Why Are Firms More Productive in the US than in the EU?*

Hugo Reichardt† Ricardo Reis‡

March 26, 2026

Abstract

When a European firm moves its headquarters to the United States, when a US firm expands by hiring across many states, or when a Japanese firm raises all its funding from US investors, they are revealing the sources of what makes the United States productive. This paper builds a model in which firms make location choices across multiple dimensions—where to hire production workers, where to place sales staff, where to raise funding, where to place the headquarters—and shows how these choices reveal the relative productivity of regions along each dimension. It constructs a new global database that tracks where firms locate their activities, and it develops an econometric strategy, including a new selection correction method, to estimate the facets of a region's relative productivity. Applying the framework to the US-EU comparison, it finds that the US productivity advantage is primarily accounted for by two factors: substantially higher efficiency of the US's funding markets and higher costs of operating across the EU's fragmented internal market.

*We are grateful to many colleagues for inspiring discussions and to the research assistance of Nicholas Tokay. This work was supported by the UK Research and Innovation grant number EP/Eo25039/1. First draft: February 2026.

†CREi and Barcelona School of Economics. hreichardt@crei.cat

‡Department of Economics, London School of Economics. r.a.reis@lse.ac.uk

1 Introduction

Between 2007 and 2024, labor productivity grew 18% more in the United States (US) than in the European Union (EU) (Draghi, 2024). Quantifying what lies behind this gap, whether it is skill shortages (Draghi, 2024), management practices (Bloom et al., 2012), adoption of technologies (Schivardi and Schmitz, 2020), existence of an integrated single market (Letta, 2024), or a weaker financing environment (Lagarde, 2024), is as difficult as it is important. A widely-reported symptom of this widening gap in competitiveness is that some successful European start-ups have relocated their headquarters (HQ) to the US. Which firms did so, when, and which of their workers and activities changed should contain information on what these firms perceive as making the US a more attractive place to do business than the EU.

This paper uses revealed preference by firms to identify the relative importance of different dimensions of the productivity differences across regions. Building on the Ricardian insight that trade flows reveal comparative advantage, we show that firms' location decisions for their activities reveal the relative efficiency of regions. Further, the extent to which firms concentrate activity within their home country/state reveals the extent and role of within-region market integration. Finally, when firms relocate their HQ, they reveal the attractiveness of a region in commanding operations in the other regions where it is located. Combining all of these firm choices, we can decompose the productivity differences across regions into different dimensions and components.

We take three steps to operationalize this approach. First, we compile a new global database that tracks where firms hire production workers and sales workers, where they obtain funding from, and where they are headquartered. Second, we build a model of multi-region firms making decisions on where to hire production workers, where to place sales and management staff, where to raise external finance, and where to locate their HQ. Third, we develop a strategy to quantitatively estimate the relative efficiency, integration, and attractiveness of regions and to aggregate them into comparable differences in productivity by using our database and guided by our model.

Applying our method to the EU-US comparison, we find the US advantage is almost entirely explained by two factors: higher efficiency of US funding markets, and higher fixed costs of operating across EU member states relative to US states. In contrast, differences between the US and the EU in labor and

sales productivity are small.

Outline. In Section 2, we describe a new global database on firms' location choices regarding their production workers, their sales personnel, their financing, and their HQ. Standard data on multinational activity rarely disaggregates the geographic allocation of a firm's inputs across regions and input types. We overcome this deficiency by combining data from two main sources: LinkedIn and Crunchbase. The first reveals where a firm's workers are located from their personal histories in the LinkedIn database. The second reveals the firm's sources of external funding.

Next, we use historical snapshots from the Internet Archive's Wayback Machine to produce data on firms relocating HQ across the world. The resulting database shows that firms have moved from Europe to the US at much higher rates than the reverse and that there is strong positive selection into moving from Europe to the United States.

Section 3 lays out a model of firms' choices on the sourcing of various inputs—production, sales, and finance—from different regions. In the intensive margin, holding fixed the set of locations that the firm operates in, the share of inputs sourced from each region reveals the firms' perception of each region's efficiency per unit cost for each type of inputs. In the extensive margin, firms decide from which region to source their inputs by trading off the decrease in marginal cost that comes from being able to shift inputs across one more region against the fixed cost of adding that region to its operations. The share of firms operating in multiple states within a region reveals how integrated production is within that region. Finally, firms choose whether to relocate their HQ for a cost, taking into account that a region is attractive to locate the HQ if it allows for a small penalty in the efficiency of operating in other regions, and a high expected growth rate of overall efficiency. All combined, the productivity of firms in one region decomposes into its relative advantages across these different activities and layers of choices.

Section 4 develops an efficient estimation strategy for these different dimensions of productivity. We show that the relative attractiveness of regions can be estimated using linear regressions over firms' worker shares across regions. We develop a new method to correct for selection on the extensive margin (where to locate) when estimating regions' efficiency in the intensive margin (how much to source from each region). Given these estimates from the intensive margin,

we calculate the potential profits that a firm obtains under each possible set of locations chosen on the extensive margin, and weight these against the fixed costs of adding another region to source from. We show that using a conditional logit regression over the share of firms that choose to operate in multiple regions identifies these costs. Given the intensive and extensive margin estimates, we calculate the profits for a firm from having its HQ in one region. We show that the matrix of probabilities of moving from one region to another and the size of those who move, are moments that identify the benefits and costs of moving HQ.

Section 5 applies our framework to the comparison of Europe and the United States. Starting with the intensive margin, we find that there is little difference in the efficiency, or costs, of employing workers in either region. There is, however, a large US relative advantage when it comes to financing. These sharp estimates are driven by the observation that firms that have their HQ in neither Europe nor the United States but employ workers in these two regions, on average have similar shares of workers in both. When they get funding from both regions, instead, they obtain a much larger share of funding in the United States.

Turning to the extensive margin, we find that the costs of opening operations in a region outside of its union are not too different for EU or US firms. However, US firms are much more likely to operate in multiple US states than European firms to operate in multiple EU countries. The model and data therefore identify a large difference in the integration of US vis-a-vis EU markets, and this holds across production, sales or funding.

Finally, in the data, more firms move to the US from the EU, than the other way around. Moreover, the firms that move are those that operate in more regions to start with and that grow faster once they move. The model therefore identifies the US as a more attractive region to place the HQ for a company. The difference in attractiveness is again dominated by finance. While firms that move from Europe to the United States relocate some workers there, they shift a significantly larger share of their funding toward US financial markets.

Section 6 combines these estimates into an overall measure of the productivity difference between Europe and the United States. We use the model to calculate a series of counterfactuals, where Europe closes the productivity gap to the US across the different dimensions. This allows us to quantify which avenues for reform might lead to a larger payoff. The counterfactuals suggest that

creating an integrated capital market would remove almost all of the gap. They also show that reforms that close the gap come with a shift in the distribution of firm sizes in Europe toward larger firms.

Section 7 concludes.

Literature review. On our question of what lies behind the US-EU productivity differences, there is a vast literature that focuses on one possible cause and tries to quantify its relevance: see, for instance, (Gutiérrez and Philippon, 2018; Bloom et al., 2012; Schoefer, 2025; Bergeaud, 2024). Our contribution is to consider multiple factors at once, and quantify their relative importance.

A rich literature uses firm outcomes to learn about productivity differences across countries (Hsieh and Klenow, 2009; Bartelsman et al., 2013). Our approach is quite different. Instead of measuring the allocation of inputs across firms within a region, we measure the firm's allocation of inputs across regions.

These questions fall into the general study of why some countries are richer and more productive than others. The typical approach is to measure the quantity and quality of inputs available within each region and their allocation across firms. This paper offers a somewhat different approach. It asks instead *within* firms, where they *choose* to get their inputs *across* regions.

Our approach of using firms' choices of location is inspired by the quantitative trade literature that uses trade flows to infer patterns of comparative advantage (Costinot et al., 2012; Levchenko and Zhang, 2016) in Ricardian models of trade (Dornbusch et al., 1977; Eaton and Kortum, 2002a). Instead of trade across regions and firms, we use data on the location of workers within a firm.

Our model brings together some of the key ingredients from the literature on trade (Eaton and Kortum, 2002a), management and sales (Lucas, 1978; Arkolakis, 2010), and finance (Markowitz, 1952). Our contribution is to explore some of the common features between these models and to integrate them to generate an empirical strategy that is common across these three choices.

