

Startup Types, Structural Policy and Macroeconomic Performance in Europe^{*}

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

February 14, 2022

Abstract

How much can policymakers improve macroeconomic outcomes by encouraging the entry of high-performance startups? We construct a novel and comprehensive data set on almost 1,5 million startup firms in ten European countries. We then apply cluster analysis to identify distinct startup types and trace their development through early life. Three new stylized facts about entrepreneurial startup strategies transpire. First, we uncover five well-separated clusters of startups, which we label *Basic*, *Large*, *Capital intensive*, *Cash intensive*, and *High leverage*. Second, these startup types are consistently present across countries, industries, and cohorts. Third, startup types are associated with specific life-cycle performance in terms of productivity, employment generation, and exit rates. Feeding these empirical results into an agnostic firm dynamics model, we 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 reaped through policies that benefit high-performance startups (such as large and capital intensive types) while discouraging the entry of underperforming startups (such as high-leverage 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 8th vintage of the CompNet data set, and Tommaso Bighelli, Çağatay Bircan, Javier Miranda (discussant), and seminar participants at University College London, University of Regensburg, Danmarks Nationalbank, University of Bristol, Hong Kong University of Science and Technology, the National University of Singapore, 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 lacklustre macroeconomic performance as reflected in anaemic productivity growth and 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 (Romer and Romer, 2010); innovation subsidies (Bloom et al., 2002; Akcigit et al., 2016); structural reforms to reduce labor market distortions (Nickell, 1997); and measures to deepen financial markets (Lelarge et al., 2010; Bach, 2014; Calomiris et al., 2017).

This paper investigates an entirely different policy lever, one that has remained largely unexplored: influencing the composition of startup firms. It is by now well established that there are large ex-ante differences between newly created firms (see the references below). This evidence suggests that startups can be thought of as belonging to certain “types”, which to a large extent determine their life-cycle performance. It follows that structural policies that can successfully shift the mix of startup types that enter the economy, may have important macroeconomic impacts. Perhaps not surprisingly, some governments have therefore started to focus their policy efforts on specific startup types. For example, the UK’s tax-relief programs for startup investors limit support to smaller firms with assets and employees below a certain threshold.¹ Policy makers often motivate these targeted policies as an attempt to help high-potential startups to survive but it is likely that they also have in mind promoting more entry of these types. For example, in May 2020 the UK government launched a Future Fund to support the “most innovative businesses” and “top-performing start-ups with huge economic potential” that will “create high-skilled jobs”.²

The idea of improving the composition of new firm cohorts *ex ante*—as opposed to “fix-

¹In the case of the Enterprise Investment Scheme (EIS) the asset limit is GBP 15 million and the full-time employee limit is 250 while for the Seed Enterprise Investment Scheme (SEIS) these limits are GBP 200,000 and 25 full-time employees.

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

ing” established generations ex post—appears attractive due to its forward-looking nature. Rates of firm entry and exit are high, typically around ten percent annually. This means that the majority of firms that will be producing twenty years from now is yet to be born, while many of the current firms will no longer exist then.³ Indeed, startups have been documented to be key drivers of job creation and productivity growth, see Haltiwanger et al. (2013) and Haltiwanger et al. (2001). Yet, in practice, government efforts to create jobs and boost economic dynamism by stimulating entrepreneurship often fail (Lerner, 2009). These observations raise important questions: Can targeted startup policies improve macroeconomic performance structurally 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 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, coming 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. The data contain information on almost 1,5 million European startups and allow us to draw a direct link between micro and macro outcomes along various key dimensions. Indeed, the data provide unique cross-country panel observations with representative coverage of the full startup population, which 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 which is limited to employment) and databases that poorly capture firm entry and exit (such as Orbis, cf.

³For instance, in the US Longitudinal Business Database it can be observed 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 or even higher.

Bajgar et al. (2020)).

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

An important 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 to deal with latent heterogeneity, see Bonhomme et al. (2021). 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 the data, we classify based on five important choice variables in the initial year of startup: employment; the capital-to-labor ratio; total assets; the leverage ratio; and the cash-to-assets ratio. The practical advantage is that these variables are widely observed from tax information from the very beginning of the firm life cycle and hence can readily be used to differentiate startups and facilitate targeted policies.

The clustering algorithm gives rise to an endogenous classification of firms into five types in the baseline (in an extension we consider six and seven types). It turns out that this clustering captures the majority of the empirical variation, between 50 and 70 percent. Moreover, the outcome of the clustering analysis is remarkably robust across countries. In almost all countries in our sample, we obtain the following five firm clusters: (i) large; (ii) capital intensive; (iii) high-leverage; (iv) cash-intensive; and (v) basic. Using a “meta-

clustering” analysis we verify that each of these types is similar 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. Also, we find that the distribution of startups across the five types is quite stable across countries and economic sectors. Finally, we track these five startup types for over a decade and document the life cycle profiles of the types in terms of the key choice variables. We find that the initial cross-type differences are highly persistent. All these findings confirm that the clustering algorithm robustly captures fundamentally different firm types.

Based on the clustering outcomes, we then document differences in performance. We find large and persistent differences in employment and productivity across the various types, implying that 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 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 and total factor productivity and is significantly more likely to exit in any given year. By contrast, startup types that stand out either because of their capital-intensive production technology or because of their relatively high cash-to-assets ratios, invariably boast higher productivity levels and lower exit probabilities.

Having documented these structural and persistent differences between startup types, we use the 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.⁴ The differentiation of tax rates alters the incentives for different firm types to enter the economy and hence affects the startup mix itself. We proceed to compute the effects of such budget-neutral tax reforms on the macro economy. We find that policy can significantly increase aggregate productivity and employment by tilting the mix of startup types. The

⁴While this particular policy instrument suits us due to 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.

productivity effects are smaller than for employment because there exists less heterogeneity in productivity across startup types than in employment levels.

We also show which kinds of startups should be encouraged and which ones 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 consequences of startup type-specific policies. Finally, we use our results to shed light on the effects that existing policies may have on the composition of startups and the resulting macroeconomic implications.

Related literature

We build on an emerging literature that documents the importance of ex-ante heterogeneity for firm performance over the life cycle.⁵ We contribute to this literature in three important ways. First, we treat ex-ante heterogeneity explicitly as a multi-dimensional object, whereas most existing studies focus on just one dimension of heterogeneity, such as firm size.⁶ We characterize firms jointly by five key choice variables at startup and we consider several performance measures (labor productivity, TFP, profitability, exit probability). Second, the quality of our data allows us to provide a micro-to-macro mapping along these dimensions and to compare this mapping across countries. This allows us to paint a rich and novel characterization of the European startup landscape. Third, the clustering algorithm classifies firms into types based on observables in the first year of operation. This makes the approach relatively straightforward to implement by policy makers. By contrast, other studies have not made firm types observable (Albert and Caggese, 2021; Sterk et al., 2021).

An added advantage of our analysis is that we obtain empirical estimates for the elasticity

⁵For example, Schoar (2010); Hurst and Pugsley (2011); Andersen and Nielsen (2012); Guzman and Stern (2015); Belenzon, Chatterji and Daley (2017); Sedláček and Sterk (2017); Sterk, Sedláček and Pugsley (2021).

⁶Other 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).

of firm entry and that we 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). 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.

We also contribute to a branch of the public finance literature that studies the effects of public policy on firm-level outcomes (Bloom et al., 2002; Zwick and Mahon, 2017; Liu and Mao, 2019; Benzarti and Harju, 2020). Much of this literature has focused on how taxation and other policies affect the entry, performance, and exit of firms but has not considered the effects on the startup mix. We are—to the best of our knowledge—the first to provide such an analysis, which usefully complements existing impact estimations.

Moreover, we add to the literature on the micro origins of aggregate productivity (Hsieh and Klenow, 2009; Restuccia and Rogerson, 2013; 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 understand the drivers of differences in aggregate productivity across countries and industries, as well as changes over time. Instead of focusing on the “start-up 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 start-up 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., 2021) and business accelerators (Gonzalez-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 birth.

2 Theoretical framework

In this section, we present a theoretical framework on which we will base the policy counterfactuals, to be presented in Section 5.3. The framework also offers clear guidance on which empirical statistics are needed to answer the questions at hand.

2.1 The model

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

Incumbent firms. There is an endogenous mass of heterogeneous firms. An individual firm, indexed by i , operates a production function given by $f_i(x, z)$, where x is a vector of endogenously chosen firm-level inputs, e.g. capital and labor, and z is a vector of exogenous variables, e.g. exogenous TFP shocks. 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 over time or with the age of the firm.

The firm also faces a demand curve $q_i(p, x, z)$ which depends on its price p and potentially on other choice variables (e.g. the firm might expand its demand via advertising, a form of intangible capital accumulation) as well as shocks (e.g. 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, e.g. it may depend on the firm’s type. Finally, we allow for an arbitrary set of constraints, which may be firm specific.

Time is discrete. 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:

$$\mathbb{E}V_{i,j} = \max \mathbb{E} \sum_{t=0}^{\infty} \Lambda^t \pi_{i,j}(x_{i,t}, z_{i,t}; \tau_j)$$

where \mathbb{E} is the expectations operator, Λ is the firm's discount factor which for simplicity we assume to be common across firms, $\pi_{i,j}(\cdot)$ is the pre-tax profit function implied by the production and demand structure, and τ_j a policy variable that is taken as given by firms and which may be type dependent. To fix ideas, it is instructive to think of τ_j as a schedule of tax rates, although the model applies more generally.

A firm exits (for ever) if and only if the firm no longer has positive value, that is if and only if $\mathbb{E}V < 0$.

Entrants. For any type j , there is 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 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 $\theta_{i,j}$ 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 $\tilde{\mathbb{E}}_j V$ 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 $\tilde{\mathbb{E}}_j V \geq \theta_{i,j}$. That is, only those entrepreneurs with a sufficiently low entry cost actually start a firm. We can now express the actual number of entrants of type j , n_j , as:

$$\ln n_j = g_j(\ln \tilde{\mathbb{E}}_j V), \quad (1)$$

where the function g_j depends on the joint distribution of the entry costs, $\theta_{i,j}$, and where we have expressed the variables in logs for convenience. The above equation pins down how entry in each type fluctuates with expected firm values.⁷

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 discount rates are taken as given for any individual country. 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 (e.g., 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 labour productivity among firms of age a and type j , and $\omega_{a,j}$ is the employment share of these firms.⁸

2.2 Effects of a policy change on startup composition

Based on the above framework, we now present a simple formula to compute the compositional effects of a change in taxes or subsidies. Let dT denote the *direct* effect of a policy change on the present value of the profits of an individual firm. In case the policy is a tax, dT is the change in the present-value tax bill absent any behavioral change by the firm.

⁷The function g_j depends on the firm type but is constant over time. That is, we assume that the distribution of entry costs and the number of potential entrants may vary across types but are time invariant.

⁸In turn, $y_{a,j}$ is measured in the micro data as the employment-weighted average labor productivity of individual firms.

Taking a first-order approximation of Equation (1) gives the percentage change in the number of entrants in type j :

$$\frac{dn_j}{n_j} = -\varepsilon_j \tilde{\mathbb{E}}_j \left[\frac{dT}{V} \right]. \quad (2)$$

The equation states that the percentage change in the number of entrants equals the product of two objects. The first, $\varepsilon_j \equiv g'_j(\ln \tilde{\mathbb{E}}_j V)$, is the elasticity of the number of entrants of type j with respect to the ex-ante expected firm value. The second, $\tilde{\mathbb{E}}_j \left[\frac{dT}{V} \right]$, is the average direct effect of the policy change on the present-value profits of firms of type j , expressed as a fraction of the firm value. 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 in 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, the percentage change in the firm value is simply the change in the tax rate, i.e. $\tilde{\mathbb{E}}_j \left[\frac{dT}{V} \right] = d\tau_j^{CI}$, where $d\tau_j^{CI}$ is the change in the corporate income tax rate (potentially differentiated by firm type). More generally, $\tilde{\mathbb{E}}_j \left[\frac{dT}{V} \right]$ is straightforward to compute, provided that one can evaluate the direct implications of a policy change for the tax payment of a certain type and can observe profits in the data, as we do.

Given the change in the number of entrants, the counterfactual firm and employment shares of the different types in the incoming cohort, and hence the $\omega_{a,j}$'s, are straightforward to compute. This in turn allows us to calculate a counterfactual time path for 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 need 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 3.4.

2.3 Effects on post-entry behavior

Note that to evaluate the formula in Equation (2), one does 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 value are equal to zero. This property makes the formula particularly convenient to apply in practice to evaluate the compositional implications of any change policy.

That said, effects of tax changes on the post-entry behavior of firms of course do have macroeconomic implications which do not work via the startup composition channel. These effects depend on the specifics of the policy. For example, a large literature has studied how various kinds of taxes and subsidies affect firm-level outcomes and what are the macroeconomic implications.⁹ 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 such studies. Indeed, to evaluate the overall macroeconomic effects of a proposed tax reform, one can simply add outcomes of our composition calculation to those from existing studies on post-entry behavior.¹⁰

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.

⁹See, for example, Akcigit et al. (2016); Zwick and Mahon (2017); Akcigit et al. (2018); Liu and Mao (2019); Benzarti and Harju (2020).

¹⁰In our concrete policy application we will consider a corporate income tax, which does not 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 e.g. Sedláček and Sterk (2019).

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 the Competitiveness Research Network (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 competitiveness data set for 20 European countries. In order 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 that 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 that 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.¹¹ 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).

