Startup Types and Macroeconomic Performance in Europe*

Ralph De Haas[†], Vincent Sterk[‡], and Neeltje Van Horen[§]

June 6, 2022

Abstract

Can policymakers improve macroeconomic performance by encouraging the entry of high-performance startups? To answer this question, we construct a novel and comprehensive data set on 1.3 million startups in ten European countries. We apply cluster analysis to identify distinct startup types and trace their development over time. Three stylized facts transpire. First, we uncover five well-separated startup clusters that are consistently present across countries, industries, and cohorts. We label these *Basic*, *Large*, *Capital-intensive*, *Cash-intensive*, and *High-leverage*. Second, the initial differences between these startup types are highly persistent over time. Third, each startup type displays a characteristic life cycle in terms of productivity, employment generation, and exit rates. We feed these empirical results into an agnostic firm dynamics model to quantify how much structural policy could improve macroeconomic performance by shifting the composition of startups. We find that substantial gains in aggregate employment and productivity can be made through policies that benefit high-performance startups (such as capital-intensive ones) while discouraging the entry of underperforming firms (such as highly leveraged ones).

JEL classification: D22; D24; G32; L11; L25; L26; O47

Keywords: Startups, firm entry, productivity, corporate tax, entrepreneurship, cluster analysis

---

*We thank Filippo di Mauro for his enthusiastic support of this project, Victoria Marino for excellent research assistance, Marco Matani and Tommaso Bighelli for their expert help with running our code as part of constructing the 8$^{\text{th}}$ vintage of the CompNet dataset, and Tommaso Bighelli, Çağatay Bircan, Javier Miranda (discussant), and seminar participants at University College London, Institute for International Economic Studies (Stockholm), University of Regensburg, Danmarks Nationalbank, University of Bristol, University of Essex, Hong Kong University of Science and Technology, National University of Singapore, Hebrew University of Jerusalem, Toulouse School of Economics, and the CompNet ProdTalk for useful comments. The views expressed in this paper are ours and not necessarily those of the Bank of England or the EBRD.

†European Bank for Reconstruction and Development, CEPR, and KU Leuven. Email: dehaasr@ebrd.com.

‡University College London and CEPR. Email: v.sterk@ucl.ac.uk. Corresponding author.

§Bank of England, CEPR and University of Amsterdam. Email: neeltje.vanhoren@bankofengland.co.uk.

1 Introduction

In many advanced economies, politicians are increasingly concerned about lackluster macroeconomic performance as reflected in anaemic productivity growth and, in some countries, low employment levels (OECD, 2015; Syverson, 2017; Akcigit and Ates, 2021). Naturally, policy-makers are looking for novel levers to structurally improve these macroeconomic outcomes. A long-standing literature has explored several directions that policy can take, including tax adjustments and innovation subsidies (Bloom et al., 2002; Akcigit et al., 2016) and structural reforms to reduce distortions in labor markets (Hopenhayn and Rogerson, 1993), financial markets (Buera and Shin, 2017), and product markets (Edmond et al., 2021).

This paper investigates an entirely different policy lever, one that has remained largely unexplored: influencing the types of new firms that are being founded. The idea of improving the composition of new firm cohorts—as opposed to "fixing" established generations—appears attractive for two reasons. First, startups have been documented as key drivers of job creation and productivity growth (Foster et al., 2001; Haltiwanger et al., 2013). Moreover, recent evidence (referenced below) suggests that ex ante heterogeneity among newly established firms helps to predict their performance later in life. It follows that structural policies that successfully shift the mix of startup types that enter the economy, may generate important macroeconomic impacts.

Second, forward-looking policies to shift the composition of startup cohorts also appear attractive because the rates of firm entry and exit are high, typically around 10 percent annually. This means the majority of firms that will be in operation twenty years from now are yet to founded, while many current firms will no longer exist by then.[1] Some governments have already begun to focus their policy efforts on specific startup types. For example, in May 2020, the UK government launched a Future Fund to support the "most innovative businesses" and "top-performing startups with huge economic potential" that

[1]For instance, the US Longitudinal Business Database shows that, in 2019, 71 percent of all firms was 20 years old or younger. The startup rate in that year was 8.2 percent. In the appendix, we show that European startup rates are comparable.

will “create high-skilled jobs”.[2] Yet, in practice, efforts to create jobs and boost economic dynamism by stimulating entrepreneurship often fail (Lerner, 2009). These observations raise important questions: Can targeted startup policies structurally improve macroeconomic performance by altering the mix of startup types? And, if sizeable gains are possible, which types of startups should be encouraged and how to identify these types? At present, there is no clear answer to any of these questions.

We tackle these questions using large-scale administrative data sets for ten European countries (Croatia, Denmark, Finland, France, Italy, Lithuania, the Netherlands, Slovenia, Spain, and Sweden). The data contain a rich set of variables, deriving from the balance sheets and income statements of individual startups. We collected these data in close collaboration with the Competitiveness Research Network (CompNet), which uses a distributed micro-data approach to generate regularly updated, micro-based, and internationally harmonized data on European firms.

Our data contain information on more than 1.3 million European startups and allow us to draw a direct link between micro and macro outcomes along various key dimensions. The data provide unique cross-country panel observations with representative coverage of the full startup population. This compares well with data sets slanted towards larger firms (such as Compustat); surveys following one specific startup cohort (such as the Kaufman Firm Survey); administrative data covering only a very limited number of variables (such as the US Longitudinal Business Database); and databases that poorly capture firm entry and exit (such as Orbis, cf. Bajgar et al. (2020)).

Our analysis of these novel data is guided by a theoretical firm dynamics framework in the tradition of Hopenhayn (1992). We show that this model can be used in a tractable way to study how policies affect the composition of startup types and thus macroeconomic outcomes. In particular, it turns out we can conduct these policy counterfactuals while remaining largely agnostic about the demand and production structure of firms. Only three

[2]https://www.gov.uk/government/news/uk-tech-firms-and-investors-brought-together-for-landmark-treasury-conference.

sets of empirical statistics are required. The first set consists of multidimensional life-cycle profiles of the various firm types. The second set of statistics contains entry elasticities with respect to (the net present value of) profits. The third is the immediate impact of a policy instrument on firm profits. Importantly, all of these statistics can be readily estimated in our data set.

A key question we face is how to classify firms into types, given that these types are not directly observed. We address this issue by using K-means clustering, an unsupervised machine learning algorithm that has recently gained popularity in the applied economics literature as a way of dealing with latent heterogeneity. Underlying our application of this method is the idea that a startup's type is revealed by a number of key choices it makes when commencing operations. Exploiting the richness of our data, we classify startups based on five important choice variables in the initial year of operation: employment; the capital-to-labor ratio; total assets; the leverage ratio; and the cash-to-assets ratio. The practical advantage of this approach is that these variables are easily observed from tax information at the very beginning of a firm's life cycle. They can therefore readily be used to differentiate startups and facilitate targeted policies.

The clustering algorithm endogenously classifies startups into five types. It turns out that this clustering captures the majority of the empirical variation. Moreover, the outcome of the clustering analysis is remarkably robust across countries. In each country in our sample, we obtain the following five startup clusters: large; capital-intensive; high-leverage; cash-intensive; and basic. Using a "meta-clustering" analysis, we verify that each of these types has similar characteristics across countries. Moreover, using Monte Carlo simulations, we check that this similarity across countries is not due to mechanical factors relating, for example, to the shape of the distribution of the clustering variables. We also find the distribution of startups across the five types is quite stable across countries and economic sectors. Finally, we track these five startup types for more than a decade and document their life-cycle profiles in terms of the main choice variables. We find that the initial cross-type differences are highly persistent. All of these findings confirm that the clustering algorithm robustly captures fundamentally different firm types.

Based on the clustering outcomes, we then document heterogeneity in performance. We find large and persistent differences in employment and productivity across the various types, implying a change in the composition of startups can potentially have large macroeconomic effects. A number of salient patterns emerge clearly. In particular, the performance of the high-leverage startup type tends to be consistently poor. Even when we compare startups within the same country and economic sector, the cluster of highly leveraged firms displays substantially lower labor productivity and total factor productivity (TFP) and these firms are significantly more likely to exit in any given year. By contrast, startups that either stand out because of their capital-intensive production technology or because of their relatively high cash-to-assets ratio, typically boast higher productivity levels and lower exit probabilities.

Having documented these structural and persistent differences between startup types, we use our model to conduct a number of counterfactual exercises to quantify the "policy space" for improving macroeconomic performance via the startup composition channel. Concretely, we consider a budget-neutral differentiation of corporate income tax rates by startup type.[3] The differentiation of tax rates alters the incentives for different firm types to enter the economy and thus affects the startup mix itself. We proceed to compute the effects of such budget-neutral tax reforms on the macro economy, and find policy can significantly increase aggregate productivity and employment by tilting the mix of startup types. The productivity effects are smaller than the employment ones because there exists less heterogeneity in productivity than in employment levels across startup types.

We also show what kinds of startups should be encouraged and which should be discouraged in order to achieve maximum policy impact (given a certain intensity of the policy intervention). Furthermore, we analyze the trade-offs that policy-makers face when there are multiple macro objectives, and we use the data to shed light on the broader welfare

[3]While this particular policy instrument suits us because of its simplicity, other policies have the same effect on the startup composition as long as they influence the profits of the various types in the same way.

consequences of policies that differentiate between startup types. Finally, we use our results to shed light on the effects that existing policies may have on the composition of startups, and therefore on macroeconomic performance.

Related literature

We build on an emerging literature that documents the importance of ex ante heterogeneity for firm performance over their life cycle. This heterogeneity manifests itself in vastly different growth expectations among new entrepreneurs (Hurst and Pugsley, 2011), predictability of firm performance based on observable characteristics of entrepreneurs and businesses at the moment of startup (Schoar, 2010; Guzman and Stern, 2015; Belenzon et al., 2017; Guzman and Li, 2019) as well as strong cohort effects over the business cycle (Sedláček and Sterk, 2017). Quantitatively, ex ante heterogeneity accounts for the majority of heterogeneity in firm-level employment, as found by Sterk et al. (2021), who estimate an employment process using micro data.

We contribute to this literature in four important ways. First and foremost, we are the first to analyze how policy can be designed to exploit ex ante heterogeneity to improve macroeconomic outcomes. Second, we treat ex ante heterogeneity explicitly as a multidimensional object, whereas most existing studies focus on just one dimension of heterogeneity, such as firm size.[4] We jointly characterize firms by five key choice variables at startup and consider several performance measures (labor productivity, TFP, profitability, and exit probability). Third, the quality of our data enables micro-to-macro mapping along these dimensions and a comparison of this mapping across countries. This allows us to paint a rich and novel characterization of the European startup landscape. Fourth, the clustering algorithm classifies firms into types based on observables in the first year of operation.

[4] Recently, Bernard et al. (2022), using Belgian data on firm-to-firm trade, have underlined the importance of multiple attributes to explain firm-level success and, therefore, dispersion in the firm-size distribution. Other work assesses the role of specific ex ante differences across startup founders, such as their prior business experience (Lafontaine and Shaw, 2016), cognitive and non-cognitive personality traits (Levine and Rubinstein, 2017) and age (Azoulay et al., 2020).

This makes our approach straightforward to implement by policy-makers. In contrast, other studies have not made firm types observable (Albert and Caggese, 2021; Sterk et al., 2021).

Another key contribution of this paper is that we obtain empirical estimates for the elasticity of firm entry and can show that these elasticities are heterogeneous across firm types and countries. Such estimates are relatively rare in the literature, even though they are a standard input in firm dynamics models, which often assume either fixed entry (that is, a zero elasticity) or a free entry condition (that is, an infinite entry elasticity).[5] Our estimates may thus be used to impose more empirical discipline on models in the tradition of Hopenhayn (1992) and models with firm entry more generally.

Moreover, we add to the literature on the micro origins of aggregate productivity (Hsieh and Klenow, 2009; Restuccia and Rogerson, 2013; Hsieh and Klenow, 2014; Midrigan and Xu, 2014; Moll, 2014; Gopinath et al., 2017). Here, our focus on startup types and their persistent performance differences provides a new perspective on productivity dispersion between firms. It also opens up a new avenue to better understanding the drivers of differences in aggregate productivity across countries and industries, as well as changes over time. Instead of focusing on the "startup deficit"—a decline in firm entry observed in several countries over the past decades (Decker et al., 2017; Alon et al., 2018)—we highlight the importance of the *composition* of new startup cohorts for aggregate productivity and employment.

Finally, our results also complement studies that investigate the effectiveness of interventions to improve firms' performance *after* they have been established, such as consultancy services and management training (Bloom et al., 2013; Iacovone et al., 2022) and business accelerators (González-Uribe and Leatherbee, 2018; González-Uribe and Reyes, 2021). In contrast, we explore how aggregate productivity can benefit from policies that influence the composition of new startup cohorts at inception.

[5] Other estimates can be found in Sedláček and Sterk (2017) and Gutiérrez et al. (2019). They rely, however, on full-blown structural DSGE models to estimate the entry elasticity, whereas our method is agnostic. Moreover, we estimate elasticities for different startup types.

2 Theoretical framework

This section presents a theoretical framework on which we base the policy counterfactuals, presented in Section 5.2. It also offers guidance on which empirical statistics are required to answer the questions at hand.

2.1 The model

Our theoretical framework is a generalization of the firm dynamics model of Hopenhayn (1992). It allows for endogenous entry and exit of firms as well as firm-level idiosyncratic shocks and ex ante heterogeneity. We now describe this framework in more detail:

Incumbent firms. There is an endogenous mass of heterogeneous firms. Let x be a vector of firm-level choices, such as its price and capital and labor inputs, and let z be a vector of variables that are exogenous to the firm, such as aggregate prices and productivity shocks. An individual firm, indexed by i, operates a production function given by $f_i(x, z)$. We do not take a stance here on the specific functional form of the production function, which may include fixed costs. Moreover, we allow the production function to be firm-specific—it may, for instance, depend on the firm's type, and it may even vary with the age of the firm.

The firm also faces a demand curve $q_i(x, z)$ that depends on its price and potentially on other choice variables included in x (for example, the firm might expand its demand via advertising, a form of intangible capital accumulation) as well as shocks (for example, z could include demand shocks). Again, we do not take a stance on the details of the demand structure and we allow it to be firm specific. For example, it may depend on the firm's type. Finally, we allow for an arbitrary set of constraints, which may also be firm-specific.

Time is discrete and indexed by t. There is a finite number of ex ante firm types, indexed by j. A "type" refers to a set of commonalities in demand and production structures, as well as constraints, among a group of firms. Yet we also allow for heterogeneity within types reflecting firm-level shocks or initial conditions.

The firm sets its choice variables in order to maximize, at any point in time and given its constraints, the expected value of profits. The firm value is thus given by:

$$V_i = \max E \sum_{t=0}^{\infty} \Lambda^t \pi_i(x_{i,t}, z_{i,t})$$

where E is the expectations operator, Λ is the firm's discount factor, which, for simplicity, we assume to be common across firms, and $\pi(\cdot)$ is the profit function implied by $f(\cdot)$ and $q(\cdot)$. A firm exits (forever) if and only if the firm no longer has positive value, that is if and only if $V_i < 0$.[6]

Entrants. For any type j, there are a certain number of potential entrants in any given period. Potential entrants do not yet know their precise startup fundamentals, such as their production function, demand function, and initial conditions. However, they do know their type and the distribution of fundamentals within their type.

Each potential entrant then faces an entry decision. If they decide to enter, they must pay an entry cost θ_i which depends on the firm type as well as the individual potential entrant. After paying the entry cost, the entrant learns its startup fundamentals. Subsequently, the firm may be hit by shocks, for instance to its productivity or demand.

Let $\mathbb{E}_j V$ be the *ex ante* expected value of a firm of type j, that is, before paying the entry cost and before learning the startup fundamentals. Optimal decision-making implies that a firm of type j is started whenever $\mathbb{E}_j V \geq \theta_i$. That is, there exists a cutoff value on θ_i such that only potential entrants with an entry cost below the cutoff actually start a firm.[7] We can now express the actual number of entrants of type j as $n_j = g(\mathbb{E}_j V, \Gamma_j)$ where g is an increasing function in $\mathbb{E}_j V$, and where Γ_j denotes the number of potential entrants of

[6] This condition could be relaxed, since our subsequent analysis will not rely on it. For instance, one could introduce exogenous exit without changing the results.

[7] *Within* a type, all entrants have the same expectations prior to drawing the demand and production function. This implies there are no entry selection effects within types (only across types). While this modeling assumption is standard in the literature following Hopenhayn (1992), we verify it empirically in Section 5.2.1.

type j and their distribution over entry costs. Thus, the number of entrants is a function of the expected firm value and shocks to the number of potential entrants. Taking a first-order approximation (in logs) of this function gives the following expression for the percentage change in the number of entrants in type j:

$$d \ln n_j = \varepsilon_j \cdot d \ln \mathbb{E}_j V + \gamma_j \tag{1}$$

Here, $\varepsilon_j > 0$, is the elasticity of the number of entrants of type j with respect to the ex ante expected firm value, which is given by the mass of entrants at the entry cutoff, relative to the mass of entrants below the cutoff. Moreover, γ_j denotes shocks to the number and distribution of entrants of type j (i.e. shocks to Γ_j). Thus, the number of entrants of a certain type may increase either because of an increase in the expected firm value, or because of a shock to the number of entrants and/or their distribution over entry costs.[8]

Equilibrium and aggregation. All countries in our empirical application are members of the European Union, and therefore their markets for goods, labor, and firm ownership are integrated. Accordingly, we assume that prices and wages are taken as given for any individual country (below we will relax this assumption). We can compute any aggregate variable Y as:

$$Y = \sum_a \sum_j \omega_{a,j} y_{a,j},$$

where $y_{a,j}$ is the aggregate among firms of age a and type j, and $\omega_{a,j}$ is the appropriate weighting factor (for example, the firm or employment share). Both variables can be observed directly in the data, given a classification of firm types. For instance, if Y denotes aggregate labor productivity, then $y_{a,j}$ is aggregate labor productivity among firms of age a and type

[8] An alternative setup to arrive at the same result is to assume an unlimited number of entrants and to assume that the entry cost is homogeneous across entrants, but an increasing function of the number of entrants (within a type). In this case, the elasticity ε_j is determined by the shape of the entry cost function, see, for example, Gutiérrez et al. (2019).

j, and $\omega_{a,j}$ is the employment share of these firms.[9]

2.2 Effects of a policy change on startup composition

Based on the above framework, we now present a simple formula to compute the effects of a change in taxes or subsidies on the distribution of startups. Let $d\mathbb{E}_j T$ denote the *direct* effect of a policy change on the present value of the profits of an individual firm, averaged across firms within type j. For instance, $-d\mathbb{E}_j T$ could be the increase in the present-value tax bill following an increase in tax rates. While the number and distribution of *potential* startups is not directly affected by the policy, startup activity will be affected through the change in firm values. Applying Equation (2), the change in the number of startups of type j due to the policy change (again, to a first-order approximation) is therefore given by:

$$d \ln n_j = \varepsilon_j \cdot [\ln(\mathbb{E}_j V + d\mathbb{E}_j T) - \ln \mathbb{E}_j V] = \varepsilon_j \cdot \frac{d\mathbb{E}_j T}{\mathbb{E}_j V} \tag{2}$$

The formula states that the change in the number of startups of type j is the product of the entry elasticity and the percentage change in the firm value as a result of the policy change. Changes in startup composition can thus be evaluated based on this formula. Intuitively, the effects of a policy on the startup composition depends on how much the startup values of different types are affected by a policy change, and on how sensitive firm entry by the various types is to changes in firm values.

In case the policy is a corporate income tax, so that tax payments are proportional to profits, $\frac{d\mathbb{E}_j T}{\mathbb{E}_j V}$ equals the change in the tax rate. More generally, $\frac{d\mathbb{E}_j T}{\mathbb{E}_j V}$ is straightforward to compute in terms of firms' profits, provided one has data on the exposure of different firm types to the policy. We will use our data to estimate the entry elasticities (Section 5.1).

Having evaluated the formula for a given policy change, we can compute the counterfactual firm and employment shares of the different types in the incoming cohort, and thus the

[9] In turn, $y_{a,j}$ is measured in the micro data as the employment-weighted average labor productivity of individual firms.

change in the weights $\omega_{a,j}$. This in turn allows us to calculate the counterfactual level of Y, which isolates the macro implications of the compositional effects of the policy change.

In conclusion, in order to compute the effects of a policy change on startup composition, and its aggregate implications, we require three sets of statistics for each startup type: (i) life-cycle profiles for the variables of interest (for example, productivity) and the weighting factor (such as firm shares or employment shares); (ii) the elasticities of entry with respect to the firm value; (iii) the direct effect of the policy change on firm values. All can be measured in the data. However, in order to compute these statistics, we also need a method to classify firms into types. We discuss our approach to this in Section 3.3.

2.3 Effects on post-entry behavior

To evaluate the formula in Equation (2), we do not need to evaluate the effects of the policy change on the post-entry behavior of firms. This follows immediately from the Envelope Theorem, which implies that up to a first-order approximation any such effects on the firm's value equal zero. This property makes the formula particularly convenient to apply in practice when evaluating the compositional implications of any policy change.