There is a rich literature on multinationals and how they organize activity along global value chains (Antras and Chor, 2013; Muendler and Becker, 2010). Our focus is not on the chain of production activities and how trade flows between regions as a result, and our data includes many firms that would not qualify as multinationals but that do have inputs in more than one region.

2 A global database on firms' location choices

Our two main sources of data are LinkedIn and Crunchbase. We discuss each in turn, before describing other series that complement them.

2.1 LinkedIn: data on workers and firms

We use data on the near-universe of LinkedIn profiles of individuals and companies in January 2025 from Revelio Labs, a company which takes snapshots of all Linkedin profiles.

There are about 700 million individual unique LinkedIn worker profiles. They include comprehensive data on individuals' educational and employment histories, including locations, job titles, salary (as estimated by Revelio Labs), and identifiers that link to the company profiles of their employers. The data also contains 26 million unique company profiles that contain information on the location of the headquarters, the industry, the type of ownership (public or privately held), the founding year, and total employment. Linking firms and workers produces roughly 1.5 billion worker–job spells.

We compute the total number of people by job type and by location within a company in a given year. We aggregate job roles into two types: production workers (like "admin", "engineers", "operations", or "scientist"), or management and sales ("finance", "marketing", "sales"). Therefore, for firm f in region i at year t, we denote by $l_{f,i,t}$ the number of production workers at that date, and $m_{f,i,t}$ the management and sales workers. The panel starts in 2008, the first year for which coverage of the LinkedIn data is large enough. In constructing the company-level data, we weight employees by the inverse probability of them having a LinkedIn profile as estimated by Revelio. Likewise, we use sampling weights for firms so that they match the distributions across size and sector in the census.

To our knowledge, LinkedIn is the only global and large-scale source of data on how firms allocate their employment across the world. Data on the activity of multinationals typically comes from financial statements of listed firms. These statements only rarely report how employment is distributed globally. Further, these data are only available for large public firms and are subject to tax-driven distortions in reporting. The LinkedIn database covers a much larger set of firms, and employment can be broken down by both location and

job type, which is essential for our identification strategy.

At the same time, there is selection into the sample. For instance, we might expect that less qualified workers and smaller firms in less developed countries are less likely to have LinkedIn profiles. Some of this selection is dealt with by sample weights. More importantly, the identification of the different parameters comes from firms that operate across multiple regions and, for a subset of the parameters, from firms that move their headquarters. Missing out on small firms with a single location has little impact on our method, even if there are many of them. Moreover, the firms that drive the productivity differences between the EU and the US are in sectors that we expect to be well covered by LinkedIn, most prominently in information and communication technologies.

2.2 Crunchbase: data on funding

Crunchbase is a large database with information on companies across the world. It focuses on start-ups and the technology sector, and contains detailed information on their HQ location, total employment, year founded, founders, and the total amount of external funding received broken down by source and location. We obtained a snapshot of the universe of this database in July of 2024, when there were more than 4 million businesses in the sample.

The Crunchbase data contain the company's social media usernames, allowing us to match it accurately with the LinkedIn data. We observe the LinkedIn username for 74 percent of Crunchbase firms. Among those, some included LinkedIn usernames in an older format that cannot be easily matched to the corresponding new usernames in the LinkedIn database. Ultimately, we matched 2.1 out of 3.4 million Crunchbase firms to their respective LinkedIn data.

Table 1 reports summary statistics for the LinkedIn data, the Crunchbase data, and the matched sample. It shows that LinkedIn contains more firms and somewhat younger firms than Crunchbase. In general, firms in Crunchbase seem more positively selected than firms in LinkedIn. We also find that LinkedIn firms that appear in Crunchbase are more likely to be (US) tech firms, suggesting that Crunchbase is more biased towards (American) tech.

For the purpose of our exercise, the selection into our sample is not necessarily undesirable. First, these firms are predominantly the fast-growing ones that obtain funding from multiple sources, and so they are the ones that will be more informative for our purposes. Further, since our identification will come

from firms operating in multiple regions, it is an advantage that our database covers large firms well.

TABLE 1: Summary statistics for LinkedIn and Crunchbase data

	LinkedIn	Crunchbase	LI ∩ CB
Firms			
Total number (millions)	17.9	3.4	2.1
Median age (years)	10.0	17.0	17.0
% w/ workers in multiple regions	15.1		37.0
% tech	31.2		38.0
% w/ funding information		9.6	8.9
% US HQ'ed	29.8	44.9	43.4
% EU27+UK HQ'ed	31.2	27.1	26.8
Employees			
Total number (millions)	367.5		193.4
% production	57.1		57.6
% management/sales	33.5		31.6
% in firms w/ workers in multiple regions	76.6		89.5
% in US HQ'ed firms	33.2		40.6
% in EU27+UK HQ'ed firms	27.5		27.4

Notes: This table reports summary statistics for our two main data sources, LinkedIn and Crunchbase, and the intersection of the two, LI∩CB. Summary statistics on employment are computed using sample weights provided by Revelio Labs.

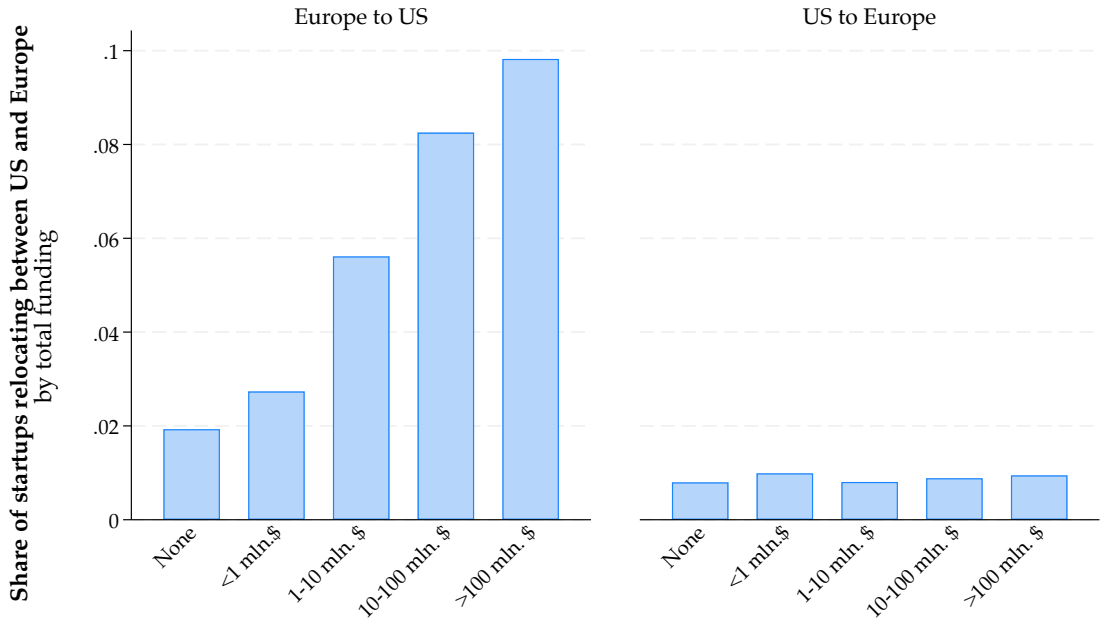
2.3 Wayback Machine: identifying movers

While the LinkedIn and Crunchbase data snapshots that we have include information on the firm's current HQ location, there is no direct observation on whether that was different in the past. In turn, commonly used datasets like Orbis update location when there is a change in the legal entity but do not keep the past locations on record.

The main source we use to document HQ migration rates is a snapshot of the Crunchbase data in December 2013, which included almost 200 thousand companies. Around eighty percent of the firms in this snapshot could be matched to a firm in the 2024 Crunchbase snapshot, using the Crunchbase usernames to match firms accurately. We complemented this database with scraped and parsed 2014 snapshots from individual Crunchbase profiles saved in the Wayback Machine, an internet archive. Finally, we obtained a database of 173,712

firms in Crunchbase for which we observe the HQ locations in both 2024 and 2014.

FIGURE 1: HQ migration rates between the US and Europe



Notes: This figure shows the share of firms that moved headquarter from Europe to the US and from the US to Europe between 2014 and 2024, based on the matched Crunchbase snapshots. The rates are broken down by the level of the funding the firm had received by 2024. Europe is classified as the 27 European Union countries and the UK.