To collect harmonized cross-country data on European startups we attached additional code to the standard instructions sent to national authorities in early 2021 in preparation for the eighth CompNet vintage.¹² Our code extracted data on all startups, defined as firms that enter the data set for the first time (see also Section 3.2). This gives us a unique cross-country and cross-sector baseline panel of all 867,789 startups established during the calendar years 2010-18.¹³ We also construct an alternative, unbalanced panel that contains

¹¹See Lopez-Garcia and di Mauro (2015) for more details on the CompNet project.

¹²We 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).

¹³We drop firms that are not observed in each year between the start of their operations and their exit from the data set.

all startups born during 2002-19, with coverage varying across countries and years. This unbalanced panel contains information on a total of 1,450,547 startups. Hereafter we refer to the group of startups that are born in a particular country and year as a ‘cohort’. We aggregate the data in each country and year at either the macro (economy-wide) level or at the 1-digit NACE Rev.2 industry level.¹⁴ For each country-industry-year 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.¹⁵ We retrieve the data also split by startup type. The classification of startups will be discussed in Section 3.3.

3.2 Startups in the CompNet database

Using the CompNet data, we aim to identify all firms that start their operations in a particular country and year (that is, a startup cohort). To do so, we first take all firms that enter the data set for the first time in a particular year. We define this year as the start year as it is the first time that a firm is 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 grow and develop in terms of their employment generation, productivity, and survival—all areas on which cross-country evidence remains scant.

¹⁴These 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.

¹⁵We use two productivity measures. The first is labor productivity, defined as real value per employee. The second is Total Factor Productivity (TFP), which is based on a production function estimated via the two-step control function approach of Ackerberg et al. (2015) and implemented at the 2-digit sectoral level.

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

In Appendix Figures A1, A2, and 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 datasets is very similar although in a few cases—such as the Netherlands and Sweden—there are gaps in the average *level* of reported employees. Lastly, Figure A3 compares exit rates (firm death) over time in both data sets. We see that the five-year cumulative exit rate is comparable at 45% (56%) in the Compnet (Eurostat) dataset. Moreover, the trends as firms age are very similar in both cases, although in France (Sweden) Compnet tends to underreport (overreport) firm exit. Overall, we conclude that the firm population that underpins the statistical moments we get 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 (see e.g. 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.¹⁷ 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.

The clustering algorithm allocates each individual firm i into one of $j = 1, \dots, k$ 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 whose mean is closest. Based on that initial categorization, new group means are determined

¹⁷See, for example, Bonhomme et al. (2021) who study the case of labour supply with unobserved worker types.

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.¹⁸ 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 that with five clusters one can capture around three fifths of the total variation in the five clustering variables. Beyond $k = 5$, the η^2 statistic still increases but levels off. The data therefore suggest that our start-ups 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).¹⁹ 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, total factor productivity, 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.²⁰

¹⁸The statistic has a similar interpretation to the R^2 often reported in regression analysis.

¹⁹All monetary variables in CompNet are denominated in thousands of Euros. Also, as is standard in the literature, we let the clustering be based on z-scores as opposed to raw variables, in order to eliminate arbitrary scaling effects.

²⁰Earlier theoretical and empirical 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 (?Bustamante and D'Acunto, 2021; Farinha et al., 2019; Derrien et al., 2020). 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).

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 first-stage clustering procedure as the observations.²¹

The four panels in Figure 1 visually summarize the outcomes 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 one 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 appear to be very similar. Moreover, for the meta clustering, we obtain $\eta^2 = 0.963$. That is, variation between clusters explains the vast majority of the overall data variation, leaving only a very small contribution for cross-country variation within meta clusters.²²

²¹That is, in the meta-clustering procedure the unit of observation are 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.

²²Another indication of the similarity is that we see 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.

One might be concerned that there are mechanical reasons why we observe very similar clusters arising in different countries, or that the similarity is a coincidence. To investigate this, we conduct a Monte Carlo experiment for the meta clustering. Specifically, we consider a large number of 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. distributed so in truth no meta clusters exist.²³ 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 0.963 observed in the real world data. The experiment thus supports our interpretation that the cluster outcomes observed in the data for different countries are indeed remarkably similar.

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 startups

Table 1 summarizes the results of our clustering analysis. We label the five archetypal startup clusters that emerge *Basic* (49 percent of all start-ups); *Large* (4 percent); *Capital intensive* (7 per cent); *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, as compared with an average of only 3 people in the other categories. Likewise, *Cash intensive* startups hold

²³We assume log-normal distributions for these draws, as the cluster variables are non-negative.

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. 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 are less widespread in Croatia, Denmark, and Lithuania. Figure A6 in the Appendix shows how this country-specific variation in the distribution of startup types persists over time.

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 a bit overrepresented among manufacturing firms but underrepresented among 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 varies across countries as well (see Figure A7 in the Appendix).

Lastly, Figure 4 shows that the distribution of startups across the five startup types (and aggregated over all countries) is relatively persistent during the period 2010-19. We nevertheless observe a small decline in the share of highly-leveraged and basic startups, while the share of cash-rich startups steadily increases.

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 A8 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 like Croatia. Likewise, the last panel of A8 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 lives. We first focus on age profiles in the choice variables that we fed 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 start-up 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 twelve 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 about 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 the 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, however, 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 points 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. As a result, there is also some convergence in the average capital intensity of these large firms over time.

Next, we are interested in how the 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 firms in a given country, one-digit industry, startup year (cohort), age and type.²⁴ The analysis uses the full (unbalanced) sample and we include dummies to indicate the four main startup types, 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 performance that might correlate with startup type.²⁵

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 total factor productivity (column 2) as well as a lower likelihood of early exit (column 3). Especially *Large* startups are considerably less likely to exit within the first decade after commencing operations. This can also be seen in the Kaplan-Meier survival curves in Appendix Figure A9. 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

²⁴For example, an observation could refer to all Spanish ICT firms in the “high leverage” category born in 2005, at age seven.

²⁵In particular, country x cohort FE absorb all time-invariant characteristics common to startups born in a specific country and year; industry x cohort FE control for time-invariant traits common to all startups in a specific industry and born in a given year; and country x industry FE absorb all time-invariant variation that characterizes startups in a specific industry and country. Lastly, 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 birth years *and* that depend on a firm’s age.

startup types (column 4).²⁶

4.3 Startup types: Profiles

This section provides a short profile or vignette 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 is the existence of 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 macro economy, such as the aggregate labour 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.

Right from the start capital-intensive firms use a large amount of 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 only partially—converges to 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 high-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

²⁶Appendix Table A2 replicates this analysis using a sub-sample of firms of between five and eight years old. The results line up closely with those in Table 2, indicating that the performance divergence between startup types is not solely driven by transient differences in early life.

strong persistence in these startups' reliance on capital in the production process (Figure 5). This not only reflects the industries that capital-intensive startups tend to operate in (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 that 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, such as highly educated workers (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 start-up population has remained relatively stable. This suggests that the global trend towards automation documented in the literature (Acemoglu and Restrepo, 2018; Graetz and Michaels, 2018) is broad-based and not exclusively driven by an increase in capital-intensive startups.

4.3.2 Large startups

About 4 percent of all startups belong to the *Large* cluster. 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 finding is consistent with empirical findings in Sterk et al. (2021) who use U.S. data to document the importance of ex-ante heterogeneity for firm-level employment.

Size heterogeneity may derive from differences in TFP combined with decreasing returns

to scale (which implies that more productive firms choose a higher optimal size), or from heterogeneity in returns-to-scale, with some production processes being scaled up relatively easily. A final possibility is that the size heterogeneity derives from differences in the scalability of demand for firms' products. Hottman et al. (2016) provide empirical evidence that differences in the scale of demand is a quantitatively important source of heterogeneity across firms.

Again, our results shed light on the life-cycle structure and performance of heterogeneous firms, this time along the size dimension. Large firms are relatively common in construction, transport and logistics, and manufacturing. Large startups are relatively productive and, importantly, this productivity premium is very persistent over time. They scale up their work force even further during the first twelve 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 total factor productivity 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 and better access to external funding—make large startups the least likely to exit during the first twelve years of operations (see also Figure A9). 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).²⁷

²⁷Branstetter 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.

4.3.3 High-leverage startups

About 14 percent of all startups belong to the *High-leverage* cluster. These firms stand out because they begin operations by taking on substantial amounts of debt relative to the size of their balance sheet. These high-leverage start-ups begin operations with an average leverage ratio of 116 percent, considerably more than the other start-up clusters.²⁸ While high-leverage start-ups 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 that they may be financially quite vulnerable.

The survival curves in Figure A9 show that high-leverage startups have the lowest 12-year survival probability. This aligns well with an earlier literature documenting a possibly causal link between firm leverage and exit probability.²⁹ Dinlersoz et al. (2019) show that young private firms in the U.S. 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 shown previously by Lemmon et al. (2008) on the basis of Compustat data (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 before these firms become profitable at all.³⁰ This cluster of star-

²⁸Owners 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)).

²⁹See, 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).

³⁰Rajan and Zingales (1995) provide cross-country evidence of a negative correlation between firm-level 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).

tups 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 particularly 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 start-up the information asymmetry between banks and firms is highest. Firms with high-quality investment projects, who know that 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 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 her enterprise goes bankrupt).

4.3.4 Cash-intensive startups

About a 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 start-up, these firms hold almost half of their assets in the form of cash, about five times as much as firms in the other start-up 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 relatively 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 hence 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 twelve 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, total factor productivity, and profitability and consequently also pay somewhat higher wages. These patterns are in line with Begenauf 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 may 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 in order 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 especially because these firms tend to own few tangible assets that could serve as collateral, or possibly because these businesses are relatively non-standard, making it more difficult to obtain bank credit.

4.3.5 Basic startups

Lastly, just below 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 start operations, these firms employ on average four people; operate with a capital intensity of 8.6 percent; have a cash ratio of 12 per cent; and a leverage ratio of 22 percent. Basic

start-ups 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

In this section, we perform counterfactual policy exercises to study how changes in taxation can improve specific macroeconomic outcomes by changing the composition of new startup cohorts, and what such tax changes might look like. We do so based on the model framework developed in Section 2. As explained in that section, to quantify the startup composition effects of tax policy, we need the life-cycle profiles as presented in the previous section. In addition we need a second set of statistics: the entry elasticities for each startup type. Below 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 using a linearization of Equation (1):

$$\ln V_j = \beta_0 + \beta_1 \ln n_j + u_j$$

where $\ln V_j$ is the ex-post average logged value of type j . $u_j = \ln V_j - \ln \widehat{\mathbb{E}V}_j$ is a residual that captures the difference between ex-ante and ex-post values arising from ex-post aggregate shocks. This residual is orthogonal to n_j since ex-post shocks are not known at the time of entry. Note that $\beta_1 = \frac{1}{\varepsilon_j}$ is the inverse of the entry elasticity that we can estimate via Ordinary Least Squares.

A final question is how to measure V_j in the data. We do so using the life-cycle profile for profits and exit rates, and the number of entrants for each country-industry-startup type-

cohort in our data set.³¹ We assume a discount factor $\Lambda = 0.96$, corresponding to an annual discount rate of about 4 percent. We only use cohorts that we can observe for at least seven years (that is, those born before 2011) and drop all observations beyond age seven to create a balanced panel. For age eight and onward, we assume that profits and the year-on-year exit rate remain fixed.³²

Table 3 presents the results of the entry regressions. Column (1) shows the estimate of the entry elasticity without conditioning on type or country. On average, the elasticity is about $1/0.648=1.5$. That is, an increase in expected firm value by one percent pushes up the number of startups by 1.5 percent. Column (2) differentiates the estimate by startup type. This shows that while the elasticity estimate has the same order of magnitude for the various startup types, there is nevertheless substantial heterogeneity. The entry elasticity varies between 0.98 for large startups and 2.1 for high-leverage startups. Finally, column (3) shows that the entry elasticity is relatively comparable across countries, although it appears to be particularly high in Spain.

5.2 Counterfactual exercises

We can now proceed to conduct 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 compute the composition channel we do not necessarily 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 of different types as a fraction of the firm value. Thus, the following results apply to a wide range of policy changes in the taxation and subsidization of firms.

³¹We construct profits by multiplying profits (EBIT) per unit of revenue and revenues, at the country-industry-startup type-cohort level.

³²That 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.

To fix ideas, however, we will 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 value are simply equal to the tax rate itself. Further, 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.³³ This generally implies that if some 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 types, 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 re-scale them to ensure that 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 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 behaviour 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 age for which we have estimated the life-cycle profile. We compute the macro outcomes using the life-cycle profiles while accounting for the exit rate of each type (which we also estimate from the data).

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

³³To be precise, they are revenue neutral excluding potential tax revenue effects due to changes in post-entry behaviour.

standard deviation of the implied tax rate changes across firms.³⁴ The idea is that larger policy changes may bring about larger 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 6 summarizes the results of our policy experiment. Panel A shows the policy space for employment. The horizontal axis measures the intensity of the potential policy change (that is, the standard deviation of tax rate changes). Moving from left to right, warmer colors indicate stronger corporate tax rate differentiation. The solid line plots the policy upperbound. This is the largest possible aggregate employment increase given a certain policy intensity. For instance, for a policy intensity of 0.1, aggregate employment can be improved by up to 6 percent. Figure 6 shows that large improvements in aggregate employment are possible by tailoring tax rates to encourage entry of particular startup types. At the same time, however, other policy changes with the same intensity could reduce aggregate employment by a similar amount (up to an amount indicated by the dashed line).

Panel B of Figure 6 again plots the policy space, but this time for aggregate labor productivity. Substantial gains are again possible, although in percentage terms the effects are smaller than for employment. This mostly reflects that there exists less heterogeneity in productivity across startup types than in employment.

