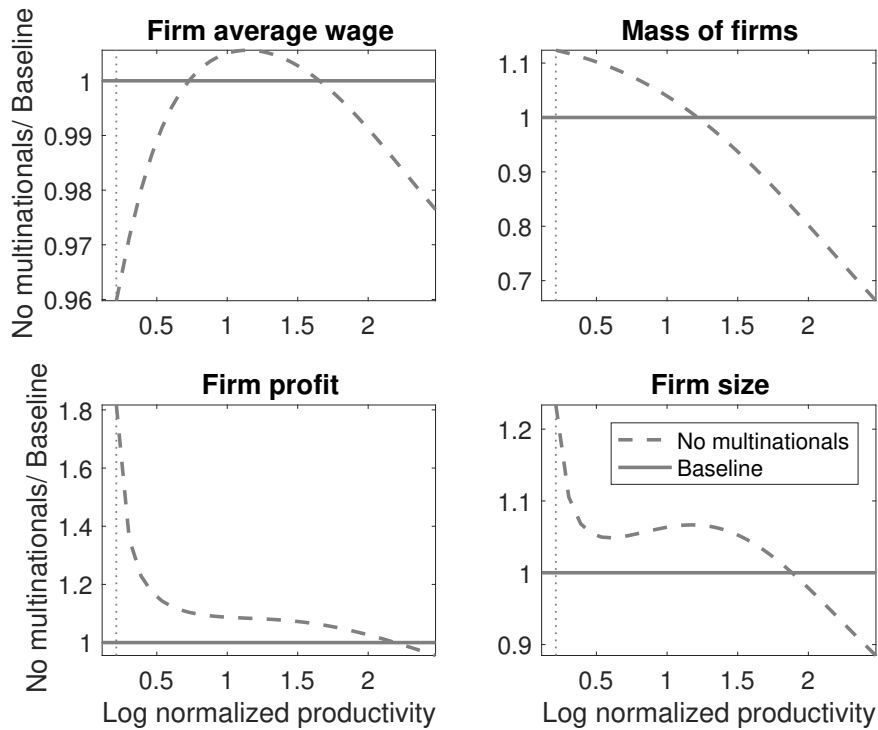


Figure 10: Heterogeneous impact of restricting multinational entry

Notes: Top left panel plots firm-level average wage in the counterfactual economy without multinationals relative to the baseline. Top right panel plots mass of firms in the counterfactual economy without multinationals relative to that in baseline. Bottom left panel plots firm profit in the counterfactual economy without multinationals relative to firm profit the baseline. Bottom right panel plots firm size in the counterfactual economy without multinationals relative to that in the baseline. Note that in each case, the variables of interest are not defined for productivity levels such that there are active firms in the counterfactual economy but no active firms in the baseline economy.

As measured by the Gini coefficient, wage inequality is largely unchanged between the baseline and counterfactual economies. However this insensitivity masks heterogeneous impacts across workers at firms of different productivity. This heterogeneity is illustrated in the top left panel of Figure 10, which plots firm-level average wages by firm productivity in the counterfactual economy with no multinational affiliates relative to the baseline. At

the lowest and highest productivity firms, wages are lower in the absence of multinational affiliates. But for firms of intermediate productivity, average wages are higher.

To understand why the impact of restricting multinational entry on firm-level average wages differs across the productivity distribution, it is useful to consider equation (12), the expression for the wage for a worker at firm of productivity p with outside option $q \leq p$. The wage is made up of three terms: the employing firm's productivity, the relevant outside option, and the discount for the option value of climbing the job ladder within the current match. A key force leading to heterogeneous differences in average wages across firms of different productivity is heterogeneous shifts in the distribution of outside options.¹⁷




Figure 11: Shifts in outside option distribution for firms of different productivity

Notes: Figure plots distribution of outside options for workers employed at firms of three different productivity levels, in the baseline economy with multinationals, and in the counterfactual economy with restricted multinational entry. The firm productivity level in each case is where the CDF equals 1, since outside options are less than or equal to the productivity of the employing firm. The vertical portion of the offer distribution in each case is where productivity is at the relevant \underline{p} , i.e. the value of productivity at which no firm can attract any worker from unemployment.