Figure 1 shows that headquarters migration is strongly asymmetric: the rate of relocation from Europe to the US is several times higher than the reverse flow across all funding categories. Further, it shows that there is sharp positive selection into moving by firm quality. Among European firms with no external funding, fewer than 2% relocated to the US over the decade; among those that had raised more than $100 million, the rate exceeds 10%. Figure A.1 shows similar patterns when dividing firms by employment size. We note that technology startups are particularly overrepresented in the historical Crunchbase data. The migration rates in Figure 1 should therefore be interpreted as more representative of the tech startup population rather than the broader firm distribution.

3 A model of multinational firm behavior

We build a model of multinational firm behavior to show how firms' input choices reveal regions' relative productivity and to inform and discipline our econometric strategy.

The world is made up of many regions, indexed by $i \in I$, that are populated by many firms, indexed by f. We consider three levels of choices by a firm. First, it chooses where to locate its HQ $h_f \in I$. Then, it chooses in which regions to locate its production, management and financing activities, respectively $(I_f^y, I_f^m, I_f^v) \subset I \times I \times I$. Finally, it chooses how many production workers to hire $l_{f,i}$, managers to employ $m_{f,i}$ and funding to get $v_{f,i}$ in each of the regions. We present and solve these choices, one by one, by backward induction.

3.1 The intensive margin

This section starts by studying the intensive margin of how much to produce, sell, and borrow in each region in which there is an operation.

3.1.1 Production

Following a large literature in labor and trade (Acemoglu and Zilibotti, 2001; Eaton and Kortum, 2002a; Grossman and Rossi-Hansberg, 2008), a firm must combine a series of tasks (or intermediate goods) $y_f(\tau) \geq 0$ in the unit interval $\tau \in [0,1]$ in order to produce its final good $y_f \geq 0$. These tasks can be substituted between each other with elasticity $\sigma_y < 1$ according to the function:

$$y = \left[\int_0^1 y(\tau)^{\frac{\sigma_y - 1}{\sigma_y}} d\tau \right]^{\frac{\sigma_y}{\sigma_y - 1}} . \tag{1}$$

The combination of tasks that are potentially located in different regions of the world matches the decentralized production structures that we see in modern firms.

The firm is able to offshore these tasks around the world according to Ricardian comparative advantage. Each task can be produced by hiring labor in each region, $l_{f,i}(\tau)$ with productivity that is the product of three components. The first is a firm-specific productivity z_f^y, which is common across the regions in which it operates. The second is a region-specific productivity $a_i^y(h_f)$, com-

mon across firms, but that may depend on where the firm's headquarters is located. For instance, if the HQ and the region of production are very far apart, its productivity may be lower. The third is an i.i.d. firm-specific task-shock to productivity in that region $\varepsilon_i(\tau)$, which has mean 1. Therefore:

$$y_f(\tau) = z_f^y \sum_{i \in I_f^y} a_i^y(h_f)\varepsilon_i(\tau) l_{f,i}(\tau)\,, \tag{2}$$

where I_f^y is the set of regions where firm f operates a production location.

The firm wants to minimize the costs of producing from paying a wage w_i to each of the workers hired in a region. Following a long tradition in international trade, dating back to (Eaton and Kortum, 2002a), we assume that the shocks $\varepsilon_i(\tau)$ are drawn from a Frechét distribution with scale parameter 1 and shape parameter $\theta_y > 0$. This extreme-value assumption ensures that, in choosing the lowest-cost region to locate each of the production, the firm does so proportionately to the relative productivity of these regions. The appendix proves the following (standard) result:

Proposition 1. *The cost-minimizing choices of locating production workers are such that the labor share per region is proportional to its relative attractiveness:*

$$\frac{w_i^y l_{f,i}}{w_j^y l_{f,j}} = \left(\frac{a_i^y(h_f)/w_i^y}{a_j^y(h_f)/w_j^y}\right)^{\theta_y} \equiv \frac{\tilde{a}_i^y(h_f)}{\tilde{a}_j^y(h_f)} \tag{3}$$

where $\tilde{a}_i^y(h) \equiv (a_i^y(h)/w_i^y)^{\theta_y}$ is the wage-adjusted productivity of region i's production workers.

Taking these choices into account, firm f's marginal cost for a given set of production locations I_f^y is:

$$c_f^y \propto \frac{1}{z_f^y}\left(\sum_{i \in I_f^y} \tilde{a}_i^y(h_f)\right)^{-\frac{1}{\theta_y}}. \tag{4}$$

Marginal costs depend on firm productivity z_f^y, region productivity $\tilde{a}^y(h_f) \equiv \{\tilde{a}_i^y(h_f)\}_i$, and the set of production locations I_f^y. Each of these components lowers marginal cost, since higher firm or region productivity allows the firms to hire fewer workers to produce an extra unit of output, while increasing the number of regions yields a lowest-cost producer for some of the tasks.

3.1.2 Management and sales

Following the literature on market penetration and management (Lucas, 1978; Arkolakis, 2010), we model the firm as having to hire a salesforce in order to sell in different regions subject to diminishing returns in the management of those workers.

A firm must set up a trading post in a region to sell in it and staff it with management and sales workers $m_{f,i}$. In each region, there is a continuum of identical districts (in the real line), and the larger is its salesforce, the larger the mass of districts $s_{f,i}$ in the region it is able to sell. Reflecting a limited span of control, there are diminishing returns to managing more workers across a larger mass of districts. Finally, in each district, the firm sets a monopoly price to sell an amount $q_{f,i}$ of its service.

We make two standard assumptions on the functional form for demand and supply. First, the demand for the service in a district has a constant price elasticity $\sigma_m > 1$:

$$q_{f,i} = z_f^q \, p_{f,i}^{-\sigma_m} \,, \tag{5}$$

where z_y^q captures the popularity or quality of the service. This demand system implies that the monopolist optimally charges a constant markup over its marginal cost in each district.

Second, the mass of districts reached by the firm in a region is given by the diminishing returns function:

$$s_{f,i} = z_f^s a_i^m(h_f) m_{f,i}^{\theta_m} k_f^{\psi} \,. \tag{6}$$

There are again two dimensions to productivity, now in management and sales: one that is firm-specific z_f^s, and one that is region-specific $a_i^m(h_f)$ and which depends on the headquarters' location. The extent of the span of control problem is captured by the positive parameter $\theta_m \in (0,1)$. Finally, the capital of the firm k_f is used for market with diminishing returns given by $\psi > 0$.

Choosing how many sales workers to hire in each region to maximize profits after paying wages w_i^m leads to the following result, proven in the appendix:

Proposition 2. *The profit-maximizing choices of sales and management workers are*

such that the labor share per region is proportional to its relative attractiveness:

$$\frac{w_i^m m_{f,i}}{w_i^m m_{f,j}} = \left(\frac{a_i^m(h_f)/(w_i^m)^{\theta_m}}{a_j^m(h_f)/(w_j^m)^{\theta_m}} \right)^{\frac{1}{1-\theta_m}} \equiv \frac{\tilde{a}_i^m(h_f)}{\tilde{a}_j^m(h_f)}, \tag{7}$$

where $\tilde{a}_i^m(h) \equiv (a_i^m(h)/(w_i^m)^{\theta_m})^{1/(1-\theta_m)}$ *is the wage-adjusted productivity of region i's sales and management workers.*

The appendix also shows that the operating profits (profits net of the cost of capital and fixed costs) are given by:

$$\Pi_f^o \propto \left(\left(z_f^y \right)^{\sigma_m - 1} z_f^q z_f^s \right)^{\frac{1}{1-\theta_m}} \left(\sum_{i \in I_f^y} \tilde{a}_i^y(h_f) \right)^{\frac{\sigma_m - 1}{\theta_y(1-\theta_m)}} \left(\sum_{i \in I_f^m} \tilde{a}_i^m(h_f) \right) k_f^{\psi/(1-\theta_m)} \tag{8}$$

This equation shows that the firm productivity in production, sales, and, management, (z_f^y, z_f^q, z_f^s), and the productivity of the regions the firm operates production and management locations, (I_f^y, I_f^m), together determine the operating profits of the firm. For a fixed marginal cost of production, more productivity in sales and management raises the ability of the firm to sell more in each district and across more districts in each market. More locations raise profits by expanding sales. Firms' idiosyncratic productivity is complementary to the number of locations in raising profits.