That said, tax changes do, of course, influence the post-entry behavior of firms, and therefore have macroeconomic effects through channels other than that of the startup composition. These channels depend on the specifics of the policy. For example, a large literature has studied how various kinds of taxes and subsidies affect firm-level and, ultimately, macroeconomic outcomes.[10] Such studies provide useful guidance to policy-makers and researchers as to which macroeconomic effects tax-specific reforms may have. Our composition formula is complementary to this literature. Indeed, to evaluate the overall macroeconomic effects of a proposed tax reform, one can simply add the outcomes of our composition calculation to those from existing studies on post-entry behavior.[11]

[10] See, for example, Akcigit et al. (2016); Zwick and Mahon (2017); Akcigit et al. (2018); Liu and Mao (2019); Benzarti and Harju (2021).

[11] In our concrete policy application in Section 5, we consider a corporate income tax, which does not

2.4 Welfare effects

At the margin, a change in startup composition has no effects on the welfare of entrepreneurs. To appreciate this point, note that, to the marginal entrant at the cutoff, the expected value of starting a firm exactly equals the entry cost. In other words, the marginal entrant is indifferent between entering and not entering.

Yet, there may well be welfare effects that go beyond this. For instance, there are likely externalities on other stakeholders, such as workers, consumers, or the taxpayer. Quantifying these effects requires the full specification of a model, including assumptions on market structures, frictions, and preferences. Rather than taking this route, we instead conduct a positive analysis that focuses on macroeconomic quantities rather than welfare. Yet, we will consider a wide range of outcomes, including profits and wages, to obtain a sense of the broader welfare effects.

2.5 Equilibrium effects

So far, we have abstracted from equilibrium effects. Arguably, this is appropriate when analyzing policy changes in individual countries that are part of a large economic union such as like the EU. Yet one may also consider a policy jointly implemented in all countries simultaneously. In that case, there may be equilibrium feedback effects.

We now extend our framework to incorporate such effects. The margin of equilibrium adjustment we consider is the labor market, as wages are a main component of firm costs, and we will later study aggregate employment effects. Nonetheless, the analysis could be extended to account for equilibrium effects in other markets along the same lines. Let w denote the real wage per worker, which adjusts to clear the labor market.[12] Equation (2)

necessarily have any direct effect on post-entry behavior. Intuitively, when the government takes a certain share of firm profits, via a corporate income tax, the profit maximization problem itself remains unchanged. See, for example, Sedláček and Sterk (2019).

[12]Equivalently, we could fix goods prices and let nominal wages adjust, or fix nominal wages and let goods prices adjust.

then extends to:

$$d \ln n_j = \varepsilon_j \frac{d\mathbb{E}_j T}{\mathbb{E}_j V} + \varepsilon_j \gamma_j d \ln w. \tag{3}$$

The second term on the right-hand side captures the effect of a change in the wage on the number of startups of type j, where γ_j is the elasticity of the expected firm value with respect to w. This elasticity is equal to the present value of wage payments relative to the present value of profits.

We assume that labor supply is a function of the wage: $L^{supply} = L^{supply}(w)$.[13] Taking logs and differentiating, we obtain $dL^{supply} = \kappa \cdot dw$, where κ is the Frisch elasticity implied by the labor supply function. Given a certain change in the policy T and the wage w, we can now evaluate the new number of firms of each type and aggregate to compute the change in total labor demand. Labor market clearing implies:

$$dL^{demand} = dL^{supply} = \kappa \cdot dw. \tag{4}$$

3 Data and clustering

Having presented the model and the implied set of sufficient statistics to analyze the compositional effects of policies, we now turn to our data and measurement.

3.1 The CompNet database

We carry out a cross-country analysis of startups based on confidential administrative micro-level sources at the national level. These data were collected in close collaboration with CompNet, which was founded in 2012 by the European Central Bank and is currently hosted by the Halle Institute for Economic Research. CompNet provides its members and external researchers with a regularly updated, micro-based, and internationally harmonized compet-

[13] We abstract from income effects on labor supply since wage adjustments are a transfer between firm owners and workers, both of which are households. The labor supply schedule can also be micro-founded explicitly with Greenwood-Hercowitz-Huffman preferences.

itiveness data set for 20 European countries. To preserve confidentiality at the level of individual firms, and to improve cross-country comparisons, CompNet uses a "distributed micro-data approach" as developed by Bartelsman et al. (2004). This means data get annually updated by sending standardized code to national statistical agencies and central banks. These organizations then run this code on the confidential firm-level data they maintain and aggregate it up to the sector-year level in a standardized fashion. The data are subsequently returned to CompNet with key statistical moments that describe the distribution of a large number of firm characteristics at the sector-year level for each country.[14] The data set contains information about all firms in all private non-financial industries. Particular care is taken to ensure a high level of cross-country consistency to allow for international comparisons and the identification of idiosyncratic country effects (CompNet, 2018).

3.2 Startups in the CompNet database

To collect harmonized cross-country data on European startups, we embedded additional code in the standard instructions sent to national authorities in early 2021 in preparation for the eighth CompNet vintage.[15] Our code extracted data on all firms that commence their operations in a particular country and year (that is, a startup cohort). We thus define the start year as the year in which a firm is first economically active. We also observe each firm's formal registration year and drop observations if one or several of the following conditions hold: the lag between firm registration and actual startup is more than four years; registration occurs *after* the actual start year (this only happens in a handful of cases); the firm has more than 50 employees at the time of startup.[16] Once a firm enters our data set, we can track it for several years. This allows us to construct comparable data on how firms

[14]See Lopez-Garcia and di Mauro (2015) for more details on the CompNet project.

[15]We cannot use data for Germany, Poland, and the Slovak Republic as the national data sources exclude firms with fewer than 20 employees (10 in the case of Poland).

[16]Bayard et al. (2018) match new Employer Identification Numbers with employer records in the US business register. They define firm startup (firm age is zero) as the time the first payroll is observed in the Longitudinal Business Database. While recent applications account for the bulk of firm inceptions within a year, there is a long tail of startups that begin operations only several quarters after initial registration.

grow and develop in terms of their employment generation, productivity, and survival—all areas on which cross-country evidence remains scant.

This approach results in a unique cross-country and cross-sector panel of all startups established during the calendar years 2002–2019. Table A1 presents the coverage per country.[17] The panel contains information on a total of 1,345,489 startups. We aggregate the data in each country and year at either the macro (economy-wide) level or at the one-digit NACE Rev.2 industry level.[18] For each country-industry-year-cohort cell, our data set contains various firm characteristics, including average number of employees, average capital intensity, average cash ratio, average leverage ratio, as well as several productivity metrics.[19] Table A2 provides a detailed description of the main variables. All monetary variables are PPP-adjusted and real variables are deflated with sectoral price indices. We also retrieve the data split by startup type. This classification of startups will be discussed in Section 3.3.

In Appendix Figures A1–A3, we compare our startup population to Eurostat's Business Demography Statistics on startups (while excluding sole proprietorships for consistency). Figure A1 shows that startup rates (the number of startups in a year as a fraction of the total firm population) are very comparable in our CompNet-based data set and the aggregate data collected and published by Eurostat. The same holds, by and large, when we compare average employment growth during the first five years after startup (Figure A2). In most countries, trend growth in both data sets is very similar. In a few cases—such as the Netherlands and Sweden—there are gaps in the average *level* of reported employees. In those cases, a comparison with other countries suggests the Eurostat data are anomalous, rather than the CompNet data. Finally, Figure A3 compares exit rates (firm death) over time in both data sets. The five-year cumulative exit rate is comparable at 45 percent (56

[17]We drop firms that are not observed in each year between starting operations and exiting the data set.

[18]These industries are administration; construction; hospitality; ICT; manufacturing; professionals; trade; and transport. Due to confidentiality reasons, cells with fewer than ten underlying firm observations return empty. We also produced results at the more granular two-digit NACE industry level, but this approach created too many empty cells.

[19]We use two productivity measures. The first is labor productivity, defined as real value per employee. The second is TFP, which is based on a production function estimated via the two-step control function of Ackerberg et al. (2015) and implemented at the two-digit sectoral level.

percent) in the CompNet (Eurostat) data set. Moreover, the trends as firms age are very similar in both cases, although in Sweden CompNet tends to overreport firm exit. Overall, we conclude that the firm population that underpins the statistical moments we derive from CompNet is representative of the firm population in our sample countries as reported by Eurostat.

3.3 Identifying startup types

We use K-means cluster analysis (Everitt et al., 2011) as a data-driven approach to categorize firms according to their startup strategy. K-means clustering is a type of unsupervised machine learning that has recently gained traction in the applied economics literature as a way to study empirical settings with latent heterogeneity, see for example Bonhomme et al. (2022). In our application, the heterogeneity is in firm startup types, and the idea underlying our use of the clustering algorithm is that choices made in the very first year help to reveal the type of startup.

Let x_i be a vector of firm-level clustering variables, in practice converted into z-scores to avoid arbitrary scaling effects. The clustering algorithm allocates each individual firm i into one of $j = 1, .., k$ clusters, in order to solve $\min \sum_{j=1}^{k} \|x_i - \bar{x}_j\|^2$, where $\bar{x}_j$ is the cluster mean. The algorithm begins with k seed values as the initial group means. Each observation is then assigned to the group with the closest mean. Based on that initial categorization, new group means are determined and these iterative steps continue until no observations switch groups.

We experiment with different k's by calculating the statistic $\eta^2 \equiv 1 - \frac{WSS}{TSS}$, where $WSS \equiv \sum_{j=1}^{k} \|x_i - \bar{x}_j\|^2$ is the within-cluster sum of squares and $TSS \equiv \sum_{j=1}^{k} \|x_i - \bar{x}\|^2$ is the total sum of squares, with $\bar{x}$ being the unconditional mean across all observations.[20] We let k vary between 1 and 10, which is visualised in the scree plots in Appendix Figure A4. At $k = 5$, the η^2 statistic is around 0.6, which means five clusters capture around three-fifths of

[20] The statistic has a similar interpretation to the R^2 often reported in regression analysis.

the total variation in the clustering variables. Beyond $k = 5$, the η^2 statistic still increases but levels off. The data therefore suggest that our startups are optimally clustered into five well-separated (non-overlapping) clusters, each representing a distinct startup strategy.

We let the cluster algorithm group startups based on five important endogenous variables that entrepreneurs decide on when starting a business: the initial number of employees; real total assets; capital intensity (amount of real fixed assets per employee); cash to total assets; and leverage (total debt to total assets).[21] We therefore cluster using variables that are decided at the moment of startup (but can be adjusted during the life of the firm) rather than outcome variables such as labor productivity, TFP, or value added. We choose these five variables because the existing literature has either directly or indirectly identified them as key startup decision parameters and because the underlying CompNet microdata are complete and of high quality across all our sample countries.[22]

3.4 Comparing clustering results across countries

We implement the cluster analysis using a separate micro data set for each individual country. There is therefore no a priori reason for the clustering outcomes to be similar across countries. Indeed, very different kinds of clusters may arise in different contexts. Moreover, even if the type of clusters would turn out to be similar, their shares might vary widely across countries.

To assess the similarity of clusters across countries, we run a second-stage "meta-clustering" analysis, which groups comparable clusters from different countries. This also provides us with an objective procedure to assign common names to similar clusters. Specifically, we repeat the clustering analysis while taking the cluster centers derived from each country's

[21] All monetary variables in CompNet are denominated in thousands of euros.

[22] Earlier work has assessed the role of startups' scale as measured by total assets or employees (Albuquerque and Hopenhayn, 2004; Kerr and Nanda, 2010; Buera et al., 2011) or set-up costs (Derrien et al., 2020)); liquidity and cash holdings (Bolton et al., 2019); and leverage and use of bank credit (Robb and Robinson, 2014; Farinha et al., 2199; Derrien et al., 2020; Bustamante and D'Acunto, 2021). We also use individual firms' capital intensity (fixed assets per employee) as a clustering variable because heterogeneity in production functions (and the resulting variation in the elasticity of substitution between capital and labor) is expressed directly in different choices regarding capital-to-labor ratios (Oberfield and Raval, 2021).

first-stage clustering procedure as the observations.[23]

The four panels in Figure 1 visually summarize the outcome of this meta-clustering procedure. Different meta-clusters are indicated with different colours. The panels clearly show five clusters arising in all countries. For example, the first panel contains a red meta cluster of capital-intensive startup clusters. Each of the individual red circles indicates a country-level cluster of startups that stand out nationally because of their exceptionally high capital intensity. Such a distinct capital-intensive cluster emerges in each country, thus allowing the meta algorithm to bunch them together in a single meta-cluster.

Importantly, the variation between clusters within countries is much greater than the variation between countries in the same meta-cluster. In other words, the clusters arising in different countries turn out to be very similar. For the meta-clustering we obtain $\eta^2 = 0.96$ and this indicates that variation between clusters explains the vast majority of the overall data variation. This leaves only a very small contribution for cross-country variation within meta-clusters.[24]

One might be concerned there are mechanical reasons for the very similar clusters in different countries, or that this similarity is a coincidence. To investigate this, we conduct a Monte Carlo experiment for the meta-clustering. We consider many random draws for the cluster variables, with means and standard deviations as observed in the data. However, in the experiment, these draws are i.i.d. so that no meta-clusters exist.[25] We repeat this experiment many times and each time compute η^2. Appendix Figure A5 shows that these η^2 statistics are much lower in the experiment than the 0.96 observed in the real-world data. The Monte Carlo experiment therefore supports our interpretation that the cluster outcomes observed in the data are indeed remarkably similar across countries.

[23]That is, in the meta-clustering procedure, the units of observation are the z-scores of the first-stage cluster centers, averaged across years and industries. The z-scores are computed within each country to allow for measurement differences and institutional variation across countries.

[24]Another indication of the similarity of clusters across countries is that, in each country, all clusters fall into different meta-clusters. This is not mechanically the case. It could have been the case that the meta-clustering would assign multiple clusters in a certain country to the same meta-cluster.

[25]We assume log-normal distributions for these draws, as the cluster variables are non-negative.

4 An anatomy of startup types

This section presents several novel stylized facts that follow directly from our clustering analysis. A first key result is the presence of five distinct startup types across countries, industries, and cohorts (Section 4.1). A second important finding is that each of these startup types has a recognizable life cycle in terms of firm traits and performance measures (Section 4.2). Lastly, Section 4.3 sketches a short summary profile of each startup type.

4.1 Five types of startup

Table 1 summarizes the results of our clustering analysis. We label the five archetypal startup clusters that emerge *Basic* (49 percent of all startups); *Large* (4 percent); *Capital-intensive* (7 percent); *Cash-intensive* (26 percent); and *High-leverage* (14 percent). These labels reflect the key dimension along which a startup type clearly differentiates itself. For example, large startups employ, on average, 20 people when they begin operations, compared with an average of only three people in the other categories. Likewise, cash-intensive startups hold, on average, 54 percent of their assets as cash when they commence operations, whereas the average is just 12 percent for other types.

Figure 2 shows that, while the clustering algorithm yields the same five startup types in each country, their local prevalence differs somewhat. For example, cash-intensive startups are relatively ubiquitous in Italy but less so in France. Likewise, highly leveraged startups are relatively common in France but less widespread in Croatia, Denmark, and Lithuania.

Figure 3 views the composition of the startup population through a sectoral lens and shows that the five startup types also emerge within each main economic sector. The exact distribution over the five types nevertheless differs somewhat between sectors. For example, large startups are slightly overrepresented among manufacturing firms but underrepresented among information and communications technology (ICT) firms and enterprises that provide professional services. Cash-intensive firms are overrepresented in the ICT industry and professional services, while high-leverage startups are relatively common in the hospitality sector. This distribution of startup types by industry also varies across countries (see Figure A6 in the Appendix).

Finally, Figure 4 shows that the distribution of startups across the five types (and aggregated over all countries) is quite persistent during the period 2010–2019. We nevertheless observe a small decline in the share of highly leveraged and basic startups, while the share of cash-rich startups steadily increases during the protracted period of monetary easing in the wake of the global financial crisis.

Each of the five startup types also consistently displays its defining characteristic across countries. For example, as can be gleaned from the first panel of Figure A7 in the Appendix, large startups are indeed consistently larger (in terms of total assets) than the four other startup types. In relative terms, this difference is particularly striking in Italy but less so in a country such as Croatia. Likewise, the last panel of Figure A7 shows, for example, that high-leverage startups are indeed the most leveraged group of startups in each country, and this is particularly so in Denmark, Lithuania, and the Netherlands.

4.2 Startup types: Early life cycles

With the clustering results in hand, we now describe key patterns in the development of startups during the first 12 years of their existence. We first focus on the choice variables that we fed into the clustering algorithm. For each of these variables, we run age-specific regression specifications in which we regress the clustering variable on dummy variables for four startup types (we omit the *Basic* type as the base group) as well as country, cohort, and industry fixed effects. The sample is the full panel data set at the one-digit industry level. We run a separate regression for each age group (one-year-old firms, two-year-old firms, etc.) and plot the successive coefficients for the startup type dummies in Figure 5.

A number of salient patterns stand out. First, while *Cash-intensive* startups, by definition, begin their operations with substantially more cash (relative to total assets) than *Basic* and other startup types, they quickly reduce this cash intensity over time as they invest in other assets. Yet, even after 12 years, the cash intensity of this startup type remains about 10 percentage points higher than that of other startups. Second, *Large* startups not only start out with significantly more employees, this size gap vis-à-vis other startup types widens further during the first decade. While large startups employ, on average, 20 more people than basic startups when they commence operations, this difference increases to about 30 employees after 12 years. Third, we find clear evidence of convergence in leverage ratios across startup types. In particular, *High-leverage* startups are about 50 percentage points more leveraged than basic startups when they begin operations (even within the same country and industry). That excess leverage is reduced quickly during the first decade of operations so that, after 12 years, the difference has shrunk to just 5 percentage points. Fourth, we also find (partial) convergence in terms of startups' capital intensity. *Capital-intense* startups begin production with an almost 50 percentage point higher capital-to-employee ratio. Over time, however, they quickly reduce this gap to about 15 percentage points. Fifth and finally, we see that large startups are not only bigger in terms of staff numbers but also in terms of total assets. Yet, while large startups gradually expand their staff numbers even further, relative to other startup types, they slightly reduce the size of their balance sheet relative to other startup types in the same country and industry. There is some convergence in the average capital intensity of these large firms over time.

Next, we are interested in how different types of startups perform as they grow older. To look into this in a systematic way, Table 2 reports panel regressions for several performance measures. Observations refer to cell averages for all new firms in a given country, one-digit industry, startup year (cohort), age, and type.[26] The analysis uses the full sample and we include dummies to indicate the four main startup types, again using the *Basic* type as the excluded category. We saturate the specifications with an exhaustive set of interactive fixed effects (FE) to flexibly control for various unobservable drivers of firm-level

[26]For example, an observation could refer to the average productivity of Spanish ICT firms in the *High-leverage* category established in 2005, at age seven.

performance that might correlate with startup type. In particular, country-cohort FE absorb all time-invariant characteristics common to startups established in a specific country and year; industry-cohort FE control for time-invariant traits common to all startups in a specific industry and established in a given year; and country-industry FE absorb all time-invariant variation that characterizes startups in a specific industry and country. Finally, we add interactions between startups' age and their country, industry, and cohort. This allows us to flexibly control for startup traits that are specific to countries, sectors, or establishment years *and* that depend on a firm's age.[27]

Panel A of Table 2 reveals some interesting patterns. First, compared to firms with a basic startup strategy, startups with a more differentiated strategy tend to outperform in terms of higher labor productivity (column 1) and TFP (column 2), as well as a lower likelihood of early exit (column 3). In particular, *Large* startups are considerably less likely to exit within the first decade after commencing operations. The latter category (as well as cash-intensive firms) also operates at an above-average profit margin (column 5)—even relative to other startups in the same country, sector, and industry—notwithstanding the fact that these firms pay substantially higher wages (column 4). An important exception are the highly leveraged firms. This cluster of startups consistently and strongly underperforms in terms of labor productivity and TFP. Highly leveraged firms also tend to operate with a lower profit margin (column 5) even though they pay lower wages compared to the other startup types (column 4).

In Panel B of Table 2, we replicate this analysis using the sub-sample of firms of between five and eight years old. The results line up closely with those in Panel A, indicating that the performance divergence between startup types is not solely driven by transient differences in early life.

[27]Throughout the paper, we provide robust but unclustered standard errors as we use data for the full population of startups across sectors and countries rather than for a random sample of sectors or countries (Abadie et al., 2017). All our results are robust to clustering at the sector-country-cohort level to account for possible correlation of model errors over time at that level.

4.3 Startup types: Profiles

This section provides a short profile for each startup type in order to consolidate and interpret the results from our empirical exploration so far in light of the literature.