Next, the two panels on the left of Figure 7 plot the corporate tax rate policies associated with the two policy upperbounds. Likewise, the two panels on the right show the associated firm shares. The upper left panel shows that, in order to achieve the productivity upperbound, taxes should be lowered for cash intensive and capital intensive types. This would

³⁴More generally, the policy intensity could be measured as the change in the average present-value tax bill of each type. Since we interpret the policy as a corporate income tax, this is equal to the change in the tax rate.

encourage entry among these types and increase their share in the total startup population. At the same time, in order to keep the policy revenue neutral, tax rates should be raised for the other types, in particular the large startups. The upper right panel shows the associated shares of the types in the startup population. These shares do not move one-for-one with the tax changes since the elasticities of entry are allowed to vary across types. Not surprisingly, the share of capital-intensive startups increases at tax schemes implementing the policy upperbound. The shares of basic and high-leverage firms tend to decline, whereas the shares of cash intensive and large firms tend to remain constant.

The lower left panel of Figure 7 shows the tax rates associated with the employment frontier. In order to boost aggregate employment, tax rates should (unsurprisingly) be lowered for the large type. Taxes generally rise for the other types but, interestingly, the tax rise for the basic type is relatively modest. Considering the associated firm shares, shown in the lower right panel, the share of basic and large startups increases, whereas the shares of cash intensive and high-leverage firms tend to decline.

To analyze the trade-offs between aggregate employment and productivity, Figure 8 provides a scatter plot of the change in both variables under each potential policy with a standard deviation of the change in tax rates of up to 0.2. The figure suggests that, generally, there is not much of a trade-off between increasing aggregate employment and productivity, as the two outcomes tend to be positively correlated, with a correlation coefficient of 0.56.

Table A1 provides insights into trade-offs among a wider set of aggregates, by computing the correlations between outcomes among all policies in the aforementioned policy space. All correlations are positive, suggesting that there are not necessarily strong trade-offs between increasing aggregate employment, labor productivity, TFP, the number of firms, and firm profits. Importantly, however, the correlation between employment and TFP is relatively low (0.12), suggesting that there is a relatively strong trade-off between these two outcomes.

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.³⁵ Likewise, the U.S. 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' 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, many countries also 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

³⁵See 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 up to EUR 25,000.

with a particular funding structure. These include special credit guarantees and loans for startups, for example.³⁶ In other cases, tax rebates depend on a startup’s physical location. In the UK, startups qualify for a 100 per cent tax exemption for five years if they are located in one of 21 enterprise zones while in China preferential tax rates apply to firms located in the country’s western provinces. China has also encouraged foreign direct investment in the manufacturing sector through lower tax rates for foreign-owned firms in so-called Special Economic Zones and Economic and Technology Development Zones.

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, British startups qualify for R&D tax relief if they have fewer than 500 employees and revenues of less than EUR 100m, or own less than EUR 86m in assets.

The above discussion shows that many countries already operate corporate tax systems that in some way or form differentiate, often implicitly, between types of startups.³⁷ Moreover, the resulting variation in effective tax rates is often substantial and reflective of particular 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

³⁶Other measures, such as the UK Enterprise Investment Scheme (EIS) and Seed Enterprise Investment Scheme (SEIS) help startups to attract equity funding by offering tax reliefs to individual investors who buy new shares in startups.

³⁷Because 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).

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. Hence, there are potential macroeconomic gains (or losses) to be had from policies that affect 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 that these macro gains/losses can be very substantial quantitatively. Moreover, our results shed light on how policy could be changed to achieve these gains. That is, which types should be encouraged (such as capital-intensive startups) and which types discouraged (such as high-leverage startups). In short, policy makers aiming to improve macroeconomic performance would be well advised to consider 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 macro 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 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

- Acemoglu, Daron and Pascual Restrepo (2018) “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108 (6), 1488–1542.
- Ackerberg, D.A., K. Caves, and G. 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,” Technical report, National Bureau of Economic Research.
- Akcigit, Ufuk, Douglas Hanley, and Stefanie Stantcheva (2016) “Optimal Taxation and R&D Policies,” Technical report, National Bureau of Economic Research.
- Albert, Christoph and Andrea Caggesse (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.
- Andersen, Steffen and Kasper Meisner Nielsen (2012) “Ability or Finances as Constraints on Entrepreneurship? Evidence from Survival Rates in a Natural Experiment,” *Review of Financial Studies*, 25 (12), 3684–3710.
- 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.

- Bach, Laurent (2014) “Are Small Businesses Worthy of Financial Aid? Evidence from a French Targeted Credit Program,” *Review of Finance*, 18 (3), 877–919.
- 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 Papers 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*.
- 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 (2020) “Using Payroll Tax Variation to Unpack the Black Box of Firm-Level Production,” *Journal of the European Economic Association*.
- Bergner, S., R. Brautigam, M. Evers, and C. Spengel (2017) “The Use of SME Tax Incentives in the European Union,” *ZEW Discussion Paper 17-006*.
- 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, Stephane, Stephane Lamadon, and Elena Manresa (2021) “Discretatizing Unobserved Heterogeneity,” *Econometrica*.
- 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.
- Bustamante, Maria Cecilia and Francesco D’Acunto (2021) “Banks’ Screening and the Leverage of Newly Founded Firms,” Technical report.
- Calomiris, Charles W, Mauricio Larrain, José Liberti, and Jason Sturgess (2017) “How Collateral Laws Shape Lending and Sectoral Activity,” *Journal of Financial Economics*, 123 (1), 163–188.
- 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 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, ebnem 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.

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

Farinha, Luísa, Sónia Félix, João AC Santos et al. (2019) “Bank Funding and the Survival of Start-ups,” Technical report.

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

Gonzalez-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, Sebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sánchez (2017) “Capital Allocation and Productivity in South Europe,” *Quarterly Journal of Economics*, 132 (4), 1915–1967.

Graetz, Georg and Guy Michaels (2018) “Robots at Work,” *Review of Economics and Statistics*,

- tics*, 100 (5), 753–768.
- Guzman, Jorge and Scott Stern (2015) “Where is Silicon Valley?,” *Science*, 347 (6222), 606–609.
- Haltiwanger, J., L. Foster, and C.J. Krizan (2001) “Aggregate Productivity Growth: Lessons from Microeconomic Evidence,” in Edward Dean, Michael Harper and Charles Hulten ed. *New Directions in Productivity Analysis*, 303–372: University of Chicago Press.
- 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.
- 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.
- Hurst, Erik and Benjamin Wild Pugsley (2011) “What do Small Businesses Do?,” *Brookings Papers on Economic Activity*, 43 (2), 73–142.
- Iacovone, Leonardo, William Maloney, and David McKenzie (2021) “Improving Management in Colombian Firms Through Individual and Group Consulting,” *Review of Economic Studies*.
- 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 Be-

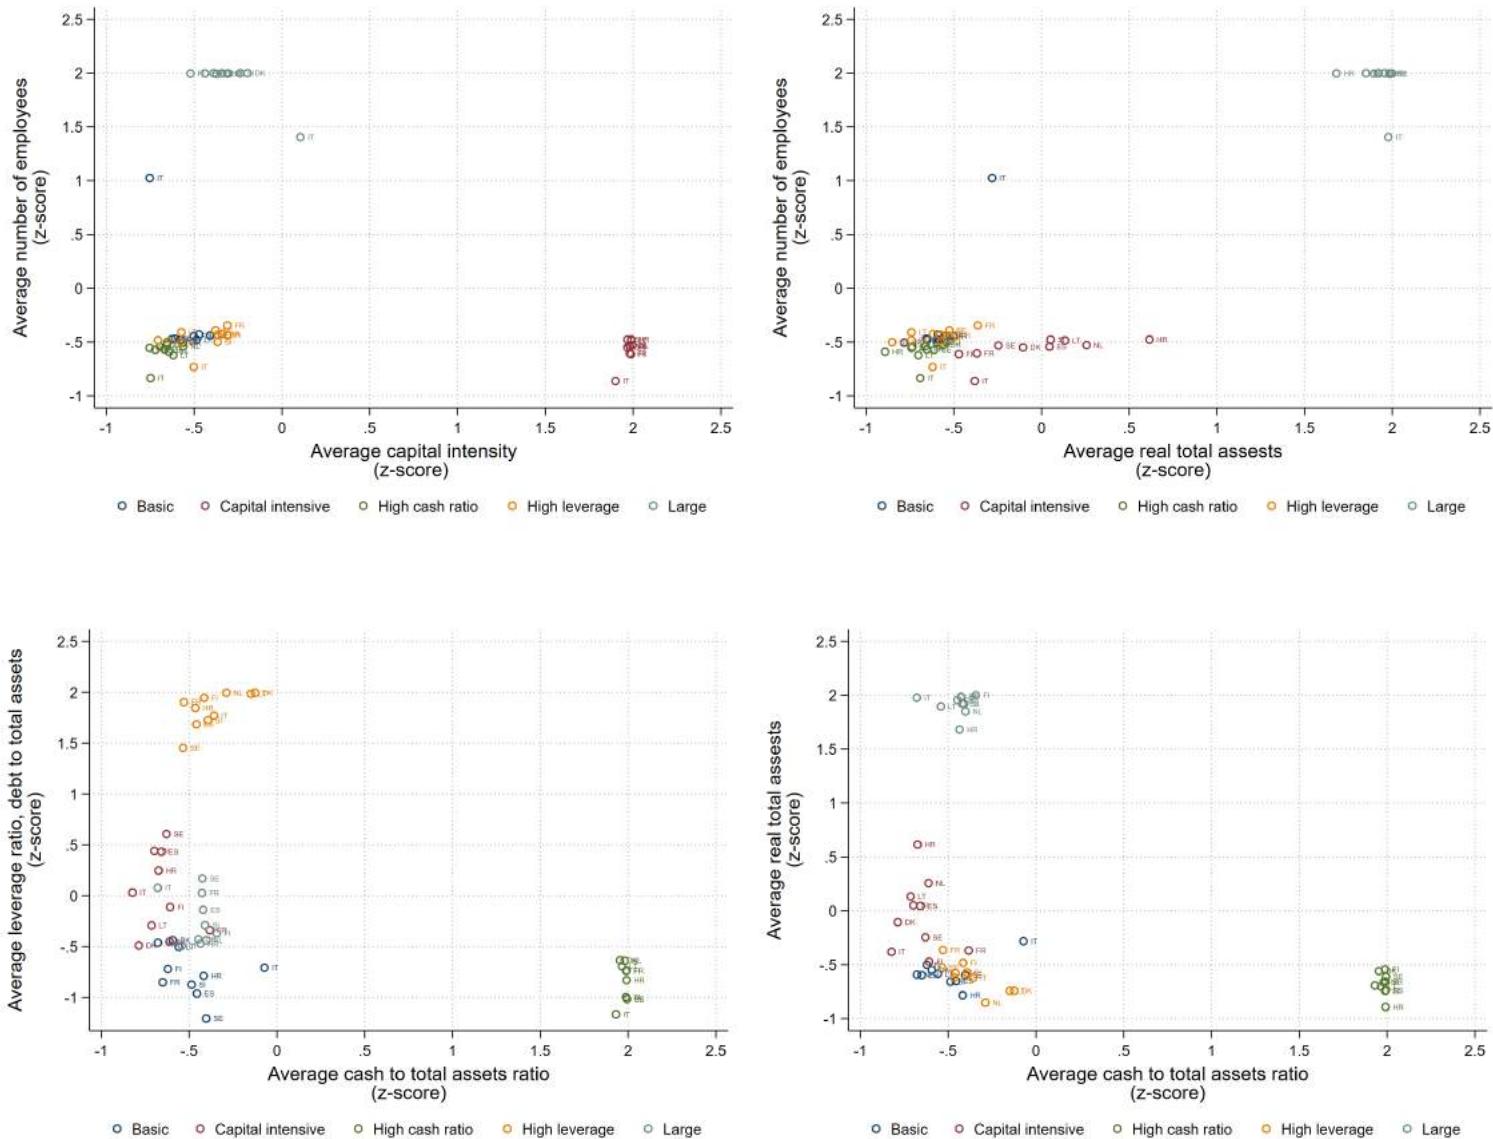
- havior: An Examination of Plant Exit and Investment Decisions,” *Review of Financial Studies*, 10 (3), 767–803.
- Lafontaine, Francine and Kathryn Shaw (2016) “Serial Entrepreneurship: Learning by Doing?,” *Journal of Labor Economics*, 34 (S2), S217–S254.
- Lang, Larry, Eli Ofek, and RenéM Stulz (1996) “Leverage, Investment, and Firm Growth,” *Journal of financial Economics*, 40 (1), 3–29.
- Lelarge, Claire, David Sraer, and David Thesmar (2010) *8. Entrepreneurship and Credit Constraints*: University of Chicago Press.
- 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*: 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, P. and F. 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.

- Nickell, Stephen (1997) "Unemployment and Labor Market Rigidities: Europe versus North America," *Journal of Economic Perspectives*, 11 (3), 55–74.
- 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_I2.
- 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."
- Robb, Alicia M. and David T. Robinson (2014) "The Capital Structure Decisions of New Firms," *Review of Financial Studies*, 27 (1), 153–179.
- Romer, Christina D and David H Romer (2010) "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, 100 (3), 763–801.
- 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.