3.1.3 Financing

The firm's capital k_f is provided by investors from around the world: $k_f = \sum_{i \in I_f^v} v_{f,i}$. We assume that the more financing the firm gets from region i the higher the cost $r_{f,i}$ is:

$$r_{f,i} = \frac{v_{f,i}^{\theta_v}}{z_f^v \, a_i^v(h_f)} \tag{9}$$

where z_f^v is the firm-specific access to external finance, $a_i^v(h_f)$ are region-specific financing conditions, which depend on HQ location, and θ_v captures the curvature of financing costs. To ensure an interior solution, we assume $1 + \theta_v > \phi/(1-\theta_m)$.

Following a long literature in financial frictions, the appendix provides a micro-foundation for the upward sloping supply of funding from each region.

The financiers in a region receive only a fraction of the excess return paid by the firm, while the rest is lost to intermediation costs. Moreover, the venture capitalists perceive that these returns are risky because the many dimensions of productivity have an unobserved random component, and will invest less the higher is the variance. Finally, different regions have different amounts of capital to lend out. All combined this generates the supply function above, where $a_i^v(h_f)$ is a composite of region-specific intermediation costs, expertise in assessing a lower risk of the firm, and depth of capital markets with more abundant capital. In the case of classic mean-variance financiers the exponent would be one, but generalizing to $\theta_v > 0$ captures different effective risk aversion by the funders of the firm.

Since the firm is risk neutral, and can perfectly substitute the source of financing across regions, it will borrow to equate the cost $r_{f,i}$ across regions. The amount of external financing from each region is then given by the following result.

Proposition 3. *The equilibrium shares of funding from each region are:*

$$\frac{v_{f,i}}{v_{f,j}} = \left(\frac{a_i^v(h_f)}{a_j^v(h_f)}\right)^{1/\theta_v} \equiv \frac{\tilde{a}_i^v(h_f)}{\tilde{a}_j^v(h_f)}, \tag{10}$$

where $\tilde{a}_i^v(h) = a_i^v(h)^{1/\theta_v}$ is the cost-adjusted productivity of region i's venture capital environment.

The optimal allocation of the firm's financing across regions implies that the interest rate as a function of total borrowing is

$$r_f = \frac{1}{z_f^v}\left(\sum_{i \in I_f^v} \tilde{a}_i^v(h_f)\right)^{\theta_v} k^{\theta_v}$$

Higher productivity in borrowing, $\tilde{a}^v(h_f) \equiv \left\{\tilde{a}_i^v(h_f)\right\}_i$, raises profits by lowering the marginal interest rate at which the firm borrows. Moreover, borrowing from more regions, I_f^v, allows the firm to tap into the capital markets of different regions and therefore borrow at a lower costs from each one. These comparative advantages are complementary with each other, and with the ones in sales and production.

3.1.4 Taking stock

Having derived the optimal choices of the firm in the intensive margin, we can now express the profits of the firm (before entry costs) as a function of the firm's idiosyncratic productivity and the productivity of the regions it operates in.

Proposition 4. *The profits of the firm (before entry costs) are:*

$$\Pi\left(I_f, \tilde{a}(h_f), z_f\right) \propto z_f \left(\sum_{i \in I_f^y} \tilde{a}_i^y(h_f)\right)^{\beta^y} \left(\sum_{i \in I_f^m} \tilde{a}_i^m(h_f)\right)^{\beta^m} \left(\sum_{i \in I_f^v} \tilde{a}_i^v(h_f)\right)^{\beta^v} \tag{11}$$

where the exponents are given by:

$$\begin{aligned}
\beta_y &= \frac{\sigma_m - 1}{\theta_y(1-\theta_m)} \cdot \frac{1+\theta_v}{1+\theta_v - \frac{\psi}{1-\theta_m}}, \\
\beta_m &= \frac{1+\theta_v}{1+\theta_v - \frac{\psi}{1-\theta_m}}, \\
\beta_v &= \frac{\theta_v \frac{\psi}{1-\theta_m}}{1+\theta_v - \frac{\psi}{1-\theta_m}},
\end{aligned}$$

and the overall idiosyncratic productivity of the firm is:

$$z_f = \left[\left((z_f^y)^{\sigma_m - 1} z_f^q z_f^s\right)^{\frac{1}{1-\theta_m}}\right]^{\beta_m} (z_f^v)^{\beta_v/\theta_v}.$$

The result in Proposition 4 clarifies the benefits from adding more regions in each dimension. From there, we describe how firms trade off this benefit against the costs.

3.2 The extensive margin

The firm chooses which markets to locate its operations in $I_f = (I_f^y, I_f^m, I_f^v)$.

We assume that opening a location in region i for production entails a fixed entry cost $e_i^y(h_f) \geq 0$. These costs capture the difficulties of adapting to the local culture, adjusting production processes to regional regulations, and navigating the local markets in which to hire the workers. If region i has a very open market, in the sense of this cost being lower, then firms are more likely to take advantage of its productivity advantages and move production there.

The same applies to sales and to financing, with associated fixed costs $e_i^m(h_f)$ and $e_i^v(h_f)$. Allowing these fixed costs to differ for production, management, and financing decisions allows the model to distinguish the extent of barriers to mobility of firms across these three dimensions. These need not be the same.

The costs depend on the location of the HQ. If region i is closer to the HQ region h_f, the costs may well be lower as a result of cultural, legal, or other similarities, in the spirit of gravity equations. More pointedly, if we consider a country as an agglomeration of regions, then the combination of these fixed costs within that country will capture how integrated its internal market is. The vector of fixed costs per HQ region $e(h_f) = \left\{e_i^y(h_f), e_i^m(h_f), e_i^v(h_f)\right\}_i$ therefore capture the fragmentation of that HQ location.

The firm weighs the benefit of more locations, raising profits, against the higher fragmentation costs. We further assume that each firm pays a firm-specific random mean-zero shock $\varepsilon^I(I_f)$ that is specific to the set of locations it has chosen for the three dimensions of its operations. These capture the heterogeneity in firm choices. Therefore, the firm's extensive margin choice is:

$$\max_{I_f^y, I_f^m, I_f^v \in \mathcal{P}(I)\setminus\emptyset} \Pi(I_f, \tilde{a}(h_f), z_f) - \sum_{i \in I_f^y} e_i^y(h_f) - \sum_{i \in I_f^m} e_i^m(h_f) - \sum_{i \in I_f^v} e_i^v(h_f) - \varepsilon_f(I_f)\,. \tag{12}$$

We make a simplifying assumption to this problem: the shocks $\varepsilon_f(I_f))$ are drawn from a Gumbel (Extreme Value Type 1) distribution with mean zero and shape parameter ζ^I across firms. This leads to a familiar result from discrete choice models:

Proposition 5. *The probability that firm f uses locations I_f, conditional on the location of its headquarters h_f and its productivity z_f is*

$$\Pr(I_f \mid \tilde{a}(h_f), e(h_f), z_f) \propto \exp\left(\frac{\Pi(I_f, \tilde{a}(h_f), z_f) - \sum_{d\in\{y,m,v\}} \sum_{i \in I_f^d} e_i^d(h_f)}{\zeta^I}\right), \tag{13}$$

For a fixed set of location costs, this optimal behavior generates a distribution of firms across locations. Inversely, given an observed empirical distribution, one can estimate the costs. Observing that a region has more firms from a certain headquarters operating in it, the model explains this in part because the productivity benefits of this region are especially high, or because the costs of having a location in it are especially low.

The benefits of operating in a location are complementary across the three dimensions. For instance, a firm which produces in many locations will have a lower marginal costs and a higher profit from selling in an additional location. The fixed costs, instead, are additively separable across regions and across activities. Yet, we might think that opening production in one location lowers the costs of having sales in that location, perhaps because of transportation costs. Likewise, it may be that opening in one European country for a US company, may lower the costs of opening in a second EU location. As our focus on the EU-US differences pushes us to consider the ICT sector, which mostly produces services distributed online, this may not be such a limitation. More importantly, our model is rich in the complementarities in benefits from opening in different locations, so all of these will simply be loaded into the profit function. It would be hard to separately identify these two dimensions of complementarity anyway, and the computational cost would be considerable.

3.3 Locating the HQ

Different HQ locations affect the relative productivities and the fragmentation costs $e(h_f)$. Since the profit function before entry costs $\Pi(I_f, a(h_f), z_f)$ is supermodular between $\tilde{a}(h_f)$ and z_f, then larger firms gain more from having their HQ located in a region where the three factors of production are more productive.