4.3.1 Capital-intensive startups

About 7 percent of all startups belong to the capital-intensive cluster. A plausible narrative for the existence of capital-intensive types concerns heterogeneity in production functions, with some firms having a particularly high production function elasticity with respect to capital. Oberfield and Raval (2021) provide evidence for heterogeneity in production elasticities and derive implications for the macroeconomy, such as the aggregate labor share of income. Our results underscore the importance of such heterogeneity, and provide further insights into the life cycles and performance of firms with particularly high capital intensities.

From the start, capital-intensive firms use many fixed assets per employee, possibly reflecting the lumpiness of initial investments (Doms and Dunne, 1998; Cooper and Haltiwanger, 2006). Over time, their capital intensity—gradually and partially—converges with that of other startup types. This suggests that, as demand grows, capital-intensive startups can scale up production by increasing the number of employees that utilize the substantial stock of machines and other assets. For example, firms may start with a core of highly skilled employees and hire additional lower-skilled staff as they grow. In some sample countries, our data indeed show a fall in wages per employee as capital-intensive startups mature.

Notwithstanding a partial adjustment process during the first decade, there is in fact a strong persistence in these startups' reliance on capital in the production process (Figure 5). This not only reflects the industries in which capital-intensive startups tend to operate (such as manufacturing and transport) but also their chosen production strategy (such as a high degree of automation). Even after more than a decade, capital-intensive startups continue to stand out—also relative to other firms in the same industry or country—by the sheer amount of fixed assets they employ. They remain about three times as capital-intensive as firms in the other startup clusters.

Empirical evidence has shown that capital-intensive firms are prone to complement valuable fixed assets with high-quality human capital (Doms et al., 1997). In line with this, we observe that capital-intensive startups pay higher wages than most other firm types and, unsurprisingly, boast higher labor productivity. They also operate at high levels of TFP, suggesting that they combine physical and human capital in a relatively productive way. As a result, the medium-term survival prospects for this type of startup are relatively good. Over time, the share of capital-intensive firms in the startup population has remained relatively stable.

4.3.2 Large startups

About 4 percent of all startups belong to the *Large* cluster. Large firms are relatively common in construction, transport and logistics, and manufacturing. These firms commence operations with both a large number of employees (more than five times as many as other startups) and a substantial asset base, reflecting their use of a scalable production technology and access to a scalable demand base. This is consistent with empirical findings in Sterk et al. (2021), who use US data to document the importance of ex ante heterogeneity for firm-level employment.

Size heterogeneity may derive from TFP differences combined with decreasing returns to scale (which implies that more productive firms choose a larger optimal size) or from heterogeneity in returns to scale, with some production processes being scaled up relatively easily. A final possibility is that size heterogeneity reflects differences in the scalability of the demand for firms' products. Hottman et al. (2016) show that differences in the scale of demand is a quantitatively important source of heterogeneity across firms.

Our results again shed light on the life-cycle structure and performance of heterogeneous firms, this time along the size dimension. Large startups are relatively productive and, importantly, this productivity premium is very persistent over time. They scale up their workforce even further during the first 12 years of their existence (thus diverging even more from other firm types in terms of employee numbers).

While large startups tend to operate on the basis of relatively modest profit margins, they display high TFP when compared to both basic and high-leverage startups—at levels similar to capital-intensive and cash-intensive startups. This good productivity performance also allows large startups to pay relatively high wages. This is in line with an extensive literature documenting a firm-size wage premium (Brown and Medoff, 1989; Oi and Idson, 1999; Troske, 1999). We show that such a premium already exists when comparing very young firms in the same country and industry.

Even though they are relatively highly leveraged, the overall advantages of size—such as more diversification—make large startups the least likely to exit during the first 12 years of operations. This is in keeping with an established literature documenting a positive relationship between firm size in early life and subsequent survival (Dunne et al., 1988; Geroski et al., 2010).[28]

4.3.3 High-leverage startups

About 14 percent of all startups belong to the *High-leverage* cluster. These firms begin operations by taking on substantial amounts of debt relative to the size of their balance sheet (their average leverage ratio is 116 percent).[29] While high-leverage startups are present in all sectors, they are relatively common in accommodation and food services, and, to a lesser extent, administrative and support services.

Interestingly, during their early life, these firms immediately deleverage, suggesting they may be financially quite vulnerable. High-leverage startups have the lowest 12-year survival probability. This aligns well with an earlier literature documenting a possibly causal link

[28]Branstetter et al. (2014) show how firm-entry deregulation in Portugal caused additional firm formation but mostly of smaller and less productive firms that were less likely to survive during the first two years.

[29]Owners of startups often finance their venture with personal loans, so that total borrowing can exceed the firm's assets (Robb and Robinson (2014) and Bustamante and D'Acunto (2021)).

between firm leverage and exit probability.[30] Dinlersoz et al. (2019) show that young private firms in the US tend to be highly leveraged but deleverage as they age. We find a similar pattern within the high-leverage cluster of startups. At the same time, we also show that, while firms in this cluster deleverage over time, their leverage remains persistently above that of other young firms. A similar pattern has been demonstrated previously by Lemmon et al. (2008) on the basis of Compustat data for publicly listed firms (and not conditioning on firm types). The authors document that, while leverage ratios exhibit substantial convergence, they are also remarkably stable over long periods of time. We also find clear evidence for both such a transitory and a permanent component in leverage ratios in our cross-country startup data.

High-leverage startups exhibit relatively low levels of productivity, and profitability and it often takes several years for them to become profitable at all.[31] This cluster of startups consequently pays considerably lower wages than all other startup types, even within the same industry or country. There are several reasons why highly leveraged firms tend to underperform in terms of productivity, profitability, and survival probability. One option is that these firms are costly to establish (for example, because an intensive marketing campaign is needed) and these fixed set-up costs are funded with bank debt (Derrien et al., 2020).

An alternative explanation for why startups with low initial leverage subsequently perform better is based on the insight that, at the time of startup, the information asymmetry between banks and firms is highest. Firms with high-quality investment projects, who know they soon will become very profitable, may not want to delay investment but instead settle for accessing (relatively little) credit early on (Bustamante and D'Acunto, 2021). A third possibility is that firms with high initial leverage are run by subsistence entrepreneurs with a

[30]See, for example, Chevalier (1995) and Kovenock and Phillips (1997). Zingales (1998) shows for the US trucking industry that initial leverage (at the beginning of a deregulation period) increased the probability of subsequent default. An important channel is highly leveraged firms' impaired ability to invest, as also shown by Lang et al. (1996).

[31]Rajan and Zingales (1995) provide cross-country evidence of a negative correlation between profitability and leverage. This is in line with pecking-order theory, which posits that profitable firms prefer to use their internal funds (retained earnings) over external debt (Myers and Majluf, 1984).

high "utility value" of being an entrepreneur but with limited business skills (Schoar, 2010). While such firms are typically only moderately profitable, subsistence entrepreneurs may nevertheless raise debt if banks accept personal collateral (so that the business owner is personally liable in case their enterprise goes bankrupt).

4.3.4 Cash-intensive startups

About one-quarter of all startups belong to the *Cash-intensive* cluster: these are typically smaller startups with a high ratio of cash to total assets. At startup, these firms hold almost half of their assets in the form of cash, about five times as much as firms in the other startup categories. They are more common in the administrative, ICT, and professional services sectors, and typically do not grow much over time. Because these firms keep a large buffer of liquid assets, the proportion of fixed assets such as machinery is initially relatively low (importantly, this holds even *within* sectors). Cash-intensive startups are therefore also firms with a quite low capital intensity.

Over time, cash-intensive firms reduce their liquidity as they gradually convert some cash holdings into both tangible and intangible fixed assets. Although their cash ratio consequently comes down, a persistent gap remains vis-à-vis the other startup clusters. Even after 12 years, the cash intensity of this startup type remains about 10 percentage points higher than that of other young firms. Interestingly, the performance of cash-intensive startups is excellent, also compared to other enterprises in the same sector. They have high levels of labor productivity, TFP, and profitability, and consequently also pay somewhat higher wages. These patterns are in line with Begenau and Palazzo (2021), who show that public firms with a high expected productivity growth, and hence a bigger need for future investment, operate with a higher initial cash-to-assets ratio.

Many cash-intensive startups are small services-oriented firms run by highly skilled individuals. The services they provide typically require few machines or other fixed assets and are only to a limited extent scalable. Few additional employees are hired over time. The results suggest that, rather than producing homogeneous goods at scale for a mass market, many cash-intensive startups instead sell to a profitable niche market of heterogeneous consumers by implementing a differentiated strategy based on targeted advertising and detailed sales advice (Johnson and Myatt, 2006). These startups do not need to borrow much to commence operations, and therefore consistently display the lowest leverage ratio of all clusters. This also suggests that they use (internal) cash buffers rather than (future) access to credit as a precaution against demand shocks. This strategy may be optimal because these firms own few tangible assets that could serve as collateral, or because they are relatively non-standard, making it difficult to obtain bank credit.

4.3.5 Basic startups

Finally, just fewer than half of all startups make up the *Basic* cluster. This largest startup cluster looks average across all the dimensions discussed so far (Table 1). When they begin operations, these firms employ, on average, four people; operate with a capital intensity of 8.6 percent; have a cash ratio of 12 percent; and a leverage ratio of 23 percent. Basic startups conduct operations at a level of labor productivity and TFP that is lower than that of all other startup types, except for the highly leveraged ones. Productivity growth is relatively muted, too. After a decade, basic startups still operate at relatively low TFP, labor productivity, and profitability levels.

5 Policy experiments

We now perform counterfactual policy exercises to study how changes in taxation can improve specific macroeconomic outcomes by changing the composition of new startup cohorts. We do this using the model framework developed in Section 2. As explained in that section, to quantify how tax policy affects startup composition, we need the life-cycle profile of each startup type. In addition, we need a second set of statistics: the entry elasticities for each type. We first discuss how we estimate these elasticities and then present our policy exercises.

5.1 Estimating entry elasticities

To estimate entry elasticities, we run a regression based on the entry condition, Equation (2). A challenge, however, is that we do not observe the ex ante expected firm value $\ln \mathbb{E}_j V$ but rather the ex post average realization $\ln V_j$. The difference between the two is because of common shocks that occur *after* entry. If we were to estimate Equation (2) based on ex post realizations, the residual would contain $\ln V_j - \ln \mathbb{E}_j V$, giving rise to correlation between the residual and the right-hand side variable, $\ln V_j$. This would introduce a bias. A second issue is that, under a null hypothesis of free entry (a baseline assumption in the theoretical literature), the entry elasticity is infinite, making the regression specification ill-defined. To circumvent these problems, we rearrange Equation (2) as:

$$d \ln V_j = \frac{1}{\varepsilon_j} d \ln n_j + \xi_j + u_j$$

where $\xi_j \equiv \frac{\gamma_j}{\varepsilon_j}$ is a fixed effect due to entry shocks, for which we control in the regression. Moreover, $u_j \equiv d \ln V_j - d \ln \mathbb{E} V_j$, which is orthogonal to n_j since ex post shocks are not known at the time of entry, and thus cannot affect the number of entrants, n_j.

We use observations at the country(c)-industry(i)-type(j)-cohort(t) level, with corresponding indices between brackets. We control for shocks to the number and distribution of entrants at the country-industry-type level, as well as at the country-industry-year level by using interactive fixed effects at these levels. The former is important as the number of potential entrants may naturally vary by country, industry, and type. The latter captures that, in a specific industry in a given country, the number of potential entrants may fluctuate over time.[32] This gives rise to the following regression specification, which we estimate via

[32] Implicit in our specification is the assumption that, within a country-industry cell, shocks over time to the number of *potential* startups are common across types.

ordinary least squares (OLS):

$$\ln V_{c,i,j,t} = \beta_0 + \beta_1 \ln n_{c,i,j,t} + \xi_{c,i,j} + \xi_{c,i,t} + u_{c,i,j,t}.$$

We note that $\frac{1}{\beta_1}$ is the entry elasticity and that, under free entry (that is, an infinite elasticity), it holds that $\beta_1 = 0$.

A final question is how to measure $\ln V_j$ in the data. We do so using each startup type's life-cycle profile for profits and exit rates, and the number of entrants for each country-industry-startup type-cohort in our data set.[33] We assume a discount factor $\Lambda = 0.96$, corresponding to an annual discount rate of about 4 percent. We only use cohorts that we observe for at least seven years (that is, those established before 2011) and drop observations beyond age seven. For age eight and onward, we assume that profits and the year-on-year exit rate remain fixed.[34]

Figure 6 plots the estimates for the inverse entry elasticity, β_1. The vertical dashed line denotes an estimate from a regression that does not condition on startup type. This overall estimate is 0.495, implying an entry elasticity of around two. This estimate is significantly different from zero at the 1 percent level, thus rejecting the traditional free-entry assumption.[35]

The figure also shows the estimates of the inverse elasticity by firm type, illustrating substantial heterogeneity. Basic startups show the highest entry elasticity, about 5.3, whereas large startups have an entry elasticity of around 1. It is perhaps not unexpected that basic, unremarkable startups react relatively flexibly to changes in economic conditions, as the creation of these startups is, for the most part, uncomplicated. Other types may require

[33] We reconstruct profits by multiplying profits (Earnings Before Interest and Tax (EBIT)) per unit of revenue with revenues, at the country-industry-startup type-cohort level.

[34] That is, we compute the firm value at age zero as: $\ln V_0 = \sum_{a=0}^{6} \Lambda^a s_{0,a}\pi_a + s_{0,7}\pi_7 \sum_{a=7}^{\infty} \Lambda^a s_{6,7}^a = \sum_{a=0}^{6} \Lambda^a s_{0,a}\pi_a + \Lambda^7 s_{0,7}\frac{\pi_7}{1-\Lambda s_{6,7}}$ where $s_{k,l}$ is the survival rate between age k and $l \geq k$. In practice, the component of the value beyond age seven is relatively unimportant for the overall startup value, due to both discounting and the high exit rates in young age groups.

[35] Our overall estimate is in the ballpark of those from estimated structural DSGE models. Gutiérrez et al. (2019) find an estimate of around 1.5, whereas Sedláček and Sterk (2017) find an estimate of around 5.5.

more specific expertise, which fewer potential entrepreneurs possess.

5.2 Counterfactual experiments

We can now proceed with our policy experiments. Specifically, we consider the aggregate effects of changes in tax/subsidy policies via the startup composition channel. As explained in Section 2, to quantify the composition channel, we do not need to take a stance on the precise nature of the tax reform. All that matters for this channel is the direct effect of the reform on the tax payments of firms, as a fraction of the firm value, for the different startup types. Thus, the following results apply to a wide range of policy changes in the taxation and/or subsidization of startup firms.

To fix ideas, however, we present the results as a change in corporate income taxation, differentiated by startup type. This is without loss of generality, but makes the magnitude of the tax reform easy to interpret since, in this case, the tax payments as a fraction of the firm's value simply equal the tax rate itself. Furthermore, we do not quantify the effects on post-entry behavior of firms, since this *does* depend on the specifics of the tax reform, and has been widely studied in a large literature, unlike the composition effects.

To impose discipline, we further restrict ourselves to policies that are revenue-neutral.[36] This implies that, if some startup types are to be taxed less, other types need to be taxed more. We also impose that the change in corporate income rate is capped at 20 percentage points for each type. Aside from these restrictions, we search the entire space of possible tax rate changes differentiated by type, and evaluate the macroeconomic implications for any admissible policy.

Our policy experiments are based on a numerical Monte Carlo procedure. Specifically, we consider a large number of uniform random draws for the tax policies and rescale them to ensure they are (post-reform) budget neutral. Then, based on the estimated entry elasticities, we evaluate the change in the number of firms within each type. This allows us

[36] To be precise, they are revenue-neutral excluding potential tax revenue effects due to changes in post-entry behavior.

to construct counterfactuals for the share $\omega_{a,j}$. Based on this, we then calculate counterfactual macroeconomic aggregates, exploiting the aforementioned result that the post-entry behavior of firms is not affected by the corporate income tax. This enables us to compute macroeconomic aggregates using the life-cycle profiles estimated from the data. We compute the macro outcomes for firms up to age 12, the oldest for which we have estimated the life-cycle profile, while accounting for the exit rate of each type (which we also estimate from the data). Figure A8 summarizes the life-cycle profiles we use in the analysis.[37]

Once we have evaluated the macro outcomes for the entire space of admissible policies, we group the policies into bins according to their "policy intensity", as measured by the absolute change in the tax rate, averaged across firms. The idea is that larger policy changes may bring about greater macro-economic effects, but may also face stronger political resistance. It is therefore useful to study the policy space conditional on a certain intensity of the policy. For each bin, we compute policies that bring about the largest positive and the largest negative change in a macro variable of interest (such as aggregate productivity or employment). We label this the "policy frontier" and present the tax-rate changes that generate the policy upper bound.

Figure 7 summarizes the results of our policy experiment. Panel A presents the policy space for aggregate labor productivity. The horizontal axis measures the intensity of the potential policy change (that is, the average absolute change in the tax rate). Moving from left to right, warmer colors indicate stronger corporate tax-rate differentiation. The solid (dashed) line plots the policy upper bound (lower bound). This is the largest possible aggregate employment increase (decrease) given a certain policy intensity.

The figure shows that substantial macroeconomic gains are possible. For instance, for a policy intensity of 0.1 (10 percentage point change in tax rate, on average), aggregate labor productivity can be increased by more than 6 percent by shifting the distribution of startup types. At the same time, substantial losses are also possible for policies that shift

[37] We also consider a version that includes firms older than 12 in the computations, which we achieve by extrapolating the life-cycle profiles. The results are very similar to our baseline findings.

the distribution away from high-productivity startup types.

Panel B of Figure 7 again plots the policy space, this time for aggregate employment. Large employment gains are possible, more than 12 percent for a policy intensity of 0.1. At the same time, even higher employment losses are possible for policy changes that move the distribution of startups away from high-employment types.

Next, the two panels on the left of Figure 8 plot the corporate tax-rate policies associated with the two policy upper bounds. Likewise, the two panels on the right show the associated firm shares. The upper left panel shows that, in order to achieve the productivity upper bound, taxes should be lowered for capital-intensive types (potentially turning the existing tax into a subsidy), financed by an increase in tax rates for the other startup types. The upper right panel shows that, as a result, the share of capital-intensive startups increases dramatically, whereas the share of basic types falls. Since the former have much higher levels of labor productivity than the latter, aggregate labor productivity increases.

The lower left panel of Figure 8 presents the tax rates associated with the employment frontier. In order to boost aggregate employment, tax rates should (unsurprisingly) be lowered for the large type (and, to some extent, the capital-intensive type). This tax cut is then financed by an increase in tax rates for the cash-intensive types (which not only employ fewer people but also have a fairly low entry elasticity, making them a relatively attractive target to tax). As a result of the policy change, the share of large startups increases markedly, whereas the share of cash-intensive startups falls. Since the former type of startup employs more workers than the latter, aggregate employment increases.

To analyze policy trade-offs, Figure 9 provides scatter plots that each depict pairs of changes in aggregate or average macroeconomic outcomes resulting from each potential policy (again, up to a policy intensity of 0.2). The top left panel suggests that, generally, there is not much of a trade-off (nor complementarity) between increasing aggregate employment and labor productivity as the two outcomes tend to be uncorrelated. Yet, there exist specific policy reforms that simultaneously boost (or reduce) aggregate labor productivity and employment. Interestingly, there is a fairly strong negative trade-off between aggregate employment and TFP (top middle panel). On the other hand, policies associated with higher aggregate labor productivity also tend to increase TFP (top right panel).

Policymakers may also be interested in effects on workers. The middle row of Figure 9 shows trade-offs between average wages and aggregate outcomes. Policies boosting aggregate employment through a different startup mix tend to be detrimental to average wages, whereas policies that increase aggregate productivity are associated with higher wages, as one might expect.

The bottom row of Figure 9 provides insight into the effects on consumers by plotting profit margins. We observe that policies that increase employment tend to reduce profit margins, which may be to the benefit of consumers. In contrast, policies to increase productivity tend to increase profit margins, suggesting the productivity gains are, at least in part, reaped by firm owners as opposed to consumers.

5.2.1 Selection effects within types?

The macroeconomic effects of the policies analyzed above are due to selection effects between different types of startups. As explained previously, in the model there is no entry selection within types. To test whether such entry selection can indeed be ruled out empirically, we run age-partitioned regressions for all startups and where observations are at the country-sector-startup type-cohort level (for example, Spanish construction firms of the cash-rich startup type that are six years old).

For each age, we regress an outcome variable—such as labor productivity or employment—on the (log) number of firms in a cohort to estimate the effect of cohort size on later startup performance. We fully saturate the regressions with interactive fixed effects at the country-sector-type, country-sector-cohort, country-type-cohort, and sector-type-cohort levels. If there is positive selection, cohorts with more startups should exhibit inferior performance later in life (when absorbing all variation that may correlate with cohort size through the use of the aforementioned battery of interactive fixed effects).

Figure A9 shows the results for firms up to age 10. The estimated effects are mostly insignificant, economically small, and do not trend in any particular direction. This indicates that—in line with our model—there is little within-type selection at the entry stage.

