THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITE DE TOULOUSE

Délivré par l'Université Toulouse Capitole

École doctorale : **Sciences de Gestion**

Présentée et soutenue publiquement par

Henri LUOMARANTA

le **9/9/2019**

Essays in Firm Dynamics, Ownership and Aggregate Effects

Discipline : **Sciences de Gestion**

Spécialité : **Comptabilité et contrôle**

Unité de recherche : **CRM (UMR 5303 CNRS)**

Directeur de thèse : **Milo BIANCHI**

JURY

Rapporteurs M, Matti, KELOHARJU, Professeur, AALTO UNIVERSITY
M, Angelo,SECCHI, Professeur, UNIVERSITE PARIS 1

Suffragants M, Matti, KELOHARJU, Professeur, AALTO UNIVERSITY
M, Angelo,SECCHI, Professeur, UNIVERSITE PARIS 1
Mme, Silvia, ROSSETTO, Professeur des universites, UNIVERSITE TOULOUSE 1
M, Milo, BIANCHI, maitre de conferences, UNIVERSITE TOULOUSE 1

Acknowledgements

First of all, I would like to thank my supervisor, Milo Bianchi, who has made this thesis possible. Your rigorous approach to research and curiosity has been both educational and inspiring. I also thank for the flexibility, patience, and constant support during this long process. I hope we can continue to work on many interesting projects in the future.

I thank my referees and committee members for reviewing this work. Matti Keloharju, Angelo Secchi, and Silvia Rossetto, I am grateful for your valuable contributions. I have received helpful and detailed comments that will improve and further my research.

Most of this work has been developed at the Finnish statistical office and has been inspired by the issues shared by economists and international statistical community. I am grateful for my employer for many interesting assignments and for the opportunity to link academic research and official statistics in a useful way. At Statistics Finland, Reetta Moilanen deserves heartfelt thanks, having strongly supported both this thesis and my professional development during the past years. Similarly, I am grateful for Mari Ylä-Jarkko, Hannele Orjala, Satu Elho and Saara Roine for supporting my research activities over the years. I also thank my past and present colleagues Jarkko Niemistö, Olli-Jussi Sonni, Lauri Saarinen, Salla Simola, Susanna Tåg, Eetu Toivanen for your support and useful comments.

I have had the pleasure to work on many cutting-edge projects with some remarkable colleagues, all exploring new data sources and registers to be transformed into useful new knowledge. I have especially been inspired by Peter Bøegh Nielsen, Samuli Rikama, and Boro Nikič, you have all provided intellectual leadership and supported this work.

A very special thanks should be directed to Paolo Fornaro, who has shared much of the research anxiety and the few highlights with me.

I started thinking about PhD studies in Paris, where I have been inspired by some exceptional individuals. Seyhun, Pankaj, Alejandro, Claire, Sabina, Sabrina, Mahreen, Marius, Thibaut, Martha, Guilherme, Robert and Diana. A major acknowledgement needs to be directed your way.

On the more immaterial level, I acknowledge Ville, Sampo and Johannes for your contributions during Junkua discussions and Enni for your invaluable support. Finally I thank my extended family Harry, Seija, Jarkko, Magda, Tuomas, Laura, Lauri, Aava, Hilla, Elia, I will always try to make you proud.

Abstract

Administrative registers maintained by statistical offices on vastly heterogeneous firms have much untapped potential to reveal details on sources of productivity of firms and economies alike.

It has been proposed that firm-level shocks can go a long way in explaining aggregate fluctuations. Based on novel monthly frequency data, idiosyncratic shocks are able to explain a sizable share of the Finnish economic fluctuations, providing support to the granular hypothesis.

The global financial crisis of 2007-2008 has challenged the field of economic forecasting, and nowcasting has become an active field. This thesis shows that the information content of firm-level sales and truck traffic can be used for nowcasting GDP figures, by using a specific mixture of machine learning algorithms.

The agency problem lies at the heart of much of economic theory. Based on a unique dataset linking owners, CEOs and firms, and exploiting plausibly exogenous variations in the separation of ownership and control, agency costs seem to be an important determinant of firm productivity. Furthermore, the effect appear strongest in medium-sized firms.

Enterprise group structures might have important implications on the voluminous literature on firm size, as large share of SME employment can be attributed to affiliates of large business groups. Within firm variation suggests that enterprise group affiliation has heterogeneous impacts depending on size, having strong positive impact on productivity of small firms, and negative impact on their growth. In terms of aggregate job creation, it is found that the independent small firms have contributed the most.

The results in this thesis underline the benefits of paying attention to samples encompassing the total population of firms in order shape more comprehensive policies. Researchers should continue to explore the potential of rich administrative data sources at statistical offices and strive to strengthen the ties with the official data producers.

Contents

Chapter 1

Introduction

1.1 Firm Dynamics, Ownership and Aggregate Effects

The link between micro-level behavior and aggregate outcomes has drawn the attention of economists for a long time. For example, modern large scale macroeconomic frameworks (such as DSGE) are based on so-called microfoundation, i.e. they model the optimization process of firms and individuals to derive the aggregate equilibrium conditions. Recently, the interest around the possible effects of firm-level performance on the economy has risen, with the works of Gabaix (2011) and Acemoglu et al. (2012). Gabaix (2011) formulates the granular hypothesis, i.e. he considers the size of firms as the key transmission mechanism of microeconomic (firm-level) shocks, due to the fact that the value added of many modern economies are characterized by having a fat-tailed distribution of firms (in terms of size), where the value added of few big companies accounts for a large fraction of the GDP, and diversification due to the large number of enterprises does not eliminate the impact of idiosyncratic disturbances.

Finland is a country that can provide the upper boundary of how large the granular effects can be. The granular hypothesis in Finland is the topic of the second chapter in this thesis, based on Fornaro and Luomaranta (2018). The results underline that micro level data can be an extremely useful source of information to understand aggregate fluctuations. From that fact, one can presume that such data can also potentially provide accurate predictions of the state of the economy in a timely fashion. *Nowcasting* aggregate economic variables using disaggregated microdata as predictors is the topic of the third chapter in this thesis.

Ownership and control are key characteristics of a firm. The instances where ownership is separated from control can create conflicts of interest and tension between the owner (principal) and the manager (agent). The so-called agency problem has been at the heart of much of the corporate finance literature since at least Berle and Means (1932) and Jensen and Mecling (1976). At the same time, while substantial work has been developed to investigate how firms' decisions (say, investment and financing) are shaped by agency conflicts, direct measures of agency costs are difficult to obtain. Extending the analysis beyond typical samples of large public firms, the fourth chapter in the thesis shows that the agency costs are important determinants of firm productivity of small firms, and are stronger in medium-sized private firms. The results highlight the importance of studying comprehensive datasets encompassing the entire universe of firms.

Firms vary along multiple dimensions in economically meaningful ways. In particular, firm heterogeneity has implications on capacity to sustain jobs, or on wealth generation. Size and ownership links between firms are the focus in the last two chapters of the thesis. I study how activities of business groups affect productivity and job creation within individual firms, drawing links to the aggregate effects. Again, I demonstrate that the results are different depending on which part of the firm population is analyzed, which is an observation that further underlines the importance of studying total populations. Within firm variation suggests that enterprise group affiliation has strong positive impact on productivity of small firms, and negative impact on their growth. In terms of aggregate job creation, it is found that the independent small firms have contributed the most, after breaking down job flows by business group membership status, age and size categories.

The results in the last two chapters provide further indication that almost any empirical study interested in firms should also consider their complex ownership structures.

The rest of this introductory chapter briefly summarizes the different chapters.

1.2 Aggregate Fluctuations and the Effect of Large Corporations: Evidence from Finnish Monthly Data

In this chapter, based on Fornaro and Luomaranta (2018), we investigate whether the granular hypothesis holds for the Finnish economy. In particular, we test if a sizable share of macroeconomic fluctuations is generated by firm-specific shocks to sales and productivity. We examine monthly firm-level data and find that the idiosyncratic shocks affecting large corporations explain around a third of business cycle fluctuations. This fact holds true both when we use the cross-sectional averages of sales and the estimated common factors, to control for common shocks. Moreover, we observe that the largest four corporations are the main drivers of this result. Finally, we detect a clear break in this relationship coinciding with the Great Recession. In particular, from 2010 onward the firm-level shocks lose their explanatory power. The findings of this paper point toward the importance of studying the granular hypothesis in a dynamic context, taking into account the possibility of breaks.

1.3 Nowcasting Finnish Real Economic Activity: a Machine Learning Approach

In this chapter, based on a joint work with Paolo Fornaro, we develop a nowcasting framework based on micro-level data in order to provide faster estimates of the Finnish monthly real economic activity indicator, the Trend Indicator of Output (TIO), and of quarterly GDP. In particular, we rely on firm-level turnovers, which are available shortly after the end of the reference month, to form our set of predictors. We rely on combinations of nowcasts obtained from a range of statistical models and machine learning methodologies which are able to handle high-dimensional information sets. The results of our pseudo-real-time analysis indicate that simple combinations of these models provides faster estimates of the TIO and GDP, without increasing substantially the revision error. Finally, we examine the nowcasting accuracy obtained by relying on traffic data extracted from the Finnish Transport Agency website and find that using machine learning techniques in combination with this big-data source provides competitive predictions of real economic activity.

1.4 Agency Costs and Firm Productivity

In this chapter, based on a joint work with Milo Bianchi, we explore how the separation between ownership and control affects firm productivity. Using Finnish administrative data on the universe of limited liability firms, we document a substantial increase in firm productivity when the CEO obtains majority ownership or when the majority owner becomes the CEO. We exploit plausibly exogenous variations to ownership and control structures, induced for example by changes in the CEO spouse's health status. Extending the analysis beyond typical samples of large public firms, we show that our effects are stronger in medium-sized private firms. We also investigate possible mechanisms and provide suggestive evidence that increased ownership boosts CEO's effort at work.

1.5 How Business Group Affiliation Improves Productivity of Small Firms: Evidence from Finnish Administrative Data

I inspect how joining a business group impacts firm productivity and job growth. Rich administrative data on the universe of Finnish limited liability firms reveals that joining a business group increases productivity, and decreases job growth within firms. This is driven by small firms that are mainly in the service sector. I provide suggestive evidence of mechanisms. I document changes in (key) employees, decrease in cost of capital and risk levels, and significant transfers of financial resources. Based on the results, the role of business groups in the economy might be most relevant in the context of small firms, which are usually not analyzed due to lack of data.

1.6 Job Creation and the Role of Dependencies

In this chapter, based on a joint work with Paolo Fornaro, we contribute to the extensive literature on the relationship between firm size and job creation, by examining the role of dependencies between enterprises. Using Finnish monthly data encompassing the population of Finnish private businesses, we calculate the gross job creation and destruction, together with the net job creation, for different size classes and industries. Importantly, we divide firms into a dependent (i.e. owned, at least partially, by a mother company) and independent category. We find that independent companies have shown a considerably higher net job creation, regardless of their size class. Once we control for age, we find that independent firms exhibit higher net job creation rates during the early years of their existence, but lower ones when they become older.

Chapter 2

Aggregate Fluctuations and the Effect of Large Corporations: Evidence from Finnish Monthly Data [1]

Paolo Fornaro*, Henri Luomaranta**

*Research Institute of the Finnish Economy, Finland

**Statistics Finland and TSM, University of Toulouse Capitole, Toulouse, France

Abstract

We investigate the effect of corporation-level shocks on the Finnish economy during the last 16 years. In particular, we test for the existence of the granular hypothesis, i.e. that a sizable share of macroeconomic fluctuations are generated by microeconomic shocks to large companies. We construct a dataset containing enterprise groups monthly sales and we find that the idiosyncratic shocks to large corporations explain around one third of business cycle fluctuations. This holds true both when we use the cross-sectional averages of sales and the estimated common factors to control for common shocks. Moreover, we observe that the largest four corporations in the dataset are the main drivers of this result. We also detect a significant break in this relationship with the Great Recession. In particular, after that period the corporation-level shocks lose their explanatory power. The findings of this paper point toward the importance of studying the granular hypothesis in a dynamic context, taking into account the possibility of breaks.

JEL Classifcation Code: C22, C55, E32

Keywords: Business Cycles, Granular Residual, Business Groups

[1] This chapter is based on a published article Fornaro and Luomaranta (2018) in Economic Modelling.

2.1 Introduction

The origins of business cycle fluctuations have been one of the most debated and explored topics in macroeconomics. Traditional frameworks, such as the real business cycle model of Kydland and Prescott (1982), identify the main shocks as being economy wide (e.g. aggregate productivity shocks) and having somewhat mysterious origins. Subsequently, Long and Plosser (1983) and the later literature (see, e.g., Horvath, 2000, and Conley and Dupor, 2003), consider sectoral disturbances in order to explain the fluctuations of the aggregate economy. In his seminal paper, Jovanovic (1982) proposed a model where microeconomic shocks are capable to generate business cycle variations.

Recently, the interest around the possible effects of firm-level performance on the economy has risen, with the works of Gabaix (2011) and Acemoglu et al. (2012). Gabaix (2011) formulates the granular hypothesis, i.e. he considers the size of firms as the key transmission mechanism of microeconomic (firm-level) shocks. Many modern economies are characterized by having a fat-tailed distribution of firms (in terms of size), where the value added of few big companies accounts for a large fraction of the GDP, and diversification due to the large number of enterprises does not eliminate the impact of idiosyncratic disturbances. An alternative explanation is provided by Acemoglu et al. (2012), who identify linkages between firms as being the main transmission mechanism. Related to this work, Frijters and Antić (2016) develop a theoretical model that connects downturns with the collapse of trade networks, through endogenous trade cycles.

di Giovanni et al. (2014) examine the universe of French firms and their annual sales growth, finding that the firm-level component is important in explaining aggregate fluctuations, and that this is mainly due to the firm linkages. Stella (2015) adopts similar methods as Foerster et al. (2011), to examine the granular hypothesis using quarterly U.S. firms' sales data. In particular, Stella (2015) uses a dynamic factor model to estimate the firms' idiosyncratic shocks and finds that the granular hypothesis does not hold for the U.S. economy.

In this work, we test the granular hypothesis of Gabaix (2011) for Finland, using monthly enterprise group data. Naturally, there have been multiple studies on the Finnish macroeconomy, especially with respect to the real business cycle. One notable example is Gorodnichenko et al. (2012), where the authors develop a dynamic general equilibrium model and they highlight how the collapse in the trade relationship with the Soviet Union contributed to the dramatic Finnish recession of 1991-1993 . Kuismanen and Kämppi (2010) study how fiscal policy shocks have affected Finnish real economic activity, finding that an increase in government spending leads to a crowding out effect on the private sector. Moreover there have been multiple works focused on the macroeconomic modeling of the Finnish economy (see, e.g., Tarkka, 1985, Lehmus, 2009 and Kilponen et al., 2016).

Despite intensive macroeconomic research, there have been relatively few studies relating microeconomic shocks to aggregate fluctuations, even though the Finnish economy seems to be one of the most extreme cases of granularity. In fact, many would argue that Finland is (or at least was) a single-firm economy, where Nokia activities represent an overwhelming share of GDP. For instance, Ali-Yrkkö et al. (2010) have shown that Nokia's production between the end of the 90's and the Great Recession has accounted for more than 2% of Finnish GDP and 10% of Finnish exports, with a peak of more than 20% of total exports during the 2000–2006 period. These considerations might explain why there has been relatively small research effort on analyzing the impact of large firms' shocks: when looking at the Finnish economy, researchers might focus on the role of Nokia and disregard the other large business groups. Given the importance of Nokia in terms of sheer size, during its golden years, this view might be justified. However, Nokia has faced a substantial drop in size and one can argue that the study of microeconomic shocks should extend to a wider group of companies. To reinforce this point, in Figure 2.1, we compute the sales herfindahl index for the top-57[2] enterprise groups in the Finnish economy. It is defined as the squared sum of the sales to GDP ratio of the companies we are interested in, and it can be interpreted as the degree of concentration for these enterprises. A higher index implies that these groups are accountable for a larger share of aggregate output. In Figure 2.1 (a), we report the scaled herfindahl index for Finland, together with the deflated monthly output measured by the Trend Indicator of Output (TIO). In Figure 2.1 (b) we report the sales herfindahl using the whole set of firms in our data against the same index obtained by excluding the top-4 corporations (in terms of average monthly sales) from the calculations.

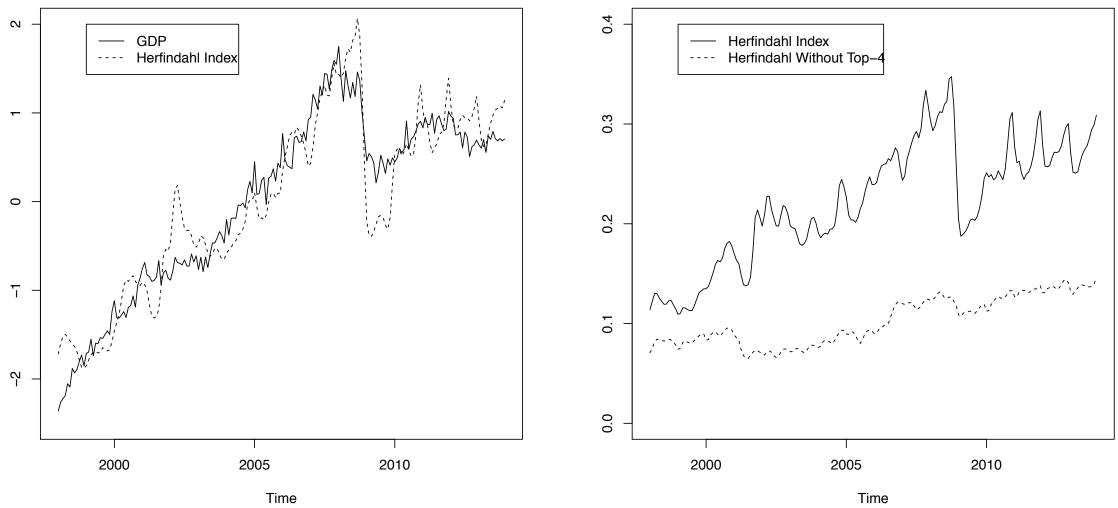

(a) Herfindahl index and Finnish monthly output

(b) Herfindahl with and without the top-4 corporations

Figure 2.1: The herfindahl index for the top-57 Finnish corporations

These plots give us some valuable preliminary insights about the dynamics of the Finnish economy in the last 16 years. First of all, the herfindahl index has not been stable over time, but has actually shown a substantial procyclicality (the correlation coefficient between TIO and the herfindahl index is 0.92). This supports a possible granular hypothesis for the Finnish economy, where aggregate fluctuations are heavily influenced by the success of large companies. The procyclicality of enterprise groups' sales is paired with the findings of Moscarini and Postel-Vinay (2012). In particular, they find that in the U.S. large firms have experienced a strong procyclicality in employment creation and

[2]We initially selected the top 100 Finnish firms by sales value. We then removed corporations which experienced extreme fluctuations or had data issues, ending up with 57 companies

destruction at business cycle frequencies compared to smaller companies.

Figure 2.1 (b) is also very informative. The solid line indicates the herfindahl index computed using the sales of the largest 57 Finnish enterprise groups, while the dashed line represents the same measure computed excluding the top-4 corporations from the data. While the two indices are fairly close during the late 1990's and early 2000's, we see a dramatic increase in the spread between them during the mid-2000's up until the Great Recession. Moreover, it is fairly evident that the herfindahl index shows a much more stable behavior when we remove the four largest corporations from the data, and that the top-4 companies in our dataset have been heavily affected by the Great Recession of 2008-2009. This suggests that few, very large, companies can be considered as a key factor in the Finnish economy and that their performance have been tightly linked with the business cycle fluctuations.

In this paper, we consider the approach of Gabaix (2011) using Finnish corporation-level sales. In particular, we estimate the granular residual and consider its impact on Finnish economic activity. One new key contribution of this research is the use of monthly data. Microeconomic level shocks are likely to have a large effect in the short run, but their impact on aggregate fluctuations might be attenuated when considering lower frequencies such as yearly data (which are commonly considered in previous work due to their availability). For example, a strike in an enterprise during a month can have a substantial effect on the aggregate output for that period, but might disappear when considering the whole year, due to the effect of temporal aggregation. Another advantage of this type of data is that it allows us to analyze the relationship of interest on a fairly short time span, without incurring in small samples problems. This means that we can verify the granular hypothesis on different subsamples and examine if events such as the Great Recession have affected it. In addition to the econometric analysis, we provide a short narrative where we examine how key events affecting large firms might have impacted the Finnish economy, using public sources. For simplicity, we restrict our analysis to the largest 57 Finnish companies.

Another important distinction from previous research is that we group together firms that belong to the same corporation, to better represent their influence on the Finnish market. Modern enterprise groups include hundreds of subsidiaries and disregarding them from the calculations would lead us to underestimate the actual influence of a company on the economy through its controlled firms. The empirical literature on internal capital markets and resource allocation suggests that projects and plants under common ownership of a parent company, have correlated investments and capital allocation (see the survey in Maksimovic and Phillips, 2013, and, for recent empirical evidence, Giroud and Mueller, 2015). Small and medium enterprises (SME) account for a large share of the gross value added of many economies (the EU28 average is 57.8% of the GDP). Even though most of these firms are independent, the gross value added generated by small and medium enterprises belonging to a large corporation, in Finland, accounts for around 50% of the value added produced by SMEs (see Airaksinen et al., 2015). This feature is present in many economies: one notable example is Germany, where the gross value added of dependent SMEs is around 43% of the total value added of small companies. These considerations should give a fairly clear idea of how important the dependences between large corporations and small firms are.

In addition to the monthly empirical analysis of the largest Finnish corporations, we use quarterly data obtained from the public reports of Nokia to verify how this individual corporation has affected the Finnish economy. The focus on Nokia is natural, given its renown importance in Finland, and can be an appealing research direction also for other countries where a small number of firms dominate the economy. One potentially interesting case could be South Korea, where Samsung and Hyundai sales account for 22% of GDP (see di Giovanni et al., 2014).

Our results show that the granular residual computed as in Gabaix (2011) is useful in explaining Finnish output fluctuations, accounting for around one third of the variation in monthly economic activity. To check the robustness of our findings, we use a dynamic factor model to compute the common component underlying our firm-level sales data. Moreover, we compute microeconomic shocks using labor productivity (relying on domestic employment) to see if our results are driven by the use of sales. In both cases, we verify that our findings are robust. We identify a clear break in the granular hypothesis with the Great Recession of 2008-2009 . While the granular residual accounts for a large share of output variation over the whole period up to the Great Recession, its explanatory power is greatly reduced in the years between 2010 and 2013. We find that removing the top-4 corporations changes the results dramatically. Notably, the explanatory power of the granular residual deteriorates substantially throughout the sample. We also decompose the variance of enterprise group-level shocks into a granular and a linkage component, finding that the granular component dominates throughout the sample, with the notable exception of the Great Recession. In the narrative analysis, using public sources, we cover a number of episodes where large Finnish companies have experienced extremely good (or poor) performance, and find that the aggregate economy was influenced by individual corporations.

The rest of the paper is organized as follows. In Section 2.2, we discuss the methodology underlying the estimation of the granular residual. In Sections 2.3 and 2.4 , we describe the data and report the main empirical results. Section 2.5 includes a brief narrative of the main events which have impacted the largest Finnish firms, with a special focus onto Nokia. Finally, Section 2.6 concludes.

2.2 Methodology

We start by introducing the methodology formulated in Gabaix (2011), in order to examine the granular hypothesis. While the original paper considers productivity, defined as the ratio between sales and the number of employees, we are skeptical in using this measure in our main analysis due to data limitations. However, we will examine productivity data as a robustness check. We will discuss this issue in more detail in Section 2.3.

To obtain the estimate of the granular residual, we compute the deviation of sales growth from the cross-sectional average $\bar{g}_t$. We use the resulting series to compute a weighted sum where the weights are determined by the enterprise group size. The granular residual is then given by:

$$\Gamma_t = \sum_{i=1}^{K} \frac{S_{i,t-1}}{Y_{t-1}} (g_{it} - \bar{g}_t), \tag{2.1}$$

where K is the number of companies, $S_{i,t-1}$ denotes the sales of enterprise group i at time $t-1$ and Y_{t-1} is the Finnish GDP. We compute $\bar{g}_t$ from both the top-57 Finnish corporations and from a much wider cross-section which includes more than 500 companies.

As an alternative to (2.1), we also compute the granular residuals by using the deviation of growth rate of sales from industry specific averages $\bar{g}_{I_i,t}$, where I_i indicates the industry in which firm i is active. The resulting formula is:

$$\Gamma_t^* = \sum_{i=1}^{K} \frac{S_{i,t-1}}{Y_{t-1}} (g_{it} - \bar{g}_{I_i,t}). \tag{2.2}$$

Once we obtain an estimate for the granular residual, we examine a model where we regress the year-on-year growth of the Finnish monthly economic activity indicator (TIO), denoted by y, on the granular residual and its lags:

$$y_t = \beta_0 + \beta(L)\mathbf{\Gamma}_t + u_t. \tag{2.3}$$

Here u_t is the error term, which we assume normally distributed and uncorrelated with the regressors, and $\mathbf{\Gamma}_t = [\Gamma_t, \Gamma_{t-1}, \cdots, \Gamma_{t-p}]$ includes the current granular residual computed using (2.1) or (2.2) and its p lags, which we select by using the Bayesian Information Criterion (BIC). Given that we are in a time series setting, the error term in (2.3) can be serially correlated. Therefore, we use the misspecification robust standard errors, adopting the Newey and West (1987) estimator, to conduct robust inference on the regression parameters.

We evaluate the explanatory power of $\mathbf{\Gamma}$ using the adjusted-R^2:

$$\text{Adjusted-}R^2 = 1 - (1 - R^2)\frac{T-1}{T-k-1},$$

where k is the dimension of $\mathbf{\Gamma_t}$ and T is the length of the time series.

As mentioned above, simple regressions do not manage to overcome the correlation-causality issue. To examine the robustness of our results, we provide a short narrative similar to the one of Gabaix (2011), where we examine large spikes (in absolute terms) in the TIO and consider shocks to Finnish enterprise groups as possible cause of these dramatic fluctuations. Related to this, in Section 2.4.4 we consider a short analysis of Nokia's implied contribution to the Finnish GDP growth using quarterly labor productivity data, obtained from the company quarterly reports. These contributions are computed using the measure

$$C_{Nokia,t} = \widehat{\beta}_{prod} dprod_{Nokia,t} \frac{sales_{Nokia,t-1}}{GDP_{t-1}} \tag{2.4}$$

where $prod_{Nokia,t}$ is the labor productivity of Nokia at time t, as measured in Gabaix (2011), $dprod_{Nokia,t}$ is its year-on-year growth rate and $\widehat{\beta}_{prod}$ is the coefficient obtained by regressing quarterly GDP year-on-year growth onto $dprod_{Nokia}$ and an intercept.

2.3 Data Description

For the first part of the analysis, we use enterprise group level data from Statistics Finland (the national statistical office). The main focus is on the monthly sales of the 57 most important corporations in Finland, over the years 1998–2013, which we refer to as top-57. The second data that is analyzed includes the sales of all the enterprises which have more than 250 employees on average, over the time span 1998–2013. This latter data is mainly used for robustness. Notice that the latter data is not aggregated to represent business groups.