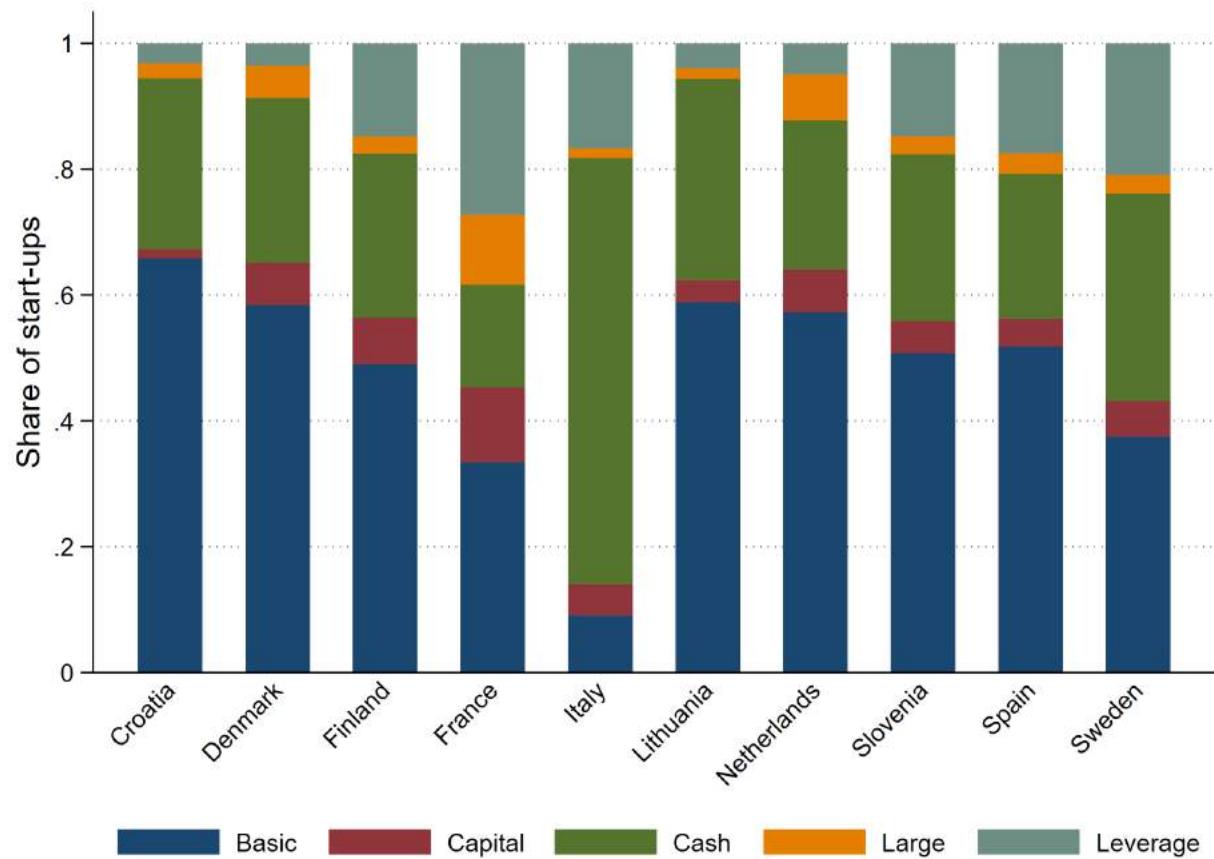
Tables and Figures

Figure 1: Meta clustering of startup types



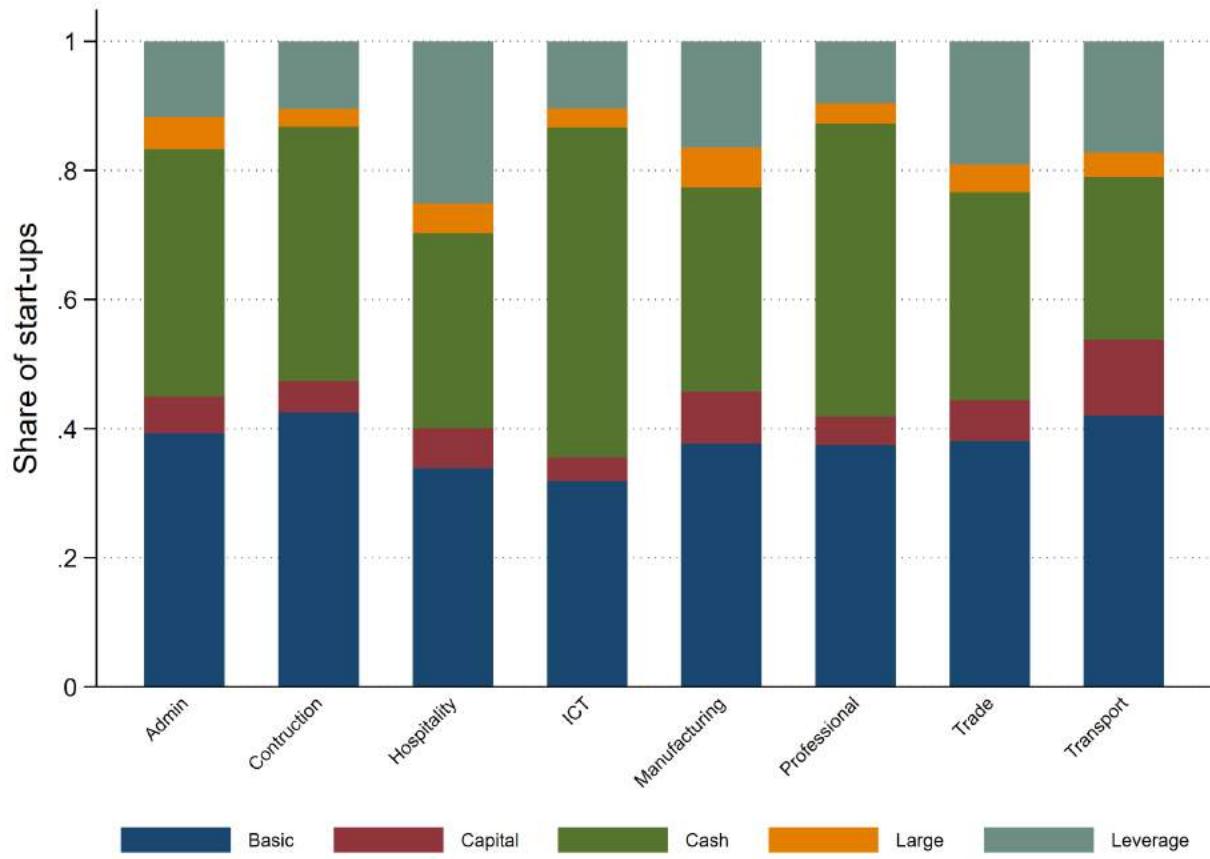
Notes: The four panels in this figure summarize the meta-clustering procedure. Different meta-clusters are indicated with different colours. The meta clustering groups comparable clusters from different countries by taking the cluster centers derived from each country's first-stage clustering procedure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

Figure 2: Distribution of startup types by country



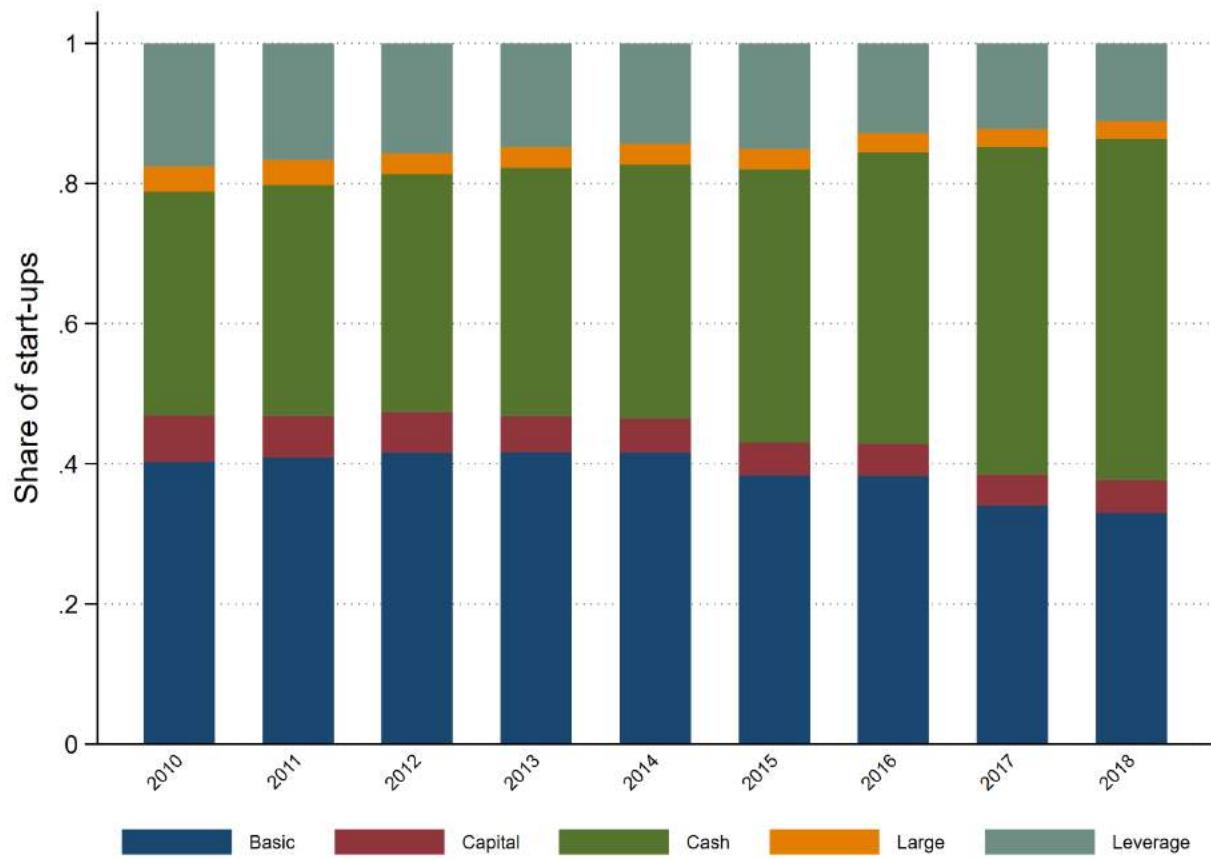
Notes: This figure shows for individual countries the distribution of the startup population across the five startup types. The startup population consists of the 2010-18 cohorts.

Figure 3: Distribution of startup types by industry



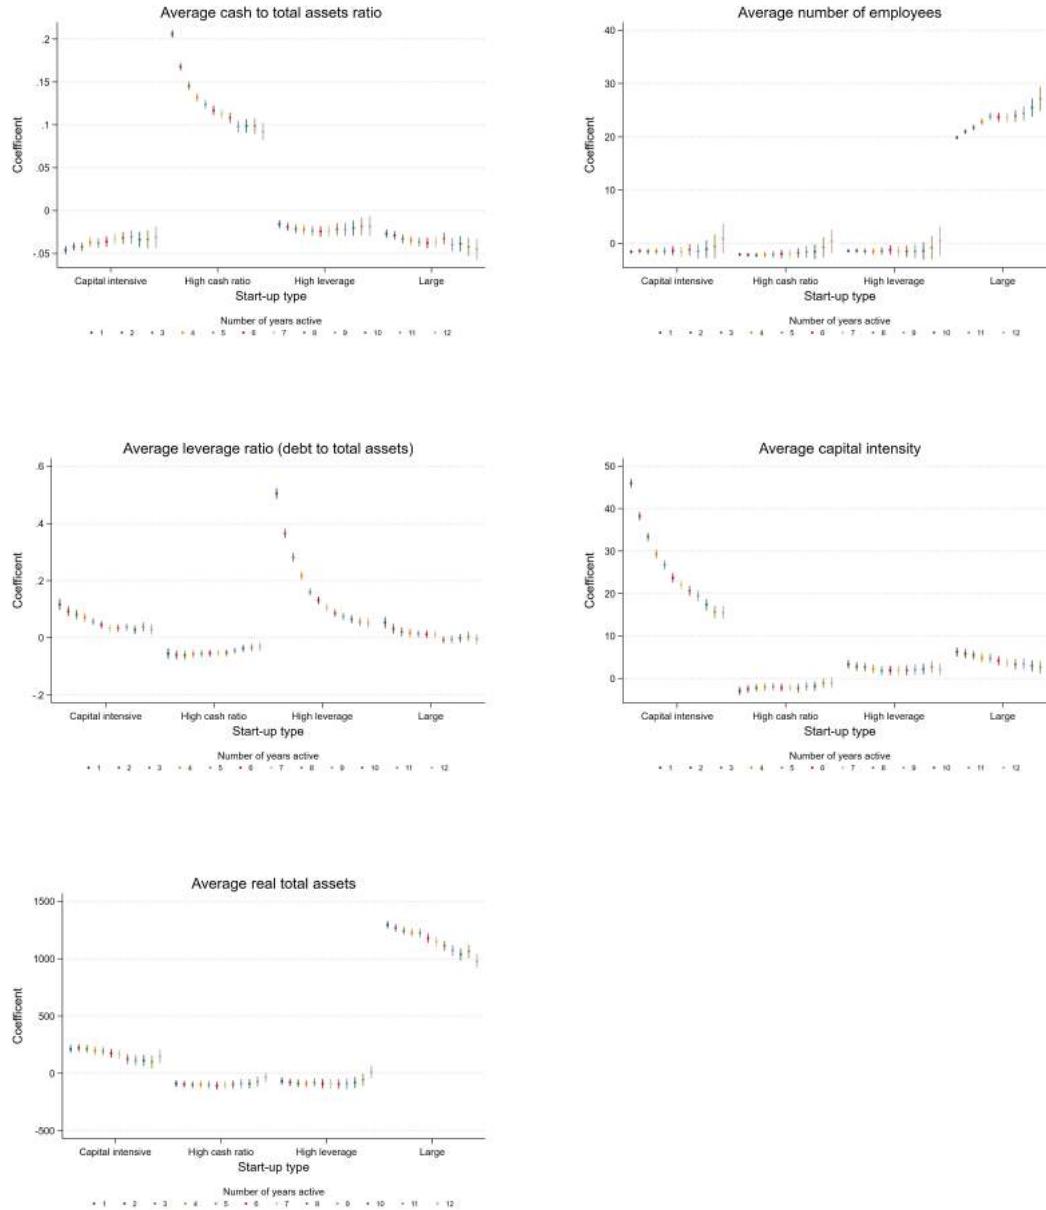
Notes: This figure shows for one-digit industries the distribution of the startup population across the five startup types. The startup population consists of the 2010-18 cohorts.