These shifts are illustrated in Figure 11, which plots the baseline and counterfactual distribution of outside options for workers employed at firms with specific low, medium, and high productivity levels. For workers at low productivity firms, average wages are lower in the absence of multinationals, due a lower value of \underline{p} . For workers at intermediate-productivity firms, this factor is at work, but it is outweighed by the fact that over the productivity

¹⁷The option value discount is *lower* in the counterfactual than in the baseline for workers with an outside option greater than \underline{p} . This tends to increase rather than decrease wages in the counterfactual.

interval where there are active firms, the distribution of outside options in the absence of multinationals is actually better than in the baseline economy. This is because in the absence of multinationals, there are more low- and medium-productivity firms, and these firms find it optimal to post more vacancies since their vacancy yield is higher (see equation (21)). Finally, for workers at high-productivity firms, the distribution of outside options in the absence of multinationals is worse than when multinational affiliates are present, as there are fewer firms high up on the job ladder, and they post fewer vacancies. This is because the vacancy yield is lower for high-productivity firms in an environment with lower unemployment.¹⁸

On the firm side, the top right panel of Figure 10 illustrates that in the absence of multinationals, there are more low-productivity firms, and fewer high-productivity firms than in the baseline (this follows directly from the productivity distribution). The bottom right panel shows that low-productivity firms are bigger, and the highest productivity firms are smaller in the absence of multinationals than in the baseline. This is again because vacancy yield moves in opposite directions for low- and high-productivity firms. Finally, the bottom left panel of Figure 10 illustrates that there are heterogeneous effects on profits across the productivity distribution. In the absence of multinationals, profits are higher than in the baseline at most levels of productivity. But they are lower at the very highest productivity firms, which are smaller than in the baseline equilibrium for the reasons just outlined.

6.4 Robustness: inelastic domestic firm entry

In an alternative counterfactual exercise, we fix the measure of potential entrants, \tilde{m}_D , at its level in the baseline calibration. The measure of active firms is still endogenous, since \underline{p} and therefore the measure of actual entrants adjust to satisfy the free entry condition. With this assumption of inelastic domestic firm entry, output in the counterfactual with no multinationals is 80% of its level in the baseline calibration. The same is true for the wage bill, while domestic firm profits are 111% of their level when multinationals are present. Meanwhile, resources available for domestic consumption are 83% of their baseline level. Overall, restricting multinational entry is more costly if domestic entry is inelastic, which is why we consider our baseline counterfactual a best-case scenario. Qualitatively, the heterogeneity in impacts is similar to that illustrated in Figure 10, for similar reasons. Full results are reported in the Appendix.

¹⁸See the Appendix for more details on shifts in vacancies and job-to-job mobility.

6.5 Robustness: values for capital share

In our baseline calibration, we assume that κ , the exponent on capital in the Cobb-Douglas production function, is equal to 0.25. As already noted, there is a good deal of uncertainty about the true value of this parameter, which moreover plays an important role in identifying ϕ , the share of match surplus that goes to labor. So we examine the robustness of our results to re-calibrating the model under the following alternative assumptions: $\kappa = \{1/3, 0.2, 0.1\}$. Unsurprisingly, the higher is κ , the higher is the calibrated value of ϕ . However vacancy posting costs (α), matching efficiency (A), and productivity parameters are also affected.

For each re-calibration, we perform the counterfactual where the multinational entry cost is set to infinity, and domestic firm entry is assumed elastic. Output falls more in the counterfactual relative to the baseline the lower is κ . However resources available for domestic consumption in the counterfactual are always between 90% and 92% of the baseline level. Qualitatively, the heterogeneous impacts we illustrate in Figure 10 are present for all the values of κ we consider. Full results of the re-calibrations and associated counterfactuals are reported in the Appendix.

6.6 Robustness: labor heterogeneity

In the data, we see labor sorting along the job ladder. The low skill share of employment is high at establishments low down on the job ladder, and low at establishments high up on the job ladder, while the reverse is true for high skill labor. This holds using educational attainment to measure skill, and also using cognitive scores obtained from military records for men born between 1950 and 1993. See the Appendix for details.

To examine the robustness of our findings to the assumption of heterogeneous labor, we calibrate a version of the model with three labor types: low, medium, and high skill. We categorize workers with less than primary, primary, and lower secondary education as low skill. Medium skilled are workers with upper secondary and post-secondary non-tertiary education. High skilled are workers with tertiary education.

We assume that the worker share in match surplus and the curvature of the vacancy posting cost function is the same for all skill types. We allow reservation utility, exogenous rates of match destruction, on-the-job search intensities, and the parameters of the matching function to differ across skill type. We also allow for complementarity between worker skill and firm productivity, as described in Subsection 4.13. In addition to the baseline set of moments, we target skill premia and the skill shares of employment in firms at the top and

bottom of the job ladder to identify the degree of complementarity. The sorting we see in the data implies modest complementarity between worker skill and firm productivity. The details of the calibration are reported in the Appendix.