5.2.2 Equilibrium effects

We now revisit the baseline policy experiment, while accounting for equilibrium effects, as discussed in Section 2.5. Such effects may be particularly relevant if an EU-wide policy were to be implemented. Concretely, we endogenize the real wage, compute the composition of startups using Equation (3), and impose the labor market clearing condition (4) as well as budget-neutrality. That is, we solve numerically for the (change in) the equilibrium wage that clears the labor market given the policy change.

To implement this exercise, we set $\kappa = 0.5$, which is in the ballpark of estimates for the Frisch elasticity in the literature. The elasticity of the firm value with respect to profits is given by $\gamma_j = -\frac{\mathbb{E}_j \sum_{t=0}^{\infty} \beta^t \Lambda^t w L_{t,j}}{V_{t,0}}$. That is, $-\gamma_j$ is the ratio of the present-value wage bill to the present value of profits. We can compute this based on our data, given we observe the firms' wage bills. We use the same truncation procedure for the present-value wage bill as for the firm value.

Figure A10 presents the results. When we account for equilibrium effects, the policy space for aggregate productivity turns out to be very similar to our baseline results. The space for employment effects is somewhat smaller. This is to be expected as an increase in labor demand pushes wages up, dampening the increase in aggregate labor demand.

5.3 Tax differentiation for startups in practice

Our counterfactual policy exercise illustrates how recalibrating corporate tax rates across startup types can bring about substantial macroeconomic gains in terms of higher aggregate employment and/or labor productivity. Such differentiation should, in principle, be feasible as it distinguishes startup types on the basis of a few observable characteristics. Indeed, many countries already—explicitly or implicitly—differentiate corporate tax rates, resulting in substantial variation in *effective* tax rates across different kinds of startups (as well as between startups and more mature firms).

First, some countries explicitly differentiate corporate tax rates based on firm size as measured by the number of employees and/or total taxable income. Belgium, France, Hungary, Latvia, Lithuania, Luxembourg, the Netherlands, Poland, Portugal, Slovakia, and Spain all apply reduced corporate tax rates to firms below a certain size threshold.[38] Likewise, the US federal government levies a 35 percent top rate for companies with at least USD 10 million in annual profits, while smaller firms benefit from lower rates. Other countries differentiate corporate income tax rates on the basis of the R&D intensity of firms. For example, China operates a base tax rate of 25 percent, but a considerably lower rate of 15 percent for qualifying high-tech enterprises. Qualification is based on firms' research and development (R&D) intensity as well as size thresholds (Chen et al., 2021). A few countries differentiate tax rates among startups specifically. For example, the rate that applies to Indian manufacturing startups depends directly on their annual turnover (Kalra, 2019).

Second, several countries engage in more implicit forms of tax differentiation among small businesses and startups. These include subsidies and grants, accelerated depreciation measures, investment allowances, as well as tax credits, breaks, and exemptions (OECD, 2020). Such measures can lower the effective tax rate for small and young businesses substantially (Rosenberg and Marron, 2015). In some cases, tax benefits only apply to startups with a particular funding structure. These include special credit guarantees and loans for startups, for example.[39] In other cases, tax rebates depend on a startup's physical location. In the UK, startups qualify for a 100 percent tax exemption for five years if they are located in

[38] See Bergner et al. (2017) and OECD (2021b). The largest difference between the reduced and the standard tax rate can be found in Portugal, where large businesses pay a rate of 31.5 percent, while small businesses pay a reduced rate of 17 percent on taxable income below a certain threshold.

[39] Other measures, such as the UK Enterprise Investment Scheme (EIS) and Seed Enterprise Investment Scheme (SEIS) help startups to attract equity funding by offering tax relief to individual investors who buy new shares in startups.

one of 21 enterprise zones, while in China preferential tax rates apply to firms located in the country's western provinces.

Third, many countries offer specific tax advantages to startups that are expected to contribute disproportionately to some desired policy outcome, such as job creation or productivity (often proxied by R&D investments or a similar innovation measure). R&D tax credits and allowances are particularly common (OECD, 2021a). Qualification for such tax credits typically depends on basic firm characteristics that are similar to those in our cluster analysis. For example, startups in the UK qualify for R&D tax relief if they have fewer than 500 employees and revenues of less than EUR 100 million, or own less than EUR 86 million in assets.

The above discussion illustrates how many countries already operate corporate tax systems that in some way or form differentiate, often implicitly, between types of startups.[40] The resulting variation in effective tax rates is often substantial and reflective of policy goals such as job creation or innovation-driven productivity growth. Many of these real-world examples of corporate tax differentiation can be quite easily translated to, or nested in, our more general policy exercise.

6 Conclusions

This paper has combined a large-scale administrative data set for multiple European countries with a theoretical framework to help understand the potential macroeconomic gains of tax policies designed to encourage a better mix of startups. Using unsupervised machine learning, we find that very similar clusters of startups (recognizable by the choices they make at entry) emerge in a number of European countries. Moreover, the subsequent performance of these types in terms of employment, productivity, and survival differs strongly. There are therefore potential macroeconomic gains (or losses) to be made from policies that affect

[40]Because both the type of investments and the financial structure of startups differ across industries, there also exists substantial cross-sector variation in effective tax rates (Rosenberg and Marron, 2015).

startup types differently and that therefore alter the composition of new startup cohorts.

Applying policy exercises based on the theoretical framework and the empirical results, we find these macro gains/losses can be quantitatively substantial. Moreover, our results inform how policy could be changed to achieve these macroeconomic gains. In particular, the counterfactual analysis sheds light on which types of startups should be encouraged (such as capital-intensive ones) and which types should be discouraged (such as high-leverage ones). In short, policymakers aiming to improve macroeconomic performance would benefit from considering the incoming generations of startups, and how existing policies impact the composition of these new cohorts and their contribution to aggregate employment and productivity growth.

Our empirical strategy and theoretical framework are easy to implement. This makes it a straightforward complement to standard analyses evaluating the macroeconomic effects of particular tax reforms, which typically ignore impacts on the composition of new startup cohorts. Increasingly, statistical agencies make rich administrative micro data sets on firms available for research purposes, and our methodology can therefore readily be applied to a wider set of countries. It would also be interesting to explore if an even more granular classification of startup types could be exploited to design more fine-tuned policies. Finally, it would be useful to explore to what extent differences in startup composition can account for cross-country heterogeneity in macroeconomic performance. We leave these issues for future research.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey Wooldridge (2017) "When Should You Adjust Standard Errors for Clustering?," *NBER Working Paper No. 24003*, National Bureau of Economic Research.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer (2015) "Identification Properties of

Recent Production Function Estimators," *Econometrica*, 83 (6), 2411–2451.

Akcigit, Ufuk and Sina T. Ates (2021) "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory," *American Economic Journal: Macroeconomics*, 13 (1), 257–98.

Akcigit, Ufuk, John Grigsby, Tom Nicholas, and Stefanie Stantcheva (2018) "Taxation and Innovation in the 20th Century," *NBER Working Paper No. 24982*, National Bureau of Economic Research.

Akcigit, Ufuk, Douglas Hanley, and Stefanie Stantcheva (2016) "Optimal Taxation and R&D Policies," *NBER Working Paper No. 22908*, National Bureau of Economic Research.

Albert, Christoph and Andrea Caggese (2021) "Cyclical Fluctuations, Financial Shocks, and the Entry of Fast-growing Entrepreneurial Startups," *Review of Financial Studies*, 34 (5), 2508–2548.

Albuquerque, Rui and Hugo A. Hopenhayn (2004) "Optimal Lending Contracts and Firm Dynamics," *Review of Economic Studies*, 71 (2), 285–315.

Alon, Titan, David Berger, Robert Dent, and Benjamin Pugsley (2018) "Older and Slower: The Startup Deficit's Lasting Effects on Aggregate Productivity Growth," *Journal of Monetary Economics*, 93, 68–85.

Azoulay, Pierre, Benjamin F. Jones, J. Daniel Kim, and Javier Miranda (2020) "Age and High-Growth Entrepreneurship," *American Economic Review: Insights*, 2 (1), 65–82.

Bajgar, Matej, Giuseppe Berlingieri, Sara Calligaris, Chiara Criscuolo, and Jonathan Timmis (2020) "Coverage and Representativeness of Orbis Data."

Bartelsman, Eric, John C. Haltiwanger, and Stefano Scarpetta (2004) "Microeconomic Evidence of Creative Destruction in Industrial and Developing Countries," *World Bank Policy Research Working Paper No. 3464*.

Bayard, Kimberly, Emin Dinlersoz, Timothy Dunne, John Haltiwanger, Javier Miranda, and John Stevens (2018) "Early-Stage Business Formation: An Analysis of Applications for

Employer Identification Numbers," *NBER Working Paper No. 24364*, National Bureau of Economic Research.

Begenau, Juliane and Berardino Palazzo (2021) "Firm Selection and Corporate Cash Holdings," *Journal of Financial Economics*, 139 (3), 697–718.

Belenzon, Sharon, Aaron Chatterji, and Brendan Daley (2017) "Eponymous Entrepreneurs," *American Economic Review*, 107 (6), 1638–1655.

Benzarti, Youssef and Jarkko Harju (2021) "Using Payroll Tax Variation to Unpack the Black Box of Firm-Level Production," *Journal of the European Economic Association*, 19 (5), 2737–2764.

Bergner, Soren M., Rainer Brautigam, Maria T. Evers, and Christoph Spengel (2017) "The Use of SME Tax Incentives in the European Union," *ZEW Discussion Paper 17-006*, ZEW—Leibniz Centre for European Economic Research, Mannheim.

Bernard, Andrew B., Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes (2022) "The Origins of Firm Heterogeneity: A Production Network Approach," *Journal of Political Economy*, forthcoming.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts (2013) "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, 128 (1), 1–51.

Bloom, Nick, Rachel Griffith, and John Van Reenen (2002) "Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979–1997," *Journal of Public Economics*, 85 (1), 1–31.

Bolton, Patrick, Neng Wang, and Jinqiang Yang (2019) "Investment under Uncertainty with Financial Constraints," *Journal of Economic Theory*, 184, 104912.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2022) "Discretizing Unobserved Heterogeneity," *Econometrica*, 90 (2), 625–643.

Branstetter, Lee, Francisco Lima, Lowell J. Taylor, and Ana Venâncio (2014) "Do Entry

Regulations Deter Entrepreneurship and Job Creation? Evidence from Recent Reforms in Portugal," *Economic Journal*, 124 (577), 805–832.

Brown, Charles and James Medoff (1989) "The Employer Size-Wage Effect," *Journal of Political Economy*, 97 (5), 1027–1059.

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin (2011) "Finance and Development: A Tale of Two Sectors," *American Economic Review*, 101 (5), 1964–2002.

Buera, Francisco J and Yongseok Shin (2017) "Productivity Growth and Capital Flows: The Dynamics of Reforms," *American Economic Journal: Macroeconomics*, 9 (3), 147–85.

Bustamante, Maria Cecilia and Francesco D'Acunto (2021) "Banks' Screening and the Leverage of Newly Founded Firms,"Technical report.

Chen, Zhao, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Yi Xu (2021) "Notching R&D Investment with Corporate Income Tax Cuts in China," *American Economic Review*, 111 (7), 2065–2100.

Chevalier, Judith A. (1995) "Capital Structure and Product-Market Competition: Empirical Evidence from the Supermarket Industry," *American Economic Review*, 415–435.

CompNet (2018) "Assessing the Reliability of the CompNet Micro-Aggregated Dataset for Policy Analysis and Research: Coverage, Representativeness and Cross-EU Comparability,"Technical report.

Cooper, Russell W. and John C. Haltiwanger (2006) "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*, 73 (3), 611–633.

Decker, Ryan A., John C. Haltiwanger, Ron S. Jarmin, and Javier Miranda (2017) "Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown," *American Economic Review*, 107 (5), 322–26.

Derrien, François, Jean-Stéphane Mésonnier, and Guillaume Vuillemey (2020) "Set-up Costs and the Financing of Young Firms," *Banque de France Working Paper*.

Dinlersoz, Emin, Sebnem Kalemli-Özcan, Henry Hyatt, and Veronika Penciakova (2019)

"Leverage over the Firm Life Cycle, Firm Growth, and Aggregate Fluctuations."

Doms, Mark and Timothy Dunne (1998) "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics*, 1 (2), 409–429.

Doms, Mark, Timothy Dunne, and Kenneth R. Troske (1997) "Workers, Wages, and Technology," *Quarterly Journal of Economics*, 112 (1), 253–290.

Dunne, Timothy, Mark J Roberts, and Larry Samuelson (1988) "Patterns of Firm Entry and Exit in US Manufacturing Industries," *RAND Journal of Economics*, 495–515.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu (2021) "How Costly are Markups?" *NBER Working Paper No. 24800*, National Bureau of Economic Research.

Everitt, Brian S., Sabine Landau, Morven Leese, and Daniel Stahl (2011) *Cluster Analysis*: Wiley, 5th edition.

Farinha, Luísa, Sónia Félix, and João AC Santos, "Bank Funding and the Survival of Start-ups," *Working Paper No. 201919*, Banco de Portugal, Lisbon.

Foster, Lucia, John Haltiwanger, and C.J. Krizan (2001) "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," in Edward Dean, Michael Harper and Charles Hulten ed. *New Developments in Productivity Analysis*, 303–372: University of Chicago Press.

Geroski, Paul A., José Mata, and Pedro Portugal (2010) "Founding Conditions and the Survival of New Firms," *Strategic Management Journal*, 31 (5), 510–529.

González-Uribe, Juanita and Michael Leatherbee (2018) "The Effects of Business Accelerators on Venture Performance: Evidence from Start-up Chile," *Review of Financial Studies*, 31 (4), 1566–1603.

González-Uribe, Juanita and Santiago Reyes (2021) "Identifying and Boosting "Gazelles": Evidence from Business Accelerators," *Journal of Financial Economics*, 139 (1), 260–287.

Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez (2017) "Capital Allocation and Productivity in South Europe," *Quarterly Journal* *of Economics*, 132 (4), 1915–1967.

Gutiérrez, Germán, Callum Jones, and Thomas Philippon (2019) "Entry Costs and the Macroeconomy," working paper.

Guzman, Jorge and Aishen Li (2019) "Measuring Founding Strategy," working paper.

Guzman, Jorge and Scott Stern (2015) "Where is Silicon Valley?," *Science*, 347 (6222), 606–609.

Haltiwanger, John, Ron S Jarmin, and Javier Miranda (2013) "Who Creates Jobs? Small versus Large versus Young," *Review of Economics and Statistics*, 95 (2), 347–361.

Hopenhayn, Hugo A. (1992) "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60 (5), 1127–1150.

Hopenhayn, Hugo and Richard Rogerson (1993) "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, 101 (5), 915–938.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein (2016) "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, 131 (3), 1291–1364.

Hsieh, Chang-Tai and Peter J. Klenow (2009) "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124 (4), 1403–1448.

Hsieh, Chiang-Tai and Peter J. Klenow (2014) "The Life Cycle of Plants in India and Mexico," *Quarterly Journal of Economics*, 129, 1035–1084.

Hurst, Erik and Benjamin W. Pugsley (2011) "What do Small Businesses Do?," *Brookings Papers on Economic Activity*, 43 (2), 73–142.

Iacovone, Leonardo, William Maloney, and David McKenzie (2022) "Improving Management with Individual and Group-based Consulting: Results from a Randomized Experiment in Colombia," *The Review of Economic Studies*, 89 (1), 346–371.

Johnson, Justin P. and David P. Myatt (2006) "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96 (3), 756–784.

Kalra, Aarti (2019) "Understanding the Corporate Tax Structure for Domestic Startup Companies," available at https://www.startupindia.gov.in/ 1764, Government of India.

Kerr, William R. and Ramana Nanda (2010) "Banking Deregulations, Financing Constraints, and Firm Entry Size," *Journal of the European Economic Association*, 8 (2-3), 582–593.

Kovenock, Dan and Gordon M. Phillips (1997) "Capital Structure and Product Market Behavior: An Examination of Plant Exit and Investment Decisions," *Review of Financial Studies*, 10 (3), 767–803.

Lafontaine, Francine and Kathryn Shaw (2016) "Serial Entrepreneurship: Learning by Doing?," *Journal of Labor Economics*, 34 (S2), S217–S254.

Lang, Larry, Eli Ofek, and René M. Stulz (1996) "Leverage, Investment, and Firm Growth," *Journal of Financial Economics*, 40 (1), 3–29.

Lemmon, Michael L., Michael R. Roberts, and Jaime F. Zender (2008) "Back to the Beginning: Persistence and the Cross-section of Corporate Capital Structure," *Journal of Finance*, 63 (4), 1575–1608.

Lerner, Josh (2009) *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed–and What to Do about It*: Princeton University Press.

Levine, Ross and Yona Rubinstein (2017) "Smart and Illicit: Who Becomes an Entrepreneur and Do They Earn More?," *Quarterly Journal of Economics*, 132 (2), 963–1018.

Liu, Yongzheng and Jie Mao (2019) "How do Tax Incentives affect Investment and Productivity? Firm-level Evidence from China," *American Economic Journal: Economic Policy*, 11 (3), 261–91.

Lopez-Garcia, Paloma and Filippo di Mauro (2015) "Assessing European Competitiveness: The New CompNet Microbased Database," Working Paper 1764, European Central Bank.

Midrigan, Virgiliu and Daniel Yi Xu (2014) "Finance and Misallocation: Evidence from Plant-Level Data," *American Economic Review*, 104 (2), 422–58.

Moll, Benjamin (2014) "Productivity Losses from Financial Frictions: Can Self-financing undo Capital Misallocation?," *American Economic Review*, 104 (10), 3186–3221.

Myers, Stewart C. and Nicholas S. Majluf (1984) "Corporate Financing and Investment Decisions when Firms have Information that Investors do not Have," *Journal of Financial Economics*, 13 (2), 187–221.

Oberfield, Ezra and Devesh Raval (2021) "Micro Data and Macro Technology," *Econometrica*, 89 (2), 703–732.

OECD (2015) "The Future of Productivity,"Technical report, Organisation for Economic Cooperation and Development, Paris.

——— (2020) "Financing SMEs and Entrepreneurs 2020: An OECD Scoreboard,"Technical report, Organisation for Economic Cooperation and Development, Paris. Available at https://www.oecd-ilibrary.org/.

——— (2021a) "Measuring Tax Support for R&D and Innovation,"Technical report, Organisation for Economic Cooperation and Development, Paris. Available at https://www.oecd.org/sti/rd-tax-stats.htm.

——— (2021b) "Targeted Statutory Corporate Income Tax Rate,"Technical report, Organisation for Economic Cooperation and Development, Paris. Available at https://stats.oecd.org/Index.aspx?DataSetCode=TABLE$_I I$2.

Oi, Walter Y. and Todd L. Idson (1999) "Firm Size and Wages," *Handbook of Labor Economics*, 3, 2165–2214.

Rajan, Raghuram G. and Luigi Zingales (1995) "What Do We Know About Capital Structure? Some Evidence from International Data," *Journal of Finance*, 50 (5), 1421–1460.

Restuccia, Diego and Richard Rogerson (2013) "Misallocation and Productivity," *Review of Economic Dynamics*, 16 (1), 1–10.

Robb, Alicia M. and David T. Robinson (2014) "The Capital Structure Decisions of New Firms," *Review of Financial Studies*, 27 (1), 153–179.

Rosenberg, Joseph and Donald Marron (2015) "Tax Policy and Investment by Startups and

Innovative Firms," *Tax Policy Center–Urban Institute and Brookings Institution*.

Schoar, Antoinette (2010) "The Divide between Subsistence and Transformational Entrepreneurship," *Innovation Policy and the Economy*, 10 (1), 57–81.

Sedláček, Petr and Vincent Sterk (2017) "The Growth Potential of Startups Over the Business Cycle," *American Economic Review*, 107 (10), 3182–3210.

——— (2019) "Reviving American Entrepreneurship? Tax Reform and Business Dynamism," *Journal Of Monetary Economics*, 105, 94–108.

Sterk, Vincent, Petr Sedláček, and Benjamin Pugsley (2021) "The Nature of Firm Growth," *American Economic Review*, 111 (2), 547–579.

Syverson, Chad (2017) "Challenges to Mismeasurement Explanations for the US Productivity Slowdown," *Journal of Economic Perspectives*, 31 (2), 165–86.

Troske, Kenneth R. (1999) "Evidence on the Employer Size-Wage Premium from Worker-Establishment Matched Data," *Review of Economics and Statistics*, 81 (1), 15–26.

Zingales, Luigi (1998) "Survival of the Fittest or the Fattest? Exit and Financing in the Trucking Industry," *Journal of Finance*, 53 (3), 905–938.

Zwick, Eric and James Mahon (2017) "Tax Policy and Heterogeneous Investment Behavior," *American Economic Review*, 107 (1), 217–48.