Figure 4: Distribution of startup types by cohort



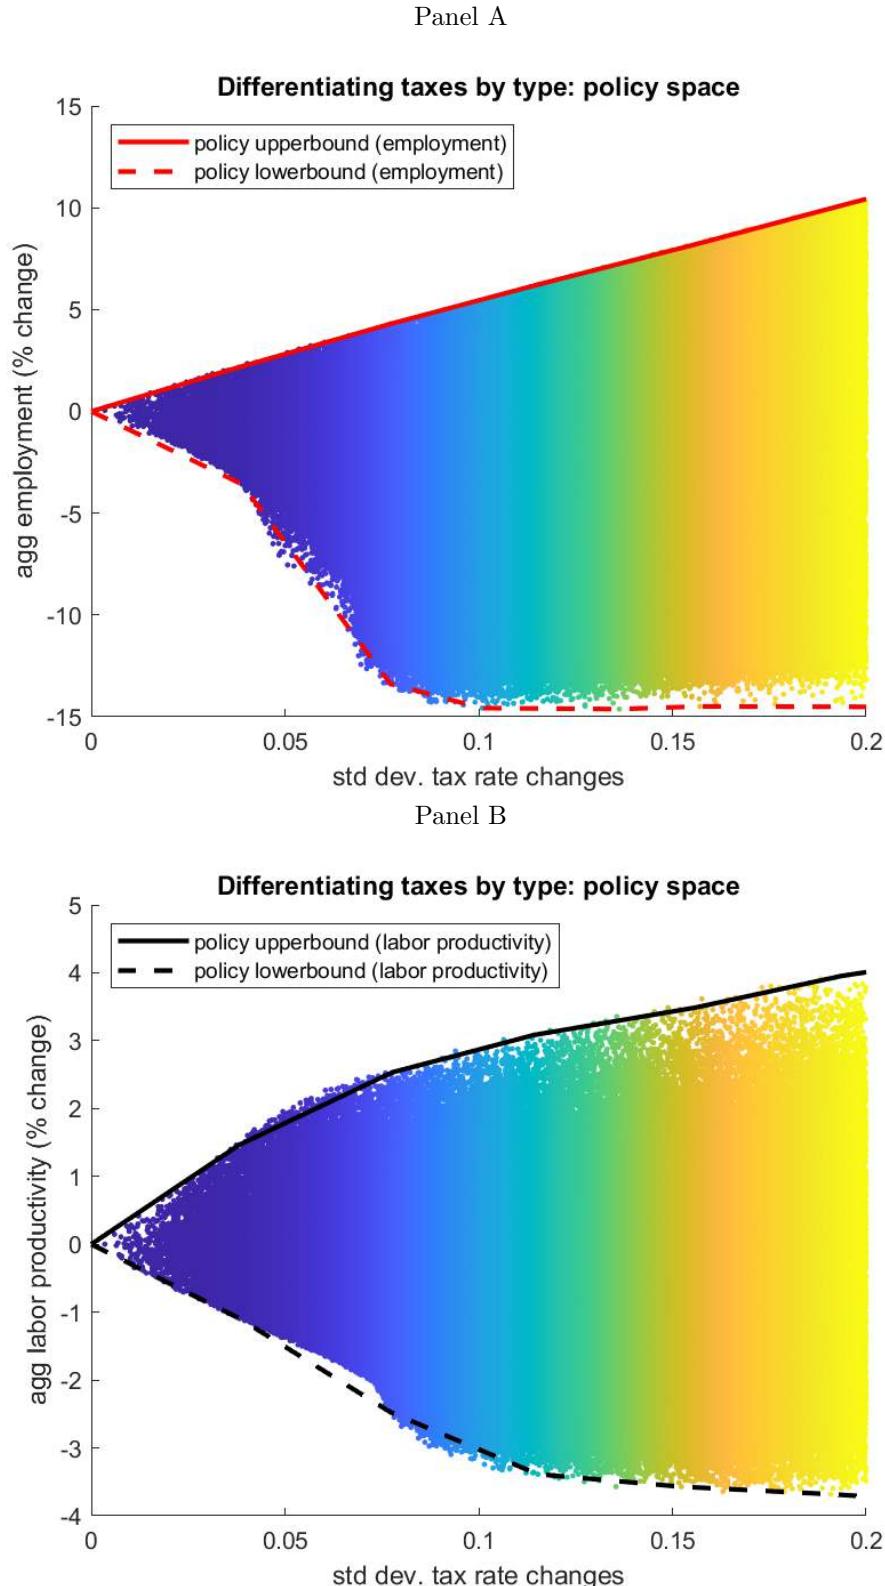
Notes: This figure shows for each cohort the distribution of the startup population across the five startup types. The startup population consists of the 2010-18 cohorts.

Figure 5: The life cycle of startup types



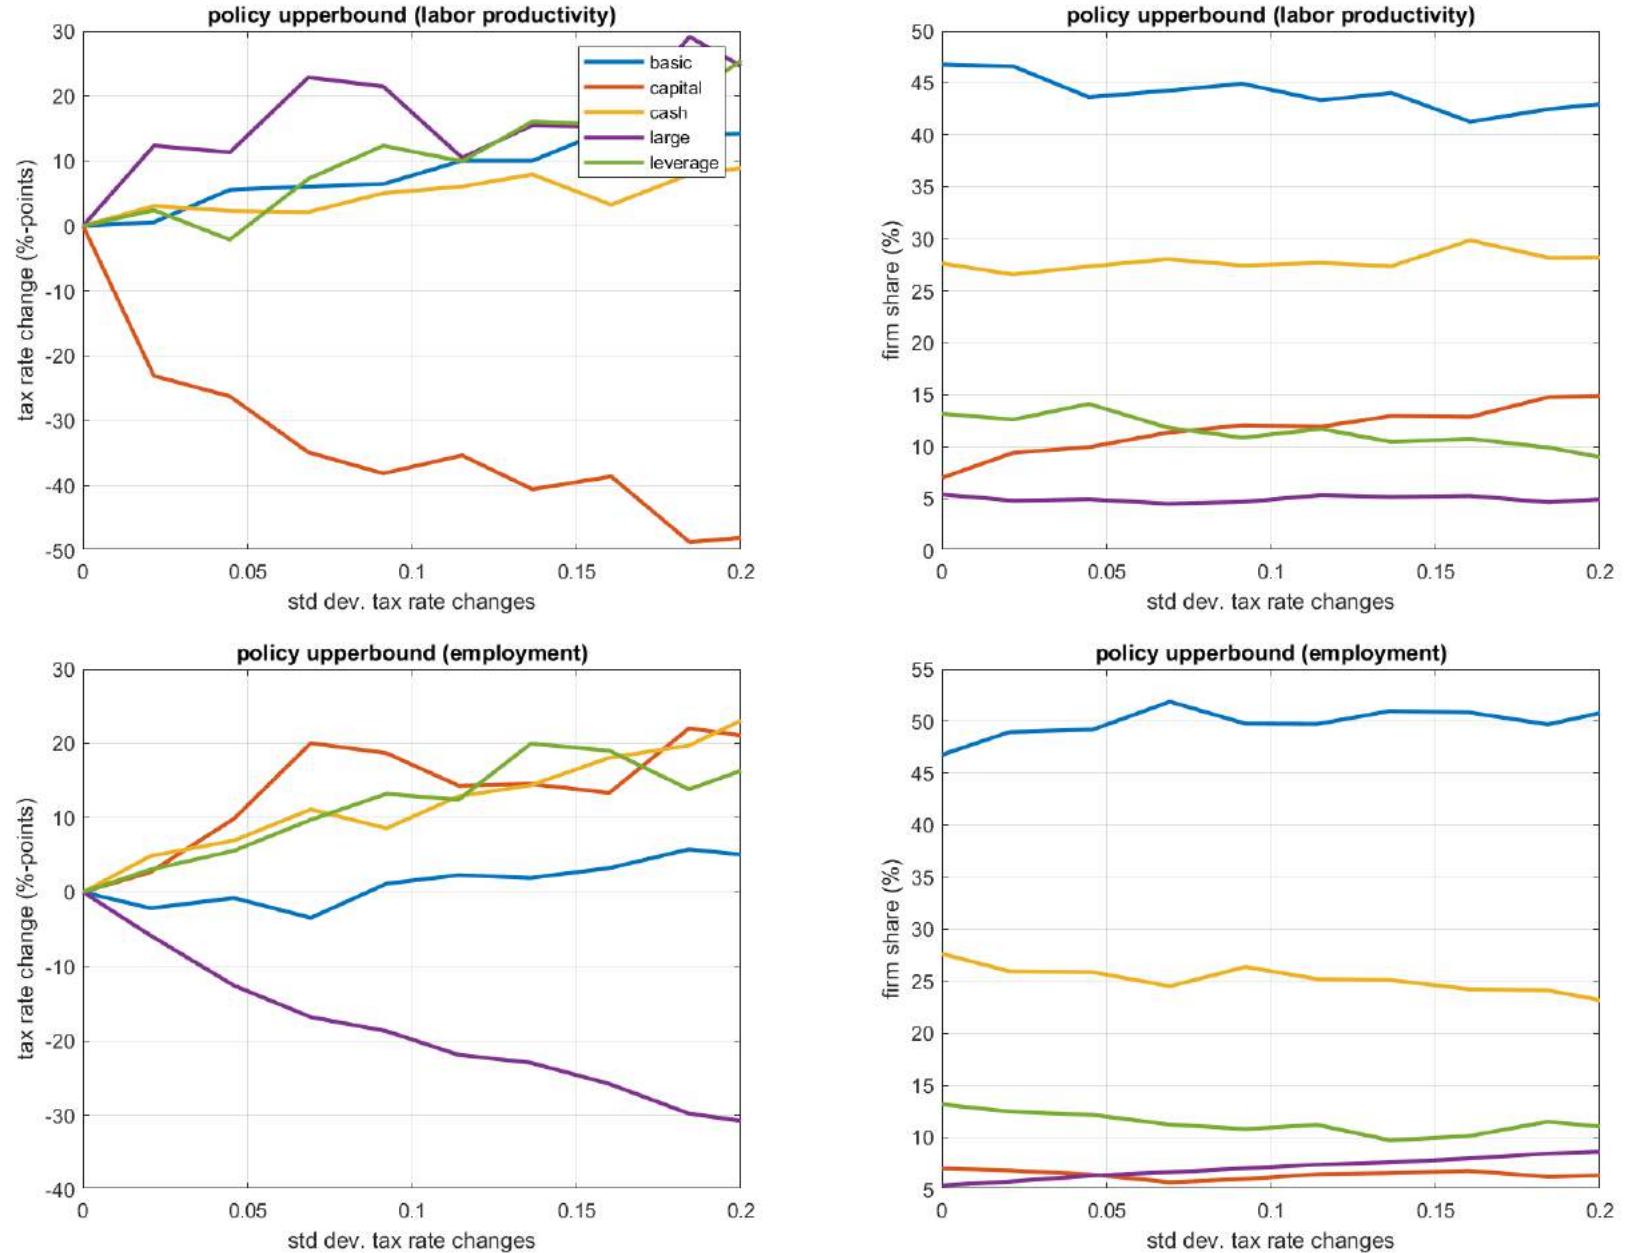
Notes: The panels in this figure summarize how the startup types develop during the first 12 years of their life in terms of the five clustering variables. Each panel corresponds to one clustering variable and plots the coefficients from 12 separate regressions where the dependent variable is this variable. Each regression is then run for an age group (age is 1, 2, ...12 years). For example, the first panel summarizes regressions in which *Average cash to total assets ratio* is regressed on dummy variables for *Start-up type* (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: Policy experiment—Tax differentiation and aggregate labor productivity and employment of young firms



Notes: This figure summarizes the policy experiment. Panel A shows the policy space for employment. The horizontal axis measures the intensity of the potential policy change as the standard deviation of tax rate changes. Warmer colors indicate stronger corporate tax rate differentiation. The solid line plots the “policy upperbound”: the largest possible aggregate employment increase given a certain policy intensity. Panel B plots the policy space for aggregate labor productivity.

Figure 7: Policy experiment—Tax differentiation and startup composition



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

Figure 8: Policy experiment—Aggregate impacts of tax differentiation on employment vs labor productivity



Notes: This scatter plot depicts the change in aggregate labor productivity and aggregate employment under each potential policy associated with a change in the standard deviation of tax rates of up to 0.2.