Collecting the productivity and entry costs into the vectors $\tilde{a}(h_f)$ and $e(h_f)$, and the scale and demand shocks into the vector z_f, then the firm's expected profit function over the random Gumbel distributed shocks are:

$$\pi_f(\tilde{a}(h_f), e(h_f), z_f) \tag{14}$$

This expected profit function net of entry costs, $\pi_f(a(h_f), e(h_f), z_f)$ is also supermodular: firms that have a higher total factor productivity—captured in the z_f vector—will benefit the most from locating their HQ in a location that gives them access to other regions for their activities at lower fixed costs $e(h_f)$

We assume that a firm has two stages in its life. In the start-up stage, the firm has a productivity draw z_f^0, and its HQ is in region ι. In the mature stage, the firm gets a new draw of productivity z_f and it may choose to relocate to a new region h_f. It is possible to expand the model to include more periods in the life of the firm and allow it to relocate more often than just once. The appendix does

so. However, in our data, we very rarely see a firm moving more than once, and we have poor data to identify whether an entrepreneur moved region in order to start his/her firm, so we cannot empirically explore these extensions.

We assume two final HQ-specific ingredients. The first is a relocation cost. If a firm changes its HQ to region $h_f \neq h_f^0$, then it must pay a cost $\mu(h_f) \geq 0$. Some regions are attractive in this sense of being easier to set up HQ there, together with whatever productivity and fragmentation advantages that they may have.

Second, we assume that the distribution of productivity draws in the mature stage of the firm's life cycle depends on the location of its HQ in the start-up stage: $F(z_f|z_f^0, h_f)$. Being born in some regions may be more conducive to learning and growing, so that for some h_f the distribution of draws of z_f is more favorable. This may capture, for instance, the existence of a hub of innovation in the region, or the ease with which its living conditions attract founders, or visa policies making it possible for HQ staff to move into the region.

Combined, these relocation costs and learning propensities capture a region's *attractiveness as a HQ location* both in terms of ease of settling in and future potential for growth.

The final ingredient, again to capture the heterogeneity in the data is that at each stage the firm draws a random shock per location, $\varepsilon_H(h_f)$, that is mean zero and iid distributed with a Gumbel distribution with scale parameter ζ^H. Combining all these ingredients, then the firm's problem is:

$$\max_{h_f^1} \left\{ \pi_f \left(a(h_f), e(h_f), z_f \right) - \mu(h_f) \mathbf{1}_{[h_f \neq \iota_f]} + \varepsilon_H(h_f^1) \right\} \tag{15}$$

This leads to the natural result, proven in the appendix:

Proposition 6. *The probability that firm f has its HQ in region h_f^1 during the mature stage of its life-cycle, conditional on its productivity z_f^1 is:*

$$\Pr(h_f \mid z_f) \propto \exp\left(\frac{\pi_f(a(h_f), e(h_f), z_f) - \mu(h_f)}{\zeta^H} \right), \tag{16}$$

where the productivity is a draw: $z_f \sim F(.|z_f^0, h_f)$.

This completes the model. An attractive HQ location not only leads to a better set of relative productivities $a(h_f)$ and fixed costs $e(h_f)$, but also to a

more favorable distribution $F(.|.,h_f)$ of future TFP, and a lower relocation cost $\mu(h_f)$.

3.4 Extensions

among the many assumptions we have made, some are inessential.

First, we ignored the use of capital in production, by assuming that it was only used for market penetration in sales. Because the scale and profits are multiplicative across the three stages of decisions we considered, this is not important. Including the use of capital in production would lead to the same conclusions, and similar expressions but with different exponents in the capital stock.

Second, we ignored tariffs and other trade costs in sales. Modeling these as iceberg costs specific to each region, they would enter the decision on how the allocation of sales staff, similarly to the productivity of the staff in that region. Like wages, they would be absorbed in our definition of $\tilde{a}_i^m(h_f)$.

Third, in financing, we did not consider earnings or collateral based financing constraints. In the appendix, we show that considering them would complicate the expression for profits, but lead to a similar expression for relative financing across regions, and a profit function that depends on the same dimensions of of productivity, with the same complementarities between them. A new element is that the net worth of the entrepreneur would now be part of the productivity of the firm, broadening the interpretation of z^f.

Fourth, we ignored different tax regimes across regions. A major source of these differences is on the taxes over the payouts to financiers, like venture capitalists or shareholders. These would be absorbed by the $\tilde{a}_i^v(h_f)$ and not change any result. Another source of these differences is the overall taxation of income and profits at the headquarters. These would be absorbed by the z_f instead.

4 Identification and estimation

This section discusses how we identify and estimate the parameters of the model using the database we built. Firms refer to LinkedIn and CrunchBase entries, and a time period is one year, from 2008 and 2025. There are 8 regions:

the United States, the 27 countries in the European Union together with the UK, and 6 other locations covering the rest of the world: Northern America excluding the US, the remaining European countries, Latin America and the Caribbean (including Mexico), Africa, Asia, and Oceania. To study the integration of the internal markets of the EU and the US, we further add the possibility for EU (and US) firms to operate in another country/state within the union.

4.1 Productivity: region heterogeneity and the HQ penalty

The model allows for productivity to vary across regions and depend on HQ location in general terms. Using the result in Proposition 1, data on the allocation of production workers by a firm would reveal the relative wage-adjusted labor productivity for all the regions that this firm operates in. These would be very noisy measures, but could be used to estimate two natural objects: the firm specific productivity, that is common across locations, and the region specific productivity, that is common across firms, but may vary with where the HQ is located. The proposition suggests a two-way fixed effects regression, over locations and firms to do so since:

$$\log w_i^y l_{f,i}^y = \underbrace{\log w_i^y l_{f,h_f}^y}_{\text{Firm FE}} + \underbrace{\log \tilde{a}_i^y(h_f) - \log \tilde{a}_{h_f}^y(h_f)}_{\text{Headquarters} \times \text{Region FE}} + \varepsilon_{f,i} \tag{17}$$

However, the headquarter-region specific term of interest, $\tilde{a}_i^y(h_f)$, is not identified. The within-firm variation identifies the productivity of region i relative to the HQ location. It does not separately identify the firm's idiosyncratic productivity and the productivity associated with the HQ location. This is because the firm fixed effects absorb all the HQ variation.

We achieve identification by assuming:

$$\log a_i^y(h_f) = \log a_i^y - \eta_{W,h_f}^y \cdot \mathbf{1}^W(i,h_f) - \eta_{O,h_f}^y \cdot \mathbf{1}^O(i,h_f), \tag{18}$$

The parameters a_i^y are the equivalent of the region fixed effects, measuring the relative productivity of a region that is common to any firm that hires production workers in that region. The effect of the HQ is now captured by only two parameters: $\eta_{W,h}^y$ and $\eta_{O,h}^y$. The indicator variable $\mathbf{1}^W(i,h_f)$ is equal to one if the location i is not the same as the HQ h_f but is still, within the same union (EU or US), while the indicator $\mathbf{1}^O(i,h_f)$ is one when $i \neq h_f$ and i is outside the union.

Therefore, relative to the benchmark productivity a_i^y, productivity will be lower by $\eta_{W,h}^y$ percent if the location is not the HQ but still within the union, and $\eta_{O,h}$ if it is outside the union. We expect these both to be less than one, and lower for the outside parameter.

The estimation equation that follows from Proposition 1 is:

$$\log w_i^y l_{f,i} = \bar{\alpha}_f + \log a_i^y - \eta_{W,h_f}^y \cdot \mathbf{1}^W(i, h_f) - \eta_{O,h_f}^y \cdot \mathbf{1}^O(i, h_f) + \epsilon_{f,i}^y, \tag{19}$$

where $\bar{\alpha}_f$ are firm fixed effects, the a_i^y are location fixed effects, but instead of HQ effects, we now have two parameters to estimate on a dummy variable of whether the location is the firm's HQ or not. These parameters capture the crux of the questions posed in the introduction. The variation across locations in the a_i^y will tell us if on average across all firms, European, American, or otherwise, putting one's production workers in a EU region comes with lower productivity than in the US. The variation across the two HQ choices in $\eta_{W,h}^y$ and $\eta_{W,h}^y$ instead will capture whether EU firms suffer a bigger decline in productivity in their production units abroad than US firms do.