It is important to notice that Gabaix (2011) focuses on firm-level productivity, while in this analysis, at least for the main part, we look at sales. This is not an unicum in the literature: both Stella (2015) and di Giovanni et al. (2014) use sales as main indicators, in order to identify microeconomic shocks. Our choice is dictated, as in the case of the two examples cited, by data availability issues. First of all, we do not have access to monthly value added, which would be the preferred measure to calculate labor productivity. However, Gabaix (2011) uses sales to compute productivity, so we can disregard this issue. A more problematic aspect is that we do not have access to the number of employees of Finnish corporations, outside of the Finnish territory. We are left with two choices: we can use the ratio of total sales to domestic employees or compute productivity as domestic sales divided by domestic workers. In both cases, we would have an extremely partial view of the actual performance of the enterprise groups we consider. The problem can be milder for the U.S., given the size of the economy, but the share of domestic activity for large Finnish multinational companies is fairly small (see Ali-Yrkkö et al., 2010, on this point). Given these considerations, we focus on sales, and leave the analysis of productivity as a robustness check, to be taken with a grain of salt. It is true that sales might have a larger risk of carrying an endogenous relationship between microeconomic shocks and aggregate fluctuations (i.e. firms' sales might be more affected by macroeconomic dynamics than productivity), but the fact that our results hold also after using a dynamic factor model, which should provide a more accurate estimate of the common components underlying our data and hence a better identification of the microeconomic shocks, gives us some reassurance in this regard.

For labor productivity we use the sales to domestic employment ratio, similar to Gabaix (2011), i.e.

$$prod_{i,t} = \frac{Sales_{i,t}}{employees_{Dom,i,t}}, \tag{2.5}$$

where $Sales_{i,t}$ refer to the total sales of the business group i in month t, and $employees_{Dom,i,t}$ is the number of domestic employees (expressed, in full time equivalents) for the same corporation. We then compute the year-on-year growth rate for this labor productivity measure and compute the granular residual as in (2.1). To analyze the role of Nokia, individually, we use quarterly productivity and net sales figures between 2003 and 2013, obtained from the the corporation's interim reports.

Statistics Finland forms the monthly turnovers by obtaining the contributions made by each enterprise to the tax authority. The mergers and acquisitions are also controlled for in the value of the comparison year (see the Appendix for a detailed explanation on the methodology). Thus, we are able to use monthly year-on-year growth rates that represent the organic growth of each company. The enterprise groups are split into four main industries of the economy using the Standard Industrial Classification (TOL 2008). [3] Below we report the number of groups belonging to each industry and how much they are accountable for the total turnovers.

Industry	Number of observations	Weight of turnovers (%)	Share of observations (%)
Finance	3	0.6	5.3
Construction	3	2.7	5.3
Manufacturing	25	66.2	43.9
Trade and Services	26	30.5	45.6

Table 2.1: Top-57 enterprises, weights by industry, January 2013

The top-57 sample is roughly evenly distributed between manufacturing (25 groups) and trade and services (26 companies). The sample also includes 3 construction and 3 finance[4] groups. By the turnover figures, the manufacturing companies are much larger, forming roughly 66.2% of the sample, while the corresponding share for trade and services is 30.5%. Construction companies account for only 2.5% of the total turnovers of the top-57 sample. In 2013, the value added generated by these 57 enterprise groups was around 34 billion Euro, accounting for roughly 17% of the Finnish GDP. In Table 2.2, we report similar descriptive statistics as in Table 2.1, for the larger dataset.

Industry	Number of observations	Weight of turnovers (%)	Share of observations (%)
Finance	18	3.7	3.5
Construction	25	2.8	4.8
Manufacturing	195	47.3	37.4
Trade and Services	283	46.2	54.3

Table 2.2: All large enterprises, weights by industry in January 2013

The second dataset contains ca. 50% of the Finnish GDP (in terms of gross value added), and includes 521 enterprises. The manufacturing industry has again the biggest weight in terms of turnovers, with 47%, but the relative importance of trade and services has now increased to 46% share of the overall sample turnovers. Construction and finance enterprises are less important, with a turnovers' share of 4.8% and 3.5%, respectively.

[3] The 5 digit statistical classification system for economic activities used in the European statistical system.

[4] The data does not include the large banks of the economy, which are problematic in terms of measuring turnovers. The corporations here are defined as "finance" by the statistical office by the value added of the headquarters, and are actual turnover producing corporations, especially after grouping the affiliated firms together.

In our analysis, we want to make sure that we capture the actual effect of a shock to a large company. If we would focus solely on individual firms (identified as single legal units in the register), we would underestimate their actual size by disregarding the potentially numerous subsidiaries. Instead, we combine the legal units to form corporation-level data through the enterprise group register information. Even though the final dataset is anonymous, the enterprise group register allows Statistics Finland to track the ownership and control relationships to achieve this systematically for the top-57 firms. Thus in our data, the mother company includes the sales of its subsidiaries, if they are relevant business units and the mother has full control over them.

2.4 Empirical Results

In this study, we use the Trend Indicator of Output (TIO) as main indicator of real economic activity. It is a monthly series constructed by Statistics Finland and provides the basis for the GDP flash estimates. In Figure 2.2, we report the TIO and the granular residual obtained using (2.2). The Finnish economy has experienced a moderate to high growth throughout the period going from the end of the 1990s until 2007, with a modest downturn around 2002-2003. However, the Great Recession had a dramatic impact on the economy, with year-on-year growth rate plunging to -10% and with economic activity growth still lagging behind toward the end of the sample.

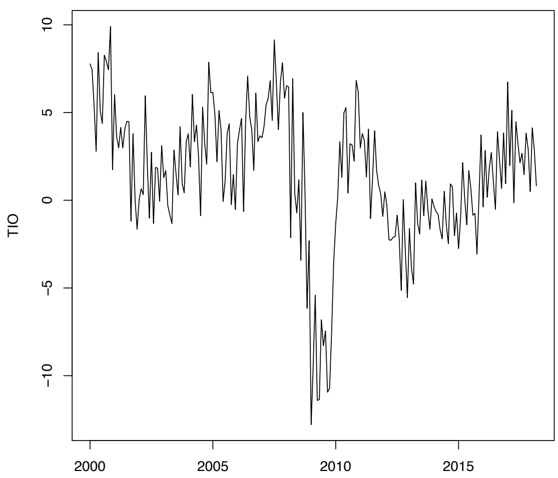

(a) Year-on-year growth rate of TIO

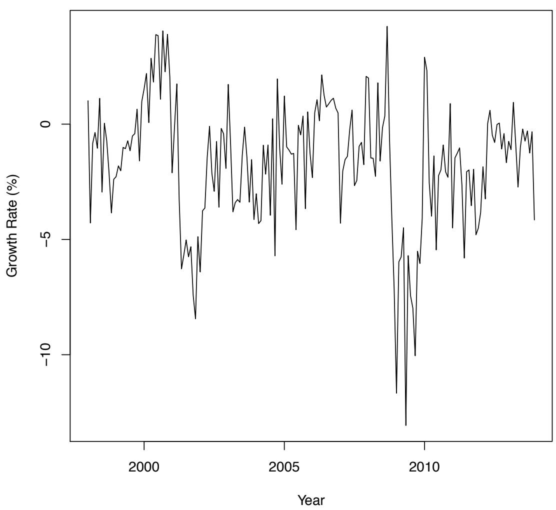

(b) The granular residual estimated using (2.2)

Figure 2.2: Finnish real monthly output and the granular residual based on the top-57 Finnish corporations

2.4.1 The granular residual in the Finnish economy

We start our empirical exercise by regressing the growth of TIO onto the granular residual computed using the full sample period available. In Table 2.3, we report results for $\bar{g}_t$ based on the top-57 companies, together with the ones obtained using the much bigger sample comprising all enterprises defined as large based on the Statistics Finland's criteria, i.e. having more than 250 employees.

	Full Sample		
	Top-57		*Large*
Model	1.	2.	3.
$\boldsymbol{\Gamma}_t$	0.821***	0.590***	0.671***
	(0.057)	(0.029)	(0.073)
$\boldsymbol{\Gamma}_{t-1}$		0.384***	
		(0.021)	
Constant	3.4***	3.7***	3.1***
	(0.2)	(0.2)	(0.3)
Observations	192	191	192
R^2	0.312	0.354	0.306
Adjusted-R^2	0.308	0.347	0.302
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01		

Table 2.3: Results of the regression of TIO onto the granular residual computed using (2.1), in the sample period going from January 1998 to December 2013. *Top*-57 and *Large* indicate how many firms were used to calculate $\bar{g}_t$. The lag length is selected based on the BIC.

Table 2.3 gives some important insights about the granularity of the Finnish economy. We find that the contemporaneous value of $\boldsymbol{\Gamma}_t$ and its first lag has a statistically significant positive effect on the TIO growth. This result follows the granular hypothesis intuition, which indicates that positive idiosyncratic shocks to large enterprise groups should benefit aggregate economic activity. Moreover, the adjusted-R^2 values (around 0.30) indicate that the granular residual is helpful in explaining a substantial part of real economic activity fluctuations, which supports the view of Finland as a granular economy.

It is interesting to notice that we get a quantitatively similar result as in Gabaix (2011), where the granular residual obtained from the top-100 US firms is able to explain around a third of GDP fluctuations. This might seem peculiar, given the smaller size of the Finnish economy, where larger corporations should be relatively more important in driving aggregate dynamics, compared to a big economy such as the U.S. one. We believe that the time frequency of our data might explain this aspect. We are dealing with monthly data, which are inherently more volatile compared to the annual ones used in Gabaix (2011), and hence harder to track. To verify this issue we compute the basic granular regression, using the contemporaneous residual and its first lag, using quarterly and annual data. The adjusted-R^2s we obtain are 0.39 and 0.44, for quarterly and annual data respectively. A similar specification in Gabaix (2011), using annual data, provides an adjusted-R^2 of 0.24, so we can see that for the Finnish case the granular residual is able to explain a substantially larger share of GDP growth. Moreover, we are going to see, in Section 4.2, that controlling for common shocks using a factor model, as suggested in Stella (2015), lowers the explanatory power of the granular residual, but the latter remains sizable. On the other hand Stella (2015) finds that, once he uses a dynamic factor model to control for common shocks, the explanatory power of the granular residual for U.S. GDP is virtually nonexistent. These considerations point toward a quantitatively larger importance of the granular residual in explaining aggregate fluctuations for the Finnish economy, compared to the U.S. case.

Using industry demeaning might give us a more appropriate estimate of the granular residual, so we re-estimate the linear regression model using formula (2.2) to compute the granular residual. In Table 2.4, we report the results using demeaning based on the top-57 firms and the larger dataset, for the full time period. These regression results give us similar findings as the ones reported in Table 2.3. The granular residual, together with its first two lags, has positive and statistically significant coefficients, and it is able to explain a considerable share of Finnish business cycle fluctuations. However, the R^2s are slightly lower than the ones in Table 2.3.

	Full Sample			
	Top-57		*Large*	
Model	1.	2.	3.	4.
$\mathbf{\Gamma}^*_t$	0.600***	0.401***	0.557***	0.448***
	(0.028)	(0.015)	(0.057)	(0.028)
$\mathbf{\Gamma}^*_{t-1}$		0.252***		0.150***
		(0.0078)		(0.016)
$\mathbf{\Gamma}^*_{t-2}$		0.236***		
		(0.0071)		
Constant	3.0***	3.5***	2.7***	2.7***
	(0.2)	(0.2)	(0.2)	(0.2)
Observations	192	190	192	191
R^2	0.239	0.333	0.252	0.259
Adjusted-R^2	0.235	0.322	0.248	0.252
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01			

Table 2.4: Results of the regression of TIO onto the granular residual computed using (2.2), in the sample period going from January 1998 to December 2013. *Top*-57 and *Large* indicate how many firms were used to calculate $\bar{g}_t$. The lag length is selected based on the BIC.

One peculiar characteristic about the Finnish economy is the presence of a single large company, Nokia, which has been shown (see Ali-Yrkkö et al., 2010) to contribute considerably to GDP growth. In the light of this information, it is interesting to see if the explanatory power of the granular residual is concentrated on few very large corporations. In Table 2.5, we report the results obtained by regressing TIO on the granular residual computed after removing the top-4 companies from the dataset. We report results for both specifications (2.1) and (2.2).

	Full Sample			
	Top-57	*Large*	*Top-57*	*Large*
Model	1.	2.	3.	4.
$\mathbf{\Gamma}_t$	0.57**	0.503**		
	(0.11)	(0.14)		
$\mathbf{\Gamma}_t^*$			0.56**	0.47**
			(0.11)	(0.14)
Constant	2.5**	2.4**	2.5***	2.4***
	(0.46)	(0.52)	(0.46)	(0.52)
Observations	192	192	192	192
R^2	0.052	0.051	0.052	0.050
Adjusted-R^2	0.047	0.045	0.047	0.045
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01			

Table 2.5: Results of the regression of TIO onto the granular residual computed using (2.1) and (2.2), in the sample period going from January 1998 to December 2013 and removing the top-4 corporations from the dataset. *Top*-57 and *Large* indicate how many firms were used to calculate $\bar{g}_t$.

Columns 1. and 2. in Table 2.5 include the results for the model based on the granular residual computed using (2.1), while 3. and 4. contain the parameter estimates and the R^2s for the industry demeaning specification. The outcomes of these regressions indicate that the four largest corporations have been a key driving factor behind the granularity of the Finnish economy. Removing them from the computation of the granular residual leads to much lower R^2s, indicating no substantial explanatory power of the microeconomic shocks (even though their regression coefficients remain statistically significant).

Finland has undergone a deep recession in the 2008-2009 period, followed by sluggish growth and further drop in GDP up until the end of our sample. It seems that the Great Recession represents a breaking point for the Finnish economy and for many of its largest enterprises (e.g. Nokia). Given this consideration, we analyze a pre-recession period from January 1998 until December 2007 and a sample covering the remaining years up to December 2013. We also examine the data from January 2010 until the end of the sample, to disregard the effects of the Great Recession. In tables 2.6 and 2.7, we report the results for regression (2.3) using the pre- and post-recession subsamples, respectively.

	Pre-2008			
	Top-57	*Large*	*Top-57*	*Large*
Model	1.	2.	3.	4.
$\mathbf{\Gamma}_t$	0.455^{***}	0.348^{***}		
	(0.007)	(0.0061)		
$\mathbf{\Gamma}_t^*$			0.423^{***}	0.313^{***}
			(0.0066)	(0.0059)
Constant	4.4^{***}	4.4^{***}	4.0^{***}	3.8^{***}
	(0.1)	(0.1)	(0.08)	(0.1)
Observations	120	120	120	120
R^2	0.218	0.214	0.205	0.194
Adjusted-R^2	0.212	0.207	0.198	0.187
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01			

Table 2.6: Results of the regression of TIO onto the granular residual computed using (2.1) and (2.2) in the sample period going from January 1998 to December 2007. *Top*-57 and *Large* indicate how many firms were used to calculate $\bar{g}_t$. The lag length is selected based on the BIC.

The results indicate a stark contrast between the pre- and post-Great Recession period, in relation to the granular hypothesis. Before 2008, we find that the granular residual is able to explain a moderate chunk of real output variations, with adjusted-R^2s consistently around 0.2. Moreover, the coefficients associated to $\mathbf{\Gamma}$ are highly significant and positive, even though lower than for the whole sample. On the other hand, as seen in Table 2.7, the Great Recession changes the results dramatically. If we include the economic decline of 2008-2009 in the analysis, we find that the granular residual is able to explain an even greater share of real activity growth, with R^2 reaching 0.4, and estimated coefficients reaching substantially higher values (being statistically significant as well), compared to the results in Table 2.6. However, by looking at the bottom panel in Table 2.7, the granular hypothesis does not seem to hold for the years after the Great Recession. The share of explained variance becomes very small for all specifications and the coefficients become negative (even though they remain statistically significant). Based on these results, it seems that the bulk of the relationship between the granular residual and the Finnish economy, during the 2008–2013 period, is generated during the Great Recession.

This break in the explanatory power of the granular residual is extremely interesting. We can only give a tentative reasoning for this phenomenon, but one of the main suspect is Nokia, specifically its drop in importance relative to the aggregate economy. In Section 2.4.4, we focus more on the role the telecommunication giant had in driving the Finnish economy. In Figure 2.4 (a), in Section 2.4.4, we depict the share of Nokia sales to GDP and notice a dramatic and steady drop after the Great Recession. Interestingly, after the large post-recession drop, there is a fairly stable period which is interrupted around 2011, where the company sales to GDP ratio restarted falling. The drop in the relatively importance of Nokia seems to be temporally linked with the loss of explanatory power of the

	Post-2008			
	Top-57	*Large*	*Top-57*	*Large*
	1.	2.	3.	4.
$\mathbf{\Gamma}_t$	0.892^{***}	0.951^{***}		
	(0.037)	(0.012)		
$\mathbf{\Gamma}_t^*$			0.455^{***}	0.778^{***}
			(0.034)	(0.02)
Constant	1.1^{**}	0.8^{**}	0.5	0.7^{*}
	(0.07)	(0.07)	(0.03)	(0.04)
Observations	72	72	72	72
R^2	0.391	0.607	0.183	0.500
Adjusted R^2	0.382	0.601	0.171	0.493
	Post-2009			
$\mathbf{\Gamma}_t$	-0.34^{**}	0.167^{**}		
	(0.03)	(0.05)		
$\mathbf{\Gamma}_t^*$			-0.372^{**}	-0.183^{**}
			(0.0163)	(0.05)
Constant	-0.1634^{**}	0.5^{**}	-0.5^{***}	0.2^{***}
	(0.01)	(0.01)	(0.05)	(0.09)
Observations	48	48	48	48
R^2	0.057	0.014	0.011	0.014
Adjusted-R^2	0.037	-0.008	-0.09	-0.008
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01			

Table 2.7: Results of the regression of TIO onto the granular residual computed using (2.1) and (2.2) in the sample period going from January 2008 to December 2013. *Top*-57 and *Large* indicate how many firms were used to calculate $\bar{g}_t$. The lag length is selected based on the BIC.

Granular residual, and this fact can offer a possible explanation in the break we found in the data. However, we have seen, in the Introduction, that the herfindahl index for the largest Finnish firms is roughly the same as in the pre-2008 period, so, statistically speaking, the Finnish economy has a similar degree of concentration before and after the break, and we might (ex-post) expect that the granular hypothesis holds throughout the sample. This can point toward the need to analyze different variables which can indicate the influence of large companies on the economy, such as investments.

In tables 2.8 and 2.9, we report the results of regressions of TIO onto the granular residual, while excluding the top-4 corporations from the computation of the latter. The evidence in these two tables shows a striking shift in the influence of the largest corporations on the Finnish economy. In the pre-recession period we find that removing the 4 largest corporations substantially eliminates the explanatory power of the granular residual, which is in line with what we found with the full sample. If we look at the post-2008 period, it seems that our results depend on how we define the granular residual. If we use the larger panel of firms to estimate the cross sectional average, then the granular residual keeps its explanatory power (albeit getting lower R^2s) even if we remove the top-4 firms. If we consider the data in the 2010-2013 period, we see that excluding the largest corporations does not change the results significantly, i.e. the granular residual is unable to explain a substantial share of economic fluctuations. Interestingly, it seems that the R^2 of certain specifications actually increases after we remove the largest corporations, even though it remains fairly low.

	Pre-2008			
	Top-57	*Large*	*Top-57*	*Large*
	1.	2.	3.	4.
$\mathbf{\Gamma}_t$	0.35***	0.28***		
	(0.024)	(0.021)		
$\mathbf{\Gamma}_t^*$			0.36**	0.28**
			0.(0.028)	(0.021)
Constant	4.14***	4.1***	4.14***	4.12***
	(0.1)	(0.2)	(0.0002)	(0.0002)
Observations	120	120	120	120
R^2	0.074	0.075	0.074	0.075
Adjusted-R^2	0.066	0.067	0.066	0.067
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01			

Table 2.8: Results of the regression of TIO onto the granular residual computed using (2.1) and (2.2), in the sample period going from January 1998 to December 2007 and removing the top-4 corporations from the set of firms. *Top*-57 and *Large* indicate how many firms were used to calculate $\bar{g}_t$. The lag length is selected based on the BIC.

To summarize, the shocks of large enterprise groups seem to have a substantial effect on real economic activity on all sample periods considered, except for the post-2009 years. Moreover, we find that the explanatory power of the granular residual is concentrated on the four largest corporations in the data, at least until the Great Recession. In the years after the economic downturn this relationship has deteriorated. In the narrative analysis (Section 2.5), we discuss how the poor performance of some of the largest Finnish companies might have affected this result.

2.4.2 Robustness check: factor analysis

As argued by Stella (2015), using the cross-sectional average of sales to control for the overall economic condition of the economy can be inappropriate. The dynamic factor model of Stock and Watson (2002a) provides an alternative approach to estimate the common component underlying the enterprise groups data. Stella (2015) shows that once the estimated factors are used to obtain the idiosyncratic shocks to firms, the granular hypothesis does not hold anymore for U.S. data.

We use the factor estimation proposed in Doz et al. (2012) to obtain the common component

	Post-2008			
	Top-57	*Large*	*Top-57*	*Large*
	1.	2.	3.	4.
$\mathbf{\Gamma}_t$	0.73	2.25**		
	(0.55)	(0.30)		
$\mathbf{\Gamma}_t^*$			0.81	2.16**
			(0.52)	(0.38)
Constant	−0.28	0.70	−0.21	0.60
	(0.83)	(0.67)	(0.75)	(0.72)
Observations	72	72	72	72
R^2	0.04	0.40	0.05	0.34
Adjusted-R^2	0.025	0.38	0.03	0.33
	Post-2009			
$\mathbf{\Gamma}_t$	−0.33**	−0.03		
	(0.03)	(0.30)		
$\mathbf{\Gamma}_t^*$			−0.80**	−0.15
			(0.10)	(0.28)
Constant	−0.16	0.40	−0.25	0.32
	(1.04)	(1.13)	(1.32)	(1.1)
Observations	48	48	48	48
R^2	0.11	0.01	0.10	0.02
Adjusted-R^2	0.09	-0.02	0.08	-0.019
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01			

Table 2.9: Results of the regression of TIO onto the granular residual computed using (2.1) and (2.2), in the sample period going from January 2008 to December 2013 and removing the top-4 corporations from the set of firms. *Top*-57 and *Large* indicate how many firms were used to calculate $\bar{g}_t$. The lag length is selected based on the BIC.

underlying our data and subsequently use it to compute the company-level shocks. In particular, denote the estimated factors, extracted from our dataset containing K enterprise groups, of dimensionality $T \times r$ as F, and the $r \times K$ loadings matrix as Λ. The common component underlying the enterprise group-level shocks is given by $C_t = F_t\Lambda$ and the granular residual is obtained by

$$\Gamma_t^F = \sum_{i=1}^{K} \frac{S_{i,t-1}}{Y_{t-1}}(g_{it} - C_t). \tag{2.6}$$

We estimate equation (2.3) using Γ^F to see if the granular hypothesis holds given the factor model specification and we analyze the relationship on both the whole sample and on the pre- and post-Great Recession periods. For the sake of readability, we report here the results of the static specification based on factors extracted from the larger dataset including all enterprise groups with more than 250 employees. The number of factors is selected by using one of the criteria formulated in Bai and Ng

(2002b), which suggests a very conservative specification with only one factor selected. For robustness, we estimate up to ten factors and find that only the first factor presents a strong correlation with real economic activity growth (0.71).

Table 2.11, confirms the results obtained so far. The granular residual seems to explain a substantial share of fluctuations of real output, even though the adjusted-R^2 values are slightly lower for the whole sample. We again find that the granular hypothesis has a clear break after the Great Recession, where the granular residual becomes statistically insignificant and the resulting R^2 is extremely low. Estimating the granular regression after using up to ten factors does change slightly our conclusions[5]. If the number of factors is fairly low (up to three), the results are qualitatively similar to the ones reported in Table 2.11, albeit with slightly different R^2s values. If we adopt ten factors, we get that the explanatory power of the resulting granular residual remains strong only for the *pre*-2008 sample, presenting low R^2s during the Great Recession and over the whole sample. The selection of the optimal number of factors is a sensitive topic, especially in a setting like the one presented in this analysis. If we use a large number of factors, we risk to include pervasive idiosyncratic shocks into the common component underlying our data, but on the other hand, it is arguable that using too few factors would lead us to underestimate the real common forces driving the firms' turnovers. Nevertheless, it is important to underline that the granular residual retains its explanatory power until the beginning of the Great Recession, regardless of the number of factors selected.

One reason between the discrepancy between our results and the ones obtained in Stella (2015) can stand in the type of data we use, i.e. firm-level versus enterprise group-level data. As the author points out, and as discussed in the previous paragraph, the statistical factor model used in the analysis of this subsection can have problems in distinguishing the common component from the propagation of idiosyncratic shocks to the rest of the firms. The use of corporations, instead of firms, in the analysis can alleviate the issue: grouping legal units together in the same mother company leads us to control for possible propagations of idiosyncratic shocks within an enterprise group. Stella (2015) uses the structural approach of Foerster et al. (2011) to filter out the possible propagations of firm-level shocks, using input-output relationships. However, because of data constraints, he has to rely on sectoral data to calibrate the input-output matrix and this might affect his results.

2.4.3 The role of linkages

While the focus of this study is on the granular hypothesis, the literature has also been interested in the effect of firm-level shocks to aggregate fluctuations through the so-called network channel (see Acemoglu et al., 2012). In practice, shocks to single enterprise groups can have substantial impact on the business cycle through the links between companies, e.g. through the intermediate goods supplier-user relationships. If these linkages would be important, we should observe a strong covariance in the company-level shocks $\epsilon_{it} = (g_{it} - \bar{g}_t)$. To examine the relative importance of the granular versus linkage channel, we follow the approach of di Giovanni et al. (2014), i.e. we decompose the

[5]To keep the analysis contained, we do not report the results for the richer factor specifications, however they are available upon request.

	Whole Sample	*Pre*-2008	*Post*-2008	*Post*-2010
	(1)	(2)	(3)	(4)
Γ_t^F	0.601***	0.384***	0.871***	-0.231
	(0.085)	(0.065)	(0.121)	(0.207)
Constant	1.304***	3.355***	−1.990***	0.788
	(0.275)	(0.203)	(0.416)	(0.522)
Observations	192	120	72	48
R^2	0.209	0.229	0.424	0.026
Adjusted-R^2	0.205	0.222	0.416	0.005
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01			

Table 2.11: Results of the regression of TIO onto the granular residual computed with (2.6) and using the dataset including the largest 57 corporations.

variance of the enterprise groups shocks (the top-57 sample) in two components. The enterprise group specific volatility can be written as $\sigma_t^2 = \sum_j^K \sum_i^K Cov(w_{j,t-1}\epsilon_{jt}, w_{i,t-1}\epsilon_{it})$. We can subsequently use the following decomposition, suggested in Carvalho and Gabaix (2013) and di Giovanni et al. (2014), for the variance of the microeconomic shocks:

$$\sigma_t^2 = \sum_i^K Var(w_{i,t-1}\epsilon_{it}) + \sum_{j \neq i}^K \sum_i^K Cov(w_{j,t-1}\epsilon_{jt}, w_{i,t-1}\epsilon_{it}), \tag{2.7}$$

where $w_{i,t-1} = \frac{S_{i,t-1}}{Y_{t-1}}$, implying that the terms in inside the variance and covariance operators are the granular residuals of the individual corporations[6]. The first summation in (2.7) is the granular component of the variance-covariance matrix of the shocks, which we denote hereafter $GRAN$. The second summation represents the covariance of the enterprise groups shocks and we label it $LINK$. Figure 2.3 presents the variance decomposition for the whole time period. To make the figure more readable, we have aggregated the series at the annual level. As we can see, the $GRAN$ component has been the dominant one for almost the entire sample, with a notable exception in 2009. This might point out that during the Great Recession the effect of corporation-level shocks has been transmitted through the linkages between firms, instead of separated individual shocks to large companies. Overall, the granular component of the variance of the microeconomic shock is predominant, which is in contrast to what has been found for the French economy in di Giovanni et al. (2014). At the monthly level, $GRAN$ accounts on average for 70% of total variance over the whole sample period, dropping to 60% during the Great Recession and with a contribution above half of total variance for almost 90% of the sample. It would be interesting to see how much this result is affected by our decision to use business groups instead of firms as the unit of observation. This could be done by replicating the analysis on firm-level data and possibly including smaller enterprises. This is, however, left for future work.