Table 1: Characteristics of startup types at time of entry

	(1)	(2)	(3)	(4)	(5)
	Number of employees	Capital intensity	Real total assets	Cash/total assets ratio	Leverage ratio
Basic	4	8.64	174.13	0.12	0.22
Capital intensive	2	91.18	407.05	0.09	0.40
Cash intensive	2	4.98	102.19	0.54	0.18
High leverage	3	13.14	135.61	0.14	1.16
Large	20	16.07	1506.81	0.13	0.34

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

Table 2: Start-up type and firm outcomes

	(1)	(2)	(3)	(4)	(5)
	Aggregate labor productivity	Aggregate TFP (GMM estimation)	Exit probability	Wage per employee	Average profit margin
Capital intensive	0.316*** (0.004)	0.048*** (0.003)	-0.062*** (0.003)	2.406*** (0.086)	0.012*** (0.001)
Cash intensive	0.030*** (0.004)	0.050*** (0.003)	-0.008*** (0.003)	0.980*** (0.078)	0.023*** (0.001)
High leverage	-0.052*** (0.004)	-0.041*** (0.003)	-0.002 (0.003)	-1.695*** (0.086)	-0.030*** (0.001)
Large	0.167*** (0.004)	0.042*** (0.004)	-0.128*** (0.003)	3.385*** (0.088)	-0.016*** (0.001)
Constant	3.327*** (0.003)	2.200*** (0.002)	0.708*** (0.002)	27.369*** (0.054)	0.043*** (0.001)
R-squared	0.902	0.978	0.630	0.909	0.612
N	26,491	19,499	28,565	27,677	27,420
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. The regressions are at the 1-digit industry level and based on the full panel of firms younger than nine years. Standard errors are in parentheses. *, **, *** indicate significance at the 10, 5 and 1 percent level, respectively.

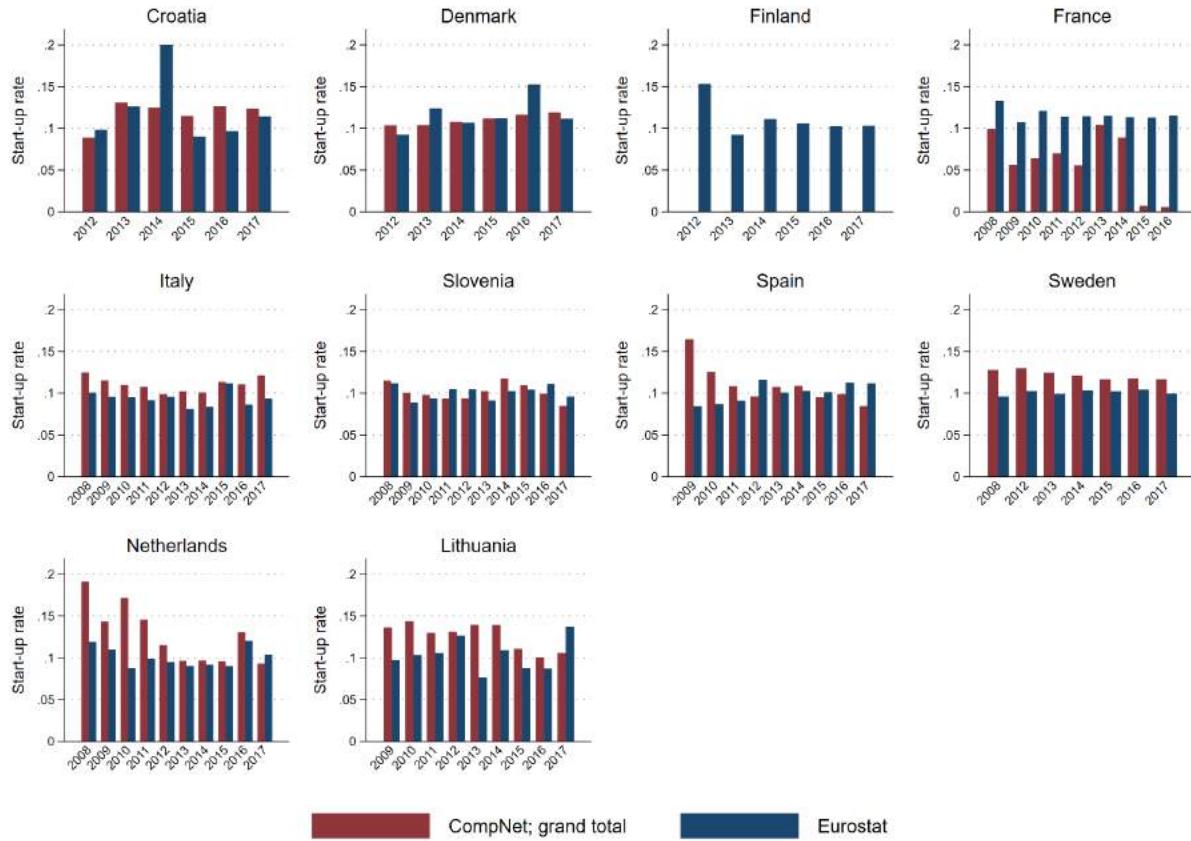
Table 3: Start-up types and entry

Dependent variable: Log total NPV	(1)	(2)	(3)
Log of number of start-ups	0.693*** (0.032)		
Basic \times Log of number of start-ups		0.645*** (0.063)	
Capital intensive \times Log of number of start-ups		0.793*** (0.062)	
High cash ratio \times Log of number of start-ups		0.699*** (0.045)	
Large \times Log of number of start-ups		0.963*** (0.079)	
Leverage \times Log of number of start-ups		0.563*** (0.053)	
Croatia \times Log of number of start-ups			0.776*** (0.105)
Denmark \times Log of number of start-ups			0.810*** (0.049)
Finland \times Log of number of start-ups			0.813*** (0.055)
Italy \times Log of number of start-ups			0.530*** (0.057)
Lithuania \times Log of number of start-ups			0.054 (0.114)
Netherlands \times Log of number of start-ups			0.776*** (0.080)
Slovenia \times Log of number of start-ups			0.613*** (0.108)
Spain \times Log of number of start-ups			0.293** (0.129)
Sweden \times Log of number of start-ups			0.744*** (0.058)
Constant	-0.812*** (0.188)	-0.561 (0.354)	-1.179*** (0.432)
Type FE	Yes	Yes	Yes
Country FE	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes
R-squared	0.594	0.600	0.610
N	1,774	1,774	1,774
Capital intensive <i>p</i> -value		0.071	
High cash ratio <i>p</i> -value		0.454	
Large <i>p</i> -value		0.001	
Leverage <i>p</i> -value		0.291	
Denmark <i>p</i> -value			0.756
Finland <i>p</i> -value			0.738
Italy <i>p</i> -value			0.037
Lithuania <i>p</i> -value			0.000
Netherlands <i>p</i> -value			0.997
Slovenia <i>p</i> -value			0.260
Spain <i>p</i> -value			0.003
Sweden <i>p</i> -value			0.787

Notes: This table shows OLS regressions at the 1-digit industry level based on the full panel of firms younger than nine years and with a start-up year prior to 2012. In column (2) *Capital intensive*, *High cash ratio*, *Large* and *Leverage* *p*-values correspond to the tests whether the interaction coefficients for these four start-up types are significantly different from the *Basic x Log number of start-up* coefficient. In column (3) *Denmark*, *Finland*, *Italy*, *Lithuania*, *Netherlands*, *Slovenia*, *Spain* and *Sweden* *p*-values correspond to the tests whether the interaction coefficients for these countries are significantly different from the *Croatia x Log number of start-up* coefficient. Standard errors are in parentheses. *, **, *** indicate significance at the 10, 5 and 1 percent level, respectively.

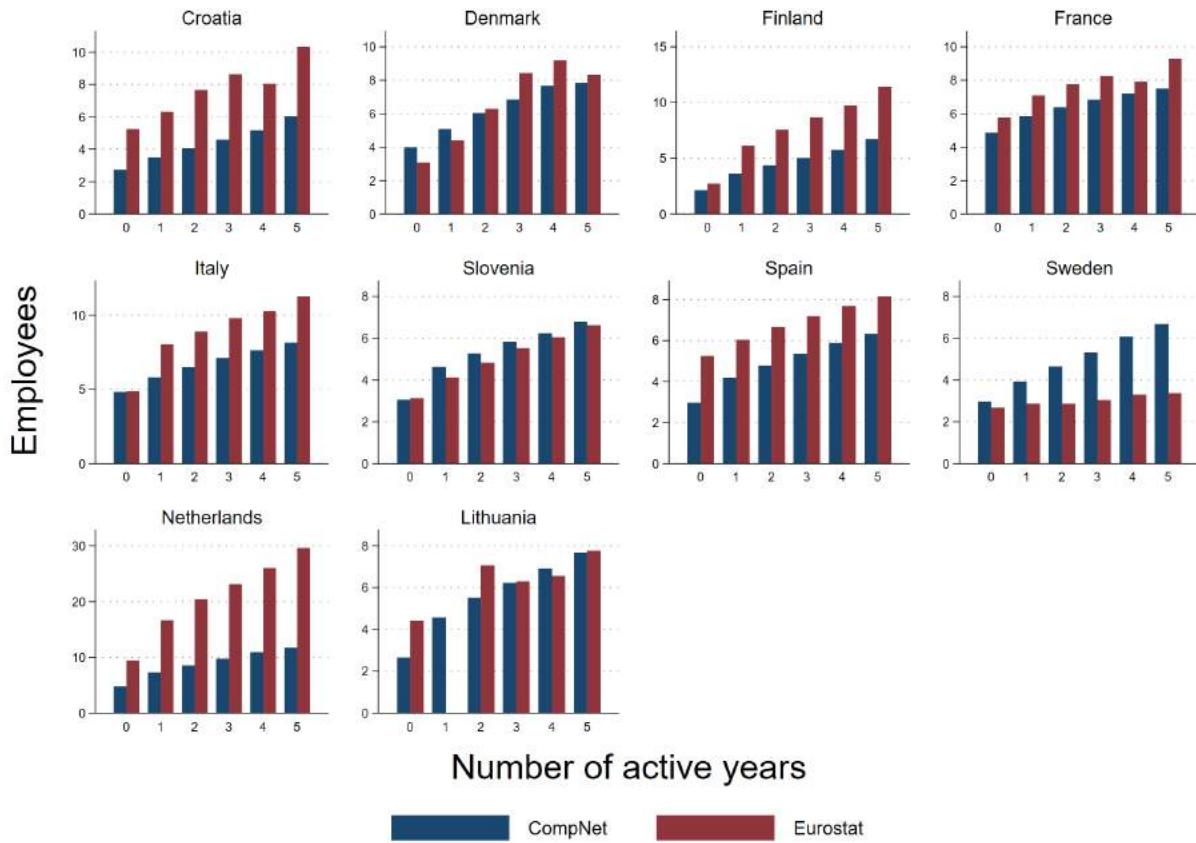
Appendices

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



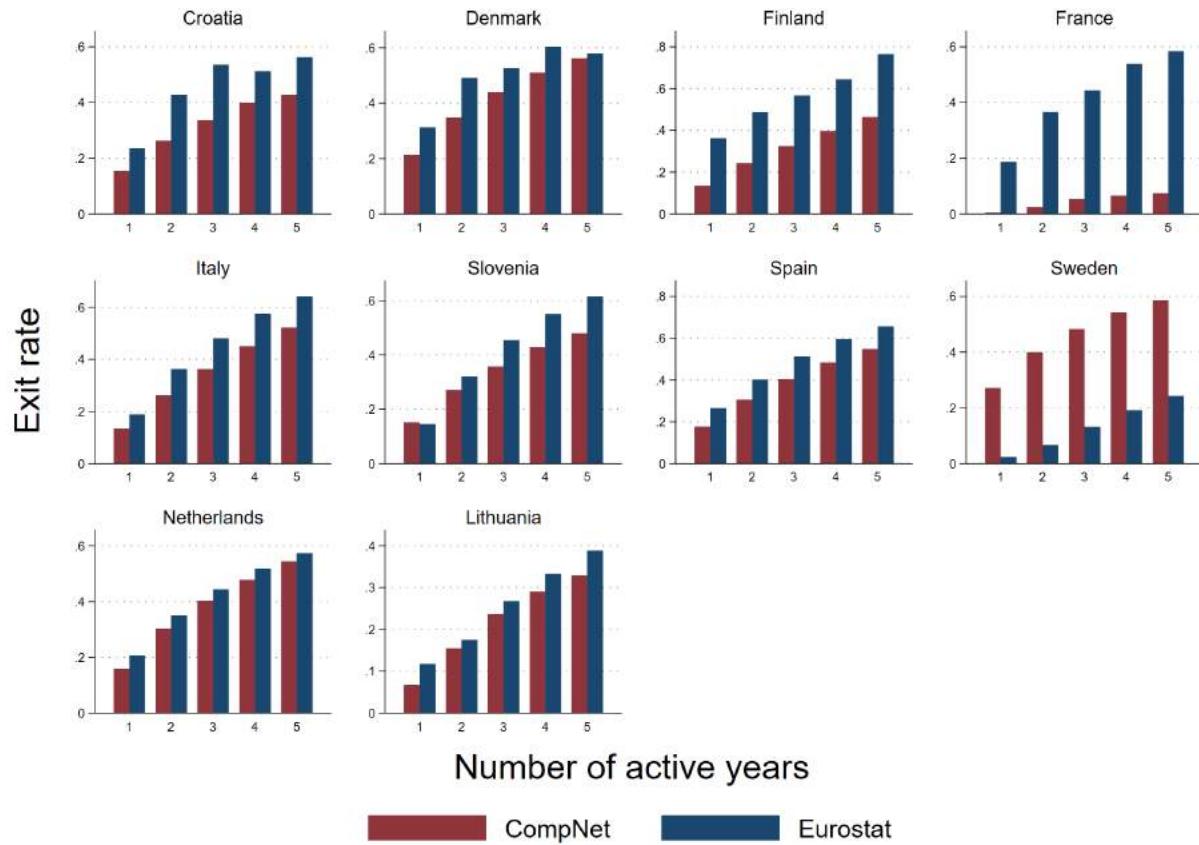
Notes: Figure comparing the total number of start-ups in the CompNet database (red bars) and births of firms as reported by Eurostat (blue bars), cohorts reported are subject to data availability. The total number of start-ups in the CompNet databases (red bars) for Finland are missing because we do not have access to this data.