Tables and Figures

Figure 1. Meta-clustering of startup types

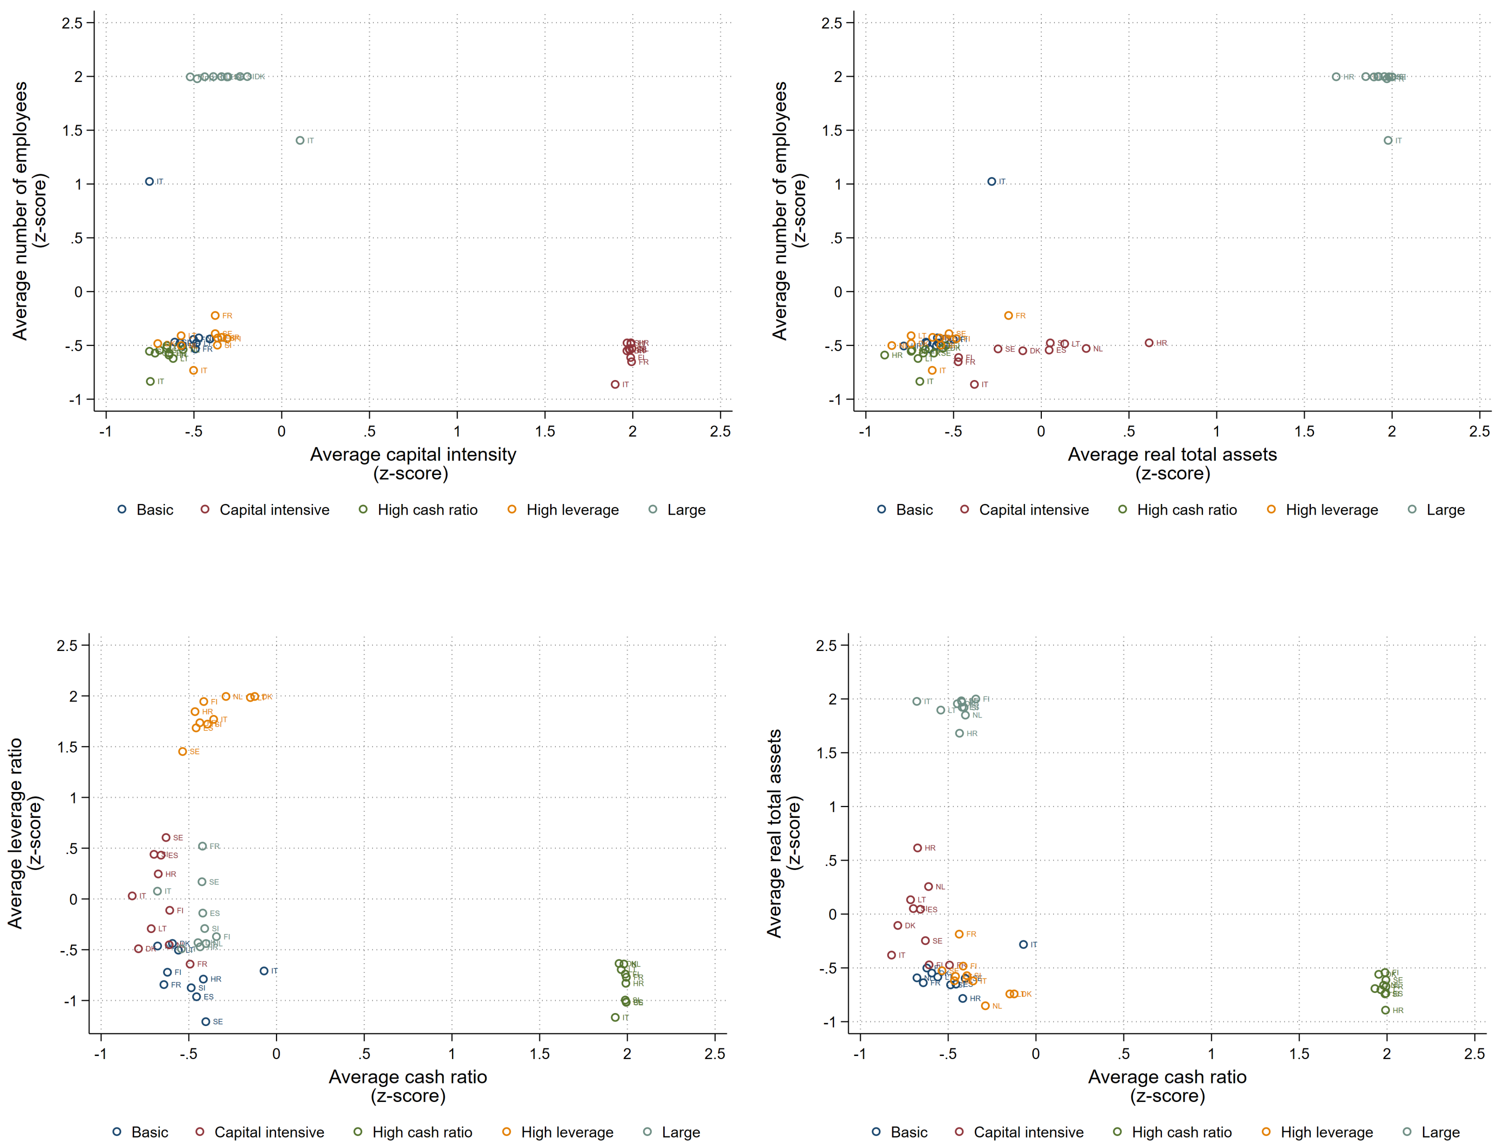

Notes: The four panels in this figure summarize the meta-clustering procedure. Different meta-clusters are denoted by different colours. The meta-clustering groups comparable clusters from different countries by taking the cluster centers derived from each country's first-stage clustering procedure as the observations. In the meta-clustering procedure, the units of observation are z-scores of the first-stage cluster centers, averaged across years and industries.

Figure 2. Distribution of startup type by country

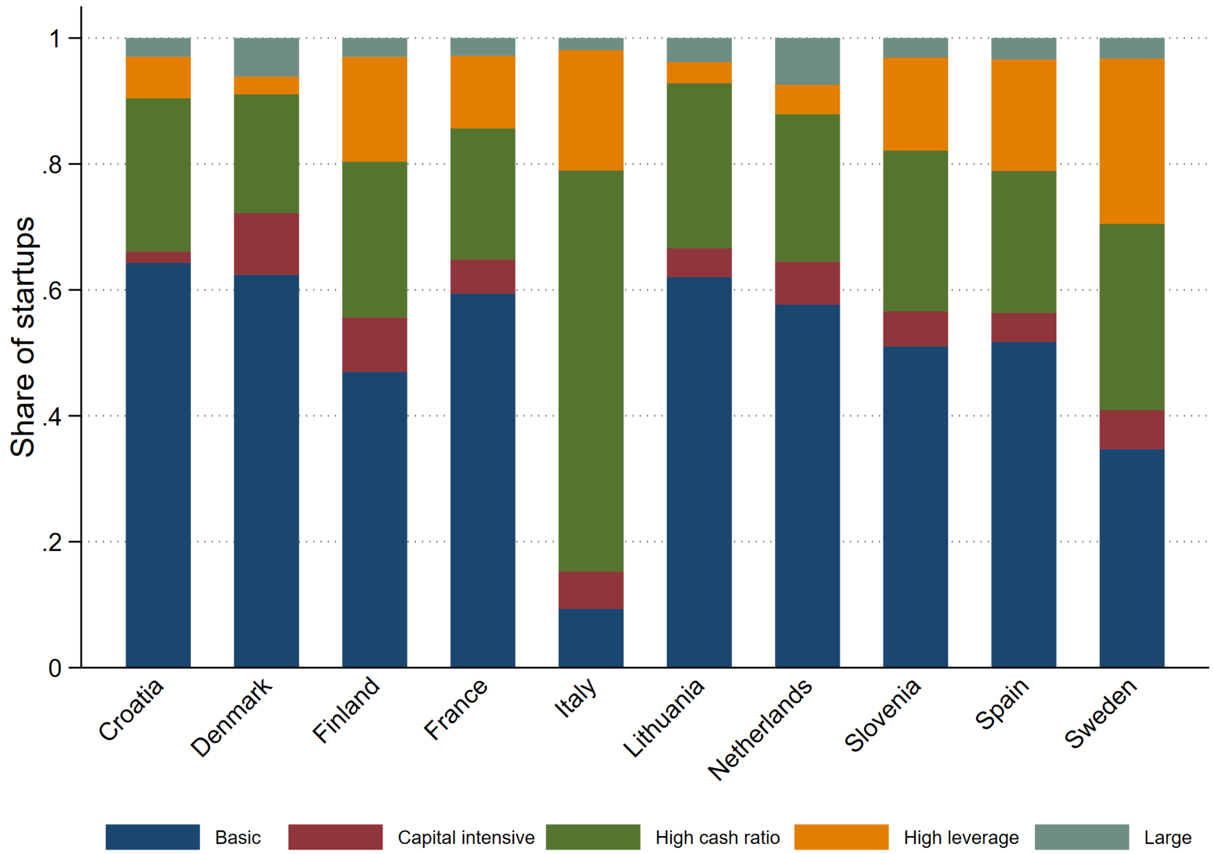

Notes: This figure illustrates the distribution of the startup population for individual countries across the five startup types. The startup population comprises all cohorts available for each country.

Figure 3. Distribution of startup type by industry

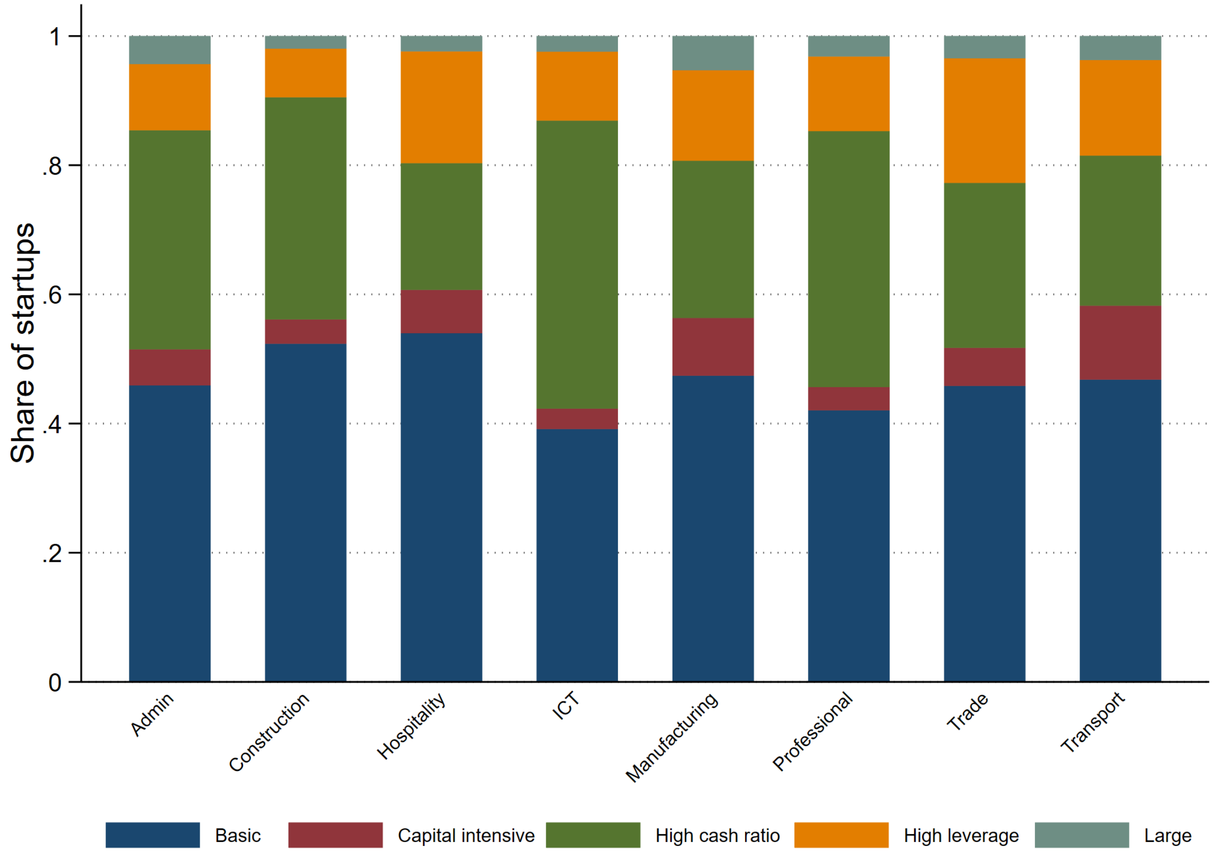

Notes: This figure illustrates the distribution of the startup population for one-digit NACE Rev.2 industries across the five startup types. The startup population comprises all cohorts available for each country.

Figure 4. Distribution of startup type by cohort

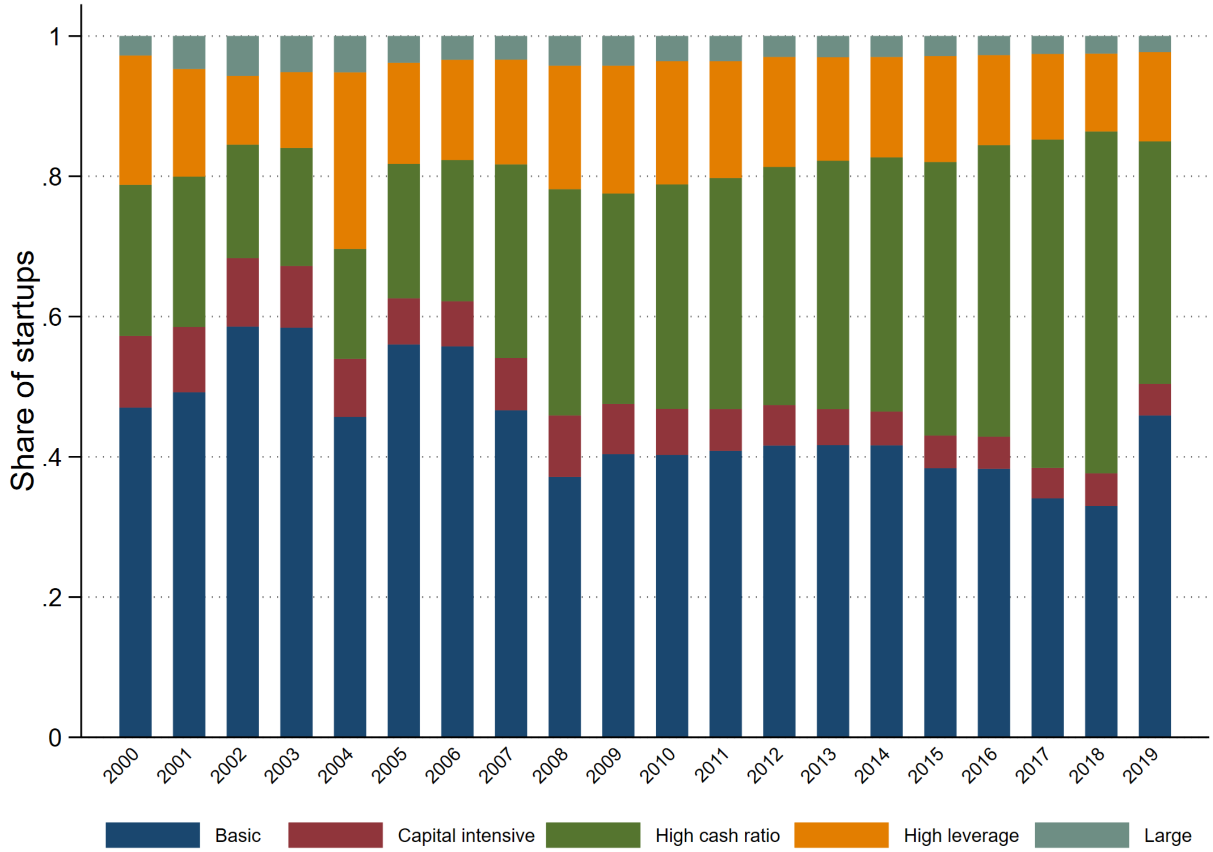

Notes: This figure illustrates the distribution of the startup population for each cohort across the five startup types.

Figure 5. The life cycle of startup types

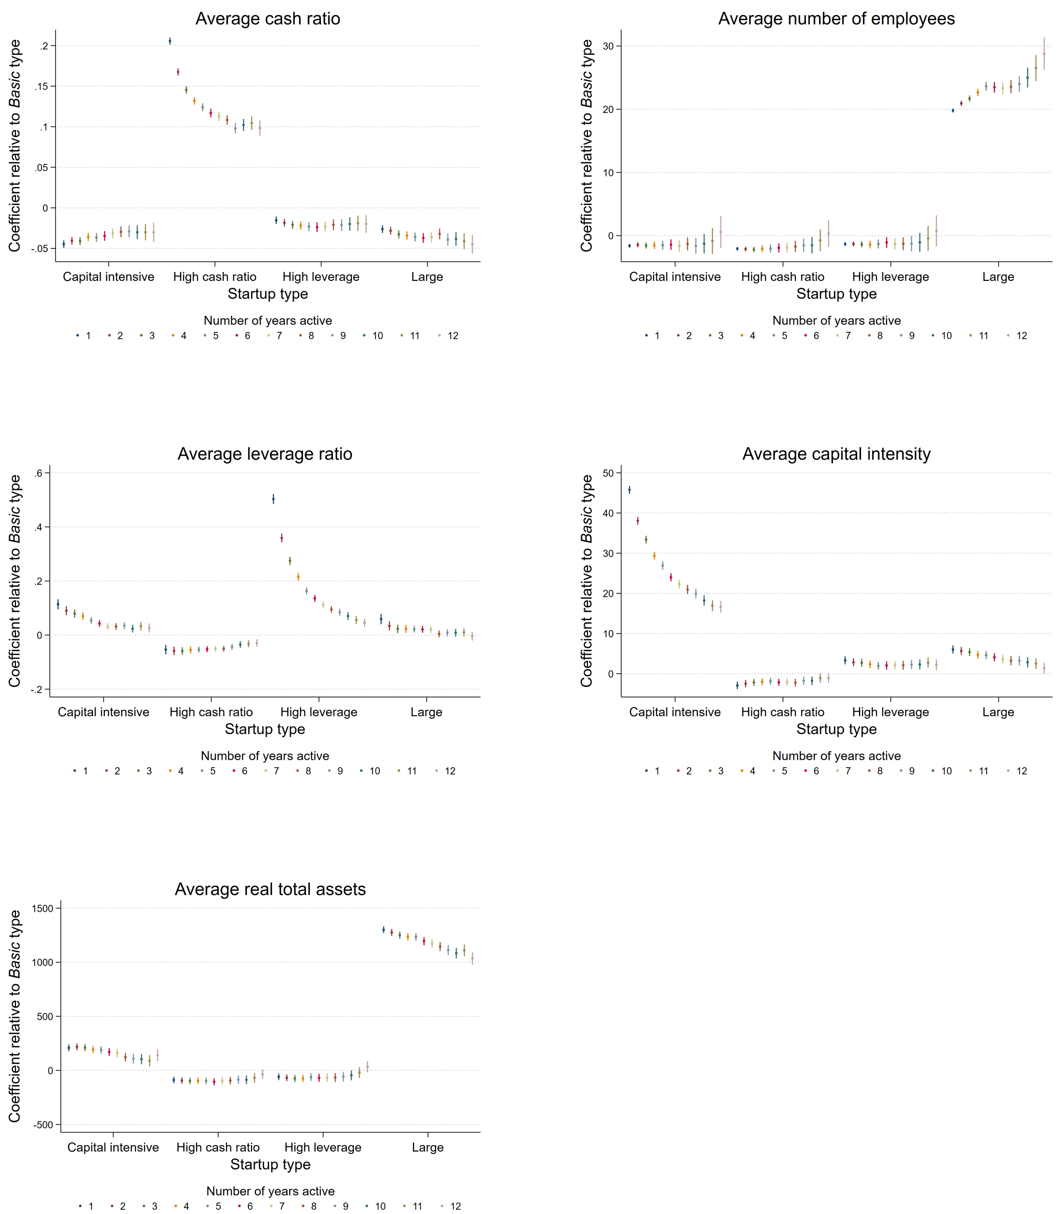

Notes: The panels in this figure summarize how the startup types develop during the first 12 years of their life in terms of the five clustering variables. Each panel corresponds to one clustering variable and plots the coefficients from 12 separate regressions where the dependent variable is this variable. Each regression is then run for an age group (age is 1, 2, ...12 years). For example, the first panel summarizes regressions in which the *Average cash to total assets* ratio is regressed on dummy variables for the startup types (the *Basic* type is omitted) as well as country, cohort, and industry fixed effects. The sample is the full panel data set at the one-digit industry level.