The fact that we use contemporaneous correlations with monthly data might lead us to disregard

[6] We have also estimated this decomposition using fixed weights, in line with di Giovanni et al. (2014). This alternative specification yields very similar results and hence it is not reported.

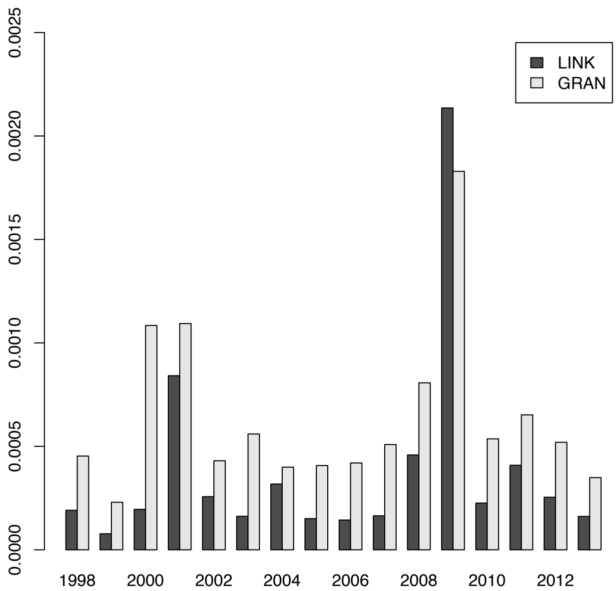

Figure 2.3: Contribution of $GRAN$ and $LINK$ component to the enterprise shocks variance.

linkages between corporations which have an effect with a lag. For example a negative shock to a firm which is supplier for a number of corporations might be transmitted after a certain number of months. To make sure that the prevalence of granularity in the variation of idiosyncratic errors is not due to the use of monthly data, we perform the analysis after aggregating the sales data at the annual level. The results are very similar to the ones reported in Figure 2.3, i.e. the $LINK$ components is predominant only in 2009. One additional interesting robustness check lies in separating our data into different industries. The fact that the contribution of the $LINK$ component becomes large during the Great Recession might indicate the spread of idiosyncratic shocks within a particular industry, such as the financial one. We estimate the decomposition in (2.7) using the data of different industries separately. We find that the trade and services industry, as well as the manufacturing one, show a similar pattern as the one in Figure 2.3. On the other hand, the decomposition for the construction and financial sectors evidences an even more prevalent role of the $GRAN$ component, even during the Great Recession.

2.4.4 The role of the size class

Next, we divide our sample of top-57 companies into three groups based on their average sales, each group containing 19 corporations. We denote the largest subset of corporations as Large (corporations 1 to 19), the second subset as Medium (20 to 38) and the last one as Small (39 to 57)[7], plus the subset containing the four largest corporations in the dataset, denoted as Giant. We then estimate the

[7]It is important to point out that even the companies in group Small and Medium are corporations with high sales and a large number of employees, in absolute terms.

following regressions

$$y_t = \beta_0 + \beta_i \mathbf{\Gamma}_{i,t} + u_t, \tag{2.8}$$

where $\mathbf{\Gamma}_{i,t}$ indicate the granular residual, and possibly its lags, of subset i at time t. Table 2.12 reports the adjusted-R^2s for the whole 1998-2013 sample and the pre- and post-recession periods.

Size Class (N. Firms)	Whole Sample	Pre-2008	Post-2008	Post-2009
Giant (4)	0.34	0.21	0.43	0.01
Large (19)	0.31	0.21	0.38	0.036
Medium (19)	0.03	0.07	0.07	0.03
Small (19)	0.002	0.09	0.16	0.003
Observations	192	120	72	48

Table 2.12: Adjusted-R^2 for different subsets of the top-57 corporations, divided by size. The number of lags used in the granular regression is based on the BIC.

Table 2.12 give us some interesting insights. It seems, as already pointed in the main analysis, that the explanatory power of the granular residual is substantially based on a very small number of large corporations. While the subsets Medium and Small do not carry any significant explanatory power, the granular residual based on the subset containing only the largest Finnish corporations is able to account for a large share of output fluctuations. Indeed, the top-4 enterprise groups seem to be the best at explaining TIO growth. The Great Recession provides an interesting shift, where all the subsets of corporations are unable to explain business cycle fluctuations after the end of the Great Recession. However, the granular residual obtained from the smaller corporations in our top-57 data is able to explain a moderate share of TIO fluctuations during the Great Recession.

These findings are also in line with what we described in Figure 2.3, where we can see that in 2009 the covariance component of the microeconomic shock dominates the granular component. Therefore, we can argue that during the Great Recession the main channel through which the company-level shocks were transmitted was through linkages, leading to a higher explanatory power of the relatively smaller companies, during this period. While this analysis is able to give us some interesting findings, it must be complemented with a more narrative approach (presented in the next section), where we can examine single episodes of granularity and their effect on economic fluctuations

2.4.5 Productivity analysis and the Nokia's case

As we mentioned before, the original study on the granular hypothesis, by Gabaix, is focused on productivity shocks. There are multiple studies in the literature that have used sales, due to data limitations, and the exercise presented here follows this practice. However, we can still calculate productivity based on domestic employment, giving us the possibility of conducting an analysis more in line with Gabaix (2011). In particular, we use (2.5) and compute the granular residual in a similar fashion as for sales, i.e. using formula (2.1). First, in Table 2.13, we report the results from regressing TIO onto the granular residual obtained from the total sales to domestic employment ratio, together with the domestic sales based productivity measure, using both a static and dynamic specification. The cross-sectional average is based on the top-57 firms.

The results in Table 2.13, columns 1. and 2., confirm the findings we obtained throughout our analysis. Firm-level shocks, based on the total sales to domestic FTEs productivity measure, present a positive and statistically significant regression coefficient, together with a fairly large R^2 value. On the other hand, once we use domestic sales, the explanatory power of the granular residual is significantly reduced, reflecting the fact that the Finnish market is fairly small, and that the largest companies are highly export oriented, thus deriving most of their value added from exports. Looking at our dataset for confirmation, we find that domestic sales share of total sales in the top-57 sample is 55%, averaged over the whole period in consideration. However, Nokia is a notable exception. Looking at the publicly available annual report for the company, in 2007, we see that the share of domestic sales to total sales is much lower. In particular, it corresponds to 0.6% in 2007, 0.94% in 2006 and 0.96% in 2005. While the average share of domestic sales for the largest companies is considerable, this measure becomes extremely small for what is widely considered historically as the most important company of the Finnish economy. As a consequence, the domestic demand cycles should have only a minor influence on the productivity dynamics of the largest firms, even though the idiosyncratic shocks (as reflected in the granular residual) appear to have a significant impact on domestic aggregate output. We interpret this evidence as lending support to the economic significance of our obtained results.

Full Sample				
	Total Sales/Emp.		*Domestic Sales/Emp.*	
Model	1.	2.	3.	4.
$\mathbf{\Gamma}^*_t$	0.59***	0.39***	0.23***	0.14*
	(0.06)	(0.08)	(0.049)	(0.06)
$\mathbf{\Gamma}^*_{t-1}$		0.28***		0.13*
		(0.08)		(0.06)
$\mathbf{\Gamma}^*_{t-2}$				
Constant	3.07***	3.21***	2.37***	2.7***
	(0.26)	(0.25)	(0.28)	(0.2)
Observations	192	190	192	191
R^2	0.34	0.38	0.10	0.12
Adjusted-R^2	0.33	0.37	0.10	0.11
Note:	Robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01			

Table 2.13: Results of the regression of TIO onto the productivity granular residual computed using (2), in the sample period going from January 1998 to December 2013. The lag length is selected based on the BIC.

We conclude the empirical analysis by looking specifically at Nokia, using publicly available data obtained from the company's interim reports. In particular, we compute quarterly labor productivity using (2.5) and we are able to obtain the total number of employees working for the company (not only the domestic ones), giving us a fairly complete view of the actual size of the group. We estimate the regression of the year-on-year growth rate of quarterly GDP onto Nokia's labor productivity growth starting from 2004, due to data limitations. The resulting R^2 is 0.25, indicating that the labor productivity of Nokia is able to explain around a quarter of GDP fluctuations during the period 2004–2013. The estimated regression coefficient is statistically significant at the 1% level and it is around 0.71. In Figure 2.4, we report the Nokia's sales to GDP ratio and the labor productivity contribution to GDP growth, computed using (2.4).

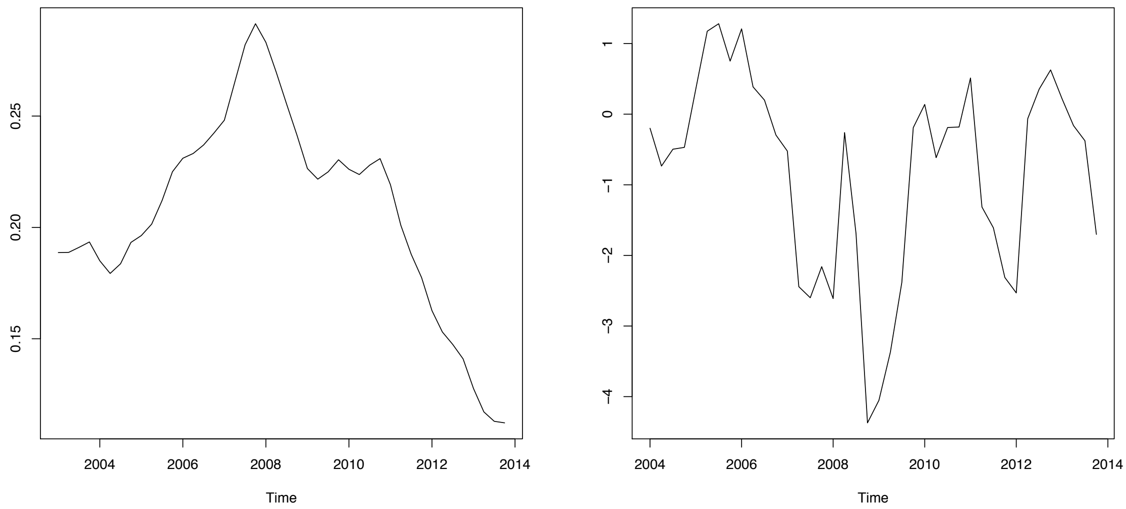

(a) Nokia's sales share of GDP

(b) GDP and implied Nokia's productivity contribution

Figure 2.4: Relationship between Nokia productivity and sales and Finnish GDP.

Figure 2.4 (a) shows how the importance of Nokia in the Finnish economy has shifted dramatically in the recent years. The sales to gdp ratio reaches its peak around 2007, where Nokia's sale represent roughly 30% of GDP. However, since 2008 the company has faced a continuous decline, with the ratio reaching its lowest values at 11%. Notice that in this analysis we are not considering the sale of the mobile division to Microsoft. Figure 2.4 (b) give us another insight on how important Nokia has been for the Finnish economy. Between the second quarter of 2005 and the first quarter of 2006, the implied contribution to GDP of Nokia's productivity is around 1.8 percentage points, accounting (on average) for around 60% of GDP growth in this period. Even more surprising is how the drop in productivity during the Great Recession seems to have affected the aggregate economy. Nokia's implied contribution during 2009 is -3.72% points on average, accounting for around 41% of the GDP drop during the year. While these figures give an idea of how large the influence of Nokia has been during the last 10 years, we have to keep in mind that there are various caveats behind this result. The productivity measure we use, while handy in terms of computation and data availability, is not the most sophisticated one and alternatives like the Solow residual can be taken into consideration.

Moreover, we are considering labor productivity without removing the cross sectional average, as we did with the granular residual and this can overestimate the correlation between Nokia's productivity and GDP growth. This is due to the lack of data on total employment for the rest of the corporations in our data (we only have domestic employment). Finally, it would be nice to use value added data and adopt a more "accounting" oriented approach as in Ali-Yrkkö et al. (2010).

2.5 Narrative analysis

In Section 2.4, we have found multiple results pointing toward the granular hypothesis holding in Finland. In this section, we look in detail at a number of episodes where output growth (or drop) has been particularly strong. We then inspect events, using public sources and data, that have affected the largest Finnish corporations around the periods of interest. We consider, as in Gabaix (2011), episodes where $|y_t| \geq 1.3\sigma_y$, where σ_y is the standard deviation of monthly real output growth (y_t). In this fashion, we obtain 50 monthly episodes and we analyze the years which include at least one of those months. Given the importance of Nokia, we split the analysis in two subsection: the first one focuses on large firms excluding Nokia, while the second part concentrates on the latter. In Table 2.14, we report some meaningful episodes to give the reader a sense of the impact of granular shocks. Unfortunately, due to confidentiality reasons, we cannot use our firm-level monthly data to investigate these episodes and, instead, we need to rely on available annual and interim reports. We cite both yearly and quarterly episodes of granularity and we specify the time frequency to which they refer to.

2.5.1 The effect of large firms' shocks

For each year included in the analysis, we describe the overall macroeconomic conditions (which we take from the corresponding Bank of Finland annual reports) and subsequently look at the performance of the largest Finnish corporations during the selected periods. If available, we rely on quarterly sales figures, otherwise we use annual sales. Because we need to work with annual and quarterly sales data, we re-estimate (2.3) using variables at lower frequency to estimate the impact of the granular residual on real economic activity. The coefficient we obtain for annual data is 1.20, while for quarterly frequency regression we obtain that the granular residual coefficient is 0.97.

1998: This year was marked by sustained growth and moderate inflation (related to the global fall in oil prices). However, output growth decelerated during the second half of the year, mainly due to a drop in exports caused by the southeast Asia and Russian crises. The decrease in export affected various sectors of the economy in dramatically different ways. For example, while the growth of exports for the manufacturing industry was 1.5%, the same figure for the electrical and electronics goods industry amounted to 24%. The months in which the growth rates were especially high are January to March, May, July and September, while the average TIO growth for the year was around 5%. One of the largest Finnish corporation of that year was Fortum, a majority state-owned energy and oil production company. This enterprise group was founded in 1998, from the merger of Neste and Imatra Voima. 1998 was not a good year for the company, mainly due to the drop in oil prices,

Year/Quarter	Company	Granular Residual %	Contribution to GDP Growth %	Explanation
1998	Fortum	-2.15	-2.58	Fall in oil energy prices (partly due to heavy rain falls)
1999.Q4	Nokia	4.5	4.3	Release of models 7110 (first to use WAP technology) and 3210
2000	Fortum	1.14	1.36	Expansion of gas trading network (in Germany and Netherlands)
2001	Nokia	-2.4	-0.6	Difficulties in commercializing the 3G technology
2004	Rautaruukki	0.11	0.14	Two large deliveries for the Sakhalin II project
2007	Nokia	3.7	4.4	Large expansion in the Indian market
2007.Q4	Neste Oil	0.48	0.46	Increase in refining margin due to various refinery shutdowns in U.S. and Europe
2009	Nokia	-2.5	-2.44	Competition due to the release of iPhone 3Gs, together with products by HTC and Samsung
2011	Nokia	-4	-4.8	Poor performance in the smartphone's market
2013	Nokia	-2.7	-2.6	Persistently poor performance leading to the sale of the Mobile Phone division to Microsoft

Table 2.14: Selected granular episodes in the Finnish economy from 1998 to 2013.

heavy rain falls which led to higher hydroelectricity production and subsequent energy price drops. These factors, together with a planned reduction in crude oil trading, led to a yearly drop in sales of around -15%. The average growth rate in sales of the Top-57 corporations was around 10% and the sales of Fortum in 1997 accounted for almost 10% of Finnish GDP, implying a granular residual for the company of -2.15% in 1998. Following the granular hypothesis, we can argue that the bad performance of Fortum during 1998 affected negatively GDP growth, which was almost 7%, by -2.58% points.

1999: The Finnish economy faced another period of sustained economic growth, albeit slightly lower than in 1998. The beginning of the year was still affected by low export demand caused by the southeast Asia and Russian crises, even though the internal demand and the success of the electronic and telecommunication sector kept output growth to high levels. The highest growth rates in TIO are found in the first three months of the year. However, many Finnish groups did not experience a positive year in terms of sales growth. A good example is UPM, a paper manufacturing company. The poor performance of UPM was mainly due to the oversupply of paper pulp during the year leading to a newsprint, magazine and fine paper price fall. The granular residual for the firm during the year is -0.4%, implying -0.48% points impact on output fluctuations for 1999.

2000: This was another extremely positive year for the Finnish economy. Together with the success of Nokia, which we discuss in the next subsection, Fortum's excellent performance provided a boost to economic activity. The yearly granular residual of Fortum is 1.14%, with a positive impact on output growth amounting to 1.36% points, i.e. almost 20% of the GDP growth for this year. There were multiple factors behind the success of the energy company: a substantial increase in crude oil prices, the increase of the international refining margin and the expansion of its gas trading network (starting operations in Germany and Netherlands). Moreover, the group experienced a 30% increase in fuel exports to North America and 45% increase in export of CityDiesel, an extremely low-sulfur diesel fuel. UPM faced a 8-day strike during April, with an estimated loss of 60 million euro over the profits for the second quarter of the year. However, this episode did not have a substantial effect over the company's performance over the year.

2004: Another good year with substantial growth, especially in the November–December period. The steel manufacturing company Rautaruukki faced a particularly positive year, with an average year-on-year demeaned growth in sales of around 6%. Various factors contributed to this successful performance: substantial increase in steel product prices in the international markets, efficiency improvements due to management decisions and two large delivery agreements for the Sakhalin II project (a gas and oil extraction facility). The implied contribution to GDP growth for the year is around 0.14% points, accounting for around 4% of output growth.

2006: Average GDP growth for the year is around 5%. Between the most successful corporations, we find the stainless steel manufacturer Outokumpu which made its best-ever operating profits in the last of quarter of the year. This was mainly due to the high increase in international demand together with productivity improvements through the commercial and production excellence programme. The implied contribution of Outokumpu during the last quarter of 2006 is around 1% points, amounting to almost 17% of GDP growth for that period. Moreover, increased demand from China, Latin America and Eastern Europe boosted two important companies in the paper manufacturing sector: Stora Enso and UPM. In particular, Stora Enso expanded in Latin America and became the only producer of coated mechanical paper in the area.

2007: Throughout the year Finland experienced sustained growth, especially during the last quarter. In this period, Nokia was one of the main driver of the good economic performance, together with the oil company Neste Oil. The latter was founded in 2005, after splitting from Fortum. The Finnish government controls 50.1% of the company. The granular residual for Neste Oil in the fourth quarter of the year is around 0.48%, implying a contribution of around 0.46% points and accounting for almost 6% of output growth. Behind the group's success was the launch of the NExBTL Renewable Diesel, which was employed by the Helsinki City Transport network. Even more importantly, Neste Oil benefited from increasing gasoline and diesel refining margins due to various refinery shutdowns in U.S. and Europe during the second half of the year.

2009: This year, together with 2008, has been marked by a profound crisis of the global economy. Finland faced an average 9% drop in monthly output throughout the year and most of the main Finnish groups had a very negative performance. For example, the yearly granular residual of Neste Oil is around -2.6% . During the April-September period the main diesel production line in the Porvoo refinery had to be shut down because of a fire leading the quarterly granular residual for Neste Oil to reach its lowest point at -2.8% (second quarter) and -3.2% (third quarter). The implied contribution of Neste Oil to GDP growth rates are 2.7% points and 3.1% points for the second and third quarter, respectively. The global crisis affected industrial production and demand in stainless steel, influencing greatly Outokumpu's activities. The company faced a great decline in sales and its granular residual for the year is -1.26%, with an implied contribution of -1.5% points. Moreover, in order to reduce fixed costs and boost profitability the group cut almost 10% of the personnel, with temporary layoffs of most employees at the Tornio factory (2350 employees), 250 workers of the Kemi mine and other 1600 employed in the steel production lines.

2010: This was a rebound year after the Great Recession. Various large firms had a positive performance, with Neste Oil having the highest granular residual in the last quarter of the year. The residual is around 1.21%, with an implied contribution to GDP growth equal to 1.17% points and accounting for 20% of the increase in GDP during the last quarter of the year.

As we can see, on multiple occasions the good performance of the Finnish economy was associated with the success of few large companies. However, especially in 1998 and 1999, we find cases where large groups have a negative contribution to output growth, even though the latter is very high. This can be explained by the fact that we have omitted, so far, the manufacturing of electronic and telecommunication products sector and its predominant firm, i.e. Nokia.

2.5.2 Nokia and the Finnish economy

In this subsection we shift the focus to Nokia. The analysis we report below is based on Nokia's annual and quarterly reports. Given the importance of this corporation, we describe in broad terms its performance in each year from 1998 to 2013, together with the main events which have affected the group.

1998: The strategic decision to focus on the telecommunications market, adopted in 1992, contributed to a strong development of the company, which became the largest mobile phone manufacturer worldwide by 1998. Net sales for the group increased by 50% from the previous year, thanks to the success of models such as Nokia 6110 and Nokia 5110 and the expansion of the GSM technology. During the first quarter of 1998, when Nokia 5110 was released, Nokia's granular residual is almost 1%, implying a contribution to output growth for the quarter of around 0.97% percentage points (14% share of GDP growth in Q1).

1999: Another great year for the company, with net sales increasing by 40% with respect to 1998. In October, the 7110 model was introduced and it proved to be extremely successful. In particular, this model was the first mobile phone allowing for internet access through the WAP technology. The granular residual for the quarter of release (fourth quarter) is 4.5% and its implied contribution to the 5.7% GDP growth for that period amounts to 4.3% points. Another model which proved to be extremely successful was the Nokia 3210, with 150 million units sold.

2000: This was another extremely successful year for Nokia, with a yearly increase in net sales of 54%. As we have seen in the previous subsection, the Finnish economy experienced a year of sustained growth with a spike of almost 8% increase in the last quarter of the year. On September 1st, Nokia announced the 3310 model, which was then released in the last quarter of the year. The granular residual for the group reaches its highest point in Q4 of 2000, amounting to around 2.6%. This implies a contribution of 2.5% points to output growth, amounting to around a 32% of the year-on-year increase in quarterly GDP. During 2000, there was another episode which exemplifies the vast influence that the corporation had on the whole Finnish economy. As reported in Helsingin Sanomat (2000), on July 28th the company released its interim report for the first half of the year. The figures indicated a record profit with a 65% improvement to the first half of the previous year. However the stock market did not respond well to the news because the company did not reach its expected performance. Nokia share went down by 21.3% and the Helsinki Stock Exchange Index was heavily impacted by this drop, with a decline of 15.9% in a single day.

2001: The period going from September to the end of the year was characterized by slow growth of Finnish output, including a TIO fall of almost -1.7% in December. Nokia experienced a substantial slowdown in the growth of sales, with a yearly increase of only 2.6% (the average top-57 growth is around 5% for the year), implying a -0.6% points contribution to annual GDP growth. The slowdown of the global economy was an important factor behind this result, however the company also pointed out difficulties in launching and commercializing the 3G technology as a cause of the low growth in sales.

2002: This was a year of fairly slow economic growth for Finland. While the mobile phones division of Nokia did not register a meaningful drop in sales, Nokia Networks had a -14% annual decline. One of the events behind this bad performance was the insolvency of the German operator Mobilcom, due to financial difficulties of the latter. In particular, as reported in CNN.com (2002), Nokia decided to write off a 300 million euro loan as part of a new financing agreement with the German operator.

2003: Another year of weak growth for Finland, especially during the first few months. Nokia experienced an overall drop in net sales, caused by a bad performance of Nokia Networks. This led to a dramatic reorganization of the company, with the creation of a multimedia division. However, mobile phone sales rose throughout the year. In the last quarter of 2003, the company experienced an increase of 22% of mobile phone units sold with respect to the previous year. The release of Nokia 1100, the most successful model in the history of the group, can give an explanation for the good performance of Nokia Mobile in this period. Interestingly, in the last quarter of 2003, Finnish economic activity started to accelerate with a growth in GDP of around 2.6%, compared to the 1.6% growth during the previous quarters. In Q4, Nokia's implied contribution to GDP growth is around 0.2% points (around 8% share).

2004: Finnish output growth was particularly high in the last quarter of the year. During the same period, Nokia experienced a moderate increase in net sales, mostly due to the success of Nokia Networks and the good performance of imaging smartphones. On the other hand, sales of standard mobile phone decreased, as a consequence of the price pressure dictated by increased competition.

2005: Nokia has a large granular residual in the first quarter of the year, amounting to an implied contribution to output growth of 0.3% points (10% share). Even more importantly, in the last quarter of the year, Nokia's share of the global mobile phone market reached 35%. In Q4 2005 the implied contribution of the corporation is around 1.3% points, accounting for almost half of the GDP growth.

2006: This year was marked by the volume record for the group, with 347 million units sold. The net sales growth for the Mobile Phones division was almost 20% over the year, mostly due to the competitiveness of the entry-level products and the success in fast-growing markets such as Latin America and China. GDP growth for that period was almost 5% on average and the implied contribution of Nokia is around 0.8% points (15% share of yearly GDP growth). During the year, the company started releasing the N-series handsets and even the 770 Internet tablet, with the aim to expand into the high-margin product market (The Economist, 2006).

2007: In April, Nokia Siemens Networks was officially formed. Nokia annual net sales increased by 24%, driven especially by the Multimedia division (producing and marketing smartphones). Despite the release of Apple's iPhone, Nokia kept its role as leader in the smartphone and mobile phones market. One factor behind the success of the company during the year was the large expansion in the Indian market, which became the second most important for the company in terms of production volume (Nokia press release, 2007). Nokia's granular residual for the year is around 3.7%, with an implied contribution to yearly GDP growth of 4.4% points, roughly half of the GDP increase for the 2007.

2008: Finnish output growth had a breaking point around June 2008, with the spread of the financial crisis from the United States. Up until then, TIO growth was around 2.7% per month on average. In the second half of the year it dropped to an average decline in output of around -0.7% (the average growth of monthly output for 2008 was around 1%). Together with the macroeconomic consequences of the Great Recession, Nokia had to a face a much stronger competition, especially with the release of Apple's iPhone 3G and the failure to adapt to the new smartphone demand (an interesting example on the feedback on Nokia Symbian smartphone can be found in Helsingin Sanomat, 2013). This led to a poor performance for the year, with a drop in the yearly sales of -1%. The implied contribution to the average growth of TIO during 2008 is around -3.6% points.