When we use this extended model to perform the counterfactual exercise of restricting multinational entry (the elastic domestic entry case), the headline impact on output and resources available for domestic consumption is bigger than in the model with homogeneous labor. Income available for domestic consumption falls to 81% of the level in the economy with multinationals. Unsurprisingly given complementarity between worker skill and firm productivity, restricting multinational entry reduces both skill premia and wage inequality. The reduction in wage inequality is due not only to the reduction in the skill premium: within-skill-type wage inequality also decreases. As in the model with a single skill type, the impact of restricting multinational entry is heterogeneous across firms of different productivity. Average wages decrease at the highest productivity firms when multinational entry is restricted, while they increase at firms with low and intermediate productivity. This is true both unconditionally, and conditional on labor type. However this effect is most pronounced for high-skill workers, who also experience the biggest decrease in within-group wage inequality from restricting multinational entry. Full results are in the Appendix.

7 Conclusion

Governments, at the level of countries, states, and cities, frequently provide incentives for multinational firms to locate affiliates in their jurisdictions. It is therefore important to understand the impact of multinational presence on the host economy. We focus on the impact of multinational presence through its impact on the labor market. We use a calibrated job ladder model of the Norwegian labor market to show that because multinational affiliates are more productive than domestic firms, restricting multinational entry reduces the total wage bill by reducing labor productivity, and that this outweighs the positive impact on profits of domestic firms of reduced competition in the labor market.

These aggregate results are of quantitative interest, but are not surprising in the light of standard models of selection into multinational status with a competitive labor market. Where our model of a frictional labor market delivers new insights is in its prediction that the impact of restricting multinational presence need not be uniform across workers and local firms. By leading to a thinner right tail of the firm productivity distribution, restricting multinational presence reduces labor market competition high up on the job ladder, while

increasing competition for workers low down on the job ladder. So average wages for workers at intermediate-productivity firms may even rise when multinational entry is restricted, though wages fall overall. In addition, our model predicts that unemployment may fall if multinational entry is restricted, because returns to search go down, making workers willing to accept worse jobs. As a result, though profits may increase for most domestic firms when multinational entry is restricted, they may fall for some very high-productivity domestic firms because of the tighter labor market.

Our model is consistent with the existence of a multinational wage premium, which has been extensively documented using data for many countries. Within the context of our model, this premium is fully accounted for by the greater average productivity of multinational affiliates. In our extended model with labor heterogeneity, this is true once worker skill is controlled for. In the Appendix, we report results from estimating worker-level wage regressions with a dummy variable for employment at a multinational establishment. If we include worker fixed effects, controlling for the poaching index along with other establishment characteristics almost eliminates the multinational wage premium.

Our counterfactual results are consistent with the findings of both Alfaro Ureña et al. [2021] and Setzler and Tintelnot [2021] that there is a positive impact of (instrumented) multinational presence in a local labor market on the wages of those employed at domestic firms. Setzler and Tintelnot [2021] find that the impact is bigger for high-paid workers, consistent with our result that multinational presence increases competition more for workers at the top of the job ladder than on the middle rungs.

Finally, there is a very large literature on productivity spillovers from FDI, which has mixed findings on the direction and magnitude of these spillovers, see, e.g. Iršová and Havránek [2013]. Our finding that the labor market impact of multinational presence may differ across firms could potentially explain these mixed findings. If the heterogeneous labor market impact of multinational affiliates are not appropriately controlled for, it may be hard to consistently isolate productivity spillovers.

References

Alonso Alfaro Ureña, 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.
- Ragnhild Balsvik and Stefanie A. Haller. Picking lemons or picking cherries? Domestic and foreign acquisitions in Norwegian manufacturing. *Scandinavian Journal of Economics*, 112(2):361–387, 2010.
- Kenneth Burdett and Dale T. Mortensen. Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2):257–273, 1998.
- 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.
- Pierre Cahuc, Pauline Carry, Franck Malherbet, and Pedro S. Martins. Employment effects of restricting fixed-term contracts: Theory and evidence. 2022.
- Allan Collard-Wexler and Jan De Loecker. Production function estimation and capital measurement error. Technical report, National Bureau of Economic Research, 2016.
- A. Kerem Coşar, Nezhir 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. *American Economic Review*, 112(12):3803–3847, 2022.
- Pablo D. Fajgelbaum. Labour market frictions, firm growth, and international trade. *The Review of Economic Studies*, 87(3):1213–1260, 2020.
- Émilien Gouin-Bonenfant. Productivity dispersion, between-firm competition, and the labor share. *Econometrica*, 90(6):2755–2793, 2022.
- 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.
- 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.
- Zuzana Iršová and Tomáš Havránek. Determinants of horizontal spillovers from FDI: Evidence from a large meta-analysis. *World Development*, 42:1–15, 2013.
- Xiao Ma, Marc-Andreas Muendler, and Alejandro Nakab. Learning by exporting and wage profiles: New evidence from Brazil. MPRA Paper 109497, University Library of Munich, Germany, August 2020.
- 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.