Figure A2: Employment growth by firm age—CompNet versus Eurostat data



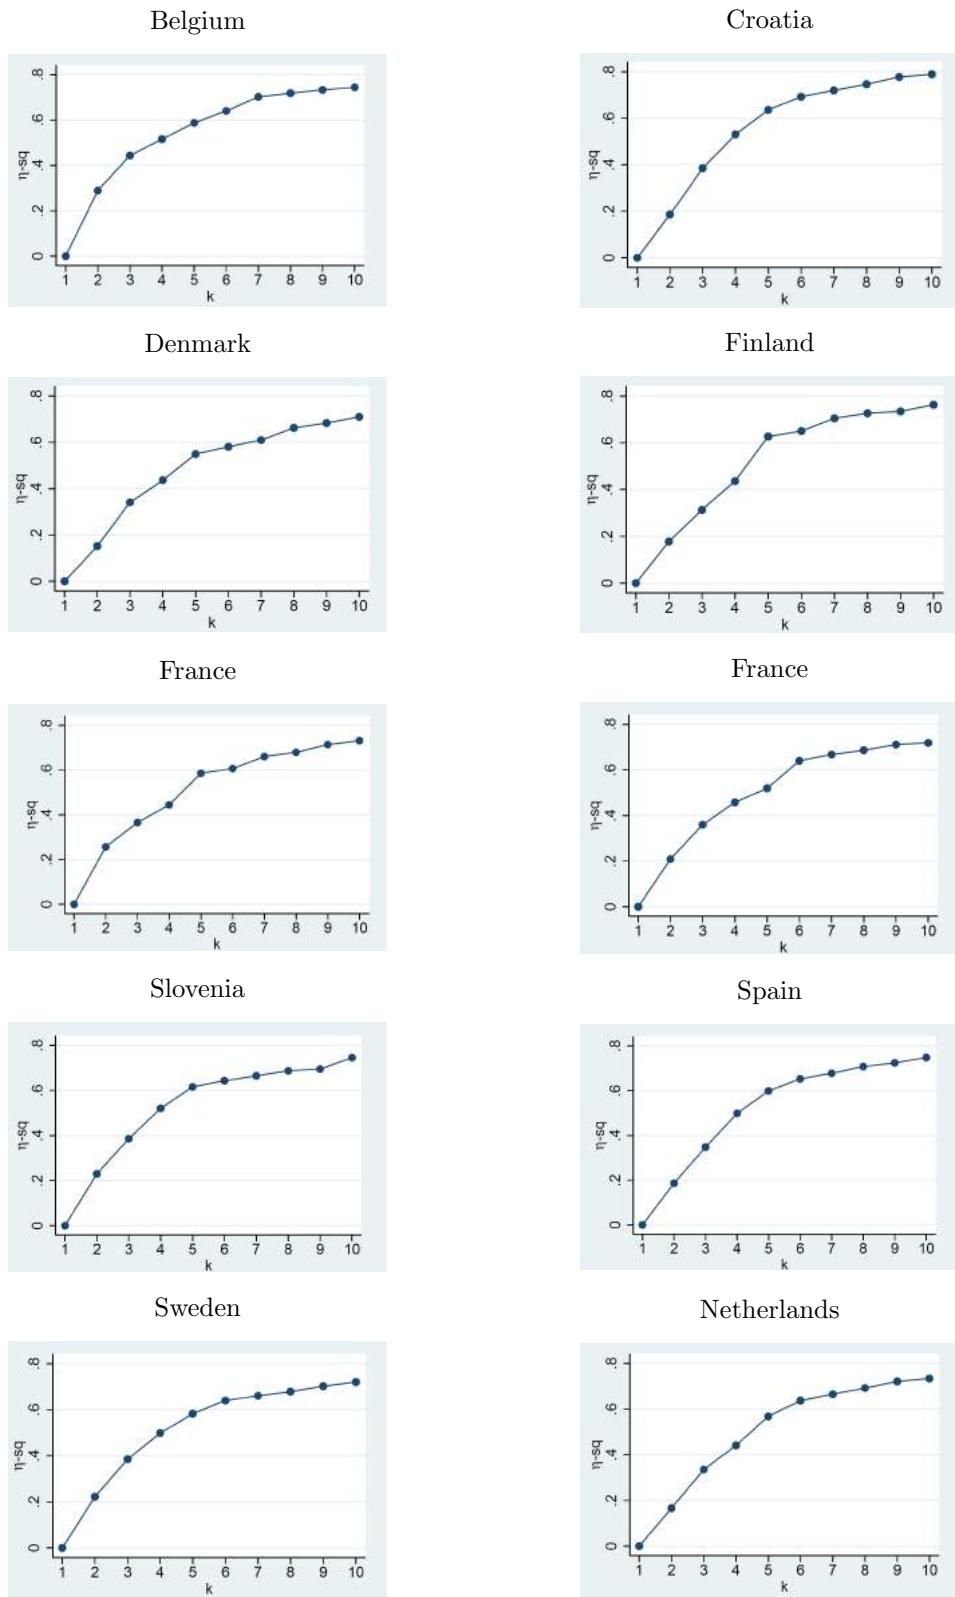
Notes: Figure comparing growth in number of persons employees by start-ups as reported by the CompNet data (blue bars) and Eurostat (red bars). For comparison purposes, we adjust the Eurostat data such sole proprietorship firms are removed and we adjust for the average number of persons employed by sole proprietorship firms. The x-axis depicts *Start-up age* which is the number of years a start-up has been active. *Employees* on the y-axis is averaged over cohorts.

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



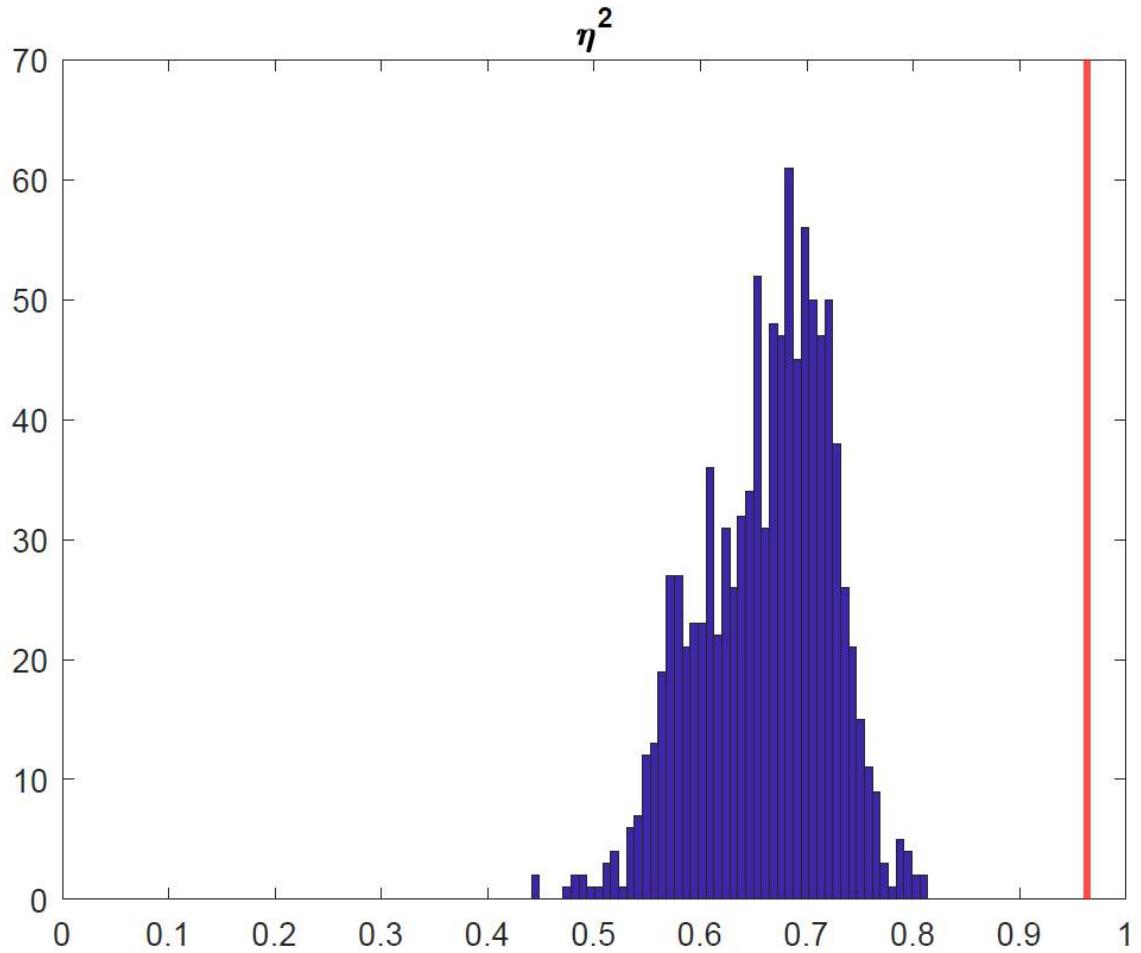
Notes: Figure comparing the cumulative exit rate of start-ups in the CompNet database (red bars) and as reported by Eurostat (blue bars). The x-axis depicts *Start-up age* which is the number of years a start-up has been active. *Exit rate* is calculated by taking an average exit rate over all cohorts for each start-up age group.

Figure A4: Scree plots



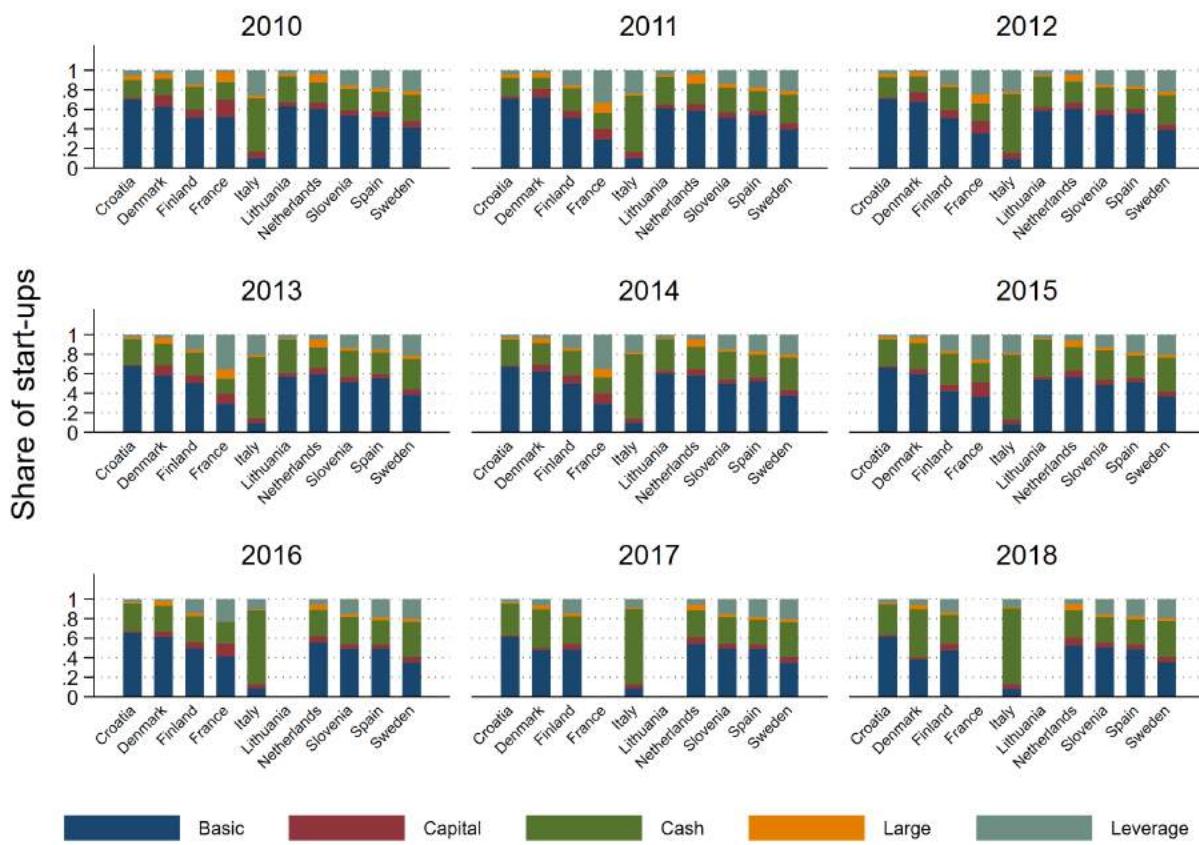
Notes: This figure shows scree plots resulting from the k -means cluster algorithm at the firm-level, for each country in the sample. On the x-axis k indicates the number of clusters. The η^2 coefficient on the y-axis measures the proportional reduction of the within sum of squares for each cluster solution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