2009: This year was marked by the dramatic drop in Finnish output (around -6.5% on average through the year) and Nokia difficulties. The granular residual for the group during 2009 is around -2.5%, with an implied contribution of -2.44% points. This corresponds to almost 40% share of GDP drop for 2009. The global economic crisis, together with the growing competition in the smartphone market, with the release of iPhone 3GS and products from LG, HTC and Samsung, were the main cause of this weak performance. During the 3rd quarter of 2009 Nokia faced its first quarterly losses since it started to report on a quarter-by-quarter basis in 1996, amounting to 913 million euro loss (The Guardian, 2009).

2010: After the Great Recession, the Finnish economy returned to positive growth from March onward. Nokia performance was still fairly weak, with net sales increasing by 3%. The cross-sectional average for the top-57 Finnish corporations for that period is around 15%, implying a -12% demeaned growth rate. The estimated contribution to annual GDP growth (which was roughly 3%) is around -3% points. It seems that the negative performance of Nokia during 2010 was a key factor holding back the rebound of the Finnish economy. A contributing factor behind this decline was the success of iPhone 4 and Samsung Galaxy S, together with problems in the supply chain with the shortages of components for mobile phones production. In particular, supply of components such as memory chips, resistors and transistors did not respond strongly to the increased demand after the Great Recession, which led to weak results for many consumer electronics manufacturers (Reuters, 2010).

2011: While this was a positive year for the Finnish economy, with a 5% output growth on average, Nokia had another extremely disappointing year with a yearly net sales drop of -18% (implying a demeaned growth of -32%) and a granular residual of around -4% (implying a -4.8% points contribution to yearly output growth). Nokia started a partnership with Microsoft to adopt Windows operating system as a primary platform for the company's smartphones. After the announcement of the partnership, Nokia share dropped to its lowest value since July 2009. Moreover, the group implemented a series of strategical and operational changes, including a large reduction in personnel and the closure of various facilities (among them, the Cluj factory in Romania, which was opened only three years earlier, Reuters, 2011).

2012: Finnish economy turned to negative year-on-year growth, with an average quarterly GDP drop of -1.4%, following the overall trend of the Euro Area. Nokia continued to struggle due to the strict competition among smartphone manufacturers, and this led to a dramatic reduction in market share. The company demeaned yearly change in sales was almost -27% and the decrease of personnel amounted to -25% with respect to the previous year. The number of smartphones sold in the last quarter of the year amounted to around 16 million units, against the 65 million Samsung devices and 27 million iPhones. As a consequence of these difficulties, Standard and Poor's eventually downgraded Nokia bonds to 'junk' status (Bloomberg, 2012). The granular residual for the company was around -5.4%, with an implied contribution of around -6% points.

2013: Both the Finnish economy and Nokia declined throughout the year. The GDP year-on-year growth was around -1.1%, while the annual granular residual of Nokia was close to -1.1%. The highest drop in GDP was in the first quarter of the year with a -2.9% year-on-year decline. Nokia's granular residual for that quarter is around -2.7% and implies almost a -2.6% points contribution (which accounts for most of the decline in GDP in Q1 2013). Nokia's year was marked by two operations: the re-acquisition of half of Nokia Siemens Network from Siemens and, more importantly, the Mobile Phone division was sold to Microsoft. Even though the announcement was made on September the 2nd, the actual deal was finalized in April 2014, which falls out of the sample period we have analyzed.

This section has outlined how the role of Nokia in the world-wide telecommunication market has dramatically shifted over the past 16 years. Since 2008, the group has lost its dominant role in the mobile phone business, due to the success of its various competitors and the rise in popularity of smartphones. After selling the Mobile division to Microsoft, in 2014, Nokia's annual sales amounted to around 13 billion Euro, against the 30 billion Euro in 2012. Interestingly, the decline of Nokia could have been the driving factor behind the post-2009 results in Table 2.7, showing that the granular residual is unable to explain a substantial amount of output growth variation. It would thus be interesting to follow up on the performance of the group in the next few years in order to track the development of its impact on Finnish output, given the company's significantly reduced size, and the apparently successful shift in it's product mix.

2.6 Conclusions

In this paper, we have examined the granular hypothesis within the Finnish economy of the last 16 years. To do this, we use a monthly dataset comprising the top-57 Finnish business groups, in terms of sales. We then use the methods suggested in Gabaix (2011) to extract the granular residual and verify its importance in explaining output fluctuations. We complement the econometric analysis with a short narrative comprising the main events which have hit the largest Finnish companies, with a special focus on Nokia. To examine individual corporations, in the narrative analysis, we use publicly available data obtained from the groups' annual and quarterly reports.

We find that the granular residual is able to explain a substantial share of the fluctuations in Finnish real economic activity. Interestingly, the Great Recession marks a strong break in the relationship, with R^2s dropping to very low levels when we analyze the period going from 2010 up to 2013. As a robustness check, we also estimate the granular residual using a factor model, as in Stella (2015), and we get similar results as in the main analysis. We also split the top-57 corporations in our dataset into four different size classes, based on their sales, and study the explanatory power of groups of different sizes. We find that the idiosyncratic shocks to the largest companies are the most important predictor, especially until the Great Recession. After 2008, the explanatory power of all corporation-level shocks drop substantially. Furthermore, we examine the explanatory power of Nokia's labor productivity, using quarterly data obtained from the company reports, and find that it accounts for around 25% of GDP fluctuations. Finally, we follow the approach of di Giovanni et al. (2014) to decompose the variance of the enterprise group shocks and find that the granular component has dominated the covariance component, except for 2009. This confirms the fact that the main transmission mechanism for microeconomic shocks has indeed been the granular one, with the network channel playing a secondary role.

In the narrative analysis, we delineate a number of episodes where large Finnish companies have experienced particularly successful (or weak) periods and examine their effect on Finnish output growth. Moreover, we dedicate a subsection to study the performance of Nokia in the last two decades and find a dramatic change in the importance of the group in the Finnish (and World) economy. From being the largest mobile phone manufacturer in the World during the mid-2000s, the company has faced a continuous decline during the years after the Great Recession, which ended with the sale of the mobile phones and smartphones division to Microsoft in 2014.

One of the most significant results of this analysis is the break in the Granular Hypothesis after the end of the Great Recession. As reported in tables 2.7 and 2.9, the impact of the shocks of large corporations onto aggregate output becomes very small from 2010 onward. At a first glance, this result can be interpreted as the effect of the decline of some large firms during the Great Recession, above all Nokia. However, the Finnish economy remains highly granular, in terms of the dominance of few large firms, as we can see from Figure 2.1 (b). While the Herfindahl index shows a large drop after 2010, it quickly reaches a similar levels as the ones of the pre-crisis period. This indicates that the sales of large corporations account for a similar share of GDP as in period before the recent economic crisis. Given that the network effect seems not to be driving our results, as shown in Section 4.3, we should seek a different explanation. It is arguable that the real impact of corporations as large and influential as Nokia goes beyond their actual size, measured by their sales. Moreover, even if a corporation plays a central role in the economy, it might not be a particularly important supplier or client for the rest of the firms in the economy (e.g. it might buy its intermediate products from foreign firms and affiliates). One additional explanation could be that the success of large corporations might indirectly influence the overall economic environment, possibly through generation of know-how or by investing into research and development outside the company. It could be interesting to carry an analysis in the fashion of Ali-Yrkkö et al. (2010), focusing on a limited number of firms and studying a wider array of economic variables and possible channels of influence of the large corporations on the economy.

While the result that the aggregate output of a small country, like Finland, might be heavily affected by the performance of few large firms can be expected, this does not reduce the practical importance of the findings of the current study. First of all, they reinforce the point that macroeconomic models and analyses of the business cycle should not discard microeconomic shocks and dynamics. Many economies have a non-normal distribution of firms, especially when considering the ownership structure that can link smaller enterprises with large corporations, and hence the standard models relying on the representative firm might be inappropriate. This also highlights important policymaking consequences, indicating that economic institutions (such as the government) should pay close attention to the performance of larger enterprises, given the potentially large effects on aggregate fluctuations.

While this analysis provides many useful insights, there are multiple possibilities to expand this research further. First of all, we focus on sales as the main indicator of a business group success, due to their monthly availability for a wide range of companies. We have also tested our results using measures of productivity such as the one in Gabaix (2011), but it can be interesting to see how relying on value added (see Ali-Yrkkö et al., 2010) would affect our findings. Moreover, we can expand the factor model analysis by employing the filtering technique of Stella (2015) and Foerster et al. (2011), to make sure that the estimated factors do not incorporate the effect of idiosyncratic shocks. However, the use of enterprise group-level data can milden this issue by grouping firms inside the same mother company. Furthermore, it would be interesting to shift the focus from the granular hypothesis to the network channel suggested in Acemoglu et al. (2012). While we touch this issue by studying the variance decomposition of the group-level shocks and examine how much the granular and link components account for their fluctuations, a more extensive analysis of the network channel can be very interesting in the light of the possible influence of Nokia onto the rest of Finnish firms. Related to this, it would be interesting to use within company relationships to examine the contribution of individual firms and establishments within a business group to the granular residual of that company. As illustrated by the case of Nokia, the large granular effects can be sourced to the success of particular products or divisions inside the firm, and developing this line of thought can indicate that the analysis of the dynamics within firms can be useful in explaining aggregate effects.

2.7 Appendix: Adjustment for legal restructuring.

In this appendix, we discuss the details of the procedure adopted by Statistics Finland to control for merger and split-offs in a set of enterprises. Assume that firm 1 is examined after an event (merger or split-off) where N firms are involved. Then the estimated sales of firm 1 one year ago is calculated by:

$$sales(firm_{1,t-12}) = \frac{sales(firm_{1,t}) * sales(firm_{1,t-12}, firm_{2,t-12}...firm_{N,t-12})}{sales(firm_{1,t}, firm_{2,t}...firm_{N,t})}$$

where t is the time periods in which the adjustment is computed, and N is the number of firms involved in a merger or split-off. The sum of the previous year sales levels in all the firms involved in the event is divided for each continuing firm weighted by their relative size at present time t. Let us go through some simple numerical examples to see how this works:

1. Assume a firm A with 2 billion euro sales in period t, that had 1 employee in t-12. Firm A acquires firm B with billion sales at time t, m and 1 billion euro one year ago. Firm A, which continues existing, will be assigned a new estimated number of sales for the comparison year, in order to make the growth rates comparable year-on-year. The comparison values of firm A is estimated as $\frac{2(1+1)}{(2+1)} = 4/3$, and the rate of change for A becomes $(2+1)/(4/3) = 2.25$ (as opposed to 3 if no correction is done)

2. Consider the situation where firm A is split into smaller units, say B and C. A has 3 billion euro sales at time $t-12$, B has 3 billion euro sales at t and C has sales for 2 billions at t. B and C

did not exist at $t-12$, so their comparison values become: $(3/3)3=3$ and $(2/3)3=2$, resulting in the rate of change for B and C to be 3/3 and 2/2 (equal to 1 for both firms). The growth rate is forced to be the same among the continuing firms after a split-off.

Chapter 3

Nowcasting Finnish Real Economic Activity: a Machine Learning Approach[1]

Paolo Fornaro*, Henri Luomaranta**
*Research Institute of the Finnish Economy, Finland
**Statistics Finland and TSM, University of Toulouse Capitole, Toulouse, France

Abstract

We develop a nowcasting framework, based on micro-level data, to provide faster estimates of the Finnish monthly real economic activity indicator, the Trend Indicator of Output (TIO), and of quarterly GDP. As main predictors we use firm-level turnovers, which are available shortly after the end of the reference month. We rely on combinations of nowcasts obtained from a range of statistical models and machine learning techniques which are able to handle high-dimensional information sets. The results of our pseudo-real-time analysis indicate that a simple nowcasts' combination based on these models provides faster estimates of the TIO and GDP, without increasing substantially the revision error. Finally, we examine the nowcasting performance obtained by relying on traffic volumes data, and find that using machine learning techniques in combination with this big-data source provides accurate predictions of real economic activity.

JEL Classifcation Code: C33, C55, E37
Keywords: Flash Estimates, Machine Learning, Micro-level Data, Nowcasting

[1]This chapter is based on an accepted article, conditional on editorial revisions, in Journal of Empirical Economics

3.1 Introduction

We live in a data-rich world. Statistical agencies, central banks, research institutes and private businesses have access to (and produce) hundreds of economic and financial indicators. The list of available data is continuously growing, with the introduction of "big data" encompassing sources such as Internet search engines, social media sites, cash registry data and many more. However, this wealth of information has not been directly translated into a faster and more accurate production of important economic statistics, such as the GDP. Instead, statistical institutes publish economic indicators with considerable lag and the initial estimates are revised considerably over time. For example, in Finland the first estimate of GDP provided by Statistics Finland is released 45 days after the end of the reference quarter (flash estimate), while the first "appropriate" version is released 60 days after the end of the quarter.

The advantages of having a timely picture of the state of the economy are multiple and concern a range of economic actors such as the central bank, the government and private investors and businesses. Therefore, nowcasting and the production of economic activity indicators in real time have been the focus of a growing econometric literature. Early works related to the tracking of economic conditions in real time are Aruoba et al. (2009), for the U.S. economy, and Altissimo et al. (2010) for the Euro Area. In these studies, the authors develop econometric frameworks with the objective to create high-frequency indicators of real economic activity. On the other hand, the nowcasting literature is interested in estimating an existing economic indicator (usually quarterly GDP growth) in real-time. Few examples drawn from the nowcasting literature are Doz et al. (2012), Evans (2005), Modugno (2013) and Aastveit and Trovik (2014), among many others. Usually, nowcasting models rely on a wide array of data, such as consumer surveys, financial variables and macroeconomic indicators, and use factor models or large Bayesian vector autoregressions to produce predictions of the variables of interest.

In this study, we combine confidential firm-level datasets and machine learning techniques, as well as traditional statistical models which can deal with large datasets, to provide faster estimates of Finnish real economic activity, both at the quarterly and monthly frequencies. The monthly series we target is the Trend Indicator of Output (TIO)[2], published by Statistics Finland 45 days after the reference month, while the quarterly series is GDP. For both series we compute nowcasts of the year-on-year growth rate. In addition, we examine the predictive power of a novel dataset based on traffic volumes' measurements, created by combining disaggregated data obtained from the Finnish Transport Authority website. The use of novel data sources, such as firm-level turnovers data and traffic measurements, in combination with the use of a wide array of machine learning techniques provides the main contribution of our study to the nowcasting literature. The use of firm-level data in providing fast estimates of real economic activity is not unique: Matheson et al. (2010) rely on qualitative responses obtained from business surveys, to obtain nowcasts of New Zealand GDP growth, while Fornaro (2016) uses a similar firm-level dataset to estimate Finnish economic activity. We expand the latter work in two main ways: firstly, we consider an additional data source, i.e. the trucks' traffic

[2]A description of this indicator is available at http://www.stat.fi/til/ktkk/index_en.html

volumes, which can be interesting with respect to the use of big data in economic forecasting and nowcasting (e.g., see Baldacci et al., 2016). Moreover, we consider a much larger array of statistical frameworks and machine learning techniques compared to Fornaro (2016), which focuses exclusively on factor models.

Before we go to the summary of our results, it can be useful to depict the timing of the statistical releases of a typical month and quarter, i.e. when the firm-level data becomes available, when the current estimates of Statistics Finland are published and when we compute our nowcasts.[3] We do this for the TIO in Figure 3.1.

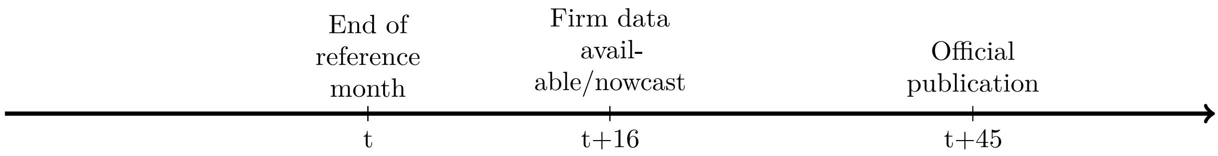

Figure 3.1: Official release and nowcast schedule for the TIO.

For the TIO, the monthly variable considered in this exercise, the release schedule and the timing of our nowcast process is fairly simple. Denoting the end of the reference month by t, the firm-level data becomes available 16 days after the end of the month $(t + 16)$ and we compute the nowcast. Statistics Finland then publishes its first version of the TIO 45 days after the end of the month $(t+45)$. The timing of the nowcast process and of the publication schedule for quarterly GDP is slightly more complicated, and we report the timeline of the events in Figure 3.2.

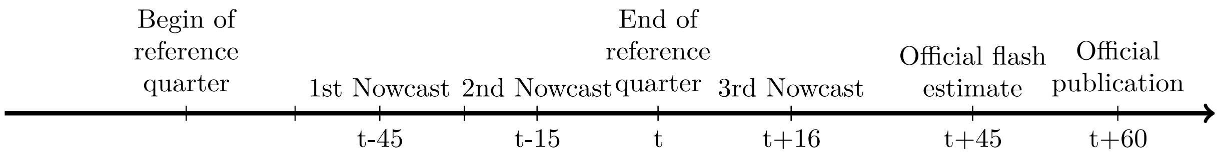

Figure 3.2: Official release and nowcast schedule for quarterly GDP.

To keep Figure 3.2 readable, we do not mention explicitly when the firm-level data for a given month of the reference quarter becomes available. However, the nowcasts are computed as soon as the firm-level data is available for the period of reference. Let's index the end of the reference quarter by t. We produce three nowcasts: the first one is done 45 days before the end of the quarter $(t - 45)$, i.e. during the second month of the reference quarter; the second nowcast is produced during the third month of the reference quarter, which correspond to 15 days before the end of the quarter $(t - 15)$. We then compute our final nowcast for the quarter 16 days after its end $(t + 16)$. Statistics Finland releases the flash estimate of GDP 45 days after the end of the reference period, while the first official figure is published at $t + 60$.

We find that our approach of combining predictions obtained by using a large set of machine learning algorithms, based on firm-level data, is able to provide accurate estimates of monthly economic

[3]In our exercise, we compute both nowcasts (predictions of a variable while the reference period is still ongoing) and backcasts (estimates referring to a period which already ended), we refer to our predictions as nowcasts, to be in line with the literature (see Banbura et al., 2011).

activity growth, with revision errors that are in line with the ones of Statistics Finland, while shortening the publication lags by 30 days. The resulting early estimates of the monthly indicator are used to compute three nowcasts of GDP year-on-year growth. The first two nowcasts provide good accuracy, even though there are some notable revision errors. However, the estimates produced after the end of the quarter are very accurate, while providing a 45 days reduction in the publication lag. Moreover, the methods we use are computationally feasible and easily automatable, making them appropriate for a real-time setting. We conduct a similar analysis using truck traffic volumes' measurements, and find satisfactory results, albeit inferior to the ones obtained from firm-level data.

The remainder of this paper is divided as follows: in Section 3.2 we discuss some of the large set of models adopted in the analysis, in Section 3.3 we describe our target indicators and data sources, and we delineate the structure of our nowcasting exercise. The results of the exercise are reported in Section 3.4, while Section 3.5 concludes.

3.2 Methodological Aspects

Given that the main contribution of this study is the use of novel data sources, we keep the description of the models adopted brief. This section does not cover comprehensively the techniques we use for two reasons: firstly, the sheer number of statistical models and machine learning techniques adopted in the exercise does not allow a thorough discussion (the full list of the techniques adopted in our study is reported in the appendix). Secondly, these techniques have been used in previous econometric or statistical studies, hence a detailed description would be superfluous. We instead try to give the basic intuition underlying some of the main classes of models used and redirect the interested readers to the original works in which the models we employ were developed or to some previous forecasting applications in which these models are adopted.

One of the most important models in our exercise is the dynamic factor model, in the form of Stock and Watson (2002a,b). The basic idea is that a handful of constructed variables, the factors, can summarize the information contained in a large dataset. Formally, the K-dimensional set of predictors for period t, X_t, follows the Stock and Watson (2002a) representation with r latent factors

$$X_t = \Lambda F_t + u_t, \tag{3.1}$$

where F_t is $r \times 1$ and Λ is the $K \times r$ matrix of factor loadings. Stock and Watson (2002a) have shown that the factors can be estimated using principal components, i.e. the r factors correspond to the first r eigenvectors of the $K \times K$ matrix XX'.[4] Factor models are especially important in our application because, in addition to the basic specifications including raw firm-level data and traffic data as predictors, we estimate specifications where we utilize latent factors (estimated via principal components) as predictors. This is done to see whether reducing the noise in our input data improves the performance of the models (for techniques which rely on dimensionality reduction we do not

[4]Alternative estimators of latent factors are presented in Forni et al. (2000) and, more recently, Doz et al. (2011). Bai and Ng (2002a) developed a series of information criteria that provide an estimate of the number of static factors r.

use factors as input variables). We have tried using the Bai and Ng (2002a) criteria to determine the optimal number of factors, but these always suggested to use a very large number of principal components (usually the upper bound). Therefore, we use three different specifications with 10, 20 and 30 principal components (using less than 10 factors leads to significantly worse results).

Another important class of models we use is shrinkage regression, in particular the ridge regression, the Lasso (Tibshirani, 1996) and the elastic-net (Zou and Hastie, 2005). The main intuition of these models is to regularize the coefficients of the predictors, in order to reduce the predictions' variance. Hastie et al. (2009) provides an in-depth review of these techniques, while De Mol et al. (2008) offers an economic forecasting application of shrinkage regressions, with an interesting comparison with principal components.

Our nowcasts are then based on a large number of machine learning techniques, which are covered extensively in Hastie et al. (2009): boosting (for forecasting applications see Bai and Ng, 2009 and Wohlrabe and Buchen, 2014), regression trees and random forests (for an application in recession forecasting, see Nyman and Ormerod, 2017), regression splines, support and relevance vector machines, neural networks (an interesting time series application that considers these three techniques is Plakandaras et al., 2015), and k-nearest neighbors (see Fernandez-Rodriguez et al., 1999 for an application in exchange rate forecasting). Finally, we include in our models set an automated ARIMA where we use principal components (extracted from the firm-level or traffic data) as external predictors.

All the models utilized in our nowcasting exercise are implemented using the `caret` package for `R`. The tuning of the models hyperparameters is done automatically, using a rolling forecasting origin technique, discussed in Hyndman and Athanasopoulos (2018). Once considering specifications with different input variables (raw data vs. sets of principal components extracted from the data), we arrive at a total of 130 models to estimate. As benchmark model, we utilize an automated ARIMA procedure, implemented using the package `forecast` for `R`, as described in Hyndman and Khandakar (2008).

3.3 Data description and empirical exercise formulation

In this section, we briefly describe the data used in our exercise and we give few details about the nowcasting process.

3.3.1 Target variables

The target variables in our exercise are quarterly GDP, and the TIO, in particular their year-on-year growth rates in real-term. The GDP is published as an early version at $t + 45$, and updated at $t + 60$. The $t + 60$ version is considered as the first official and reliable estimate of GDP, and it is the one we target. The TIO is a monthly series that describes the development of the volume of produced output in the economy. It is constructed by using early estimates of industry-level turnover indexes (not publicly available), which are appropriately weighted to form the monthly aggregate index. The TIO is published monthly at $t + 45$, and its value for the third month of a quarter is used to compute the flash estimate of GDP. We plot our target variables in Figure 3.3.

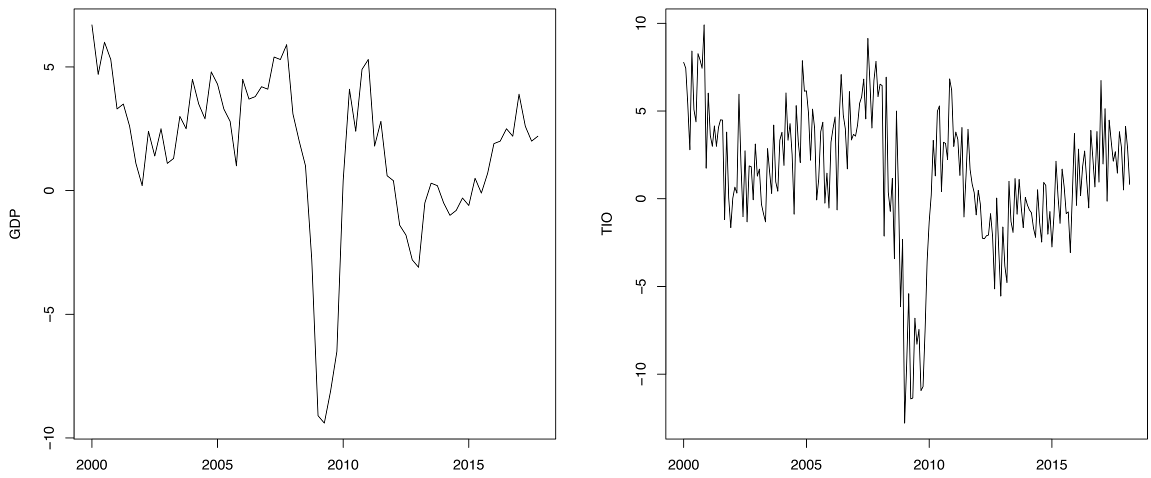

(a) GDP year-on-year growth, Q12000-Q42018 (b) TIO year-on-year growth, Jan.2000-Dec.2018

Figure 3.3: Target variables.

One aspect that is important to underline is how closely related the TIO and GDP growth are. If we aggregate the latest version TIO growth to the quarterly level, using a simple arithmetic average, we obtain a series that closely tracks GDP growth (the resulting correlation coefficient is 0.99). This demonstrates that providing a good estimate of TIO leads to a greater nowcasting accuracy of GDP.

3.3.2 Firm data

The main predictors in our nowcasting application are firm-level sales extracted from the sales inquiry, a confidential monthly survey conducted by Statistics Finland for the purpose of obtaining turnovers from the most important firms in the economy. This dataset covers around 2,000 enterprises and encompasses different industries (services, trade, construction, manufacturing), representing ca. 70% of total turnovers. The data is available soon after the end of the month of interest and a considerable share of the final data is accumulated around 15 to 20 days after the end of the reference month. Formally, Statistics Finland imposes a deadline to the firms on the 15^{th} day of the month after the reference period, but this deadline is not always met, thus our set of firms' sales does not cover the entire sample. We compute the nowcast of TIO on the 16^{th} day after the reference month and typically have 800 firms in the predictors' set. We compute the year-on-year growth of sales, starting from January 2006 until the month we want to nowcast. If the firm has reported sales by the $t + 16$ of the target month, but has missing values during the time span (i.e. the firm did not reply at some earlier date, or the firm was not included in the turnover inquiry at some point in time), we try to obtain the missing growth rates from VAT data, which should include all the firms in the economy. VAT data is available at t+52 days after the month has ended, and we can therefore use it for all the missing values occurring up to two months before the reference month. We trim the remaining firms for which we can not impute their missing values, leading us to have a balanced panel of firm-level data. The data accumulation (i.e. how new firm-level data becomes available over time) is realistically simulated by using the time stamp of the reported sales, which allows us to track what data was available by each date of a month. Furthermore, starting from January 2017 we have collected data in real-time. Overall, we reliably mimic the information available to the statistical institute at each point in time.

3.3.3 Traffic data

The other main predictors used in our exercise are the traffic volumes records obtained by the Finnish Transport Agency website.[5] This dataset contains records of individual vehicles passing through a number of measurement points (about 500) around Finland, observed through an automatic traffic monitoring system. The information recorded includes direction, speed and, importantly, vehicle class. This dataset contains numerous missing values, due to the fact that some measurement points do not have observation for certain days or months, and it is not structured. For our nowcasting analysis, we collect data for trucks' traffic volumes from January 2006 onward, and consider only measurement points in the region of Helsinki (which contains a bit less than 100 measurement points), because the data download procedure is time consuming. Focusing on this region should give us valuable information, given that it accounts for a large share of the Finnish economy. Trucks' traffic presents an interesting link with aggregate economic activity. We expect that in periods of economic growth, when trade volumes and production are increasing, we should observe a higher number of trucks' passages, in order to move goods. Of course, this does not cover the transfer of services and other types of economic activities, but it should still present some positive correlation with economic activity growth.