Figure 6. (Inverse) entry elasticities by startup type

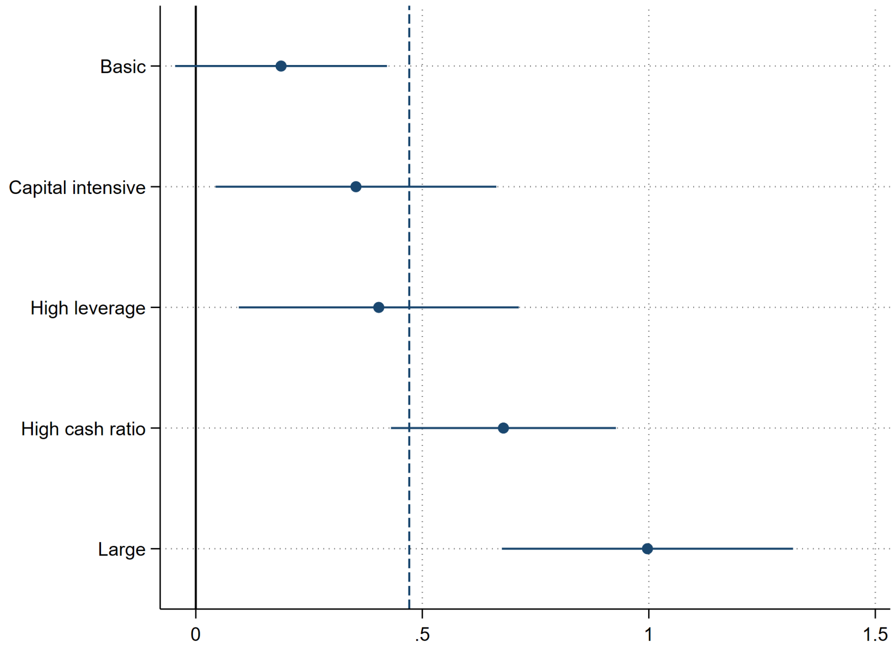

Notes: This figure shows coefficients from OLS regressions where the dependent variable is the log net present value of firm profits. The estimated coefficients represent the inverse entry elasticity, β_1, for each startup type. *Country* $\times$ *Startup type* $\times$ *Industry* and *Country* $\times$ *Industry* $\times$ *Cohort* fixed effects are included throughout. The regressions are at the one-digit NACE Rev.2 industry level. The data set only includes cohorts observed for at least seven years (those founded before 2011) and drops observations beyond age seven. For age eight and onward, profits and the year-on-year exit rate are assumed fixed. Confidence intervals are at the 95 percent level. The vertical dashed line denotes an estimate from a regression that does not condition on startup type.

Figure 7. Policy experiment: Tax differentiation and aggregate labor productivity and employment of young firms

Panel A

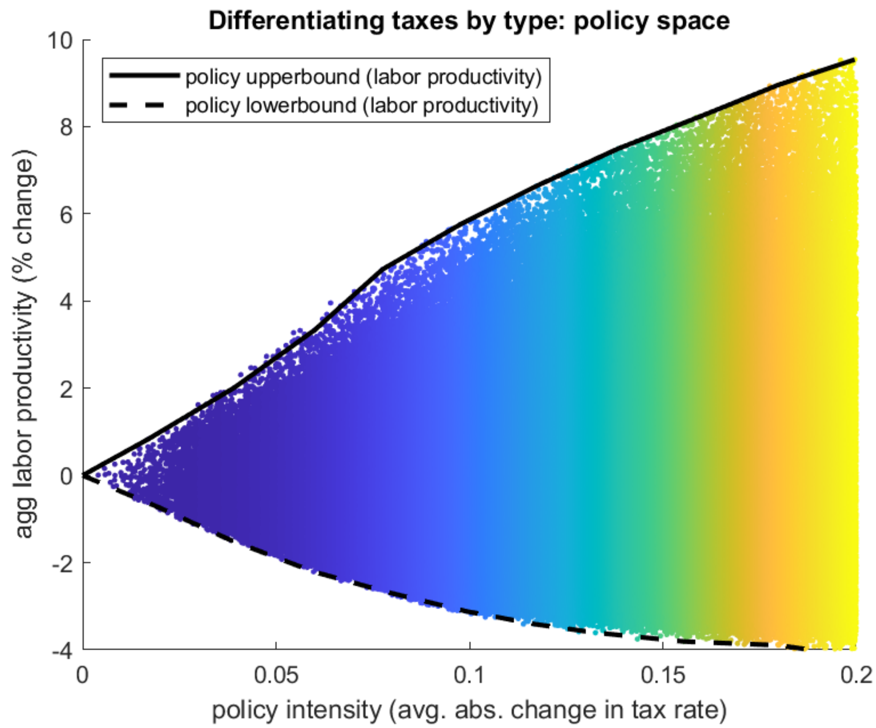

Panel B

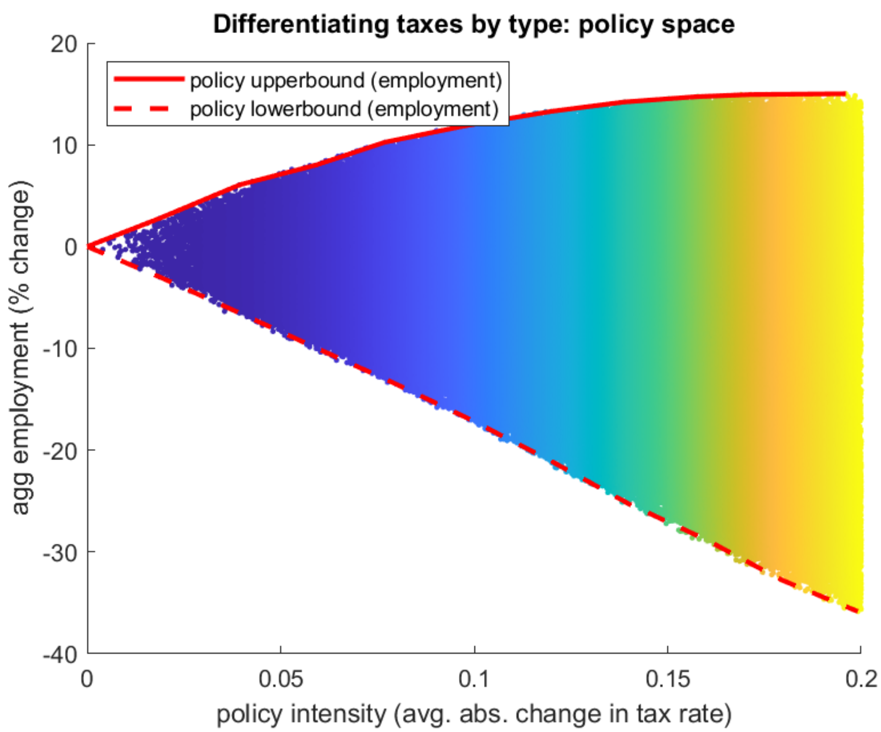

Notes: This figure summarizes the policy experiment. Panel A shows the policy space for aggregate labor productivity. The horizontal axis measures the intensity of the potential policy change as the absolute change in the tax rate, averaged across firms. Warmer colors indicate stronger corporate tax rate differentiation. The solid line plots the "policy upper bound": the largest possible aggregate labor productivity increase given a certain policy intensity. Similarly, Panel B plots the policy space for aggregate employment.

Figure 8. Policy experiment: Upper bounds of macro impacts due to tax differentiation

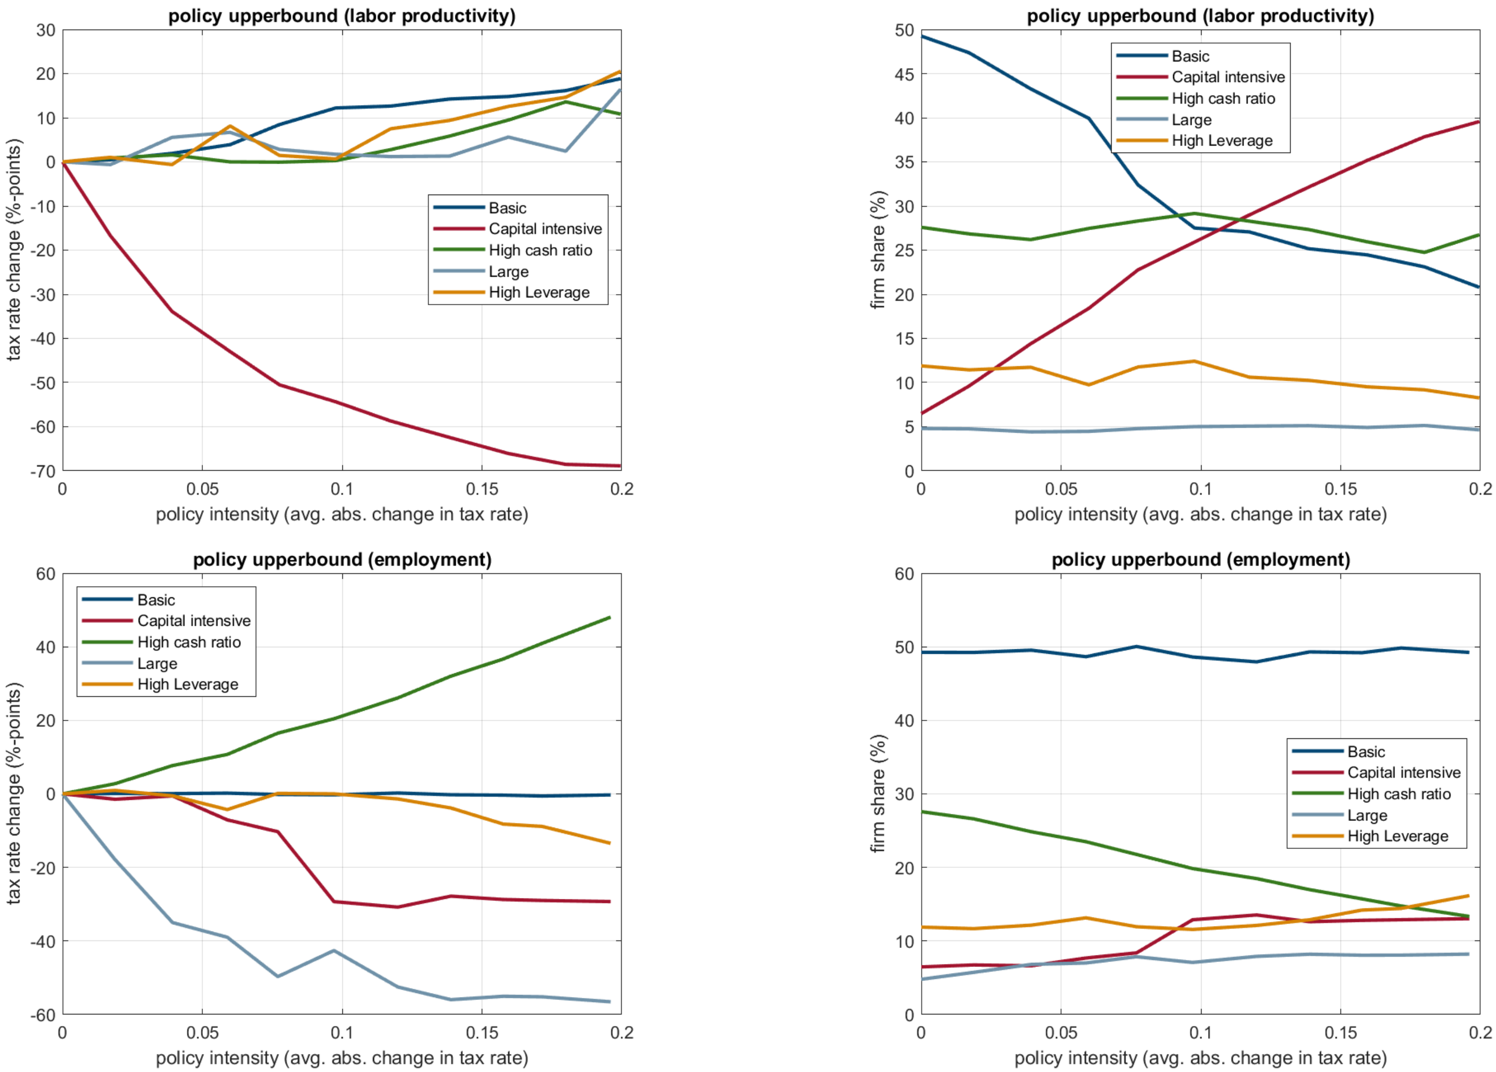

Notes: The two panels on the left (right) plot the corporate tax rate policies (startup shares) associated with the two policy upper bounds. The upper left (right) panel shows the tax rates (startup shares) associated with the labor productivity frontier. The lower left (right) panel shows the tax rates (startup shares) associated with the employment frontier.

Figure 9. Policy experiment: Macroeconomic trade-offs

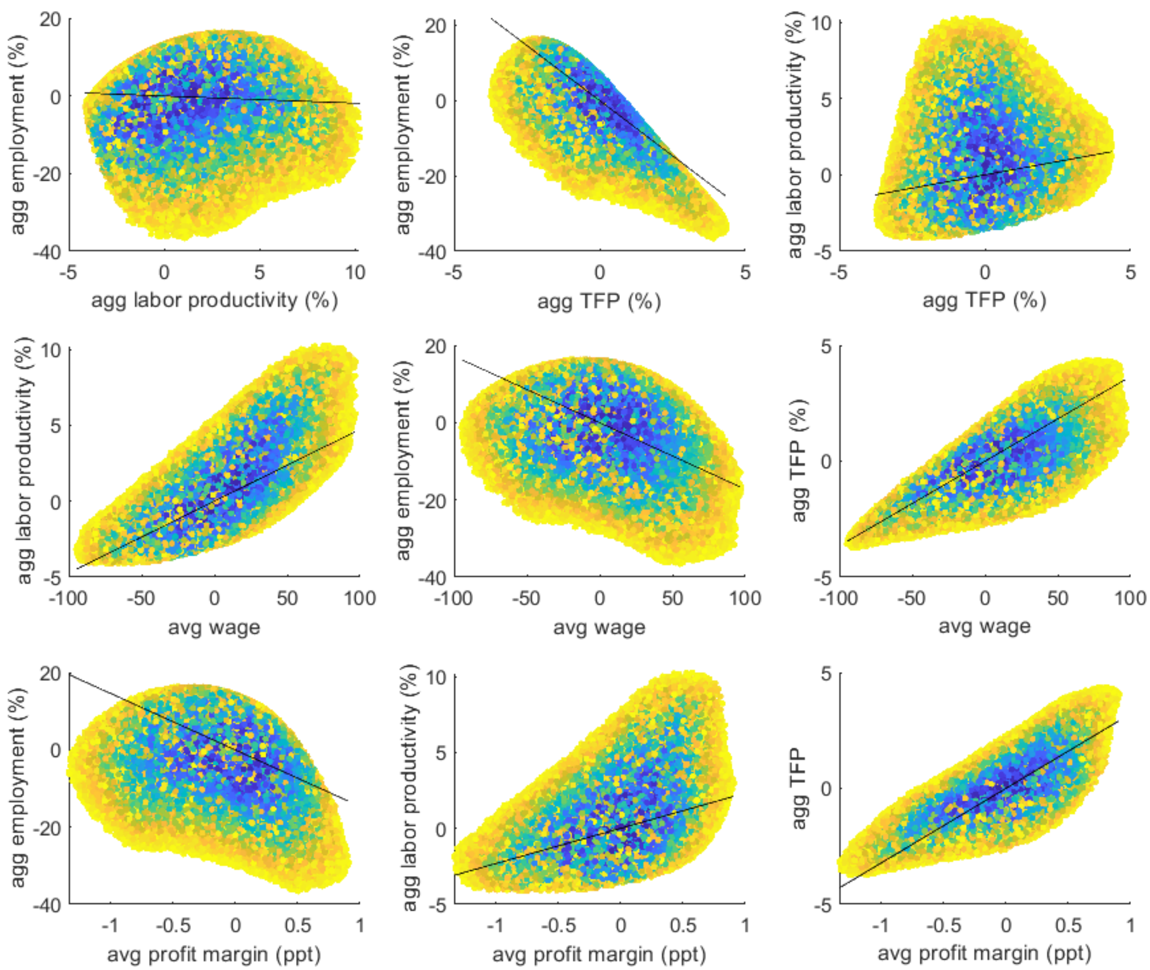

Notes: These scatter plots each depict pairs of changes in aggregate or average macroeconomic outcomes resulting from each potential policy (up to a policy intensity of 0.2). Wages per employee are in thousands of euros. Linearly fitted regression lines are shown in black.

Table 1. Characteristics of startup type at time of entry

	(1)	(2)	(3)	(4)	(5)
	Number of employees	Capital intensity	Real total assets	Cash ratio	Leverage ratio
Basic	4	8.56	166.59	0.12	0.23
Capital-intensive	2	93.18	405.12	0.09	0.41
Cash-intensive	2	4.67	92.46	0.54	0.18
High-leverage	3	12.90	122.97	0.14	1.18
Large	20	16.05	1488.30	0.13	0.34

Notes: This table presents the cross-country means of the cluster variables for each of the five startup types in the year of establishment. Means are unweighted and based on the full panel.

Table 2. Startup types and firm outcomes

Panel A: All firms

	(1)	(2)	(3)	(4)	(5)
	Aggregate labor productivity	Aggregate TFP	Average exit probability	Average wage per employee	Average profit margin
Capital intensive	0.315***	0.048***	-0.061***	2.449***	0.012***
	(0.005)	(0.004)	(0.003)	(0.094)	(0.001)
Cash rich	0.032***	0.050***	-0.007***	1.092***	0.022***
	(0.003)	(0.002)	(0.002)	(0.058)	(0.001)
High leverage	-0.045***	-0.039***	-0.000	-1.516***	-0.031***
	(0.004)	(0.003)	(0.002)	(0.064)	(0.001)
Large	0.165***	0.046***	-0.124***	3.406***	-0.018***
	(0.004)	(0.003)	(0.003)	(0.094)	(0.001)
Constant	3.353***	2.213***	0.710***	27.732***	0.041***
	(0.002)	(0.002)	(0.001)	(0.033)	(0.000)
R-squared	0.905	0.978	0.635	0.909	0.605
N	27,562	19,848	29,642	28,744	28,490

Panel B: All firms, aged 5-8

	(1)	(2)	(3)	(4)	(5)
	Aggregate labor productivity	Aggregate TFP	Average exit probability	Average wage per employee	Average profit margin
Capital intensive	0.195***	0.011*	-0.063***	0.879***	0.013***
	(0.007)	(0.006)	(0.003)	(0.121)	(0.001)
Cash rich	0.019***	0.031***	-0.023***	1.145***	0.014***
	(0.006)	(0.004)	(0.003)	(0.095)	(0.001)
High leverage	-0.027***	-0.021***	0.002	-1.394***	-0.008***
	(0.006)	(0.005)	(0.003)	(0.105)	(0.001)
Large	0.151***	0.050***	-0.141***	2.596***	-0.013***
	(0.007)	(0.007)	(0.004)	(0.154)	(0.001)
Constant	3.464***	2.247***	0.812***	30.151***	0.048***
	(0.003)	(0.003)	(0.001)	(0.052)	(0.001)
R-squared	0.906	0.981	0.616	0.930	0.654
N	8,731	6,055	9,569	9,149	9,081
Country × cohort FE	✓	✓	✓	✓	✓
Industry × cohort FE	✓	✓	✓	✓	✓
Country × Industry FE	✓	✓	✓	✓	✓
Age × Country FE	✓	✓	✓	✓	✓
Age × Industry FE	✓	✓	✓	✓	✓
Age × Cohort FE	✓	✓	✓	✓	✓

Notes: This table shows OLS regressions where the dependent variable is indicated in the column heading. Observations are at the country x sector x startup type x cohort x age level. The dependent variables in the first two columns are aggregate outcomes (i.e. employment-weighted averages) by startup type. The dependent variables in the last three columns are simple averages by startup type. Industries are defined at the one-digit NACE Rev.2 level. Regressions are based on the full panel of firms younger than nine years in Panel A and firms between age five and eight in Panel B. Standard errors are in parentheses. *, **, *** indicate significance at the 10, 5 and 1 percent level, respectively. Wages are denominated in thousands of euros.