Figure A5: Monte Carlo experiment of the meta clustering



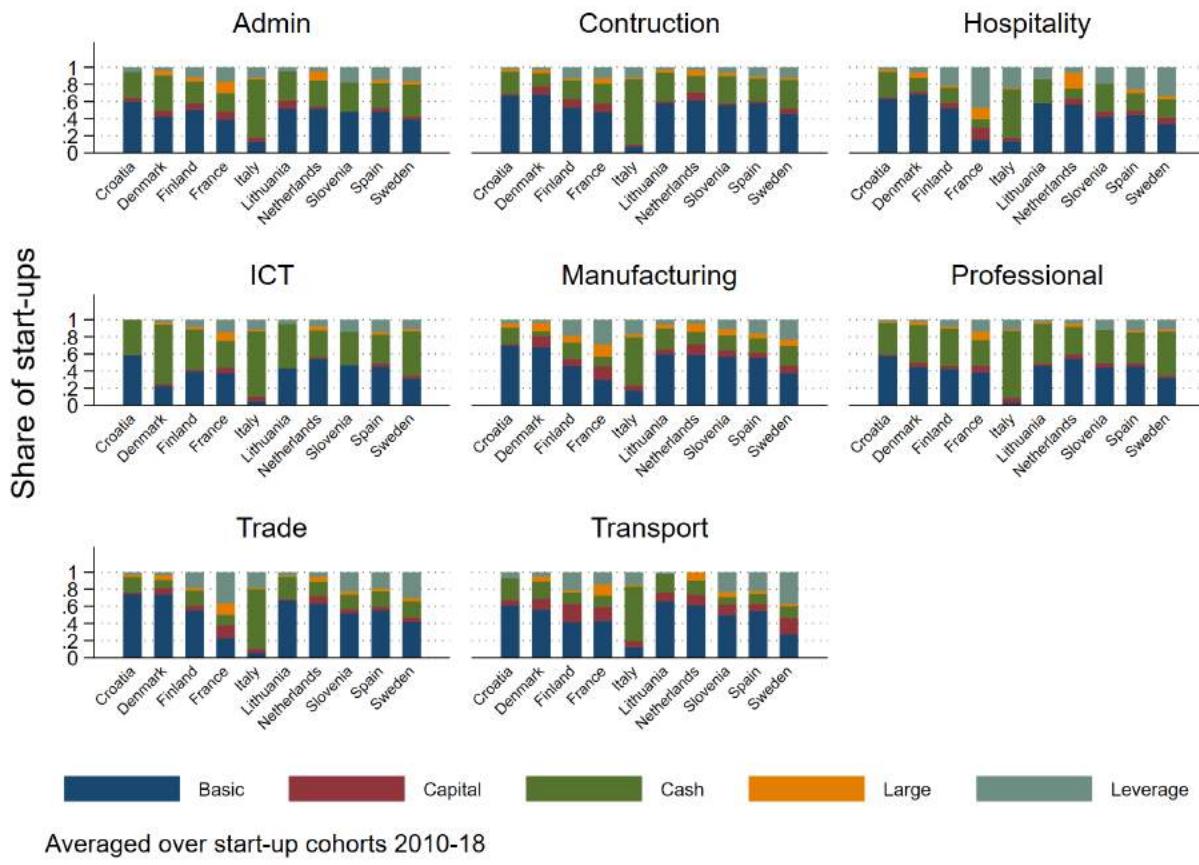
Notes: This histogram summarizes a Monte Carlo experiment consisting of a large number of random draws for the cluster variables, with means and standard deviations as observed in the data. These experimental draws are i.i.d. distributed so that no clusters exist. The experiment is repeated many times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

Figure A6: Distribution of startup types by country and by annual cohort



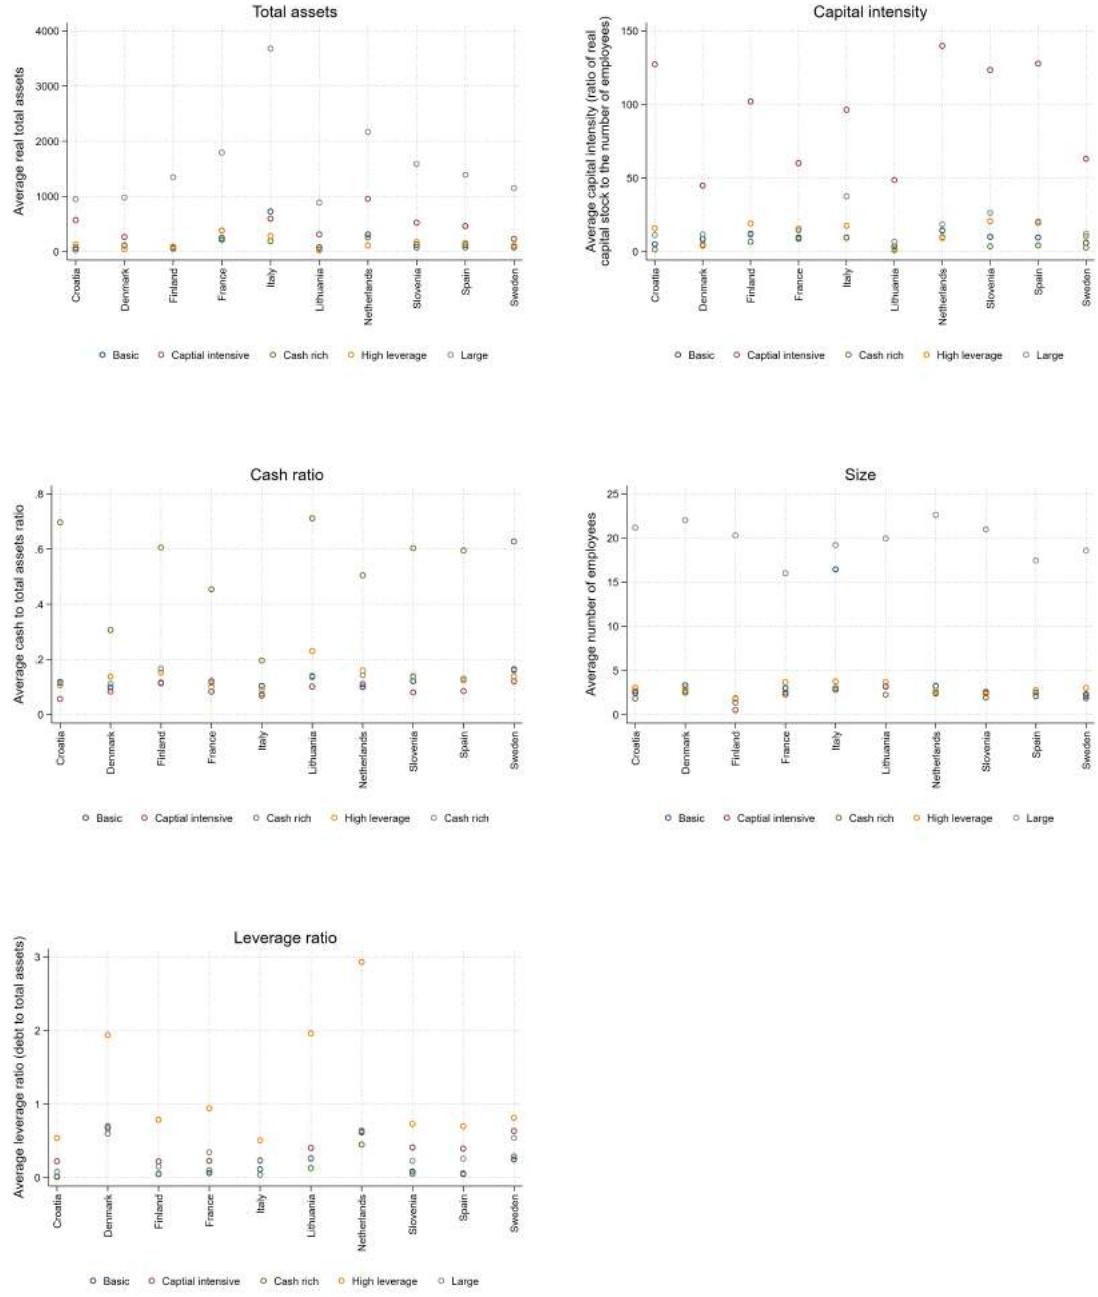
Notes: This figure shows for individual countries and cohort years the distribution of the startup population across the five startup types.

Figure A7: Distribution of startup types by industry and by country



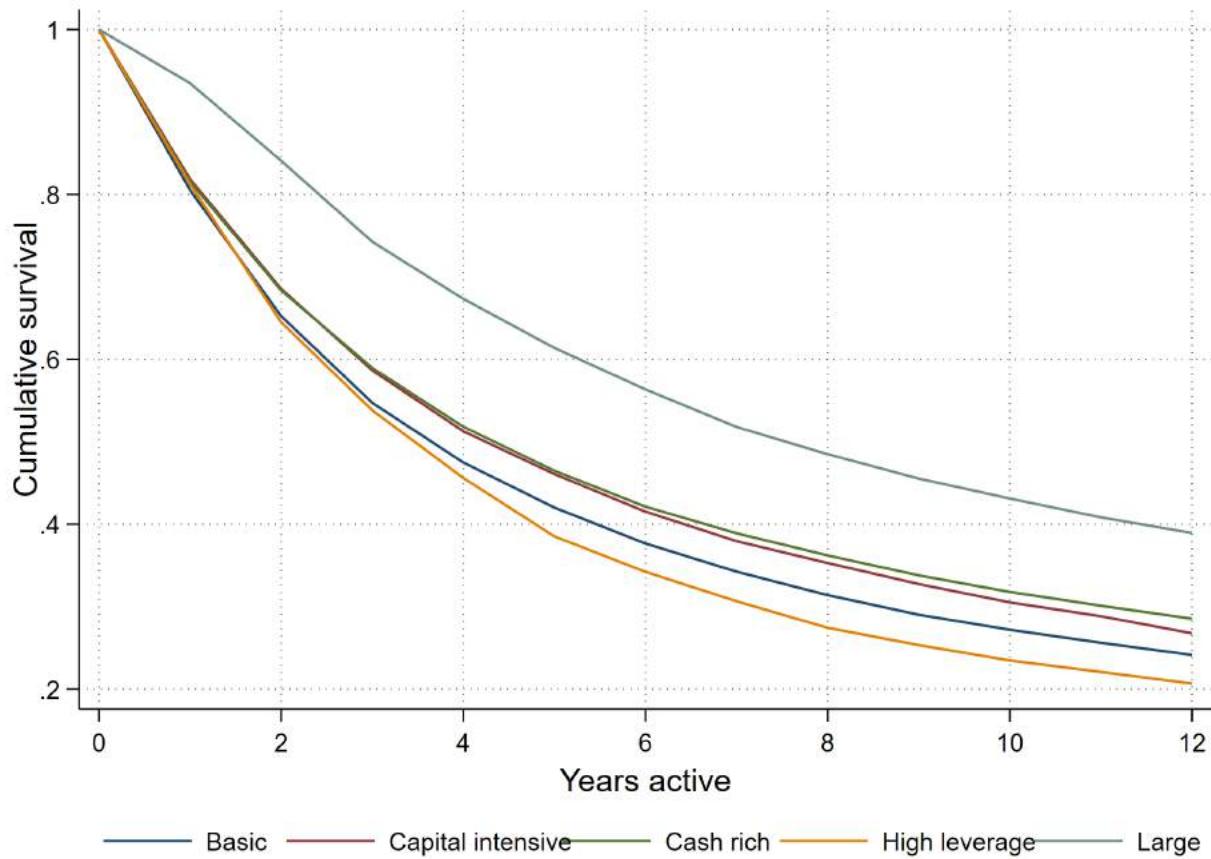
Notes: This figure shows for individual 1-digit industries and countries the distribution of the startup population across the five startup types. Shares are averaged over start-up cohorts 2010-18.

Figure A8: Startup characteristics at time of firm entry, by type and country



Notes: This figure shows for the five startup types, the country-level mean of the five cluster variables in the year of firm entry. The sample corresponds to the full panel.

Figure A9: Kaplan-Meier survival curves, by startup type



Notes: This figure shows Kaplan-Meier survival curves for the five startup types. The sample consists of the full panel for the 2000-07 startup cohorts. Due to data availability, Italy, the Netherlands and Spain are not included.

Table A1: Policy experiment: Correlation matrix

	Agg. Employment	Agg. Labor Productivity	Agg. TFP	No. Firms	Avg. Profits
Agg. Employment	1.00				
Agg. Labor Productivity	0.56	1.00			
Agg. TFP	0.12	0.30	1.00		
No. Firms	0.58	0.66	0.80	1.00	
Avg. Profits	0.72	0.61	0.74	0.98	1.00

Notes: Table showing correlation matrix between aggregate employment, aggregate labor productivity, aggregate total factor productivity, number of firms and average profits.

Table A2: Startup type and firm outcomes

	(1) Aggregate labor productivity	(2) Aggregate TFP (GMM estimation)	(3) Exit probability	(4) Wage per employee	(5) Average profit margin
Capital intensive	0.191*** (0.007)	0.011* (0.006)	-0.064*** (0.004)	0.749*** (0.129)	0.013*** (0.001)
Cash intensive	0.017*** (0.006)	0.030*** (0.005)	-0.023*** (0.003)	1.010*** (0.112)	0.014*** (0.001)
High leverage	-0.038*** (0.007)	-0.024*** (0.006)	0.003 (0.004)	-1.614*** (0.131)	-0.009*** (0.001)
Large	0.153*** (0.007)	0.043*** (0.006)	-0.144*** (0.004)	2.611*** (0.132)	-0.014*** (0.001)
Constant	3.431*** (0.004)	2.229*** (0.003)	0.809*** (0.002)	29.598*** (0.077)	0.049*** (0.001)
R-squared	0.902	0.981	0.610	0.930	0.653
N	8,259	5,903	9,092	8,677	8,607
Country × cohort FE	✓	✓	✓	✓	✓
Industry × cohort FE	✓	✓	✓	✓	✓
Country × Industry FE	✓	✓	✓	✓	✓
Age × Country FE	✓	✓	✓	✓	✓
Age × Industry FE	✓	✓	✓	✓	✓
Age × Cohort FE	✓	✓	✓	✓	✓

Notes: Table showing OLS regressions where the dependent variable is indicated in the column heading. The regressions are at the 1-digit industry level, we use the full panel and keep only those firms which have been active for between 5-8 years. Standard errors are in parentheses. *, **, *** indicate significance at the 10, 5 and 1 percent level, respectively.