As first step, we download raw daily records of traffic volume, where one file contains measurement informations for each vehicle passing through a single measurement point for a given day, from the Finnish Transport Agency website. The traffic data is then aggregated at the monthly level, by averaging the daily volume of truck traffic over the month, and we then compute year-on-year growth rates for each month and for each measurement point. We assume that our estimation of TIO is conducted around 16 days after the end of the reference month (as in the main exercise). This allows us to use the Statistics Finland's estimates of TIO for the $t-1$ month, where t represents the period we want to nowcast. In principle, the traffic data we utilize allows for nowcasts during the month of interest, given their daily frequency, but a daily estimate of TIO which relies on traffic volumes would risk to be extremely volatile, given the large variations in traffic volumes at the daily frequency. As mentioned previously, data on traffic volumes contain many missing values; in order to impute the missing observations, we rely on the regularized principal component technique illustrated in Josse and Husson (2016).

3.3.4 Nowcasting exercise formulation

In this study, we try to mimic as well as possible the conditions faced by the statistical institute in real-time. In particular, we use the original vintages of data (both for the target variables and the predictors), reflecting the information available at the time that the nowcast would have been

[5]The data is available at https://aineistot.vayla.fi/lam/rawdata.

computed. Both TIO and GDP series are revised multiple times, because of new data availability and of benchmarking.[6] Consequently, we use vintages reflecting the first estimate of TIO and adopt these initial figures as target to evaluate our nowcasts (because otherwise our nowcasts would contain errors that are not due to the lack of predictors but that are instead caused by the lack of smoothing and benchmarking). Unfortunately, the historical vintages for TIO are available only since March 2012, meaning that our nowcasting exercise does not cover some interesting periods such as the Great Recession of 2008–2009. However, we are left with more than 80 predictions to be made and the timespan going from 2012 until the end of 2018 does include periods of high growth and months of considerable output drop.

We described the timing of our exercise in the introduction, but it can be beneficial to use an example to make clear how the procedure is carried out in practice. Suppose that we want to nowcast TIO year-on-year growth for March of a given year: we would compute the nowcast on April 16th, using TIO data up to February to estimate the models and then compute the nowcast using the March firm-level sales as predictors. When estimating quarterly GDP, we do not rely directly on the GDP series but rather use TIO, which means that we do not have problems in terms of publication lag.

Now to the structure of our empirical exercise: we start to compute monthly nowcasts of the TIO from March 2012. In particular, we extract a panel of firm-level sales which starts in January 2006 and ends in March 2012 (the same goes for traffic volumes). Moreover, we include the first 12 lags of TIO as predictors, starting from the TIO growth for $t-1$ (in this case February 2012) until $t-12$ (March 2011). Notice that the lags of TIO which are included in the predictors set reflect the available information realistically, because we obtain it from the proper data vintage. We repeat this procedure for each month until December 2018, expanding the estimation window (instead of using a rolling window approach). As an example, in the case where we use r factors, extracted from the firm-level data, and the first 12 lags of TIO as predictors in a linear model, the nowcasting model for period T is given by:

$$y_T = \mathbf{y}'_{T-1}\beta_1 + \hat{F}'_T\beta_2 + \epsilon_T, \tag{3.2}$$

and the nowcast is obtained by

$$\hat{y}_T = \mathbf{y}'_{T-1}\hat{\beta}_1 + \hat{F}'_T\hat{\beta}_2, \tag{3.3}$$

where ϵ_T is a normally distributed idiosyncratic error, $\mathbf{y}'_{T-1} = [y_{T-1} \cdots y_{T-12}]$, $\hat{F}'_T = [\hat{F}_{T,1} \cdots \hat{F}_{T,r}]$, and $\hat{\beta}_1$ and $\hat{\beta}_2$ are vectors of coefficients estimated by ordinary least squares, using data from January 2006 until $T-1$. Of course (3.2) and (3.3) can take many forms depending on the model we adopt, but the principle is similar: we first estimate the models using data until the latest month for which we have TIO values and then we use the most recent micro data information to compute the nowcast, given the estimated model parameters.

[6] Statistics Finland adjusts monthly TIO figures so that they are consistent with quarterly GDP growth estimates, once the latter become available. The same adjustment is done to quarterly GDP when yearly GDP figures are released. The practical implication of this procedure is the presence of large revisions of historical growth rates at the monthly and quarterly frequency.

Our quarterly estimate of GDP are entirely based on TIO, both the released version and our nowcasts. As we mentioned in the data description, TIO provides the basis for the initial estimate of GDP, hence it is optimal to use it as a predictor in a nowcasting exercise. We compute the GDP nowcasts differently, depending on the month in which we make the estimate. In our setting, the nowcasts for a given quarter are computed three times: during the second month of the quarter, during the third month and 16 days after the end of the quarter. In the first case, we would use the nowcast of TIO for the first month of the quarter, then estimate an automated ARIMA model to obtain the forecasts of the remaining months. If we compute the GDP nowcast during the third month, we would use the first TIO estimate made by Statistics Finland for the first month, then use our nowcast of TIO growth for the second month and then compute the 1-step ahead forecast for the third month. When we estimate GDP growth 16 days after the end of the quarter we use the TIO growth computed by Statistics Finland for the first two months and augment them with our nowcast of TIO for the last month of the quarter. Eventually, we are going to have an estimate of TIO growth for each month of the quarter of interest and we obtain GDP growth by taking a simple average over the three months (the same procedure is done by Statistics Finland to obtain the flash estimate of GDP). Denote the estimate of GDP growth for quarter q going from month $t-2$ to t as $\widehat{GDP}_{q,t}$, then our quarterly nowcast is $\widehat{GDP}_{q,t} = 1/3(\hat{y}_{t-2} + \hat{y}_{t-1} + \hat{y}_t)$ Notice that this procedure is rather similar to the one of bridge regression, which links quarterly and monthly variables via simple linear models. We have tried to estimate a linear regression of GDP growth onto the quarterly average of TIO growth, i.e. estimating the linear model $GDP_{q,t} = \beta\frac{(\hat{y}_{t-2}+\hat{y}_{t-1}+\hat{y}_t)}{3} + \epsilon_t$, but our results indicate that the simple average of TIO growth is a better predictor than using the bridge formulation.

Another practical issue we wish to mention is computational feasibility. We estimate more than 130 nowcasting models, some of which are computationally burdensome. Given that we would like to produce the nowcasts around $t+16$, using the information set available by then, we need to find some sort of compromise between having the largest spectrum of models and being able to estimate TIO quickly. In order to do that, we first select the models which produce nowcasts with historical mean error (in absolute terms) below the 20^{th} percentile.[7] Afterward, we compute a simple unweighted average of this subset of models (Stock and Watson, 2004, point out that a simple forecasts' average outperforms more complex schemes) and use this combination to form the latest nowcast. The choice of keeping models with low historical mean error is driven by the high importance, for the statistical institute, of having unbiased flash estimates. Notice, that the selection of models retained in the combination is based on the out-of-sample performance of the individual models over the whole period of the analysis (from March 2012 to December 2018). We have also computed the nowcast combinations iteratively, allowing different sets of models to be selected each period, but this leads us to give up on a number of evaluation periods, in order to compute the first combination. The results do not change substantially from the ones in the main analysis, thus they are not reported, but they are available upon request.

[7] In our exercise, this set of models includes 21 specifications, such the factor augmented automated ARIMA, regression splines, tree based regressions, ridge regressions, support vector machine, k-nearest neighbors and boosting.

3.4 Empirical results

3.4.1 Results for TIO nowcasts

As pointed out in Section 3.3, the TIO is a monthly indicator of real economic activity. Our nowcasting exercise is centered on providing fast estimates for the year-on-year growth rate of TIO, starting from March 2012 (the first month for which we have the vintage of the data) and ending in December 2018. We now provide the results for our pseudo out-of-sample analysis, starting from the nowcasts obtained from firm-level data. Specifically, we report the results of the models which provide the lowest root mean squared error (RMSE), the lowest mean error (ME), mean absolute error (MAE), and finally for the model with the lowest maximum absolute error (MaxE). In addition, we report the results for the simple nowcast combination. We plot the nowcasts obtained from the nowcast combination against the first published version of TIO, in Figure 3.4.

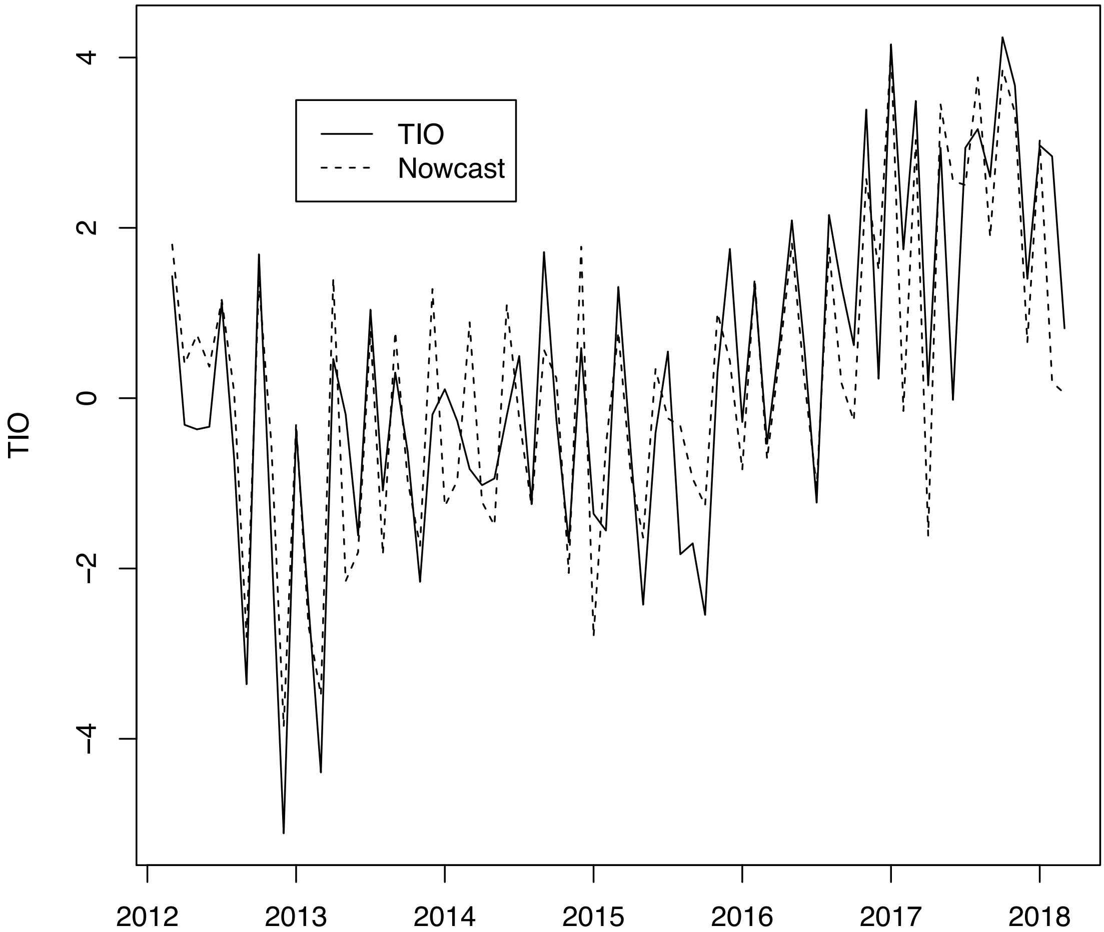

Figure 3.4: First version of TIO year-on-year growth and nowcasts combination, using the unweighted average of models selected based on low mean errors. The first version of TIO is published 45 days after the end of the reference month, while the nowcasts are computed 16 days after the end of the reference month. The set of predictors is based on firm-level turnovers.

Figure 3.4 indicates that our firm-level data offer a good basis for providing flash estimates of TIO. The nowcasts track fairly well the original series, while they provide a substantial gain in terms of publication lag (around 30 days). Next, we provide some numerical indicators of the nowcasting performance, for the models described at the beginning of this subsection. Moreover, we report the results obtained by using an automated ARIMA procedure, using the latest available TIO vintage at the time of the nowcast to estimate the model and compute the prediction.

	Lowest ME	Lowest RMSE	Lowest MAE	Lowest MaxE	Combination	ARIMA
ME	-0.00	-0.25	-0.25	-0.25	-0.01	0.11
MAE	1.06	0.75	0.75	0.75	0.78	1.36
RMSE	1.35	0.95	0.95	0.95	0.96	1.79
MaxE	4.60	2.17	2.17	2.17	2.52	5.85

Table 3.1: ME, MAE, RMSE and MaxE for different nowcasting models. Lowest ME, RMSE, MAE and MaxE indicate the models with the lowest mean error, root mean squared error, mean absolute error and max error, respectively. The Combination column contains performance measures for the simple nowcast combination based on the unweighted average of a subset of our models. The set of predictors is based on firm-level turnovers.

As we can see from Table 3.1, the nowcasting performance of our selected models is better than the one of an automated ARIMA procedure. In the first column, we report the results for the model with lowest mean error (an automated ARIMA with principal components extracted from the firm data), which shows a fairly poor performance in terms of MAE, RMSE and max error. Interestingly, the same model (a boosted generalized additive model with factors as input variables) has the best performance in terms of MAE, RMSE and MaxE, however its mean error is fairly high (indicating biased nowcasts). The simple combination of nowcasts shows very similar performance compared to the best possible model in terms of MAE and RMSE, with a slightly higher maximum error. The benefit brought by the nowcasts combination approach is the very low mean error, which means that the combination of nowcasts does not systematically undershoot or overshoot the TIO. Consequently, for the rest of this paper, e.g. when we look at the results for quarterly GDP growth, we focus on the nowcasts obtained by combining different model predictions.

The main target of our nowcasts is the first version of the TIO. This is because the later versions of this series are adjusted both for prediction errors and for additional benchmarking, meaning that we cannot be sure whether the nowcast error is due to the mistake in the prediction or because of some subsequent benchmark. However, it is still interesting to check the performance of our nowcasting framework against the final version of TIO, also because it allows us to compare our revision error against the one based on Statistics' Finland publications. We first plot the nowcasts obtained by combining the original predictions, together with the latest version of TIO. We also plot the first version of TIO against the final revision available.

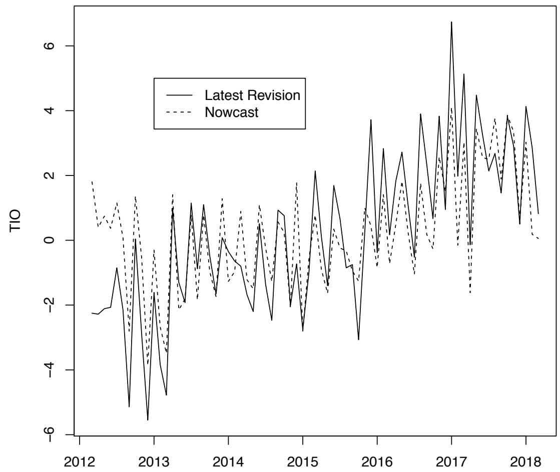

(a) TIO year-on-year growth, final version and nowcasts combination.

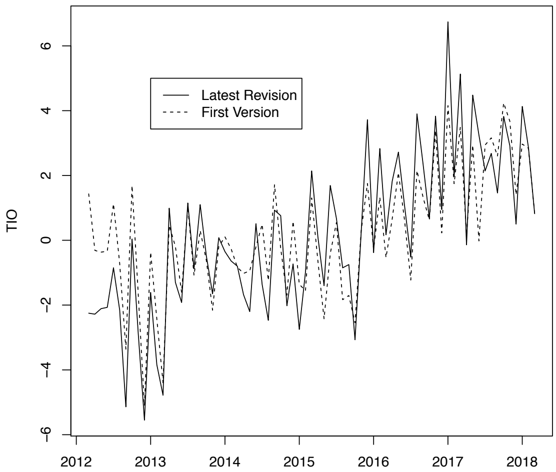

(b) TIO year-on-year growth, final version and first publication.

Figure 3.5: TIO year-on-year growth rate, first publication, final version available and nowcast. The set of predictors is based on firm-level turnovers.

Figure 3.5 (a) shows a lower nowcasting performance for our approach, which is expected, given that the TIO series we use in the estimation of our model has substantial difference from its later revisions. This can be seen from Figure 3.5 (b), where we depict the first and final version of TIO: the difference between the two series is remarkable, especially for certain periods. For example, the first official release of the year-on-year growth of TIO for June 2017 was -0.02 percentage point, which was then revised to 3.25 percentage points (interestingly, our nowcast for this month is much closer to the final value of TIO than the first release of Statistics Finland). While such extreme revisions are not common, they do show the difficulties in creating flash estimates of real economic activity. Next, in Table 3.2, we report the predictive performance measures for the nowcast combination approach, using the final value of TIO as target, even though we still use the original vintages of TIO in the estimation. We also report the same measures to evaluate the performance of the Statistics Finland's first publication.

	Combination	Statistics Finland's first
ME	-0.03	-0.03
MAE	1.15	0.95
RMSE	1.50	1.20
MaxE	4.19	3.64

Table 3.2: ME, MAE, RMSE and MaxE for the nowcast combination approach and for the Statistics Finland's first publication of TIO. The target is the latest available version of the year-on-year growth of TIO. The set of predictors is based on firm-level turnovers.

The performance measures reported in Table 3.2 confirm the fact that our nowcasting approach fares worse when it is evaluated using the latest revision of TIO. However, it is interesting to see that the predictions of our simple nowcasting combination do not show a much larger revision error compared to the first publication of Statistics Finland (which suffers from a much longer publication lag).

So far, we have evaluated the performance of nowcasts based on firm-level turnovers, the core predictors of this study. However, as mentioned before we have also constructed flash estimates based on measurements of trucks' traffic volumes. First, we report the plots of the predictions obtained by simple model combinations, where we use a similar procedure as the one explained in Section 3.4. We depict both the nowcasts against the first version of TIO and compared to the latest available revision, in Figure 3.6.

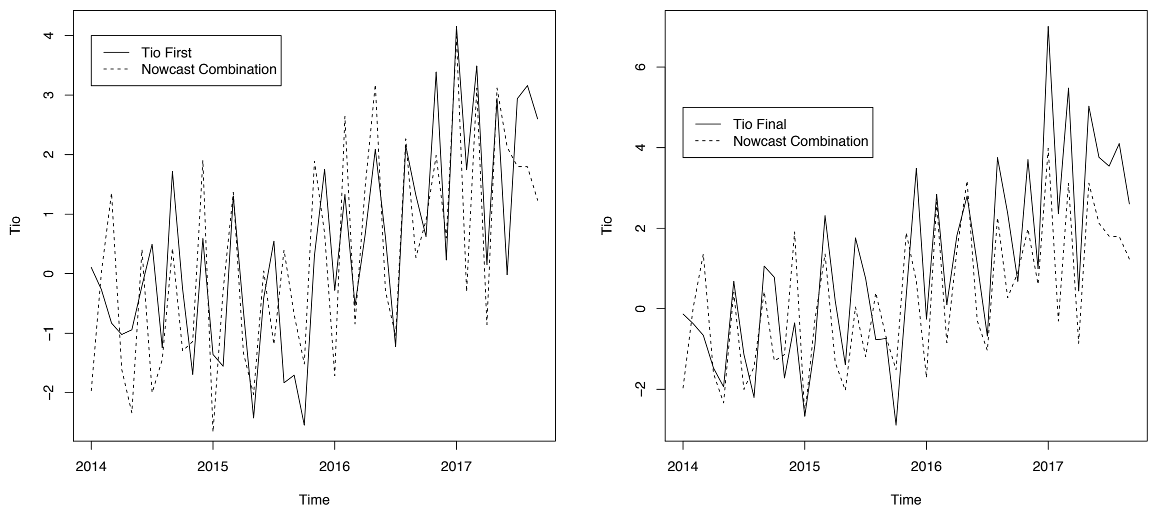

(a) TIO year-on-year growth, first version and nowcasts combination.

(b) TIO year-on-year growth, final version and nowcast combination.

Figure 3.6: TIO year-on-year growth rate, first publication, final version available and nowcasts. The set of predictors is based on trucks' traffic volumes.

While there are still some substantial nowcasting errors, it is impressive that an unstructured and peculiar data source such as traffic volumes is able to provide estimates that track economic activity fairly well. To gain a better grasp of how our approach is performing, we report the nowcast error measurements that we have used throughout the report in Table 3.3, both for the first and final version of TIO.

	Combination vs. First	Combination vs. Final	ARIMA vs. First	ARIMA vs. Final
ME	-0.07	-0.09	0.11	-0.14
MAE	0.86	1.19	1.36	1.68
RMSE	1.09	1.58	1.79	2.23
MaxE	3.16	4.44	5.85	5.76

Table 3.3: ME, MAE, RMSE and MaxE for the nowcast combination approach, evaluated using the first version of TIO growth and its latest available version.The set of predictors is based on trucks' traffic volumes.

Table 3.3 gives us some really interesting insights. With respect to the first version of TIO, the nowcasts combination based on traffic data provides slightly worse predictions, at least compared to the sales' data. However, the MAE and MaxE are fairly low, and much lower than the ones of the automated ARIMA model, indicating a satisfactory nowcasting performance. When looking at the nowcasts error with respect to the final version of TIO, the gap between the performance of traffic data based nowcasts and the predictions computed using firm-level data narrows, even though the latter remains superior. We have also checked whether merging the firm-level and traffic volumes datasets improves our nowcasts. The predictions obtained are very similar to the ones produced by using firm-level data, hence we do not report these results (they are available upon request).

To summarize the results of this subsection, we have seen that combining firm-level data with statistical models and machine learning techniques that are able to deal with large dimensional datasets provide fairly accurate nowcasts, both with respect to the first and to the final version of TIO. The good predictive performance is matched with a substantial gain in timeliness, around 30 days compared to the current publication schedule. The results for the estimates based on traffic volumes evidence the potential of this data source. While the predictions are slightly worse then the ones based on firm-level data, especially compared to the first release of TIO, the errors are not extremely large. Notably, the maximum revision error obtained from this data source is even lower than the one of the first Statistics Finland's publication. The potential real-time availability of traffic data, combined with their satisfactory nowcasting performance, indicates that it is a data source that should be studied further.

3.4.2 Results for quarterly GDP nowcasts

We now turn to the results regarding the estimation of quarterly GDP year-on-year growth, in real terms. In particular, we nowcast the $t + 60$ release of GDP, which is the first official release made by Statistics Finland. As we did for TIO, we start by plotting our nowcasts (based on the nowcast combination procedure), against the GDP year-on-year growth. We do this for the nowcasts computed during the second month of the quarter, the ones produced during the third month and finally the nowcast computed 16 days after the reference quarter. The nowcasts are provided for the period going from 2012 Q2 until 2018 Q4.

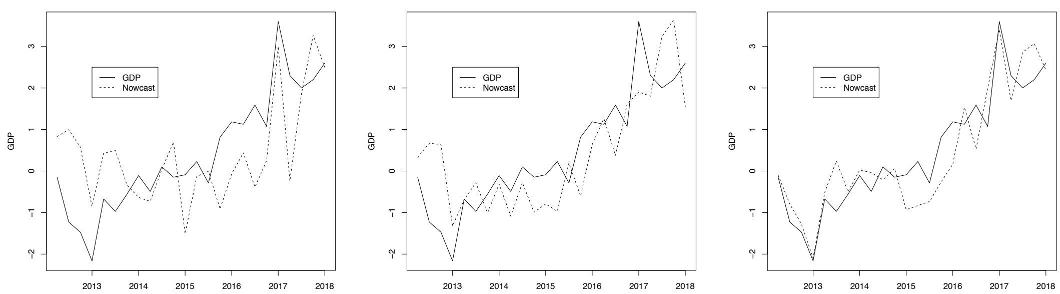

(a) Nowcasts second month (b) Nowcasts third month (c) Nowcasts 16 days after

Figure 3.7: GDP year-on-year growth rate, first publication, and the nowcasts' combination. The set of predictors is based on firm-level sales.

Figure 3.7 indicates that the estimates of TIO based on our nowcasting approach provide good predictions for GDP growth, in a timely fashion. The performance of our models seems to be particularly strong when we compute the predictions during the third month of the quarter and 16 days after the end of the quarter, providing us a 45 to 75 days reduction in the publication lag. Next, we report the nowcasting performance measures for these three sets of predictions. We also compare our results against the performance of the Statistics Finland's flash estimate of GDP, which is based on the arithmetic average of TIO year-on-year growth for the three months of the reference month, and it is published 45 days after the reference quarter. Notice that even in this application, we are using only the vintage of data which would have been available at the time the nowcasts or flash estimates were to be computed.

	Nowcast second month	Nowcast third month	Nowcasts 16 days after	StatFi Flash
ME	0.24	0.03	0.00	-0.04
MAE	0.82	0.66	0.50	0.50
RMSE	1.00	0.85	0.63	0.64
MaxE	2.13	1.86	1.15	1.45

Table 3.4: ME, MAE, RMSE and MaxE for the nowcast combination approach, evaluated using the first version of quarterly GDP year-on-year growth. The set of predictors is based on firms' sales. Nowcast second month refers to the estimates of GDP computed during the second month of the reference quarter, nowcast third month are the estimates computed during the third month of the quarter and nowcasts 16 days after are computed after the end of the reference quarter.

Looking at Table 3.4, we see that our nowcasting framework is able to predict GDP accurately. As we can expect, the performance of the models improves the later we compute the nowcasts and, from the second estimate onward. In particular, the latest estimates presents a very similar (actually slightly better) performance compared to the Statistics Finland flash estimates, providing a 30 days reduction in publication lag.

Finally, we examine the performance of the nowcasts based on traffic data. We start by depicting plots similar to the ones in Figure 3.7, i.e. we report the predictions computed during the second and third month of the reference quarter, together with the 16 days after the end of the quarter estimates.

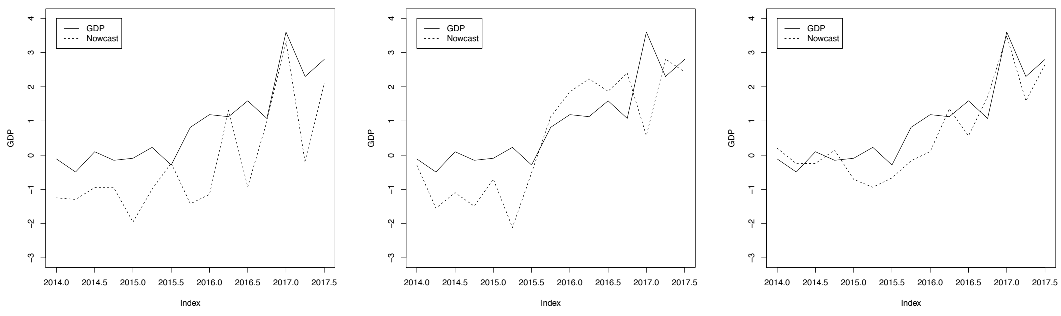

(a) Nowcasts second month. (b) Nowcasts third month. (c) Nowcasts 16 days after.