Appendices

Figure A1. Startup rates by cohort and country—CompNet versus Eurostat data

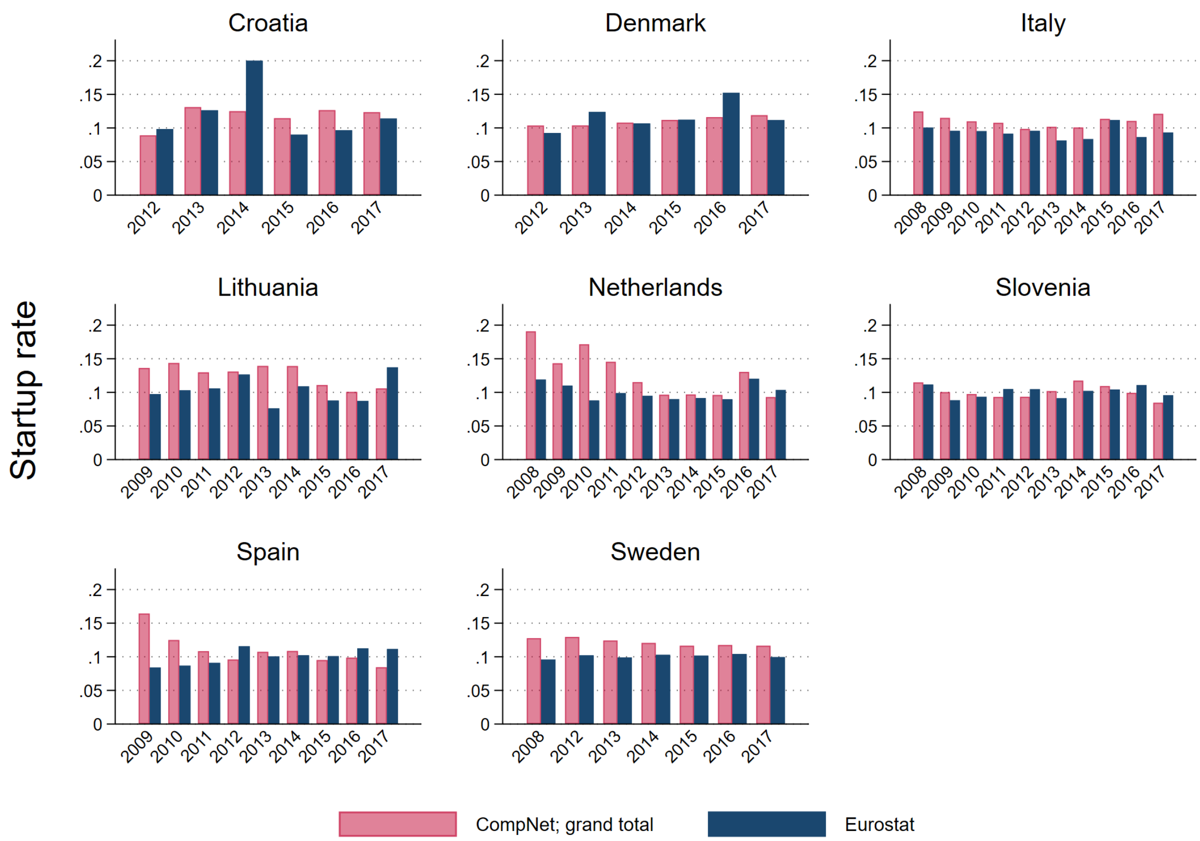

Notes: This figure compares the total number of startups in the CompNet database (pink bars) with establishment of firms in Eurostat (blue bars). Cohorts reported are subject to data availability. Finland is omitted in this chart because we cannot access the total number of startups in the CompNet database for that country. France does not appear as Eurostat does not report firm entry prior to 2008.

Figure A2. Employment by firm age—CompNet versus Eurostat data

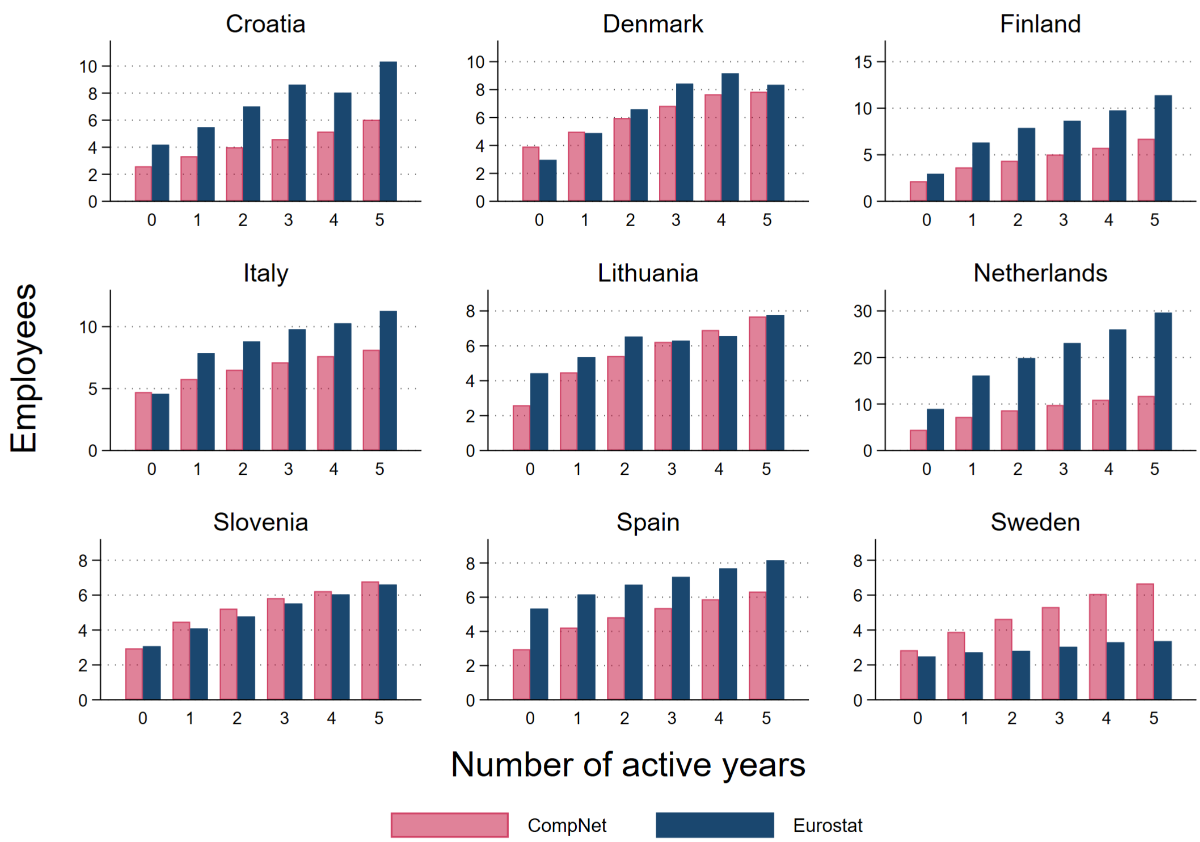

Notes: This figure compares growth in number of persons employed by startups as reported by CompNet (pink bars) and Eurostat (blue bars). CompNet and Eurostat data are matched based on the startup cohort and age group. For comparison purposes, we adjust the Eurostat data such that sole proprietorship firms are removed and we adjust for the average number of persons employed by sole proprietorship firms. The x-axis depicts *Startup age*, which is the number of years a startup has been active. *Employees* on the y-axis is averaged over cohorts. France does not appear as Eurostat does not report firm entry prior to 2008.

Figure A3. Cumulative exit rates by firm age—CompNet versus Eurostat data

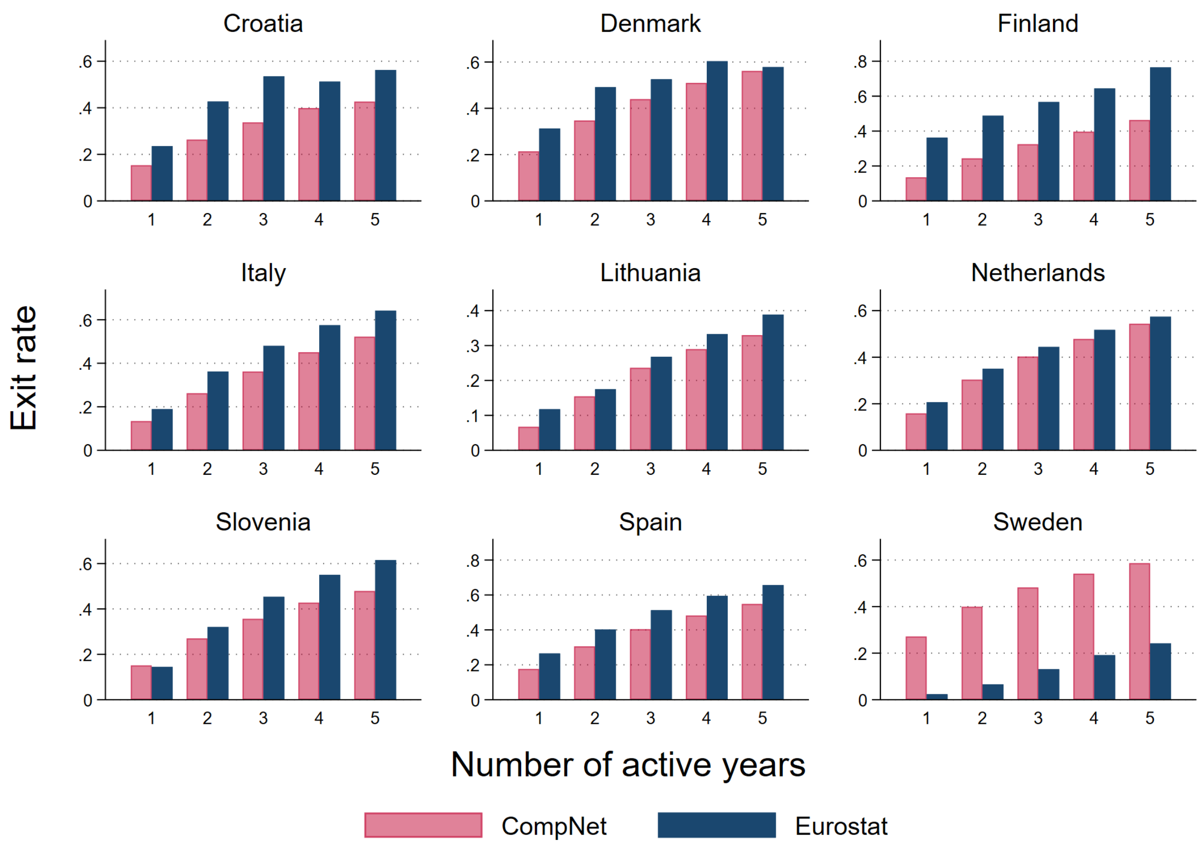

Notes: This figure compares the cumulative exit rate of startups in CompNet (pink bars) and Eurostat (blue bars). CompNet and Eurostat data are matched based on the startup cohort and age group. The x-axis depicts *Startup age*, which is the number of years a startup has been active. *Exit rate* is the average exit rate over all cohorts for each startup age group. France does not appear as Eurostat does not report firm entry prior to 2008.

Figure A4. Scree plots

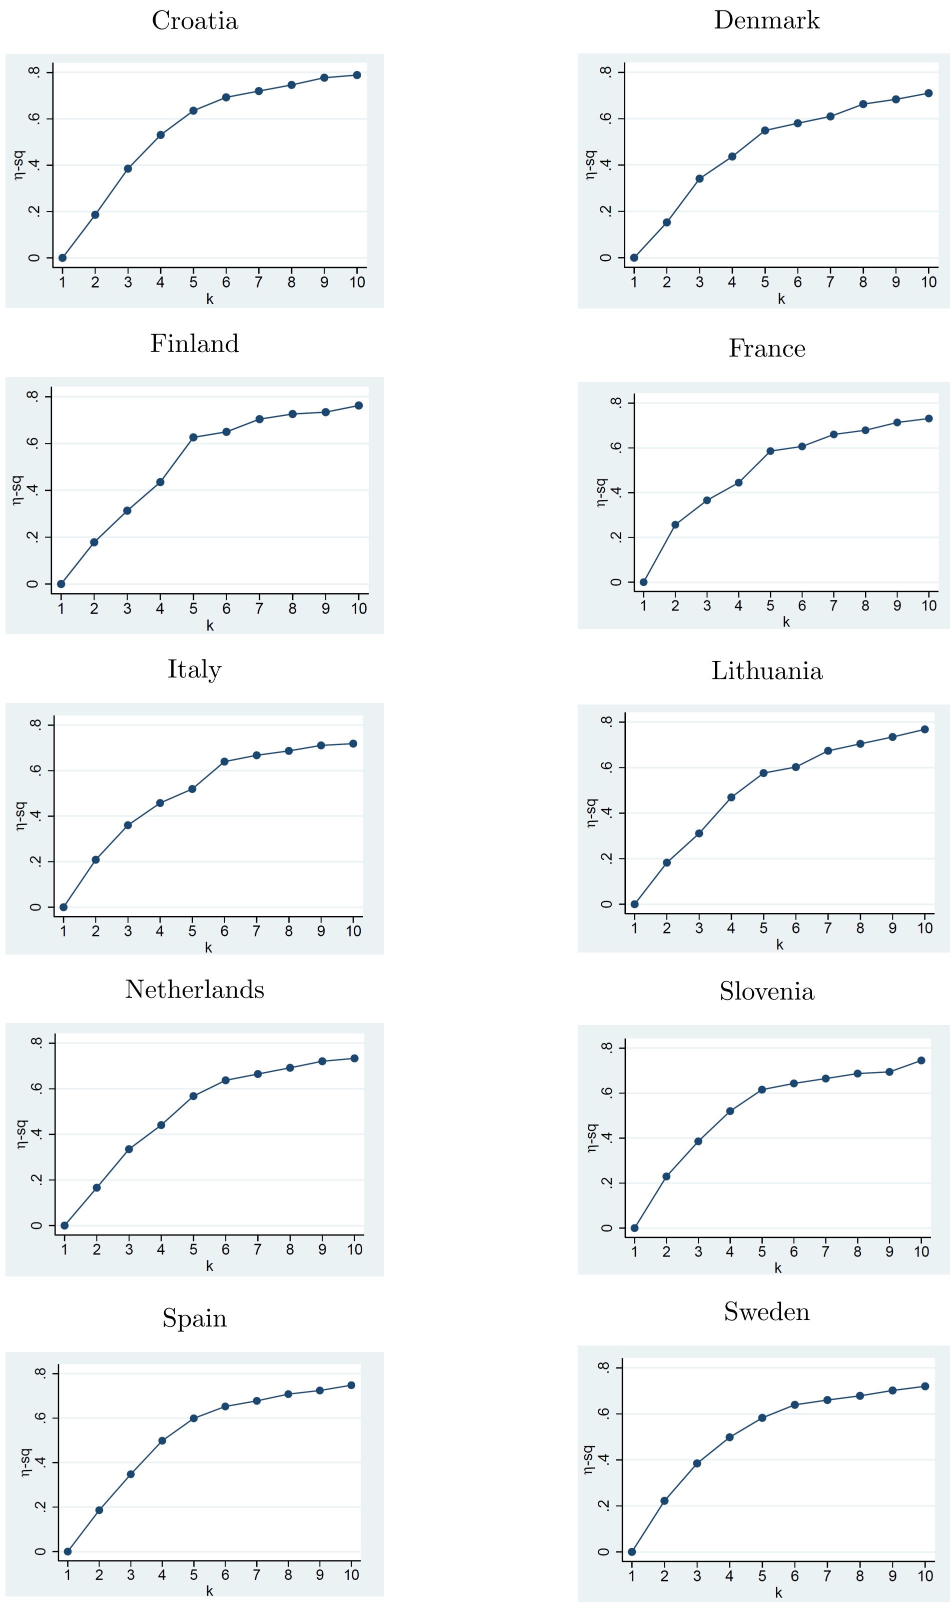

Notes: This figure shows scree plots resulting from the k-means cluster algorithm at the firm-level, for each country in the sample. On the x-axis, k indicates the number of clusters. The η^2 coefficient on the y-axis measures the proportional reduction of the within sum of squares for each cluster solution k compared with the total sum of squares.

Figure A5. Monte Carlo experiment of the meta-clustering

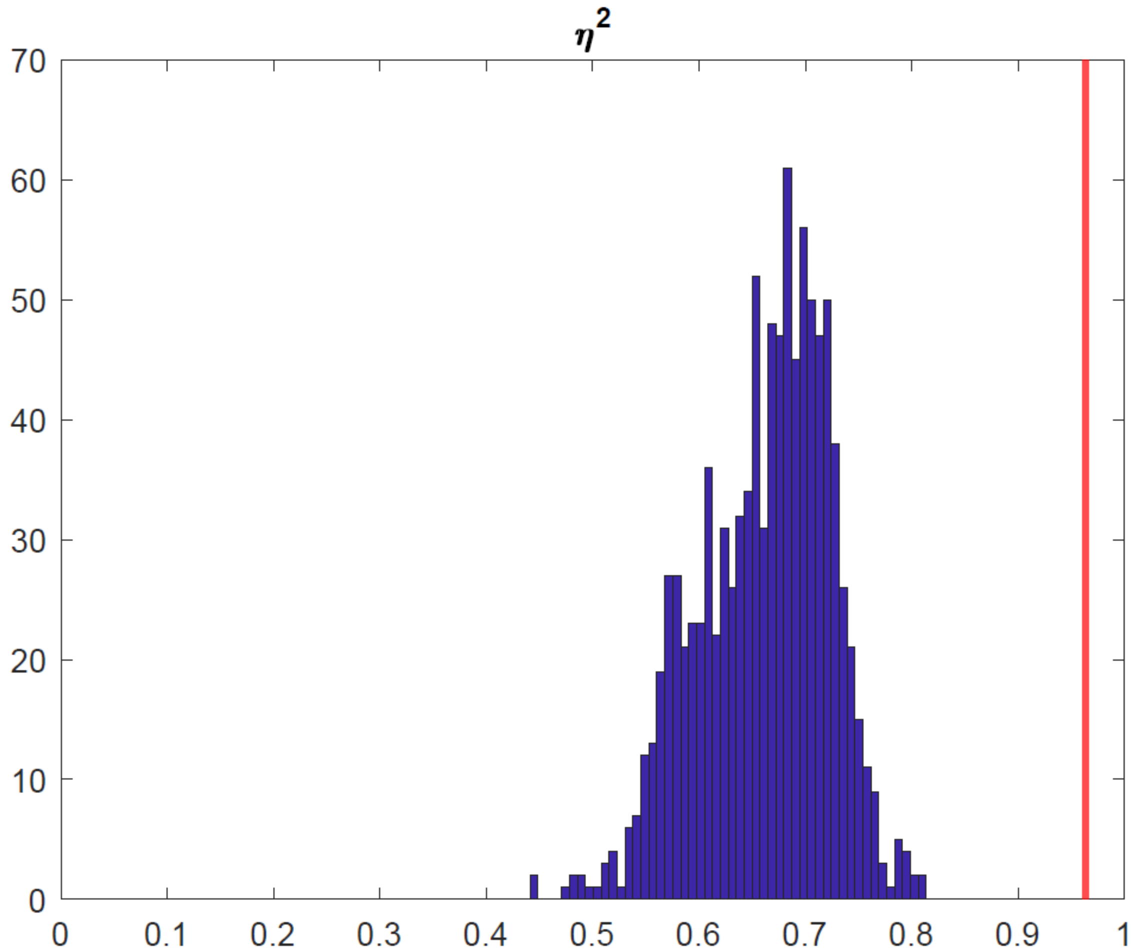

Notes: This histogram summarizes a Monte Carlo experiment consisting of a large number of random draws for the cluster variables, with means and standard deviations as observed in the data. These draws are i.i.d. so that no clusters exist in the experimental data. The experiment is repeated many times and each time η^2 is computed (blue bars). The vertical red line indicates the true η^2 statistic based on the actual data.

Figure A6. Distribution of startup type by industry and by country

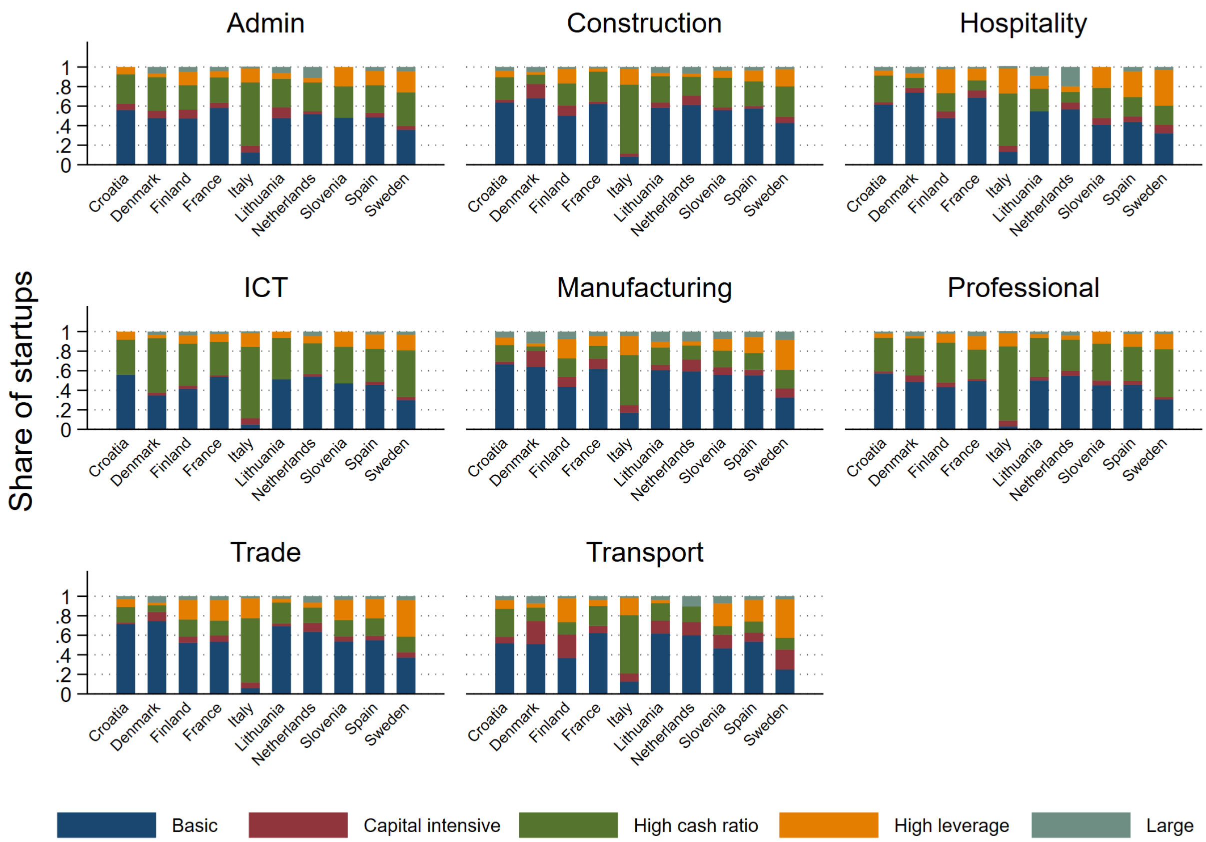

Notes: This figure presents the distribution of the startup population for individual one-digit NACE Rev.2 industries and countries across the five startup types. Shares are averaged over all startup cohorts.