To better understand the identification, take the example of an Asian firm. Comparing its choice to have more workers in the US than the EU identifies the relative productivity of these two regions. In turn comparing the share of workers that a US firm has in the US versus other locations, with the shares in and out of the EU for a EU firm identifies the respective penalty distances for the two regions.

The same arguments apply to management and funding as we make analogous assumptions about $a_i^{\cdot}(h_f)$.

4.2 Entry costs and estimation challenges

Similarly, we also assume that entry costs have a region-specific component, a penalty if the region is not the HQ but is in the union, and a (likely larger) penalty if the HQ is outside of the union:

$$e_i^{\cdot}(h) = e_i^y + \phi_{I,\mu(h)}^{\cdot} \cdot \mathbf{1}^W(i, h_f) + \phi_{\mu(h)}^O \cdot \mathbf{1}^O(i, h_f). \tag{20}$$

We can construct the profits before entry costs $\Pi(I_f^y, I_f^m, I_f^v)$ for each alternative set of locations. Proposition 4 shows how the location choices scale up prof-

its depending on the parameters β_y, β_m, β_v. This implies that we can compute profits (before entry costs) under alternative sets of locations by appropriately scaling up (or down) the observed profits with the productivity of the alternative location set (relative to the observed set). Doing so effectively controls for the firm's unobserved idiosyncratic productivity z_f.

The demand elasticity is set to $\sigma_m = 5$ (Head and Mayer, 2014). The Fréchet shape parameter for production is set to $\theta_y = 8$ (Eaton and Kortum, 2002b). The span-of-control parameter is set to $\theta_m = 0.6$ and $\psi = 0.20$. Finally, the curvature of financing costs is set to $\theta_v = 1$, the value implied by classic mean-variance portfolio theory (Markowitz, 1952). This calibration implies $\beta_y = 1.67$, $\beta_m = 1.33$, and $\beta_v = 0.33$.

Having constructed the estimates of the profit gain from adding more locations for each firm, we can then estimate the entry costs. Proposition 5 shows that the entry costs can be estimated from a conditional logit regression controlling for the number of locations—inside and outside the union—interacted by the union of the HQ region.

Intuitively, identification comes from the following variation in the data: if EU firms operate on average in fewer regions within Europe than US firms, then we would identify higher within-union fixed costs, and if they are less likely to have workers or funding outside the EU relative to US firms outside the US, this identifies a higher penalty for being a European-HQ firm.

Estimating the conditional logit regression faces a significant computational barrier as the number of alternative configurations is very large. For example, a firm headquartered in Asia, can decide whether or not to operate in 7 other regions in each of the 3 dimensions, yielding $2^{7\cdot 3} \approx 2.1 \cdot 10^6$ possibilities in $\mathcal{P}(I) \setminus \emptyset$ that one would have to evaluate the logit function over for each firm.

We address this computational problem by following McFadden (1978); Daly et al. (2014) in estimating the parameters over a subset $\mathcal{S} \subset \mathcal{P}(I) \setminus \emptyset$ of feasible location combinations. Consistency requires adding the correction term for the inclusion probability $\ln \pi(\mathcal{S} \mid j)$ to the indirect utility of each alternative j. We sample S alternatives with replacement (ensuring the chosen alternative is always included), which yields $\pi(S \mid j) = \frac{K(S)}{q_j}$ with $q_j = 1 - (1 - p_j)^{|S|}$.[1] We set p_j equal to the share of firms in the HQ region with a presence in each alternative region.

[1] Since $K(\cdot)$ does not vary across alternatives, $-\ln q_j$ is the only correction needed.

4.3 Handling selection

While we have so far discussed the estimation of the extensive and intensive margins separately, the two choices might be connected. Unlike in the model, in reality, a firm may observe its idiosyncratic productivity shock when choosing location regions on the extensive margin. Therefore, all else equal, and especially keeping entry costs fixed, its choice of locations will be tilted towards regions where its productivity is higher. A selection problem arises because we only observe the intensive margin decisions for a region if the firms chose to open in that region, biasing the estimates of the region's productivity upward for those regions that were ex-ante unlikely to be chosen (since conditional on being selected, the expectation of their idiosyncratic productivity is higher).

More formally, equation (19) shows that a firm's input intensity decisions are affected by the region's productivity a_i^y that we want to estimate, but also by the idiosyncratic shocks $\epsilon_{f,i}^y$. The firm knows its draw of $\epsilon_{f,i}^y$ but the econometrician does not. This selection means that $\mathbb{E}[\epsilon_{f,i}^l \mid i \in I_f^l, h_f]$ is not the same across the regions and so the estimates will be biased. Intuitively, if a particular region i is generally unlikely to be chosen as a production platform, then conditional on it being chosen the shock $\epsilon_{f,i}^l$ must on average be high.

As in Heckman (1979), the problem is solved by estimating $\mathbb{E}[\epsilon_{f,i}^l \mid i \in I_f^l, h_f]$ and controlling for it when estimating the intensive margin parameters. The challenge here is that this expectation in our model is a complicated function of all the parameters. We develop a new method to estimate it based on observables, which should have more applications elsewhere. By taking a first-order approximation to the control function, we derive the following:

Proposition 7. *If the shocks on the intensive margin are normally distributed, the selection function can be identified up to a first order*

$$\mathbb{E}[\epsilon \mid i \in I_f^l, h] \approx \frac{\sigma_\eta^2}{\zeta} \left(1 - \pi^l(i \mid h)\right) \mathbb{E}\left[\Delta\Pi_i(\epsilon, h) \mid i \in I_f^l, h\right]. \tag{21}$$

where: $\pi^l(i \mid h)$ is probability of including region i in data and $\Delta\Pi_i(\epsilon, h)$ is the firm's profits times the share of employment in region i (times a constant).

This equation shows that the expected value of the unobserved shock conditional on selection can be expressed as a function of (i) the probability that region i is chosen by firms with headquarters in h_f and (ii) the marginal con-

tribution of region i to firm profits. Intuitively, the more exceptional a region must be to justify entry, the larger is the expected upward bias in measured productivity conditional on entry. We use this relationship to construct a selection adjustment that is feasible to implement and substantially reduces bias in the intensive-margin estimates.

We apply this formula to correct the estimates of the intensive margin.

4.4 TFP, movers, and the HQ choice

Given estimates of the intensive margin productivities and fixed costs, evaluating the expected profit function $\pi_f(a(h_f), e(h_f), z_f)$ is conceptually straightforward but again computationally expensive because of the large set of alternatives. We do so by using a sampled estimator of the expectation within a restricted set: the appendix has more details. Importantly, while the productivities of production, sales and management enter multiplicatively, and so cannot be separately identified, the productivity of financing can be separated from the others.

Since we only consider two possible HQ location choices, the EU and the US, the conditional probit in Proposition 6 is relatively easy to compute. At the same time, with only two options, the variability in the data is limited. To compute the two moving costs, $\mu(h_f)$, we match two moments in the data: the conditional probability that firms move their HQ from the EU to the US, and vice versa. The share of movers tells us the costs of moving.

We need also to identify how the HQ affects the growth of TFP, that is the conditional distribution $F(z_f|z_f^0, h_f)$. We assume a simple linear first order Markov process for the scalar composite of the productivities of production, sales and management:

$$z_f - z_f^0 = \lambda_h - \chi_h z_f^0 + \epsilon_{f,h}, \tag{22}$$

where the parameter λ_h captures he relative benefit of having your HQ in location h in order to benefit from higher TFP growth. The parameter χ_h controls for mean reversion.