Figure 3.8: GDP year-on-year growth rate, first publication, and the nowcasts combination. The set of predictors is based on truck's traffic volumes.

Figure 3.8 confirms the promising performance of traffic data for the production of early estimates of GDP, especially for the $t+16$ nowcasts. To assess in a more formal way the performance of our nowcasts, we report the error measures as before.

	Nowcast second month	Nowcast third month	Nowcasts 16 days after	StatFi Flash
ME	0.17	0.07	-0.01	-0.04
MAE	0.83	0.66	0.51	0.50
RMSE	0.99	0.85	0.66	0.64
MaxE	2.07	1.95	1.43	1.46

Table 3.5: ME, MAE, RMSE and MaxE for the nowcast combination approach, evaluated using the first version of quarterly GDP year-on-year growth. The set of predictors is based on trucks' traffic volumes. Nowcast second month refers to the estimates of GDP computed during the second month of the reference quarter, nowcast third month are the estimates computed during the third month of the quarter and nowcasts 16 days after are computed after the end of the reference quarter.

The results of Table 3.5 confirm that the nowcasts produced using traffic date have a satisfactory predictive performance, very similar to the one based on firm-level sales. Overall, it is interesting to see that traffic data are allowing us to create precise estimates of GDP growth well before the official publication by Statistics Finland.

The quarterly results reported in this subsection highlight the ability of models based on firm-level data and traffic data to provide accurate estimates of GDP growth. Even if the very early estimates, the ones computed during the quarter of reference, exhibit substantial nowcasting errors, the performance of our framework becomes significantly better when we consider the predictions at $t+16$. While these flash estimates occur after the end of the quarter of reference, they allow for a 45 days reduction in the publication lag compared to the first official release (and 30 days reduction w.r.t. the Statistics Finland's flash estimate), which represents a substantial improvement.

3.5 Conclusions

We have examined the potential of large micro-level datasets, in combination with statistical models and machine learning techniques that are able to handle high-dimensional information sets, for the production of faster estimates of real economic activity indicators, both at the monthly and at the quarterly frequency. In particular, we have examined the nowcasting performance of firm-level data, and of trucks' traffic volumes measurements.

We find that a simple combination of the nowcasts obtained from a large set of machine learning techniques and large dimensional statistical models is able to produce accurate estimates of monthly real economic activity, or at least estimates that do not lead to a much larger revision error compared to the current official publications. While the revision errors do not increase substantially, our approach allows for a reduction in the publication lag of roughly 30 days, when considering the monthly indicator. Turning to the results related to quarterly GDP, we find that our nowcasts would produce accurate estimates of GDP growth during the third month of the reference quarter, even though there are few large errors. On the other hand, the nowcasts computed at $t+16$ do not show large revisions, or at least produce revisions that are compatible with the ones of Statistics Finland. Even though these estimates would be released after the end of the quarter, they still allow for more than a month reduction of the publication lag. Finally, it is important to underline the satisfactory performance of traffic measurements data. The potential of this source of information should be explored further, given its real-time availability.

In the Finnish setting, the traffic loop data is open to the general public, while the firm level data is collected for the purpose of official statistics production and is subject by strict confidentiality standards. However, similar data collections exist in the other statistical offices of most countries, making our proposed approach and data source an interesting possibility for data users who need timely information on the state of the economy. Statistical offices have the possibility to increase their own relevance as information producers by using this kind of novel techniques. The relatively small investments that are required are related to modeling skills (in maintaining and updating the models) and adding a few features in the existing IT systems for storing information on the models, results and source data. The users of these types of estimates should be regularly informed about the expected and realized nowcast errors and revisions in the target indicators.

3.6 Appendix: Model List

In the table below we report the list of the main model families used in this study. Notice that we do not report every specifications (which depend also on the type of input variables used), thus the number of models reported here are fewer than what is mentioned in the main text (130 specifications). For each model family we report the full name, the method name in `caret` (for variations of the same model family, we report the different `caret` names related to the model family) and a reference where the reader can find a description of the technique.

Table 3.6: List of models

Model/Technique	Name in `caret`	Reference
Factor models/principal components regression	pcr	Stock and Watson (2002a)
Independent component regression	icr	Hyvärinen. and Oja (2000)
Ridge regression	glmnet	Hastie et al. (2009), Chapter 3.4.1
Lasso	glmnet	Hastie et al. (2009), Chapter 3.4.2
Elastic-net	glmnet	Zou and Hastie (2005)
Least angle regression	lars	Hastie et al. (2009), Chapter 3.4.4
Bayes Generalized Linear Model	bayesglm	Gelman et al. (2008)
Gaussian process	gaussprLinear, gaussprPoly, gaussprRadial	Williams and Barber (1998)
Partial least squares	kernelpls, pls, simpls	Hastie et al. (2009), Chapter 3.5.2
Bagged MARS	bagEarthGCV	Hastie et al. (2009), Chapter 9.4
Regression Trees	ctree	Hastie et al. (2009), Chapter 9.2.2
Boosting	BstLm, gbm, xgbTree	Hastie et al. (2009), Chapter 10
Random forests	parRF, ranger, RRFglobal	Hastie et al. (2009), Chapter 15
Nearest-neighbors	knn, kknn	Hastie et al. (2009), Chapter 13
Neural network	pcaNNet	Hastie et al. (2009), Chapter 11
Support vector machine	svmLinear, svmPoly, svmRadial svmRadialCost, svmSigma	Hastie et al. (2009), Chapter 12
Penalized regression	penalized, rqnc	Hastie et al. (2009), Chapter 16

Chapter 4

Agency Costs and Firm Productivity

Milo Bianchi*, Henri Luomaranta**

*Toulouse School of Economics and TSM, University of Toulouse Capitole, Toulouse, France.

**Statistics Finland and TSM, University of Toulouse Capitole, Toulouse, France

Abstract

We explore how the separation between ownership and control affects firm productivity. Using Finnish administrative data on the universe of limited liability firms, we document a substantial increase in firm productivity when the CEO obtains majority ownership or when the majority owner becomes the CEO. We exploit plausibly exogenous variations to ownership and control structures, induced for example by shocks to the CEO spouse's health. Extending the analysis beyond typical samples of large public firms, we show that our effects are stronger in medium-sized private firms. We also investigate possible mechanisms and provide suggestive evidence that increased ownership boosts CEO's effort at work.

JEL Classifcation Code: G30, M12, D24, E23, L25.

Keywords: agency costs, firm productivity, CEO ownership

4.1 Introduction

How costly is the separation between ownership and control within a firm? The question is perhaps one of the most fundamental in the study of modern corporations, it has been at the heart of much of the corporate finance literature since at least Berle and Means (1932) and Jensen and Mecling (1976). At the same time, while substantial work has been developed to investigate how firms' decisions (say, investment and financing) are shaped by agency conflicts, direct measures of agency costs are difficult to obtain.

From an empirical viewpoint, a satisfactory answer to this question needs to confront (at least) two major obstacles. The first is data availability. Ideally, the question requires having detailed information on the firms' operations and outcomes, on their employees (in particular, the top management) and on their ownership structure. While both firm micro data and matched employer–employee data are increasingly available, firm ownership structure is typically observed only for listed firms. This limits substantially the scope of the analysis. Listed firms are a tiny minority of the population of firms, they may have specific ownership and control structures, they may face specific regulatory constraints, or, more generally, they may be intrinsically different from the other firms. Relying only on listed firms also makes the empirical exercise difficult as in these firms there is basically always separation between ownership and control, so it is not clear how to define a benchmark in which agency costs are minimized.

The second key challenge is endogeneity. Ownership and control structures are not randomly assigned, and they are often likely to be themselves affected by firm outcomes, or by possibly unobserved factors affecting both firm's outcomes and its governance. This makes it hard to interpret these relations as causal, and to provide clear guidance to the corporate governance policy debate.

This paper provides an estimate of agency costs by exploiting Finnish administrative data covering the universe of limited liability firms. We have access to detailed firm-level data on the firm balance sheet, a rich set of information about its employees and, importantly, the firm' ownership structure in terms of identity and holdings of its shareholders. This offers the unique opportunity to investigate issues of ownership and control in the entire population of firms, and to uncover whether agency conflicts can be even costlier outside the usual sample of listed firms. The exceptional richness of these data and its panel structure will also allow to address in a novel and we believe convincing way some issues related to the endogeneity of ownership and control structures, as we detail below.

Our setting is also interesting in terms of external validity. Finland is a country that scores very highly in terms of corporate governance; for example, it was ranked first in the world by the World Bank's Corporate Governance Index (Kaufmann (2004)). As we will see, our estimates of agency costs are quite large, and we find them remarkable especially in a setting in which, under this perspective, these costs should be minimal.

The logic of our empirical exercise is very simple. We define the person who has control over the firm's operations as the CEO (we explain below our procedure to identify the CEO among the firm's employees), and we say that there is no separation between ownership and control when the CEO is also the majority shareholder (we perform various robustness checks using other thresholds on CEO ownership). We then compare firm productivity (defined, in our baseline specification, in terms of value added per worker) when ownership and control are in the same hands relative to when they are separated.

We start with fixed-effects regressions in which we exploit variations in CEO ownership within the same firm with the same CEO. That is, we compare firm productivity within the same firm-CEO pair in years in which the CEO is the majority owner vs. years in which ownership and controls are separated. In addition to any common time trend, this specification allows to capture any time-invariant characteristic of the firm, of the CEO, and of the firm-CEO matching. In our preferred specification, we show that when the CEO is also the majority owner output per worker is approximately $1,000$ euros larger, which corresponds to a 1.9% increase in labor productivity. The effect is large, as compared for example to an average productivity growth in our sample of 0.7%.

This effect is robust to alternative definitions of our treatment and, as we show, it is related to changes in CEO ownership and not to any change in the ownership structure. The effect is also confirmed when employing alternative measures of productivity and profitability, and when performing specification tests, checking for the possibility of selection bias and for violations of parallel trends.

At the same time, a causal interpretation of these result requires that unobserved heterogeneity is time invariant within a given firm-CEO pair. This assumption may be violated if unobserved pair-specific shocks induce a change in CEO ownership and at the same time they affect future firm productivity. An ideal setting to address these concerns would be one in which the CEO has majority ownership and for exogenous reasons she has to step down as CEO while at the same time keeping her shares. This would induce an exogenous separation between ownership and control within the same firm and the same ownership structure. Our IV procedure attempts to get closer to such ideal situation by exploiting shocks to CEO ownership induced by CEO's retirement, by worsened health conditions of the CEO, and of the CEO spouse.

The CEO retirement decision may be useful as it is partly driven by reasons that are orthogonal to the future productivity of the firm (see Weisbach (1995), Denis and Denis (1995) for studies employing this instrument). At the same time, the decision is voluntary and as such may be related to unobservable confounding factors. We address this by looking at CEO changes induced by shocks to CEO health, which we measure by the amount of health benefits paid out from the Finnish health insurance scheme. Increased health benefits are associated to worsened health conditions. In a similar and somewhat more extreme way, CEO health shocks have been exploited also in Johnson et al. (1985), who use CEO death, and in Bennedsen et al. (2012), who use CEO hospitalization.

When exploiting CEO health shocks, we can allow for direct contemporaneous effects of CEO health on firm productivity as well as for the possibility that past firm performance affects current CEO health. We need however to assume that current CEO health is not directly associated to future firm productivity. In order to relax this assumption, one would like to exploit shocks that induce the CEO to resign, but are completely orthogonal to any dynamics occurring within the firm. We consider shocks to the CEO spouse health, and to make the test even sharper, we restrict to CEO spouses who are not working in the firm and have no direct effect on the firm operations.

These shocks may induce the CEO to leave and, when the CEO is also the majority shareholder, they may induce an exogenous separation between ownership and control. In order to exploit this variation, we cannot fix the firm-CEO pairs; rather, we define pairs in terms of firm and its largest owner, and estimate changes in firm productivity, within the same firm-owner pair, in years in which the owner is also the CEO vs. years in which the two are separated.

The IV estimates confirm our results, showing that firm productivity is significantly larger when ownership and control are in the same hands. Estimated coefficients are similar across specifications, and if anything larger than the OLS counterpart. This is confirmed in various robustness checks. We further investigate the validity of our instruments by performing some placebo regressions in which our instruments are used to induce changes in CEOs not associated to changes in ownership. We show that it is not a change in CEO *per se* that drives our effects, but CEO changes associated to ownership changes.

We then explore whether our effects are heterogenous across firms. In particular, we investigate whether the estimates are similar in large or in listed firms, that are the typical focus of existing studies. We show that agency costs are in fact larger in medium-sized private firms (51-250 employees). We then replicate some existing results showing that in listed firms the effect CEO ownership on productivity is inverted U-shaped, and it is overall negative. We show however that these effects cannot be found outside the sample of listed firms. We believe these results highlight the importance of exploring agency costs outside typical samples. The results one gets in our broader sample are richer, and they suggest that agency costs may be particularly severe in firms that, due to data limitations, are often excluded from corporate governance studies.

Finally, we explore some possible mechanisms trough which agency costs may affect firm productivity (see Stein (2003) for a review). We first consider variables associated to empire building such as investments, assets, capex, acquisition activities, cash holdings, leverage, dividends, and find no significant changes in these variables in relation to our treatment. We then consider variables associated to quiet life. We measure CEO's effort at work by the number of employment relations the CEO has in other firms and by the number of days the CEO has been absent from work. We show both in OLS and in IV regressions that our treatment induces the CEO to take fewer external engagements and fewer days off. While this analysis is preliminary, it suggests that the quiet life hypothesis is a plausible mechanism behind our treatment effects. When the CEO is also the owner, she exerts more effort at work.

Literature The literature has investigated the relationship between CEO ownership and firm performance mostly by focusing on subsamples of listed or very large firms. Morck et al. (1988) document an inverted U-shaped relation between CEO ownership and Tobin's Q on Fortune500 firms; a similar relation is found in McConnell and Servaes (1990) on a sample of listed firms. Lilienfeld-Toal and Ruenzi (2014) show that firms with larger CEO ownership provide larger stock market returns and suggest this is due to reduced agency conflicts. Fabisik et al. (2018) expand the sample to about

1,800 firms in the US and show that the relation between CEO ownership and Tobin's Q is negative. Ang et al. (2000) is one of the few studies investigating small private firms. They define the Jensen and Meckling's zero agency costs benchmark as a situation in which the CEO is the only owner, and show that firms further away from this benchmark are less efficient.

These estimates display substantial variation depending on the sample of firms under study and on the estimation method. Most of the literature relies on cross-sectional comparisons, while effects are hardly significant when adding firm fixed effects, possibly due to limited time-series variation in these samples (Himmelberg et al. (1999), Zhou (2001)). Relative to this literature, our data cover the universe of limited liability firms over a relatively long panel, that allows exploiting significant time-series variations. We estimate our effects not only within firms, but within firms with the same CEO or the same owner.

Our results are also related to the literature on family firms, and in particular to studies investigating how having a member of the family as CEO affects firm value. Existing results are mixed; e.g., Pérez-González (2006), Bennedsen et al. (2007), Bandiera et al. (2017) show a decrease in firm value, while Anderson and Reeb (2003) and Villalonga and Amit (2006) provide a less negative view. Relatively to this literature, we concentrate on the separation between ownership and control and, by fixing the firm-CEO pairs, we can control for the quality of the CEO and of the firm-CEO matching.

More broadly, our work provides distinct and complementary insights to several themes in the corporate governance literature. Relative to studies looking at how CEO characteristics affect firm value (e.g. Bertrand and Schoar (2003)), we keep the identity of the CEO fixed in our baseline analysis and vary her ownership share. Relative to studies on how ownership structure affect firm value (e.g. Edmans and Holderness (2017)), our focus is on CEO ownership, keeping other characteristics of the ownership structure fixed. Lastly, differently from the literature on majority vs. minority shareholders (Shleifer and Vishny (1997)), we focus on the possibility of agency conflicts between the CEO and the (majority) owner.

4.2 Data

We exploit data from the Finnish Longitudinal Owner-Employer-Employee database (FLOWN) constructed by Statistics Finland, which we match with balance sheet information from the business register. We obtain a yearly panel from 2006 to 2014 covering the universe of limited companies (*osakeyhtiö*) in the business sector. Balance sheet data provide a rich set of information on firms' characteristics, operations and performance. The matched employer-employee structure allows to have information on the employees of the firm, and in particular, as we explain below, to identify its CEO. For our purposes, the key distinctive feature of these data is the detailed information on the firms' ownership structure. The Finnish tax authority requires that firms report the identity of the 10 largest shareholders or, if there are more than 10 shareholders, of any shareholder with more than 10% of firm shares. Building on this information, Statistics Finland has identified the ultimate individual shareholders of a given firm.[1]

We exclude one-man companies and we are left with around 110,000 firms. Out of those (measured when the firm first appears in the panel), 84% are micro firms with less than 10 employees, 13% are small (10-50 employees), 2.4% are medium (51-250 employees), and 0.6% are large (>250 employees). Manufacturing firms are 36% of the sample (including construction) while the rest are services (including trade).

CEO We are interested in identifying the CEO in each firm, interpreted as the person who has control on the firm's operations. We employ the following sequential procedure, as e.g. in Queiró (2016). First, we identify a person as the CEO if he or she is explicitly defined as such among the list of employees. This is the case for 7% of the firms. For the remaining firms, we consider those employees identified as having managerial responsibilities, and say that the CEO is the manager with the highest salary. This identifies an additional 30% of the CEOs. For the remaining firms, we look at whether an active entrepreneur (as classified by the tax administration) appears in the list of employees, in which case the person is identified as the CEO.[2] This is the case for 23% of our CEOs. The remaining 41% of the CEOs are defined as the highest paid worker in the firm. As a validation test, we notice that 86% of the CEOs explicitly defined as such (our first criterion) also have the highest salary in the firm.

Ownership In terms of ownership structure, we have some information for 92% of the firms in our sample; on average, we observe 82% of the firm ownership. In our sample, 39.5% have one shareholder, the median number of shareholders is 2, and 29.9% of firms have more than 2 shareholders. In firms with more than one shareholder, the average ownership share of the largest shareholder is 41%.

In order to investigate agency costs, we say that ownership and control are in the same hands in firms where the CEO is also the majority shareholder. We define our treatment variable as the dummy *CEO Owner*, which equals one when the CEO owns more than 50% of the firm shares. This is a simple way to extend the Jensen-Meckling's zero agency costs benchmark mentioned above to situations in which firms have possibly multiple shareholders.

In our sample, the CEO is also the majority owner in 29% of the firms, and 10.5% of the firms experience a change in the treatment, in 5.6% of the cases the CEO obtains majority, and in the remaining 4.9% the CEO looses majority. As intuitive, these changes are more likely to occur in micro and small firms**.** We observe large variations in CEO ownership. Conditional on observing a positive change, the average ownership change is 50%; conditional on a negative change, the average is -43%.[3] Out of these changes, 26% are associated to a change in the majority owner and so in our treatment *CEO Owner*. Conditional on having a change in the treatment, the average ownership change is 73%

[1] Identifying the ultimate owners is complicated also by the possibility of linkages of firms and business owners via holding companies and enterprise groups. Statistics Finland has implemented a procedure to track down the individual owners behind each firm along the ownership chains. See Maliranta and Nurmi (2019) for a detailed presentation of the data.

[2] The tax administration identifies an active entrepreneur in a firm if a person owns at least 30% of the shares and receives a significant income from the firm (at least 9,663 euros in 2006).

[3] In fact, these figures are similar to observed changes in ownership of the largest shareholder (whether or not she is the CEO) for which, conditional on a positive change, the average is 43% and, conditional on a negative change, the average is -38%.

for positive changes and -76% for negative changes.

As these figures suggest, the effects we obtain from our treatment *CEO Owner* are typically not driven by small changes in CEO ownership around the 50% threshold. In fact, we do not assume any specific effect around the threshold, and we perform several robustness checks by considering alternative thresholds. We define *CEO 100*, that is a dummy equal to one when the CEO owns 100% of the firm shares; and *CEO 0*, a dummy equal to one when the CEO owns any positive fraction of firm's shares. We also consider a dummy *CEO Largest*, which equals one when the CEO is the largest (though not necessarily the majority) shareholder, as well as the continuous variable *CEO shares*, that is the fraction of shares held by the CEO.

Productivity Our main interest is to investigate how our treatment affects firm's productivity. In most of our analysis, we define productivity as labor productivity, that is, value added (in real terms) over full time equivalent units of labor.[4] The measure is constructed directly by Statistics Finland in a way that is comparable across firms and over time. The measure estimates the value of the output generated by a worker in the firm without having to take a specific stand on the firm's production function nor to estimate the value of capital in the firm, which may be problematic for some firms in our sample (e.g. micro service firms). It does not measure profit and it does not serve as a tax base, so it may be less subject to discretionary accounting practices.

We will check the robustness of our results when employing other efficiency and profitability measures (described in more details below). We will also consider productivity measures based on standard TFP estimates, and we will account for possible biases induced by the inability to observe firm level prices.

We winsorize all financial variables, including productivity measures, at the 0.25th and the 99.75th percentiles. Descriptive statistics of our variables appear in Table 4.1.

4.3 OLS estimates

4.3.1 Basic results

The first set of results are based on fixed-effects OLS regressions in which we exploit variations in CEO ownership within the same firm with the same CEO. Our basic specification is

$$y_{i,t} = \alpha_i + \beta T_{i,t} + X_{i,t}^{'}\gamma + \mu_t + \varepsilon_{i,t}, \tag{4.1}$$

where i denotes a firm-CEO pair, $y_{i,t}$ is the productivity of firm-CEO i in year t, α_i and μ_t are respectively firm-CEO and year fixed-effects, and $T_{i,t}$, is a dummy equal to one when the CEO owns more than 50% of the firm shares. Our baseline set of controls $X_{i,t}^{'}$ includes industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number

[4]Value added is defined as the value of sales minus the value of purchases, accounting for changes in stocks, other operating incomes and product taxes. An industry specific index based on 2010 prices is used to deflate the nominal value added.

of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. In specifications without CEO fixed effects, we also include CEO's education, age, years of experience within the firm and in total. Our coefficient of interest is β, which measures productivity differences within the same firm-CEO pair in years in which the CEO is the majority owner vs. years in which ownership and controls are separated.

Table 4.2 reports our estimates. In column 1, we include no control and no fixed effects and observe a negative relation between *CEO Owner* and firm productivity. Once we include our set of controls (column 2) and firm fixed effects (column 3), the relation turns positive. Our preferred specification is reported in column 4, which corresponds to equation (4.1) and includes firm-CEO fixed effects.[5] According to these estimates, *CEO Owner* is associated to around $1,000$ euros larger output per worker, that corresponds to a 1.9% increase relative to the unconditional mean. This effect is large. As comparison, the average productivity growth of private sector Finnish firms in our sample period is 0.7% per year.[6]

In column 5, we exclude observations in which the CEO loses ownership ($T_{i,t} - T_{i,t-1} = -1$) and focus on the effect of the CEO becoming a majority shareholder. Similarly, in column 6, we exclude observations in which $T_{i,t} - T_{i,t-1} = 1$ and focus on the CEO losing ownership. These estimates show no significant asymmetries between the two effects, an observation we will use again in our IV estimates.

4.3.2 Robustness

We perform a series of robustness checks, starting by alternative definitions of our treatment. As mentioned, we attach no specific value to the 50% threshold in terms of CEO ownership, and we now consider alternative thresholds. In column 1 of Table 4.3, we consider *CEO 100*, a dummy equal to one if the CEO is the only owner. In column 2, we consider *CEO 0*, a dummy equal to one if the CEO has some ownership. In column 3, we focus on *CEO Largest*, that is a dummy indicating that the CEO is the largest shareholder. In column 4, we consider the continuous variable *CEO shares,* that is the fraction of shares held by the CEO. In all these cases, the effect on firm productivity is similar to our main estimates. In particular, in column 4, we estimate that productivity increases by $1,063$ euros following an average increase of CEO ownership by 50% in positive changes, and about -43% for negative changes. The effect is 6% larger than the coefficient on *CEO Owner* in Table 4.2, which as mentioned corresponds to an average change in CEO ownership of about 74%. This may suggests some concavity in the effect of CEO ownership, but not strong enough to reject linearity.[7] In fact, we explore more explicitly non-linear effects of *CEO shares* in column 5 and find no significant effect. We will investigate again these patterns (and show they are different) in listed firms.

In order to support our interpretation, we make sure that our estimates capture the specific effect

[5] As we include firm-CEO fixed effects, we do not include controls for CEO's education, age, experience; hence the higher number of observations relative to column 3.

[6] The corresponding figure for EU countries is 0.9% and for the US is 1.13%, see the OECD's website at data.oecd.org.

[7] A similar picture emerges from the estimates in columns 1 and 2. The average change in CEO ownership associated to a positive change in *CEO 100* is 78% and it is -75% for a negative change. The average change in CEO ownership associated to a positive change in *CEO 0* is 63% and it is -62% for a negative change. Out of all changes in CEO ownership, 20% of them are associated to a change in *CEO 100* and 64% are associated to a change in *CEO 0.*

of changes in CEO ownership, as opposed to any change in the ownership structure. In column 6, we consider the dummy *Owner Change,* which equals one when the majority owner changes from period $t-1$ to t and in any subsequent period, irrespective of whether or not this is associated to a change in CEO ownership. We show no significant impact on productivity in this case, suggesting that our effects are related to changes in CEO ownership and not to any change in ownership.

In Table 4.4, we report a set of robustness checks concerning our productivity measure. In column 1, we consider gross operating surplus (GOS), defined as VA minus personnel costs per unit of labor. In column 2, we consider net profit margin, defined as net profit (VA minus personnel cost, overheads and other costs, interest and taxes) over revenues. In column 3, we consider returns on assets, defined as standard as net income over total assets. These regressions confirm that *CEO Owner* is associated to higher operating efficiency and profitability.

In column 4, we consider a standard estimate of TFP, obtained as the residual of a Cobb-Douglas in which value added is regressed over capital and labor for each 2 digit industry. In column 5, we estimate TFP by adding the firm's market share and fixed effects at the industry-year level. Controlling for industry-specific time trends is a simple way to account for possible biases due to inability to observe firm prices (see Beveren (2012) and De Loecker and Goldberg (2014) for excellent reviews). In addition, in column 6, we exclude multiproduct firms that may be subject to price shocks in different industries. Again, we observe a positive relation between *CEO Owner* and firm productivity, and our coefficient of interest barely changes across these specifications.[8]

Finally, we perform some specification tests. We start by considering possible biases due to sample selection. Importantly, our sample is not selected in the sense that at each point in time we consider the universe of firms, we do not restrict to survivors. Our fixed effects specifications in equation (4.1) may also mitigate sample selection biases (Verbeek and Nijman (1992)). As additional checks, we repeat our regressions in equation (4.1) on various *selected* samples. The results are presented in Table 4.5. In column 1, we restrict the sample to *No Exit* firms, these are firms that do not die in our sample. In column 2, we consider *Persistent* firms, defined as those firms with number of observations above the median, that is equal to 9 (that is, half of our firms are observed for the entire sample of 9 years). We repeat the same procedure in terms of firm-CEO pairs, considering in a similar way *No Exit* firm-CEOs (column 3) and *Persistent* firm-CEOs, where the median number of observations for firm-CEOs is 4 (column 4). The estimated impact of *CEO Owner* is similar across the various specifications, and not statistically different from our baseline estimates on the entire population. This further limits the concerns that our effects are biased due to sample selection.