Figure A7. Startup characteristics at the time of firm entry, by type and country

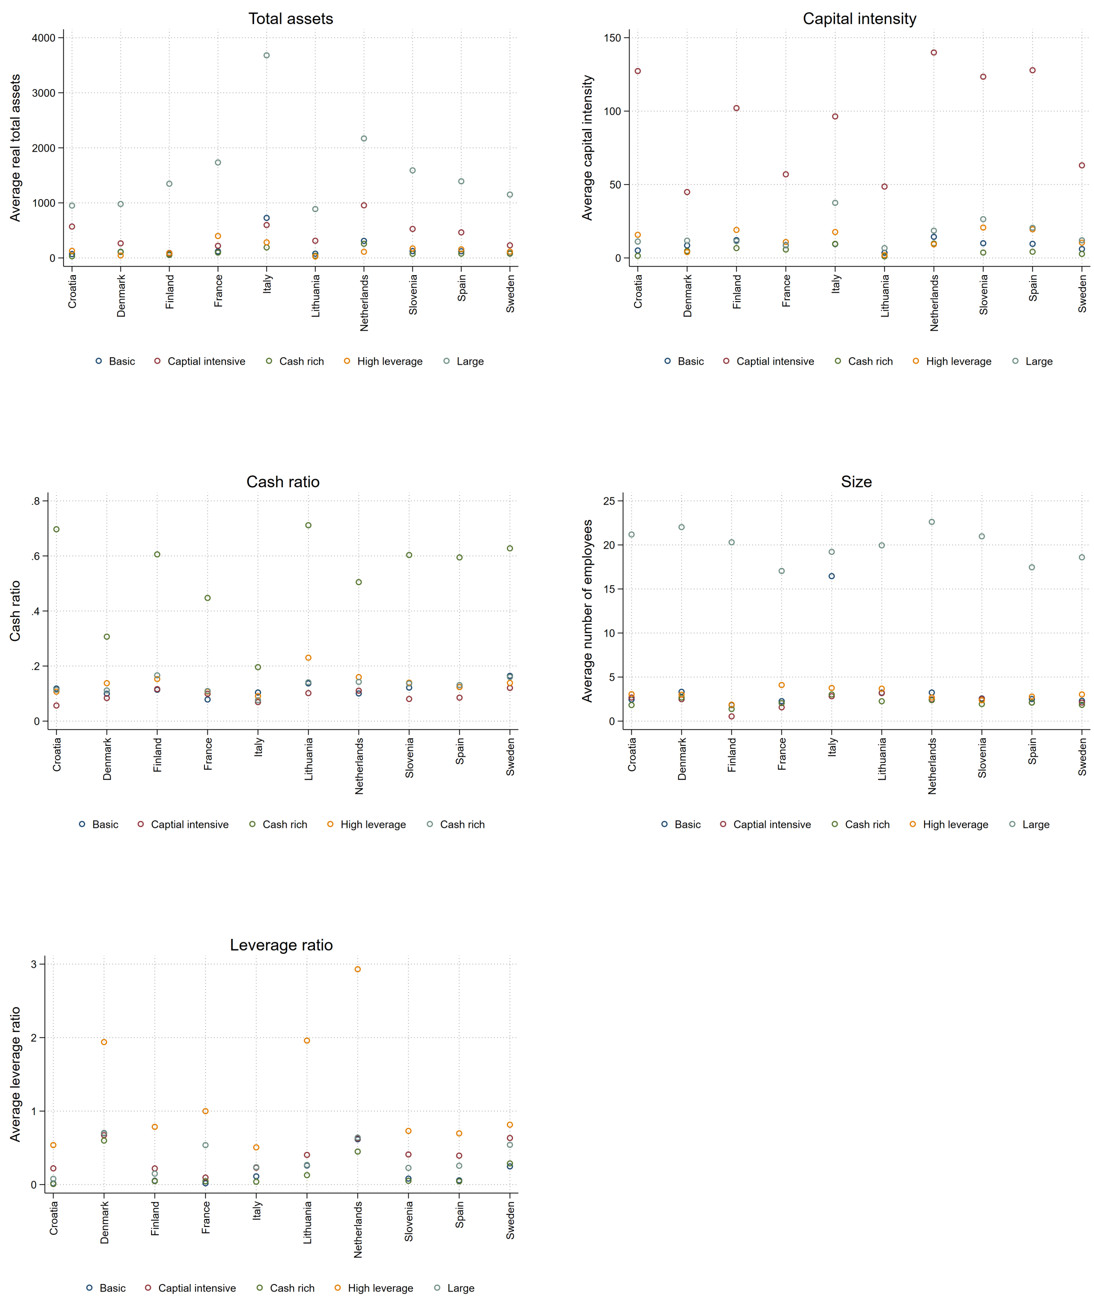

Notes: This figure presents the country-level mean of the five cluster variables for the five startup types in the year of firm entry.

Figure A8. Life-cycle profiles used in the policy experiment

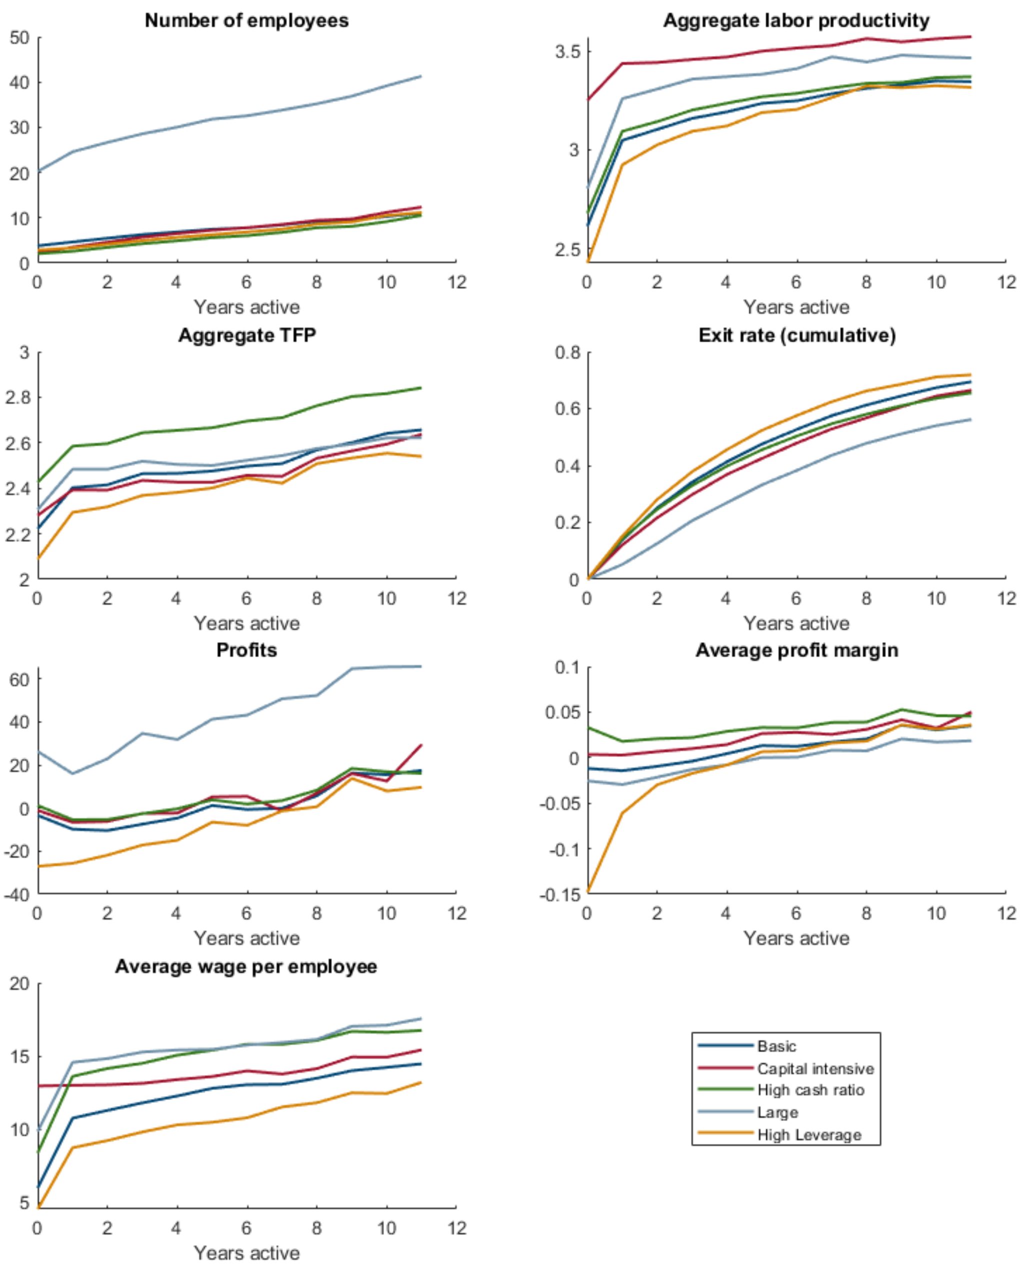

Notes: These charts summarize the development over time of the five startup types in terms of various outcomes. The results are based on panel regressions that include country, cohort, and industry fixed effects.

Figure A9. Selection at entry: Number of startups in a cohort and performance later in life

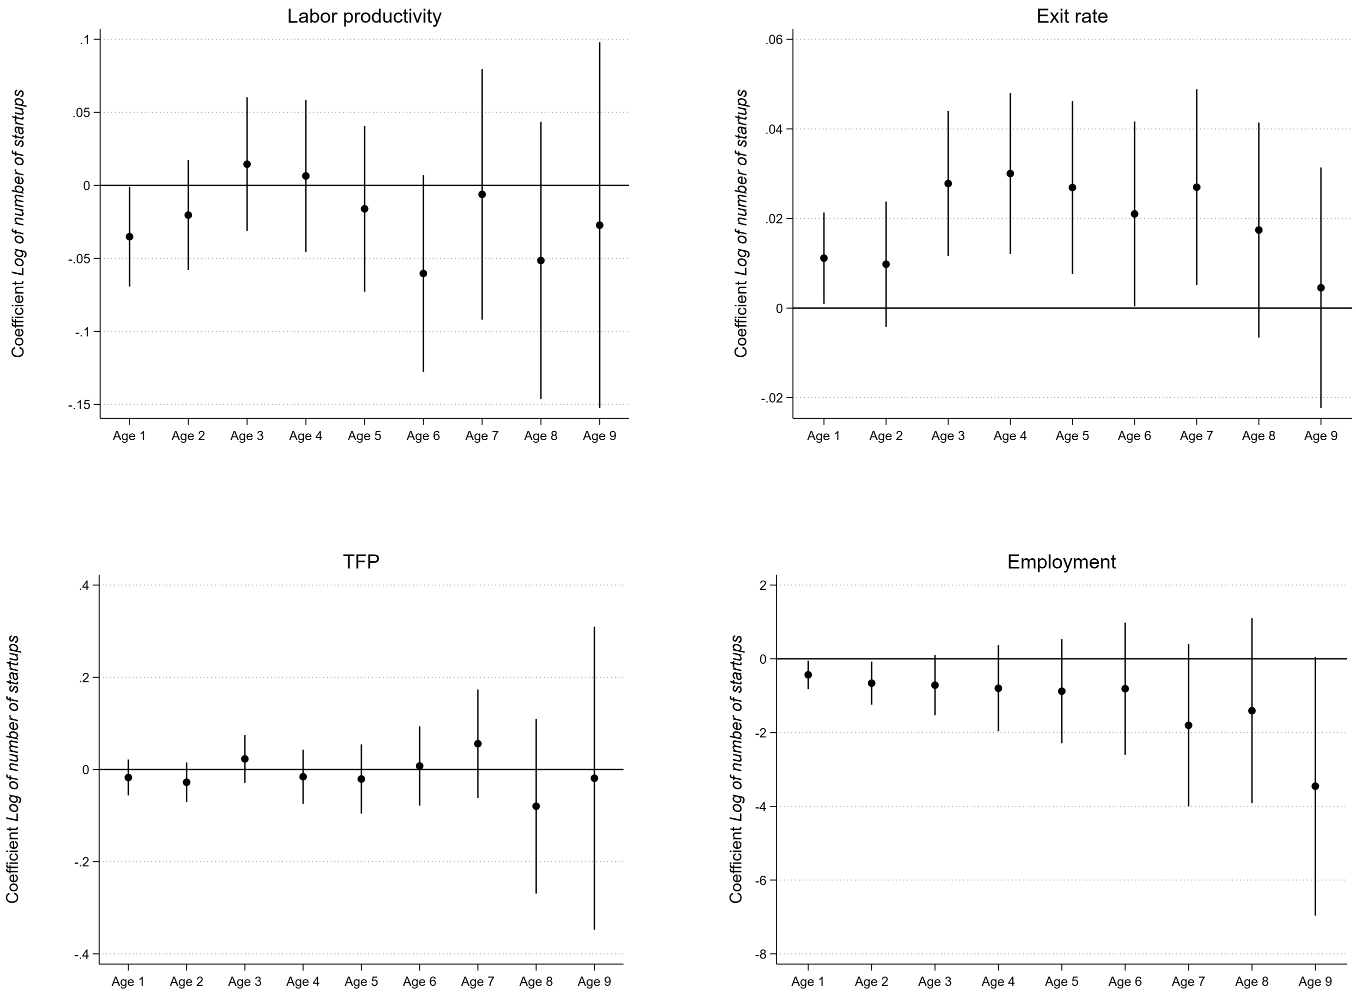

Notes: This figure summarizes OLS regressions where the dependent variable is indicated in the sub-figure titles. *Exit rate* is the cumulative exit rate defined as 1 minus the ratio of number of startups that survived until t divided by the number of startups in *t-0*. Each dot corresponds to a partitioned regression for all startups of a particular age, as indicated on the x-axis, with interactive fixed effects for industry $\times$ startup type $\times$ country; industry $\times$ country $\times$ cohort; startup type $\times$ country $\times$ cohort; and industry $\times$ startup type $\times$ cohort. The regressions are at the one-digit NACE Rev. 2 industry level. Whiskers indicate 95 percent confidence intervals.

Figure A10. Policy experiment with equilibrium adjustment: Tax differentiation and aggregate labor productivity of and employment within young firms

Panel A

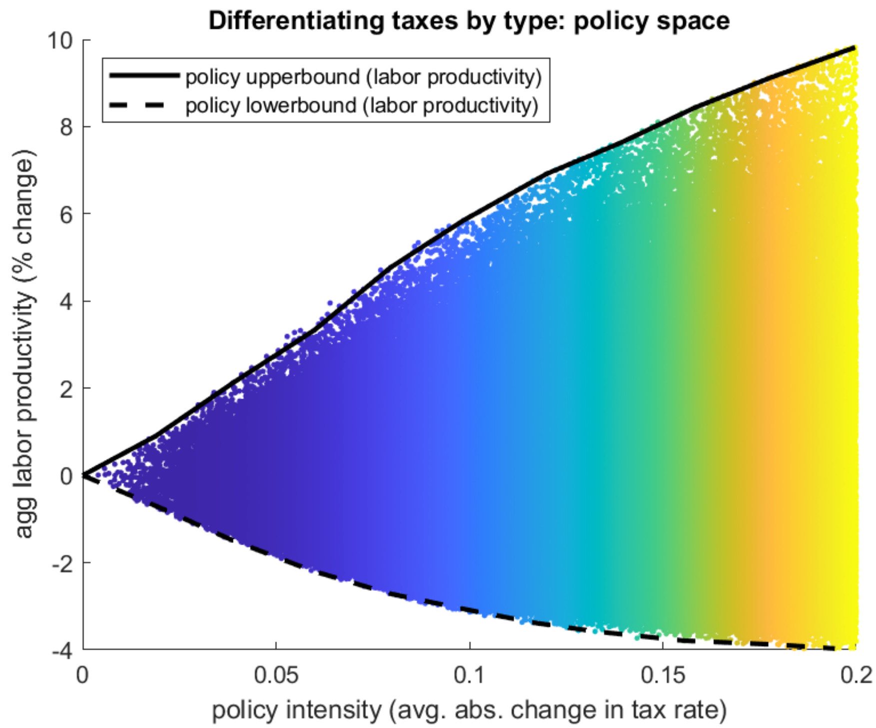

Panel B

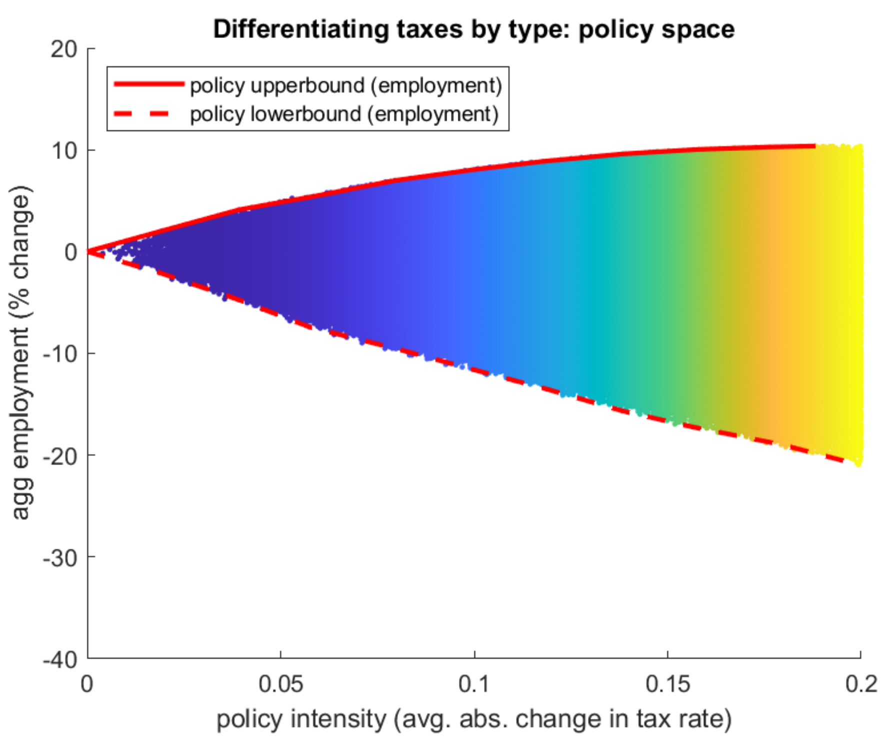

Notes: This figure summarizes the policy experiment allowing for equilibrium adjustment in the labor market. Panel A shows the policy space for aggregate labor productivity. The horizontal axis measures the intensity of the potential policy change as the absolute change in the tax rate, averaged across firms. Warmer colors indicate stronger corporate tax rate differentiation. The solid line plots the "policy upper bound": the largest possible increase in aggregate labor productivity given a certain policy intensity. Similarly, Panel B plots the policy space for aggregate employment.

Table A1. Sample composition

Country	Sample period	Number of unique startups
Croatia	2003-2019	64,760
Denmark	2002-2018	114,195
Finland	2000-2019	126,554
France	2005-2007	210,033
Italy	2007-2018	322,893
Lithuania	2001-2017	47,322
Netherlands	2008-2018	63,729
Slovenia	2006-2019	23,435
Spain	2009-2018	236,192
Sweden	2004-2019	136,376
Total		*1,345,489*

Notes: This table shows the sample composition of the full panel of startups.

Table A2. Variable definitions

Variable	Definition
Capital intensity	Average ratio of real capital stock to the number of employees
Cash ratio	Average cash to total assets ratio
Employment	Average number of employees
Exit probability	1 minus the ratio of number of firms that survived until t divided by the number of startups in t=0
Labor productivity	Logarithm of average labor productivity defined as real value-added divided by number of employees
Leverage ratio	Average ratio of debt to total assets
Number of employees	Average number of persons employed
Profit margin	Average ratio of operating profit (Earnings Before Interest and Tax (EBIT)) to revenue
Real total assets	Average real total assets (thousands of euros)
TFP	Average total factor productivity based on a GMM estimation following Ackerberg, Caves and Frazer (2015) and assuming a Cobb–Douglas production function
Wage per employee	Average wage per employee

Notes: All monetary variables are PPP-adjusted.