In the data, the four parameters $\lambda_{EU}, \lambda_{US}, \chi_{EU}, \chi US$ are identified as follows: in the cross section of EU firms for which we have information on their size in the two snapshots in time, we estimate a version of regression (22) for the

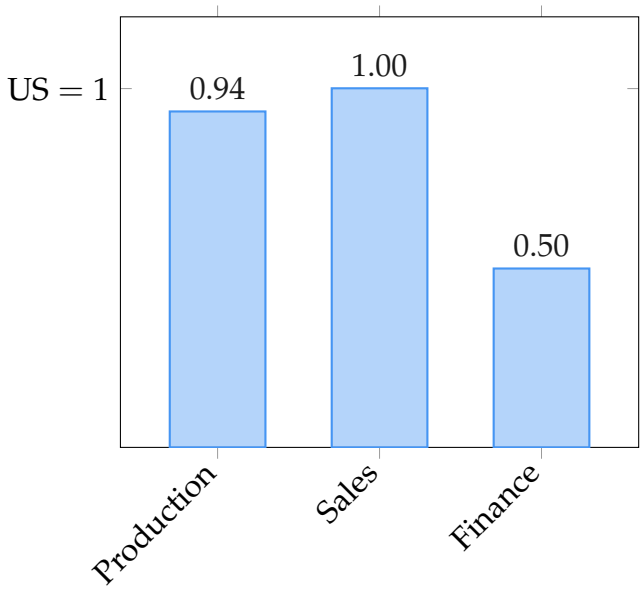
number of workers employed by the firm. Their average growth rate for the ones that keep their HQ in the EU identifies χ_{EU} while the growth in size for those who move identifies the relative $\lambda_{EU} - \lambda_{US}$. The same applied to the US firms identifies all four parameters. Finally, for the productivity of financing z^v we follow the same strategy, but using the size of external financing

5 The estimates

This section presents the paper's main estimates comparing Europe and the United States.

Our first key finding is that capital markets, rather than labor markets, are the main driver of the productivity difference between the US and the EU. Figure 2 reports the estimates of relative productivity per unit cost in the EU compared to the US for production, sales/management, and finance. The estimates show that the EU is close to parity in production, with a relative value of 0.94, and identical in sales, with a value of 1.00. In contrast, the productivity gap in finance is large: the productivity of the EU is only 0.50 of the US benchmark. These estimates are directly informed by the share of workers (or finance) that firms source from the EU relative to the US, conditional on sourcing from both in the extensive margin. Figure 2 therefore provides the first direct indication that capital-market advantages, rather than labor-market advantages, are the main driver of the overall US productivity edge.

FIGURE 2: EU's productivity per unit cost relative to the US by dimension: $\tilde{a}_i$.



Notes: This figure shows estimates of $\tilde{a}_u$ for the EU relative to the United States (which is normalized to 1). The EU contains the 27 EU states and the United Kingdom. Table B.1 presents estimates for all regions.

The second finding is that US firms can operate across states with much smaller efficiency losses than EU firms can across member countries. Table 2 reports the estimated productivity losses per unit cost when firms operate away from their headquarters. First, operating outside the union is associated with sizable losses in production and sales productivity for both US and EU firms (although somewhat larger losses for EU firms). Second, the productivity losses from operating foreign locations within the union are consistently larger for EU firms than for US firms. In other words, US firms can operate across states with much smaller efficiency losses than EU firms can across member countries. Overall, the estimates point to a more fragmented internal market in Europe, across production, sales, and finance.

TABLE 2: Productivity loss per unit costs when operating outside of HQ region

	Outside union			Foreign state within union		
Region	Production	Sales	Finance	Production	Sales	Finance
US	0.27	0.52	1.19	0.66	0.95	1.74
	(0.00)	(0.01)	(0.05)	(0.00)	(0.00)	(0.03)
EU	0.22	0.44	0.83	0.29	0.48	0.76
	(0.00)	(0.01)	(0.04)	(0.00)	(0.01)	(0.04)

Notes: This table reports estimates of productivity per unit costs when operating outside the HQ union (first three columns) and when operating outside the HQ country/state but within the union (last three columns). All values are relative to productivity in a region when a location is operated from a local HQ. Standard errors are reported in parentheses. Table B.2 reports estimates of the outside union productivity penalty for each region.

The third finding is that the US's internal market is more integrated than the European market as measured by substantially lower fixed costs of operating across the states of the union. The fixed costs of operating a production or sales location in another state within the US is close to zero, dramatically lower than the entry cost outside of the union. In contrast, the entry cost of operating foreign locations within the EU is much closer to the entry cost outside of the union. Interestingly, we find that entry costs for financing are non-negligible within the US too, even though still smaller than the EU.

TABLE 3: Fixed costs per foreign location (in mln $) by HQ region

	Outside union			Foreign state within union		
Region	Production	Sales	Finance	Production	Sales	Finance
US	0.33	0.35	0.76	0.00	0.03	0.53
	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)
EU	0.28	0.30	0.78	0.18	0.19	0.64
	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.01)

Notes: This table reports estimates of fixed costs per foreign location (in mln $) when the location is outside the HQ union and when it is outside the HQ country/state but within the union. Standard errors are reported in parentheses. Table B.3 reports estimates of the fixed costs per location outside the union for each region.

6 How to make the EU more productive?

This section will report counterfactual exercises to show the effect of closing the various dimensions of the productivity gap between the European Union and the United States.

7 Conclusion

Why is the US more productive than the EU? Is it because its production workers are more skilled (Ricardo (1817)'s comparative advantage)? Is it because US management is more efficient at organizing sales and operating in multiple locations (a hypothesis associated with Chandler, 1977)? Or is it instead because financing is more abundant or cheaper in the US (a venture capital story as in Kortum and Lerner, 2000)?

Is it because labor markets are less fluid in Europe and regulations are stricter and more intrusive, making it harder to operate in multiple EU regions in a productive way (Schoefer, 2025)? Or is it the fragmentation in EU goods and services markets, in contrast with the US's single larger homogeneous market (as in Marshall, 1890)? Finally, is it because of differences in taxes and regulations biased toward small firms that lower the effective returns from growing fast in the EU, or is it because the costs of moving to the US are lower because of visa policies or openness to immigrants?

The famous report by Draghi (2024) concluded that it was a bit of all of them. It argued that the EU is a worse business environment for firms than the US because of regulatory barriers to scaling up, a fragmented single market, a shallow capital market, and shortage of skills in technologically advanced domains.[2]

This paper takes a step forward by offering a transparent quantitative assessment of which reasons matter most. We developed a model of firm choice combining production, management/sales, and financing locations, built a global

[2] A few quotes from the introduction of the report are illustrative: "Regulatory barriers to scaling up are particularly onerous in the tech sector, especially for young companies." "The lack of a true Single Market also prevents enough companies in the wider economy from reaching sufficient size to accelerate adoption of advanced technologies." "Europe needs to make it easier for "inventors to become investors" and facilitate scaling up of successful ventures. " "A better financing environment for disruptive innovation, start-ups and scale-ups is needed as barriers to growth within the European markets are removed." "Skills shortages are acting as a barrier to innovation and technology adoption".

database on firms' inputs, and combined the two to identify different margins of relative efficiency.

Using the observed choices of location by firms, we concluded that better funding conditions in US markets were the main reason firms chose to move their HQ to the US. When they do so, the productivity penalty from operating workers far from HQ leads to some dislocation of workers to the US as well. Otherwise, all else equal, workers in the EU did not appear to be less productive than those in the US. Removing the barriers to a single market in the EU would also give a significant boost to its productivity, even if far less than closing the gap on financing.

References

ACEMOGLU, D. AND F. ZILIBOTTI (2001): "Productivity differences," *The Quarterly Journal of Economics*, 116, 563–606.

ANTRAS, P. AND D. CHOR (2013): "Organizing the Global Value Chain," *Econometrica*, 81, 2127–2204.

ARKOLAKIS, C. (2010): "Market Penetration Costs and the New Consumers Margin in International Trade," *Journal of Political Economy*, 118, 1151–1199.

BARTELSMAN, E., J. HALTIWANGER, AND S. SCARPETTA (2013): "Cross-country differences in productivity: The role of allocation and selection," *American economic review*, 103, 305–334.

BERGEAUD, A. (2024): "The past, present and future of European productivity," Tech. rep., POID wrking paper 103.

BLOOM, N., R. SADUN, AND J. V. REENEN (2012): "Americans do IT better: US multinationals and the productivity miracle," *American Economic Review*, 102, 167–201.

CHANDLER, A. D. (1977): *The Visible Hand: The Managerial Revolution in American Business*, Cambridge, MA: Harvard University Press.

COSTINOT, A., D. DONALDSON, AND I. KOMUNJER (2012): "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas," *Review of Economic Studies*, 79, 581–608.

DALY, A., S. HESS, AND T. DEKKER (2014): "Practical solutions for sampling alternatives in large-scale models," *Transportation research record*, 2429, 148–156.

DORNBUSCH, R., S. FISCHER, AND P. A. SAMUELSON (1977): "Comparative advantage, trade, and payments in a Ricardian model with a continuum of goods," *The American Economic Review*, 67, 823–839.

DRAGHI, M. (2024): "The Future of European Competitiveness," Tech. rep., European Commission.