In columns 5 and 6, we consider specifications in which, instead of firm-CEO fixed effects, we control for lagged values of the dependent variable (one lag in column 5, and three lags in column 6). These specifications are more appropriate if unobserved characteristics are not time invariant within a given firm-CEO pair, but they are instead better captured by time-varying individual-specific past productivity patterns. Estimated effects are still positive and (not significantly) smaller in size.[9]

[8] Similarly, controlling for industry-year fixed effects and market share in our baseline regressions on labor productivity has no effect on our coefficient of interest.

[9] See e.g. Guryan (2004) for a discussion on how fixed effects and lagged dependent variable specifications provide

Finally, as standard in diff-in-diff specifications, a causal interpretation of our estimates requires that treated and control units are not exposed to different trends before the treatment. In order to make sure this is the case, we consider the following regression

$$y_{i,t} = \alpha_i + \sum_{s=1}^{4} \beta_{-s}\mu_{t-s}T_{i,t} + \sum_{s=1}^{4} \beta_{+s}\mu_{t+s}T_{i,t} + X_{i,t}^{'}\gamma + \mu_t + \varepsilon_{i,t}, \tag{4.2}$$

where μ_{t-s} and μ_{t+s} correspond to years before and after the treatment and the other variables are as in (4.1). Figure 1 reports the estimated coefficients $\beta_{-4}, ..., \beta_4$ and the associated 95% confidence intervals. Within the same firm-CEO pair, there are no significant pre-treatment differences, which supports the parallel trend assumption.

If we assume that unobserved characteristics are time invariant within a given firm-CEO pair, our estimates in (4.1) can be interpreted in a causal sense. These estimates however cannot account for firm-CEO specific shocks that may induce a change in the treatment and at the same time they affect future productivity. We address these concerns in the next section.

4.4 IV estimates

A causal interpretation of our OLS estimates may be challenged for example on the basis that the CEO may have private information on the future firm productivity, and decide to acquire majority shares in anticipation of a productivity increase. More generally, changes in ownership and control structures may be correlated to unobserved pair-specific shocks that may be also correlated to future productivity. Ideally, one would like to exploit purely random changes in CEO ownership. For example, one would like to observe a firm in which the CEO has majority ownership and for exogenous reasons she has to step down as CEO while at the same time keeping her shares. The shock would exogenously generate a separation between ownership and control within the same firm and the same ownership structure. In the next analysis, we attempt to get as close as possible to such ideal situation by exploiting shocks to CEO ownership induced by CEO's retirement, by worsened health conditions of the CEO, and of the CEO spouse.

4.4.1 Instruments

Our first instrument exploits changes in the CEO due to retirement. The retirement decision is partly driven by reasons that are orthogonal to the future productivity of the firm and, in fact, it has been used by the literature to investigate the effects of CEOs on firm value (Weisbach (1995), Denis and Denis (1995)). We define the dummy *CEO Retire* that equals one if the CEO is older than the legal retirement age (63 years old) or receives pension benefits at t.

A potential issue with retirement is that its decision is voluntary and as such may be related to unobservable confounding factors. For example, a CEO may decide to retire when she expects a decline in firm productivity. We address this concern by considering a second instrument, based on shocks to

bounds for the estimated causal effect.

CEO health. For each CEO, we obtain the amount of health benefits paid out from the Finnish health insurance scheme. The scheme is mandatory and universal and it compensates the beneficiary for income losses related to health issues. As such, an increase in health benefit is due to worsened health conditions. Relative to CEO changes induced by retirement, health shocks are less likely to be driven by expected productivity shocks. The logic of the instrument generalizes, in a somewhat less extreme way, a classic approach of using CEO death as a shock (Johnson et al. (1985)) and, more recently, the approach by Bennedsen et al. (2012), who use CEO hospitalization events in Danish firms.

We use CEO health at $t-1$ as an instrument for changes in the CEO from $t-1$ to t. The validity of our instrument does *not* rely on excluding direct effects of CEO health at t on firm productivity at t, we use past health shocks to induce changes in CEO. One may also conjecture that past firm performance may affect current CEO health. If the CEO changes associated to our health shocks were driven by past firm productivity, however, we would observe a violation of parallel trends, which as shown above is not the case. A remaining issue may be that current CEO health is directly associated to future firm productivity.

In order to take this possibility into account, one should consider health shocks that induce the CEO to resign, but are completely orthogonal to any dynamics occurring within the firm. One such case is a shock to the CEO spouse health. The exceptional richness of the data allows us to recover the amount of health benefit paid to the CEO spouse, again by the national health insurance scheme. In fact, to make this test even sharper, we can restrict to CEO spouses who are not working in the firm and so have no direct effect on the firm operations. To our knowledge, this instrument is novel and in our view considerably less exposed to the above mentioned concerns.

4.4.2 Specifications

Before turning to our IV estimates, we start with an OLS estimate of

$$y_{i,t} = \alpha_i + \beta T_{i,t} + X_{i,t}^{'}\gamma + \mu_t + \varepsilon_{i,t}, \tag{4.3}$$

in which all terms are as in equation (4.1) except that we define a pair i in terms of a firm and its largest owner. In equation (4.3), the coefficient β describes what happens to firm productivity, fixing the firm and its largest owner, in years in which the owner is also the CEO vs. years in which the two are separated. While equation (4.1) exploits variations associated to the CEO becoming (or stopping being) the owner, equation (4.3) exploits variations associated to the owner becoming (or stopping being) the CEO. For the purpose of estimating agency costs, both variations should lead to similar insights. There are two reasons to focus on specification (4.3) for the next analysis. First, as further discussed below, it helps addressing the above mentioned concerns about the CEO having private information about the future profitability of the firm. Second, our instruments are shocks that may induce the CEO to leave and exploiting them requires that the CEO is not fixed in our analysis.

In order to implement our IV approach, we consider two specific features of our setting. First, our variable of interest $T_{i,t}$ is a dummy. For this reason, we first estimate a probit regression in which $T_{i,t}$

is regressed over a given instrument $Z_{i,t-1}$ and a set of controls. Then, we use the predicted $\hat{T}_{i,t}$ as an *instrument* in a standard 2SLS regression. As shown in Wooldridge (2010), this allows improving the efficiency of our estimator and obtaining an estimate of the *average* treatment effect, that is easier to compare to OLS estimates. Moreover, the procedure is robust to possible misspecifications in the probit equation and it does not require considering generated regressor issues.

A second observation is that the effect of a given instrument $Z_{i,t-1}$ on our treatment $T_{i,t}$ depends on $T_{i,t-1}$. When the CEO is the owner at $t-1$, the instrument (say, a shock to CEO health) may induce the CEO to leave and so (if anything) a negative change to the treatment, from $T_{i,t-1} = 1$ to $T_{i,t} = 0$. When the CEO is not the owner at $t-1$, the instrument may induce (if anything) a positive change in the treatment, from $T_{i,t-1} = 0$ to $T_{i,t} = 1$.

Accordingly, our IV estimates are based on the following procedure. First, we estimate the probit regression

$$T_{i,t} = \Phi(\alpha + \beta_1 Z_{i,t-1} + \beta_2 Z_{i,t-1}(1 - T_{i,t-1}) + \beta_3(1 - T_{i,t-1}) + X_{i,t}^{'}\gamma), \tag{4.4}$$

in which $Z_{i,t-1}$ is one of the above mentioned instrument and in which β_1 measures the effect of the instrument on $T_{i,t}$ when $T_{i,t-1} = 1$. This case is of particular interest, as the instrument induces a plausibly exogenous separation between ownership and control. As mentioned, we then use the predicted $\hat{T}_{i,t}$ as instrument in a 2SLS in which the first stage is a standard OLS.

4.4.3 Results

We present our results in Table 4.6. In column 1, we report OLS estimates of equation (4.3), showing that, in the same firm with the same owner, firm productivity is larger when the owner is also the CEO. As mentioned, the result is useful to address the concern with specification (4.1) that the CEO may decide to acquire ownership as she expects an increase in future profitability. In equation (4.3), instead, it is the owner who decides to become the CEO and the variation is less likely to be driven by the CEO's private information.[10] This result also confirms our estimates in Table 4.2 and it serves as a useful benchmark for the next IV estimates.

The results of our IV procedure are reported in columns 2-5. The bottom part of the table reports the probit estimates of equation (4.4), not the first stage of the 2SLS. The coefficient on $Z_{i,t-1}$ is negative, showing that our instruments have a significant impact on the treatment. If the CEO is the owner at $t-1$ and, for example, she becomes sick, she is more likely to leave and so induce a negative shock to the treatment. In column 2, the instrument is a dummy equal to one if the CEO has retired. In column 3, the instrument is the amount of health benefits received by the CEO at $t-1$ (in 10,000 euros). In column 4, the instrument is the amount of health benefits received by the CEO spouse.[11] In column 5 the sample is restricted to cases where the CEO spouse is not an employee of the firm. Results in columns 2-5 reveal a robust effect. Exploiting plausibly exogenous shocks, we show that firm productivity is significantly larger when ownership and control are in the same hands. Estimated

[10] The CEO may decide to leave as he expects future productivity to decrease, but this would go against our results.

[11] In order to keep the same sample throughout columns 2-4, we set health benefit to zero when the CEO has no spouse (that is, we make no distinction between having a spouse with no health benefits and having no spouse). Restricting our sample to CEOs with a spouse would give very similar estimates in terms both of magnitude and of standard errors.

coefficients are similar across specifications, and if anything larger than the OLS counterpart.

The validity of our instrument requires that our shocks affect firm productivity at t only thought the induced change in the treatment. In order to investigate this further, we check whether any change in CEOs induced by our instruments, whether or not it is associated to an ownership change, has a similar effect on firm productivity. If this were the case, the validity of our instrument would be questioned.

We report our results in Table 4.7, which replicates the structure of Table 4.6 but looks instead of the effect of *CEO Change*, that is a dummy equal to one if the CEO changes from $t-1$ to t and in any subsequent period. In column 1, we report OLS estimates and show no significant effect. This can be seen as a placebo test of the effect in column 1 of Table 4.6, showing that it is not a change in CEO *per se* that drives our effects on firm productivity, but CEO changes associated to ownership changes. This is confirmed in IV estimates in columns 2-5. The estimation procedure is the same as in Table 4.6, except that there is no need to consider the interaction between $Z_{i,t-1}$ and *CEO Change* at $t-1$ in the probit. The probit coefficient on $Z_{i,t-1}$ is positive, showing that our instruments indeed significantly increase the probability of having a change in the CEO. Importantly, however, these changes have no significant impact on firm productivity unless they are associated to changes in ownership, as described by our treatment. We view this as an important finding in support of the validity of our instruments.

As mentioned, our health instruments do not rely on excluding direct effects of CEO health at t on firm productivity at t. In columns 1-2 of Table 4.8, we add health at t (that is, $Z_{i,t}$) as control in our 2SLS estimates. In column 1, we see that CEO health at t has a negative impact on firm productivity at t, while in column 2 the impact of CEO spouse health is not significant. Irrespective of these effects, our estimated impact of *CEO Owner* is not affected. In our specifications, we fix the firm's largest owner, and any variation to the treatment $T_{i,t}$ is due to changes in the identity of the CEO. In this case, CEO health at time $t-1$ is not correlated to health at t, precisely because the CEO is not the same.[12]

We also consider alternative definitions of our treatment *CEO 100* and *CEO 0*, as defined in Section 3.2. Columns 3-4 present OLS estimates with firm-owner fixed effects, columns 5-6 present IV estimates employing CEO spouse health as instrument and restricting to CEO spouses not working in the firm. Results are robust and consistent in all these specifications.

Finally, in Table 4.9, we consider alternative IV specifications. In columns 1-3, we consider standard 2SLS estimates in which each instrument $Z_{i,t-1}$ is directly used in the first stage. In columns 4-6, we use probit regressions and the predicted $\hat{T}_{i,t}$ as instrument, but differently from equation (4.4) we include no interactions with $T_{i,t-1}$. As instruments, we consider CEO retire, CEO health, and CEO spouse health when the spouse is not employed in the firm. Estimated impacts of *CEO Owner* are still positive and significant, confirming the robustness of our findings.

[12] If this were not the case, we could have for example cases in which the CEO gets sick at $t-1$, she does not step down, but she rather sells her majority shares. We would observe a change in the treatment, but not a change in the CEO, which may be problematic since for a given CEO health at $t-1$ is likely to be correlated to health at t and CEO health at t may in turn affect firm productivity at t.

4.5 Interpretation

4.5.1 Representativeness of typical samples

As mentioned, most of the literature on CEO ownership focuses on samples of very large and/or listed firms. A question is whether the effects identified on those firms are representative of the population. We explore this issue in Table 4.10.

In column 1, we check whether our estimates of agency costs vary with the size of the firm. We interact our treatment indicator with the dummies Small, Medium, and Large, indicating respectively that the firm has 10-50 employees, 51-250 employees, or more than 250 employees. The omitted category are micro firms with less than 10 employees. Estimated agency costs appear largest for medium-sized firms.

In column 2, we interact our treatment with a dummy indicating whether the firm is in the service (as opposed to the manufacturing) sector. We observe that agency costs are significantly larger in manufacturing firms.

In order to highlight the effects on listed firms, we consider the continuous measure *CEO share* instead of *CEO Owner*. In listed firms, it is hardly the case that the CEO is the majority shareholder. As shown in column 3, the effect on non listed firm is positive (and not surprisingly similar to the one in Table 4.3) while the effect on listed firm is negative. That is, differently from the vast majority of firms, larger CEO ownership is associated to lower productivity in listed firms. The result is consistent with Fabisik et al. (2018) who focus on listed firms.

It has also been shown that, in listed US firms, the relation between CEO ownership and firm value is inverted U-shaped (Morck et al. (1988), McConnell and Servaes (1990)). Indeed, if we restrict to listed firms, this is the case in our sample as well (column 4). As we have shown in Table 4.3, however, such non-linearity is not so strong (in fact, the squared term is not significantly different from zero) in the broader sample including non listed firms.

These results highlight the importance of exploring agency costs outside typical samples. The results one gets in our broader sample are richer, and they suggest that agency costs may be particularly severe in firms that, due to data limitations, are often excluded from corporate governance studies.

4.5.2 Mechanisms

Finally, we would like to investigate some possible mechanisms trough which agency costs affect firm productivity. Indirectly, these results can also shed light on which types of agency costs matter the most in our setting (see e.g. Stein (2003) for a review). We distinguish in particular between empire building, according to which agency costs are driven by the manager's tendency to undertake inefficient projects (Jensen (1986)), and quiet life, according to which agency costs are driven by the manager's tendency not to put effort at work (Bertrand and Mullainathan (2003)).

We first consider variables associated to empire building. Specifically, we test whether *CEO Owner* is associated to changes in investments, assets, capex, acquisition activities, cash holdings, leverage, dividends. None of these variable appear significantly related to our treatment.

We then consider variables associated to quiet life. While direct measures of CEO's effort are hard to find, we can observe the number of employment relations the CEO has in other firms (e.g. second job, board membership, or consultancy). We can also observe the number of days the CEO has been absent from work, typically due to sick leave or for study reasons. In Table 4.11, we report our estimates fixing the firm-CEO (columns 1 and 4), fixing the firm-owner (columns 2-3 and 5-6), and the IV as in Table 4.6 with CEO spouse health as instrument, restricting to CEO spouses not working in the firm (columns 5-6). We observe that our treatment induces the CEO to take fewer engagements outside the firm and fewer days of absence from work, which is consistent with increased effort in the firm.

While this analysis is preliminary, it tends to provide support to the quiet life hypothesis as a plausible mechanism behind our treatment effects. When the CEO is also the owner, she exerts more effort at work.

4.6 Conclusion

We have shown that agency costs are an important determinant of firm productivity. This result has been established both in OLS regressions with firm-CEO or firm-owner fixed effects and in IV regressions in which we exploit health shocks of the CEO and of the CEO spouse as a source of exogenous separation between ownership and control. We believe this result is important as it establishes in a direct way the magnitude and scope of agency costs.

The possibility to exploit ownership data on the universe of limited liability firms has allowed us to estimate agency costs also in samples which had not been investigated in the previous literature. We have found agency costs to be particularly important in medium-sized private firms that are usually not the main concern for corporate governance regulation. We hope this can serve as motivation for similar data collection efforts and investigations in other countries.

4.7 Tables and Figure

Variable	Obs.	Mean	Std. Dev.
LP	566,266	53,695	41,844
CEO Owner	566,266	0.29	41.18
CEO 100	566,266	0.18	0.39
CEO 0	566,266	0.57	0.5
CEO Largest	566,266	0.39	0.49
CEO Share	566,266	0.36	0.39
GOS	565,526	15,540	35,677
Profit	565,302	-0.018	0.341
ROA	564,847	2.284	28.94
TFP	556,013	0.0009	0.51
TFP2	556,000	-4.4e-10	0.54
CEO changes	566,266	0.36	0.47
Owner changes	566,266	0.08	0.28
Dependent	566,266	0.12	0.33
HHI ownership	566,266	5,484	3,656
Workers w/ Bac	566,266	10.88	94.1
Workers w/Master	566,266	1.37	24.82
Workers w/ PhD	566,266	0.07	1.74
White Collars	566,266	5.05	62.04
Blue Collars	566,266	11.09	117.91
CEO tenure	555,431	7.04	6.74
CEO age	566,266	44.68	10.33
CEO experience	536,651	19.46	5.31
CEO job mobility	566,266	15.25	23.46
CEO retires	566,266	0.04	0.197
CEO health benefits	566,266	98.83	978.19
Spouse health benefits	566,266	94.52	890.12
Services	566,266	0.64	0.47
Micro	566,266	0.78	0.42
Small	566,266	0.18	0.39
Medium	566,266	0.03	0.18
Large	566,266	0.008	0.09
Listed	566,266	0.001	0.038
Free cash flow (1000)	566,266	121	274
Capex (1000)	566,266	64	221
Dividends	566,264	32,076	214,335
Leverage (D/E)	566,260	77,016	471,247
Investments	566,266	87,693	546,818
Acquisition activities	566,266	0.005	0.07
Assets (1000)	566,260	5,132	151,000
CEO Engagement	555,688	1.241	0.739
CEO days leave	561,715	7.60	36.31

NOTE: This table reports summary statics of all the variables used in our analysis. Minimal and maximal values cannot be reported due to confidentiality.

Table 4.1: Descriptive Statistics

Dep Variable	Labor Productivity					
	(1)	(2)	(3)	(4)	(5)	(6)
CEO Owner	-3,556	1,491	856.48	1,001	958.65	1,057
	(-33.75)***	(13.19)***	(5.98)***	(3.30)**	(2.47)**	(3.11)***
Controls	No	Yes	Yes	Yes	Yes	Yes
Fixed Effects	No	No	Firm		Firm-CEO	
Sample			All		0/1	0/-1
Number of Obs	566,266	555,431	555,431	566,260	557,372	557,438
Number of Groups			109,503	214,077	214,077	214,077
R-squared	0.002	0.17	0.01	0.01	0.01	0.01

NOTE: This table reports the results of OLS regressions. The dependent variable is labor productivity. CEO Owner is a dummy equal to one if the CEO has majority ownership in the firm. In column 3, regressions include firm and year fixed effects. In columns 4-6, regressions include firm-CEO and year fixed effects. In column 5, we exclude observations in which the CEO loses ownership. In column 6, we exclude observations in which the CEO gains ownership. In columns 4-6, controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. In addition, in column 3, controls include CEO's education, age, years of experience within the firm and in total. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.2: Main Result

Dep Variable	Labor Productivity					
	(1)	(2)	(3)	(4)	(5)	(6)
CEO 100	867.62					
	(2.60)***					
CEO 0		586.33				
		(2.84)***				
CEO Largest			1,380			
			(5.13)***			
CEO shares				1,063	1,755	
				(4.06)***	(1.91)*	
CEO shares squared					-769.81	
					(-0.81)	
Owner changes						241.52
						(0.99)
Number of Obs	566,260	566,260	566,260	566,260	566,260	566,260
Number of Groups	214,077	214,077	214,077	214,077	214,077	214,077
R-squared	0.013	0.013	0.013	0.013	0.012	0.013

NOTE: This table reports the results of OLS regressions. The dependent variable is labor productivity. CEO 100 is a dummy equal to one if the CEO has 100% ownership in the firm. CEO 0 is a dummy equal to one if the CEO has some ownership in the firm. CEO Largest is a dummy equal to one if the CEO is the largest shareholder in the firm. CEO share is the fraction of CEO ownership in the firm. Owner changes is a dummy equal to one if the majority owner in the firm changes from any previous period. All regressions include firm-CEO and year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.3: Robustness

Dep Variable	GOS	Profit	ROA	TFP	TFP2	
	(1)	(2)	(3)	(4)	(5)	(6)
CEO Owner	1007	0.006	1.26	0.014	0.015	0.018
	(3.85)***	(2.25)**	(2.73)***	(2.75)***	(2.86)***	(3.35)***
Mean Dep Var	15,540	-0.018	2.28	0.0009	-4.4e-10	
Number of Obs	565,526	565,302	564,847	556,007	556,000	513,033
Number of Groups	213,692	213,590	213,422	209,853	209,848	197,448
R-squared	0.025	0.010	0.024	0.007	0.007	0.007

NOTE: This table reports the results of OLS regressions. In column 1, the dependent variable is Gross Operating Surplus. In column 2, the dependent variable is net profit margin. In column 3, the dependent variable is Returns on Assets. In column 4, the dependent variable is TFP, obtained as the residual of a Cobb-Douglas in which value added is regressed over capital and labor for each 2 digit industry. In column 5, TFP is estimated by adding firm's market share and fixed effects for industry-years. In column 6, TFP is estimated as in column 5 but multiproduct firms are excluded. CEO Owner is a dummy equal to one if the CEO has majority ownership in the firm. All regressions include firm-CEO and year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.4: Robustness (2)

Dep Variable	Labor Productivity					
	(1)	(2)	(3)	(4)	(5)	(6)
CEO Owner	1,033	815.37	998.35	1,024	660.29	699.93
	(2.98)***	(2.09)**	(2.45)**	(2.84)***	(5.32)***	(3.90)***
LP(t-1)					0.71	0.51
					(571.81)***	(178.11)***
LP(t-2)						0.19
						(60.46)***
LP(t-3)						0.14
						(51.47)***
Sample	Firm		Firm-CEO			
	No Exit	Persistent	No Exit	Persistent		
Fixed Effects	Firm-CEO				No	
Number of Obs	418,868	264,535	255,621	310,669	308,547	121,616
Number of Groups	140,487	74,907	66,485	53,123		
R-squared	0.012	0.016	0.012	0.015	0.609	0.693

Note: This table reports the results of OLS regressions. The dependent variable is labor productivity. CEO Owner is a dummy equal to one if the CEO has majority ownership in the firm. In column 1, we restrict the sample to firms that do not die in our sample. In column 2, we restrict the sample to firms with number of observations above the median (equal to 9). In column 3, we restrict the sample to firm-CEO pairs that do not die in our sample. In column 2, we restrict the sample to firm-CEO pairs with number of observations above the median (equal to 4). In columns 5 and 6, LP(t-1)-LP(t-3) are lagged values of labor productivity with 1-3 lags. Regressions in columns 1-4 include firm-CEO and year fixed effects, regressions in column 5-6 include year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.5: Robustness (3)

Figure 4.1: Parallel Trends

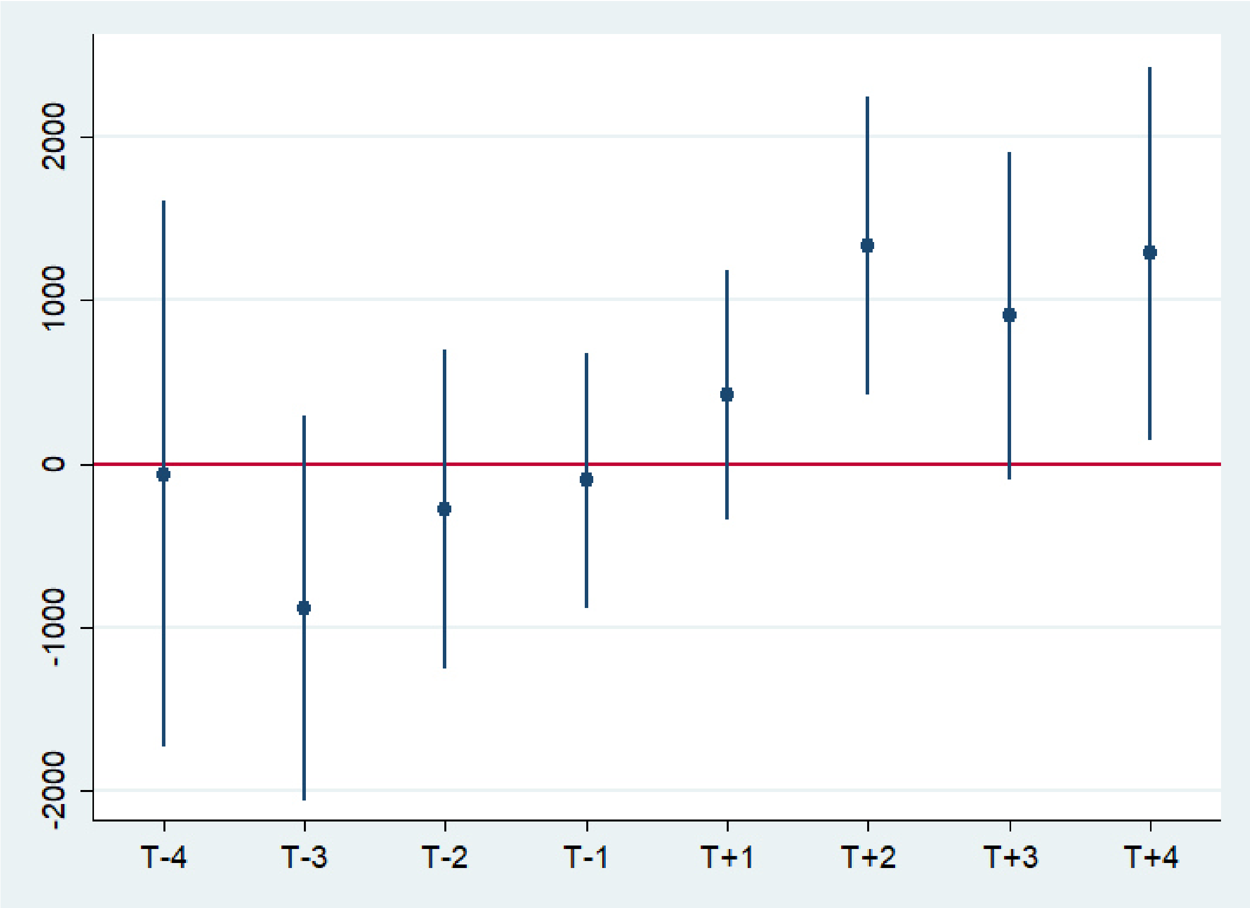

NOTE: This figure plots the estimated coefficients of equation (4.2). T-4/T-1 correspond to beta coefficients before the treatment, T+1/T+4 correspond to beta coefficients after the treatment. The bars correspond to 95% confidence intervals.