EATON, J. AND S. KORTUM (2002a): "Technology, geography, and trade," *Econometrica*, 70, 1741–1779.

——— (2002b): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.

GROSSMAN, G. M. AND E. ROSSI-HANSBERG (2008): "Trading tasks: A simple theory of offshoring," *American economic review*, 98, 1978–1997.

GUTIÉRREZ, G. AND T. PHILIPPON (2018): *How EU markets became more competitive than US markets: A study of institutional drift*, w24700, National Bureau of Economic Research New York,.

HEAD, K. AND T. MAYER (2014): "Gravity Equations: Workhorse, Toolkit, and Cookbook," in *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Amsterdam: Elsevier, vol. 4, 131–195.

HECKMAN, J. J. (1979): "Sample selection bias as a specification error," *Econometrica: Journal of the econometric society*, 153–161.

HSIEH, C.-T. AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403–1448.

KORTUM, S. AND J. LERNER (2000): "Assessing the contribution of venture capital to innovation," *RAND journal of Economics*, 674–692.

LAGARDE, C. (2024): "Follow the money: channelling savings into investment and innovation in Europe," Tech. rep., Speech.

LETTA, E. (2024): "Much More Than a Market: Speed, Security, Solidarity. Empowering the Single Market to Deliver a Sustainable Future and Prosperity for All EU Citizens," Tech. rep., Council of the European Union.

LEVCHENKO, A. A. AND J. ZHANG (2016): "The evolution of comparative advantage: Measurement and welfare implications," *Journal of Monetary Economics*, 78, 96–111.

LUCAS, R. E. (1978): "On the size distribution of business firms," *The Bell Journal of Economics*, 508–523.

MARKOWITZ, H. (1952): "Portfolio Selection," *The Journal of Finance*, 7, 77–91.

MARSHALL, A. (1890): *Principles of Economics*, London: Macmillan.

MCFADDEN, D. L. (1978): "Modelling the Choice of Residential Location," in *Spatial Interaction Theory and Residential Location*, ed. by A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull, Amsterdam: North-Holland, 75–96.

MUENDLER, M.-A. AND S. O. BECKER (2010): "Margins of multinational labor substitution," *American Economic Review*, 100, 1999–2030.

RICARDO, D. (1817): *On the Principles of Political Economy and Taxation*, London: John Murray.

SCHIVARDI, F. AND T. SCHMITZ (2020): "The IT revolution and southern Europe's two lost decades," *Journal of the European Economic Association*, 18, 2441–2486.

SCHOEFER, B. (2025): "Eurosclerosis at 40: Labor market institutions, dynamism, and European competitiveness," Tech. rep., National Bureau of Economic Research.

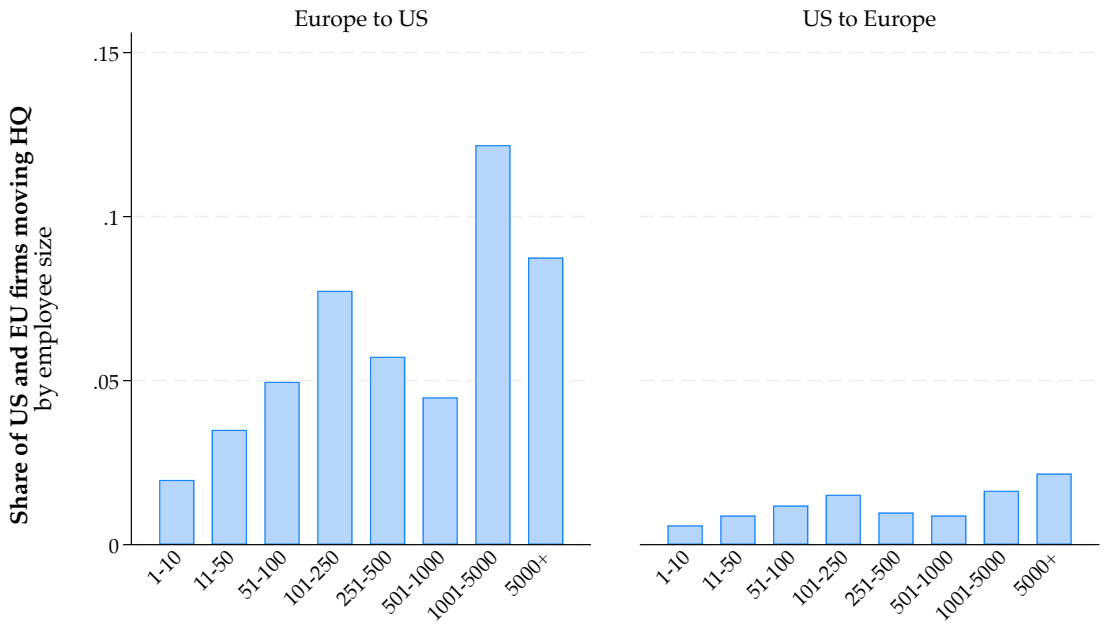
Appendix

A Figures 32

B Tables 33

A Figures

FIGURE A.1: HQ migration rates between the US and Europe



Notes: This figure shows the share of firms that moved headquarter from Europe to the US and from the US to Europe between 2014 and 2024, based on the matched Crunchbase snapshots. The rates are broken down by the number of employees in 2024. Europe is classified as the 27 European Union countries and the UK.

B Tables

TABLE B.1: Productivity per unit costs by region and dimension: $\tilde{a}_i^{\cdot}$

Region	Production	Management	Financing
United States	1.00	1.00	1.00
	(.)	(.)	(.)
EU27 and UK	0.94	1.00	0.50
	(0.01)	(0.01)	(0.01)
Africa	0.43	0.41	0.19
	(0.01)	(0.01)	(0.01)
Asia	0.89	0.73	0.51
	(0.01)	(0.01)	(0.01)
Latin America	0.46	0.42	0.14
	(0.01)	(0.01)	(0.01)
Oceania	0.33	0.33	0.24
	(0.00)	(0.00)	(0.01)
Rest of Europe	0.38	0.38	0.21
	(0.00)	(0.01)	(0.01)
Rest of North America	0.32	0.35	0.24
	(0.00)	(0.00)	(0.01)

Notes: This table reports estimates of $\tilde{a}_{iu}^{\cdot}$ for each region, relative to the United States (which is normalized to 1). Standard errors are in parentheses.

TABLE B.2: Productivity when being outside HQ union ($\exp(\eta_{O,h})$, 1 is inside HQ union)

Region	Production	Management	Financing
United States	0.27	0.52	1.19
	(0.00)	(0.01)	(0.05)
EU27 and UK	0.22	0.44	0.83
	(0.00)	(0.01)	(0.04)
Africa	0.05	0.09	0.17
	(0.00)	(0.00)	(0.03)
Asia	0.12	0.25	0.39
	(0.00)	(0.00)	(0.02)
Latin America	0.03	0.07	0.16
	(0.00)	(0.00)	(0.02)
Oceania	0.04	0.10	0.15
	(0.00)	(0.00)	(0.02)
Rest of Europe	0.11	0.21	0.12
	(0.00)	(0.00)	(0.02)
Rest of North America	0.06	0.15	0.28
	(0.00)	(0.00)	(0.03)

Notes: This table report estimates of the productivity when operating outside the HQ union (relative to the productivity in operating from that union when the HQ is in the respective union). Standard errors are reported in parentheses.

TABLE B.3: Fixed costs per foreign location (in mln $) by HQ region

Region	Production	Sales	Finance
United States	0.33	0.35	0.76
	(0.00)	(0.00)	(0.01)
EU27 and UK	0.28	0.30	0.78
	(0.00)	(0.00)	(0.01)
Africa	0.25	0.29	0.73
	(0.00)	(0.00)	(0.02)
Asia	0.27	0.31	0.68
	(0.00)	(0.00)	(0.01)
Latin America	0.28	0.31	0.80
	(0.00)	(0.00)	(0.02)
Oceania	0.23	0.26	0.73
	(0.00)	(0.00)	(0.02)
Rest of Europe	0.25	0.28	0.70
	(0.00)	(0.00)	(0.02)
Rest of North America	0.25	0.27	0.66
	(0.00)	(0.00)	(0.01)

Notes: This table reports estimates of fixed costs when operating outside the HQ union, by HQ broad region and function.

Notes: This table reports estimates of fixed costs per foreign location (in mln $) when the location is outside the HQ union. Standard errors are reported in parentheses.

References