Dep Variable	Labor Productivity				
	(1)	(2)	(3)	(4)	(5)
CEO Owner	775.43	1,795	1,674	1,656	1,744
	(4.95)***	(2.88)***	(2.64)***	(2.61)***	(2.18)**
			Probit		
Z(t-1)		-0.807	-0.515	-0.13	-0.194
		(-37.05)***	(-12.20)***	(-2.52)**	(-3.10)***
Z(t-1)*(1-T(t-1))		0.767	0.66	-0.030	0.050
		(24.66)***	(9.19)***	(-0.40)	(-0.52)
Instrument		Retire	CEO Health (10k)	Spouse Health (10k)	(not working)
Fixed Effects			Firm-Owner		
Number of Obs	555,425	367,911	367,911	367,911	290,006
Number of Groups	145,579	74,642	74,642	74,642	63,004
R-squared	0.01	0.02	0.02	0.02	0.02

NOTE: This table reports results of OLS regressions (column 1) and of Probit and IV regressions (columns 2-5). The dependent variable is labor productivity. CEO Owner is a dummy equal to one if the CEO has majority ownership in the firm. The bottom panel of columns 2-5 report probit regressions of equation (4). In column 2, the instrument is a dummy equal to one if the CEO is classified as retired in the previous period. In column 3, the instrument is the amount of health benefits received by the CEO in the previous period (in 10,000 euros). In columns 4-5, the instrument is the amount of health benefits received by the CEO spouse in the previous period (in 10,000 euros). In column 5, the sample is restricted to cases where the CEO spouse is not an employee of the firm. All regressions include firm-owner and year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.6: Exogenous Variations

Dep Variable	Labor Productivity				
	(1)	(2)	(3)	(4)	(5)
CEO Change	190.32	-5,681	-1,561	3,277	-346.7
	(1.24)	(-1.31)	(-0.35)	(0.74)	(-0.07)
			Probit		
Z(t-1)		0.322	0.157	0.105	0.067
		(27.21)***	(5.83)***	(4.38)***	(2.33)**
Instrument		Retire	CEO Health	Spouse Health	
			(10k)	(10k)	(not working)
Fixed Effects			Firm-Owner		
Number of Obs	555,425	367,921	367,921	367,921	290,016
Number of Groups	145,579	74,643	74,643	74,643	63,004
R-squared	0.01	0.01	0.02	0.02	0.02

NOTE: This table reports results of OLS regressions (column 1) and of Probit and IV regressions (columns 2-5). The dependent variable is labor productivity. CEO Change is a dummy equal to one if the CEO has changed in any previous period. The bottom panel of columns 2-5 report probit regressions as in equation (4) without interactions with T(t-1). In column 2, the instrument is a dummy equal to one if the CEO is classified as retired in the previous period. In column 3, the instrument is the amount of health benefits received by the CEO in the previous period (in 10,000 euros). In columns 4-5, the instrument is the amount of health benefits received by the CEO spouse in the previous period (in 10,000 euros). In column 5, the sample is restricted to cases where the CEO spouse is not an employee of the firm. All regressions include firm-owner and year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.7: Exogenous Variations: Placebo

Dep Variable	Labor Productivity					
	(1)	(2)	(3)	(4)	(5)	(6)
CEO Owner	1,692	1,746				
	(2.48)**	(2.03)**				
Z(t)	-867.42	-479.61				
	(1.79)*	(-0.76)				
CEO 100			727.75		1,501	
			(3.65)***		(2.05)**	
CEO 0				596.98		5,141
				(4.19)***		(1.92)*
	Probit				Probit	
Z(t-1)	-0.515	-0.194			-0.166	-0.224
	(-12.20)***	(-3.10)***			(-1.87)*	(-5.03)***
Z(t-1)*(1-T(t-1))	0.66	0.05			0.045	0.15
	(9.19)***	(0.52)			(0.37)	(2.12)**
Instrument	CEO Health (10k)	Spouse Health (not working)			Spouse Health (not working)	
Fixed Effects			Firm-Owner			
Number of Obs	367,911	290,006	555,425	555,425	312,348	312,470
Number of Groups	74,642	63,004	145,579	145,579	85,437	85,462
R-squared	0.02	0.02	0.01	0.01	0.02	0.02

NOTE: This table reports results of OLS regressions (columns 3-4) and of Probit and IV regressions (columns 1-2 and 5-6). The dependent variable is labor productivity. CEO Owner is a dummy equal to one if the CEO has majority ownership in the firm. CEO 100 is a dummy equal to one if the CEO has 100% ownership in the firm. CEO 0 is a dummy equal to one if the CEO has some ownership in the firm. The bottom panel of columns 1-2 and 5-6 report probit regressions of equation (4). In column 1, the instrument is the amount of health benefits received by the CEO in the previous period (in 10,000 euros). In columns 2,5,6, the instrument is the amount of health benefits received by the CEO spouse in the previous period (in 10,000 euros) and the sample is restricted to cases where the CEO spouse is not an employee of the firm. In columns 1 and 2, Z(t) correspond to the amount of health benefits received the current period. All regressions include firm-owner and year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.8: Exogenous Variations: Robustness

Dep Variable	Labor Productivity					
	(1)	(2)	(3)	(4)	(5)	(6)
CEO Owner	1,773	1,534	1,566	6,393	5,574	5,749
	(2.41)**	(2.02)**	(1.51)	(5.60)***	(4.84)**	(4.80)***
		First Stage			Probit	
Z(t-1)	-0.146	-0.077	-0.01	-0.28	0.269	-0.123
	(-18.15)***	(-6.45)***	(-0.78)	(-19.38)***	(8.90)***	(-3.34)***
Z(t-1)*(1-T(t-1))	0 .158	0.101	0.007			
	(18.19)***	(7.09)***	(0.44)			
Instrument	Retire	CEO Health	Spouse	Retire	CEO Health	Spouse
Fixed Effects				Firm-Owner		
Number of Obs	367,921	367,921	290,016	367,911	389,668	290,006
Number of Groups	74,643	74,643	63,004	74,642	74,642	63,003
R-squared	0.02	0.02	0.02	0.01	0.02	0.02

NOTE:This table reports results of IV and Probit regressions. The dependent variable is labor productivity. CEO Owner is a dummy equal to one if the CEO has majority ownership in the firm. The bottom panel of columns 1-3 report first stage OLS regressions. The bottom panel of columns 4-6 report probit regressions as in equation (4) without interactions with T(t-1). In columns 1 and 4, the instrument is a dummy equal to one if the CEO is classified as retired in the previous period. In column 2 and 5, the instrument is the amount of health benefits received by the CEO in the previous period (in 10,000 euros). In columns 3 and 6, the instrument is the amount of health benefits received by the CEO spouse in the previous period (in 10,000 euros) and the sample is restricted to cases where the CEO spouse is not an employee of the firm. All regressions include firm-owner and year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.9: Exogenous Variations: Robustness (2)

Dep Variable	Labor Productivity				
	(1)	(2)	(3)	(4)	(5)
Treat	906.93	2,004	657.99	1,027	460,801
	(2.13)**	(3.53)***	(1.63)	(2.67)***	(2.01)**
Treat*Small	1,401		2,017		
	(2.59)***		(2.98)***		
Treat*Medium	6,761		5,829		
	(3.31)***		(2.41)**		
Treat*Large	1474.00		9,180		
	(0.19)		(0.6)		
Treat*Services		-1,249			
		(-1.98)**			
Treat*Listed				-196,271	
				(-1.69)*	
Treat*Treat					-1,787,041
					(-2.15)**
Treat	CEO Owner		CEO Shares		
Sample	All				Listed
Number of Obs	313,789	313,789	313,789	313,789	839
Number of Groups	112,875	112,875	112,875	112,875	308
R-squared	0.019	0.018	0.019	0.018	0.098

NOTE: This table reports the results of OLS regressions. The dependent variable is labor productivity. In columns 1 and 2, Treat is CEO Owner, that is a dummy equal to one if the CEO has majority ownership in the firm. In columns 3-5, Treat is CEO Shares, that is the fraction of CEO ownership in the firm. Small is a dummy equal to one if the firm has 10-50 employees, Medium is a dummy equal to one if the firm has 51-250 employees, Large is a dummy equal to one if the firm has more than 250 employees. Services is a dummy equal to one if the firm is in the service sector. Listed is a dummy equal to one if the firm is listed. In column 5, the sample is restricted to listed firms. All regressions include firm-CEO and year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.10: Effects by Size and Industry

Dep Variable	CEO Engagements			CEO Days Off		
	(1)	(2)	(3)	(4)	(5)	(6)
CEO Owner	-0.01	-0.31	-0.07	-1.42	-13.24	-4.63
	(-1.82)*	(-54.59)***	(-4.12)***	(-3.75)***	(-48.39)***	(-5.46)***
Fixed Effects	Firm-CEO	Firm-Owner		Firm-CEO	Firm-Owner	
Estimates	OLS		IV	OLS		IV
Number of Obs	555,682	555,682	389,588	561,709	561,709	387,678
Number of Groups	209,292	145,593	96,391	212,669	146,551	96,110
R-squared	0.012	0.023	0.113	0.003	0.016	0.019

NOTE: This table reports results of OLS regressions (columns 1,2,4,5) and of IV regressions (columns 3 and 6). In columns 1-3, the dependent variable is the number of employment relations of the CEO in other firms. In columns 4-6, the dependent variable is the number of days of leave of the CEO. In columns 3 and 6, the instrument is the amount of health benefits received by the CEO spouse in the previous period (in 10,000 euros) and the sample is restricted to cases where the CEO spouse is not an employee of the firm. In columns 1 and 4, regressions include firm-CEO and year fixed effects. In columns 2,3,5 and 6, regressions include firm-owner and year fixed effects. Controls include industry fixed effect (2 digits), firm's age, leverage, a dummy indicating if the firm is part of a business group, the number of workers by level of education and occupation (white vs. blue collar), the HHI index of ownership concentration. Robust t-statistics are in brackets. *, ** and *** denotes significance at 10%, 5% and 1% level, respectively.

Table 4.11: Mechanisms

Chapter 5

How Business Group Affiliation Improves Productivity of Small Firms: Evidence from Finnish Administrative Data

Henri Luomaranta

Statistics Finland and TSM, University of Toulouse Capitole, Toulouse, France

Abstract

I inspect how business group affiliation impacts firm productivity and job growth. Rich administrative data on the universe of Finnish limited liability firms reveals that joining a business group increases productivity, and decreases job growth within firms. This is driven by small firms that are mainly in the service sector. I provide suggestive evidence of mechanisms. I document changes in (key) employees, decrease in cost of borrowing and risk levels, and significant transfers of financial resources. Based on the results, the role of business groups in the economy might be most relevant in the context of small firms, which are usually not analyzed due to lack of data.

JEL Classifcation Code: D22, E24, E32, L25

Keywords: dependencies;business groups; firm size; growth;productivity

5.1 Introduction

Business group is a very widespread phenomenon in modern economies, and its role is debatable. For example, in developing countries, it is an easy argument to say that business groups can alleviate capital market inefficiencies, rendering business group affiliated firms more profitable. Yet, the evidence is not conclusive across different developing countries (see, e.g. Khanna and Yafeh, 2007). In developed countries, for example in the U.S., several studies have documented the *diversification discount* (see, e.g. Schoar, 2002), where conglomerate firms' stocks are traded at a discount compared to other comparable firms.

Whether business groups are "paragons" or "parasites", as dubbed by Khanna and Palepu (2000), remains an open question. After a careful meta-analysis of the business group literature, Carney et al. (2017) concludes that business groups should be analyzed in detailed contexts or within specific legal frameworks from many angles.

For instance, it is an often overlooked fact, that large groups own numerous small affiliates, and they are usually grouped together with other small firms in economic statistics and empirical research. A recent Eurostat report (Airaksinen et al., 2015) shows that the share of dependent[1] firms' employment within the SME[2] category is substantial in several European countries, accounting for as much as half of the SMEs' employment in some economies. This is why firm size and enterprise group dependency is an important research question. Especially so, because small firms have major potential to provide positive productivity and job creation dynamics in the economy (see, e.g., Aghion and Howitt, 2009 and Neumark et al., 2011).

This paper studies productivity and growth of Finnish firms joining or leaving a business group, and demonstrates that size is a key determinant that can explain how business group affiliation is related to those variables. I show what happens to productivity and job growth when a firm joins or leaves a business group in the within firm context, and show evidence of the potential mechanisms. The sample is obtained from a comprehensive business register, which allows to examine heterogeneity of the results by size and industry. The business register captures virtually all the ownership relationships between firms. While many benchmark studies concentrate on larger listed firms (e.g. Maksimovic and Phillips, 2013) or manufacturing industry (e.g. Atalay et al., 2014) it is important to incorporate total populations of firms in the analysis, because most of the firms in a typical economy are small (in Finland, around 99% are SMEs) and most of the new jobs are created in the service industry, which is employing more people than the goods producing sector.

Disregarding the group relationship, a small dependent firm can appear exactly the same as an independent firm. Dependent firms have their own ownership structures, they are independently legally liable, and they may have considerable autonomy in decision-making. It is somewhat difficult to draw a hypothesis regarding what kind of impact dependency from a group may have on growth and

[1] Dependent from a group, including any firm that is controlled by or controls a group. Control is defined as having 50% of votes directly or indirectly

[2] Throughout the analysis, we use the EU recommendation 2003/361 for the cut-off points for defining micro (<10), small (10-49), medium (50-249), and large (>=250) size categories by using persons employed, measured in full time equivalents.

productivity. The dependent firm might have to compete for resources (see, e.g. Giroud and Mueller, 2015), or the business group as a whole might be more resilient against competition, and it might be easier to engage in wasteful investments, leading to higher growth and lower productivity. On the other hand, dependent firms can have access to different tangible or intangible resources, such as financing options or managerial talent (see, e.g., Atalay et al., 2014, who provide evidence that transfers of intangible resources are quantitatively important in vertically integrated firms). Therefore, the effect we might expect is not clear-cut.

I employ fixed effects panel regressions, and find evidence that dependency is positively related to labor productivity (2.3% increase from the unconditional sample average), and negatively related to growth rate of employees. Both relationships seem to be driven by small firms, mostly in the service sector. Given that endogeneity is an issue, and we are not able to find suitable instruments, I try to eliminate various alternative explanations in order to motivate the interpretation of the results.

Furthermore, I explore some potential mechanisms and document several subsequent changes related to firm policies and market outcomes that might explain the results. Risk level of sales and cost of borrowing decreases, CEO and management are reshuffled, and investment in high skilled staff is increased. I also document significant transfers of financial resources, notably in the form of loans from the enterprise group. I obtain some indicative evidence, that it is the human resources policies that have the most explanatory power behind the results (changes in key staff members and wage structures).

The results underline the importance of studying business group activities in the total populations of firms, and highlight how small and large firms are interlinked by ownership channel, on top of the usual input-output relationships.

Literature Business groups and affiliate performance are studied in the management literature with varying results. Recently, Carney et al. (2017) carries out an extensive meta-analysis of business group literature, and concludes that one cannot yet draw a conclusive explanation of why business groups are still prevalent in modern economies, and the economic impact should be analyzed in more detailed contexts. Productivity of firms in integrated ownership structures are more widely studied using data on large manufacturing firms (see, e.g. Maksimovic and Phillips, 2002 and Maksimovic and Phillips, 2013). Related literature has also been interested in foreign direct investments and affiliates of foreign large multinationals, and has shown superior productivity of foreign affiliates (e.g. Criscuolo, 2005). However, most of the dependent firms in the Finnish economy are domestically owned (80% of the sample) and tend to be smaller in size.

Merger and acquisitions literature is also typically focused on large and listed firms, and empirical research on mergers among small firms is relatively scarce. Recent examples I find are Arvanitis and Stucki (2015) who study Swiss small firm mergers, and Xiao (2015) who uses register based data from Sweden to study new technology firm acquisitions by business groups.

Directly related to this study is Boccara (1997), where job creation of small and medium French firms belonging to enterprise groups is computed over the 1984–1992 period, finding that the small firms which are part of a corporation exhibit higher job creation rates. Bamiatzi et al. (2014) analyses data from the U.K., and shows that business group affiliation has a positive impact on small firms especially in declining industries. A recent example studying business group affiliate performance from emerging economy is Bhaumik et al. (2017)), that suggests for example that the insurance mechanism is an important advantage of belonging to a business group for firms willing to take risk.

There is evidence from Finland on the impact of dependent firms in the SME category, and the focus is on analyzing the aggregate effects. Fornaro and Luomaranta (2016) shows that the small independent firms are behind most of the positive net job creation[3], and Fornaro and Luomaranta (2017) shows, in the light of a productivity growth decomposition model, that the dependent SMEs have larger productivity contributions than their independent counterparts. This is mostly due to more efficient reallocation of labor towards more productive firms, and due to highly productive new dependent entrants. These latter papers point towards the fact that business groups have an influential role in the Finnish economy and in its business renewal process through creative destruction among small firms.

Methodological aspects and causal interpretation Studies interested in firm performance in a business group are plagued by endogeneity issues due to our inability to randomly impose group structures on firms. For example Larrain et al. (2018) proposes a technique to establish causal relationships in the cases where firm leaves a business group, but the replication of this methodology for Finland would require more observations, since they select a very specific sample of exactly 2 industry groups. In the absence of exogenous variations, the approach of studying the within dynamics of firms where the dependency status changes, observing virtually all such cases in the economy (improving external validity), offers a partial solution. Due to richness of the data, I can control for many of the known determinants of productivity and firm outcomes, alleviating concerns for omitted variable problem (in Section 5.3.3). Furthermore, I eliminate other possible explanations more specifically in Section 5.3.2. I show what happens to firms when main owner changes, and what happens when administrative records reveals an M&A event (where physical resources are transferred), in order to see if these alternative explanations, that plausibly can happen simultaneously with the dependency status change, are driving the results. I also separately look at what happens if firm becomes a foreign affiliate, to see if it would be sufficient to only look at FDIs in our setting. These alternative explanations do not mask the impact of the main treatment of joining a business group. Another important consideration is reverse causality. We might capture the effect of small firms being more successful and therefore being acquired by other firms. An important check against this possibility is the parallel trends assumption, which turns out not to be violated. Regardless, I avoid making causal statements based on the regressions.

Definition of a business group In the European context, Eurostat's official definition for an enterprise group reads as follows.

An enterprise group is an association of enterprises bound together by legal and/or financial links and

[3]the paper is included in the next chapter of the thesis

controlled by the group head. A group of enterprises can have more than one decision-making centre, especially for policy decisions on production, sales and profits. It may centralize certain aspects of financial management and taxation. It constitutes an economic entity which is empowered to make choices, particularly concerning the units which it comprises. Eurostat, 2019.

Similar structures in the U.S. are known as conglomerates, but the slight conceptual difference is that conglomerates should operate in several industries. The main focus in this paper is the Finnish *konserni*, which is defined as its European enterprise group counterparts. To be more precise, a business group in this paper is defined as a group of firms, consisting of a mother and affiliates, where a mother has a controlling stake (over 50% of votes) in each of the affiliated firms. These kind of enterprise group structures are common, and large firms in Finland (and Europe) are typically organized as such. The focus of this paper is on changes in the business group membership status, where at one point we observe ownership links, and do not observe such links at another point in time during the firm panel. The business group membership does not imply transfers of physical assets to another firm, although this can happen, but rather transfers of voting rights. It also implies that the mother will supply a consolidated financial statements on behalf of the entire group (on top of its own). The enterprise group itself is not a legal entity or liable to pay taxes.

The remainder of this paper is organized as follows. Section 5.2 describes the data and dependent variables, Section 5.3 provides the main results, showing what happens when the dependency status changes within firms and discusses possible alternative explanations, Section 5.4 provides indicative evidence on the various mechanisms that can explain the results, Section 5.5 summarizes Fornaro and Luomaranta (2017) to highlight what the aggregate effects are in the light of an productivity growth decomposition model, and Section 5.6 concludes.

5.2 Data and dependent variables

The main data source is the Finnish structural and financial statistics database, and the sample includes the universe of limited liability firms active in Finland, spanning the years 2006 to 2014. The data includes balance sheet information that firm discloses annually. The database covers all active enterprises in the non-financial business economy (NACE Rev.2 sections B to N, excl. K). I use the deflated value added at factor cost (VA), computed by the statistical office, as the measurement for firm output. VA is calculated by deducting the costs of operating activities from the income. Costs exclude the costs related to personnel. The statistical office tries to clean the measure from the effects of transfer pricing, which is an important quality consideration in our case. Employees are converted to full time equivalents so that, an employee working half-time represents one half of a person and similarly two employees working half-time during one year represent one annual FTE. I exclude firms with less than 1 FTE (one man companies).

Labor productivity (LP) for firm i at t is

$$LP_{i,t} = VA_{i,t}/FTE_{i,t}, \tag{5.1}$$

where VA is in real terms, deflated using the implicit price deflator at 2-digit industry level, and FTE stands for full-time equivalent units of labor.

Job creation measure is obtained from a specific monthly level data source (the statistical office's short term business statistics database) where employment figures, measured in FTEs, are adjusted to represent the *organic* growth, disregarding the effects of mergers, split-offs and other legal restructuring. I aggregate these monthly observations to yearly level. The data is manually inspected for the cases of restructuring for the most important firms by the statisticians. In addition, an automated correction procedure for mergers and split-offs is adopted (Appendix B contains the details of this procedure). The file on job creation is merged with the one obtained from the structural and financial statistics database, with slightly different coverage. From the original 566,037 firm year observations I give up 30,824 observations when I use the job growth variable.

The annual growth rate is computed as:

$$Growth_{i,t} = (FTEi, t - FTE_{i,t-1})/FTE_{i,t-1} \tag{5.2}$$

Finally, I link data from the Finnish Longitudinal Employee-Employer Database (FLEED) and the Finnish Longitudinal Owner-Employee Database (FLOWN). The FLEED data includes the professional classification and salaries of firm employees, and FLOWN includes the identities of the main owners[4]. I use these two datasets to identify CEOs, executives, and share of wages paid to high-skill staff[5] for each firm[6]. More detailed data description is in the Appendix A.

Financial variables are winsorized at 0.25% and 99.75% of their empirical distributions. In general, one would like to keep as much information as possible, without biasing the results. The job growth variable is somewhat sensitive to this issue and winsorizing 0.25th and 99.75th percentile would still leave large outliers, which seem to be related to mistakes in recording restructuring events. I therefore winsorize job growth at 1% and 99% of its distribution. I define a dummy, *Dependent*, which takes value 1 if the firm belongs to a business group, and 0 otherwise. This variable serves as the main treatment of interest throughout the analysis.

5.3 Main results: how dependency status change relates to productivity and growth

The first set of regressions analyzes what relationship the dependency status change has with productivity and employment growth rates, keeping other firm characteristics fixed. The main specification

[4]Ownership is reported if a shareholder holds >=10% of shares, we observe at least the largest shareholder for 92% of the firms.

[5]I define high-skill staff as the job category of "professionals" in the ISCO-08 (major group 2)

[6]Very small firms do not always have a professional category of a CEO, because the classification is based on the main activity of the employee (in small firms, the entrepreneur does not spend most of her time in managerial duty). In cases where I fail to identify a CEO by the job category, I use the information on ownership and salaries (I check if one of the entrepreneurs is working at the firm, and rank employees based on salary). Similar procedures have been applied to Portuguese data in Queiró (2016). The details are left in the Appendix.

is

$$y_{i,t} = \alpha_i + \beta T_{i,t} + X_{i,t}^{'}\gamma + \mu_t + \varepsilon_{i,t}, \tag{5.3}$$

where i denotes the firm, $y_{i,t}$ is the productivity measure or job creation rate of firm i in year t, $T_{i,t} = 1$ if firm is dependent and $X_{i,t}^{'}$ controls for characteristics of the firm, α_i and μ_t are firm and year fixed-effects. Throughout, I control for 2-digit NACE industries, size category (indicator variables for Micro (left out), Small, Medium and Large), leverage, age and age^2.

The following table provides the main results. It also provides a simple robustness check against biases arising from the fact that we are not able to observe prices (see Beveren (2012) and De Loecker and Goldberg (2014) for reviews).

Dep Variable	LP	LP	Job growth	Job growth	LP	log(LP)	LP	Job growth
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent	25,660	17,725	-0.28	-0.11	1,254	0.01	1,245	-0.049
	(92.2)***	(59)***	(-63.6)***	(-19.35)***	(2.2)**	(1.76)*	(1.97)**	(-2.70)***
Mkt Share							4,920	
							(3.00)***	
Individual Fixed Effects	NO	NO	NO	NO	Firm	Firm	Firm	Firm
Controls	NO	YES	NO	YES	YES	YES	YES + time*industry	YES
Number of Obs	566,266	566,260	535,442	535,436	566,037	550,116	522,693	535,436
Number of Groups					110,453	107,831	106,827	103,470
R-squared	0.04	0.16	0.004	0.06	0.018	0.018	0.04	0.05

Note: Robust standard errors in parentheses.
*p<0.1; **p<0.05; ***p<0.01
In column 7, controls include interaction terms year*NACE, otherwise controls are: year, 2-digit NACE industries, size category (indicator variables for Micro (left out), Small, Medium and Large), leverage, age and age^2.

Table 5.1: Main result: how dependency status is related to productivity measures (LP, log of LP) and job growth. In regressions 1-4, I show the cross sectional relationship. In 5 I fix the firm to give a baseline result in the within firm context. To make sure that our results would be robust for alternative trimming practices, in 6 I follow e.g. Hyytinen and Maliranta (2013) who uses a strategy proposed by Mairesse and Kremp (1993) to handle outliers, removing 4.4 standard deviations above or below the input weighted industry distributions of log of labor productivity. In column 7, I use a restricted sample by removing firms with multiple plants, and control for time specific industry trends and market share (firm turnover/industry turnover), in order to eliminate possible biases arising from our inability to observe prices, and shocks arising from firm's other industries. In column 8 the Job growth is the dependent variable fixing firm.

From Table 5.1 above, we obtain significant positive relationship with labor productivity and negative relationship with job growth. Notice that the sample on job growth is slightly different, but the results on productivity are similar if I use exactly the same sample (and are available upon request). Log transformation of LP is somewhat problematic, because firms may have negative value added measures. This is caused by input costs being higher than output can generate income, and can occur naturally in rapid expansions or start-up phases, or indeed during adverse business conditions. The log-level regression in column 6 suggests 1% increase in labor productivity if the firm joins a business group. The coefficient in the 5th column translates to 2.3% increase from the sample average, and €1,254 increase in firm productivity. Column 7 has coefficient of similar size, and it includes controls for market share and industry interacted with time to control for possible industry specific shocks and price fluctuations due to demand, or due to price setting ability of the firm. Based on column 8, one can expect 4.9% decrease in job growth rate once the firm joins a business group. In economic terms, the size of the coefficients are substantial, as the average productivity growth of Finnish firms in our sample period is 0.7% per year.[7]

The role of size and industry A key element in this paper is the use of total population of firms, making it possible to inspect heterogeneities with respect to size and industries. The table below provides evidence, that it is the smallest firms in the service sector that are driving the results. I interact size categories with treatment, and industries with treatment, in order to see if there are size and industry differences in the relationships we have uncovered. I have explored heterogeneity by age without seeing a different impact on young firms.

Dep Variable	LP	Job growth	LP	Job growth
	(1)	(2)	(3)	(4)
Dependent	5,457	-0.15	-370.66	-0.009
	(5.89)***	(-6.03)***	(-0.66)	(-0.18)
Small*Dependent	-6,621	0.15		
	(-6.15)***	(5.70)***		
Medium*Dependent	-9,971	0.32		
	(-6.06)***	(7.7)***		
Large*Dependent	-19,887	0.34		
	(-4.45)***	(2.42)**		
Services*Dependent			2249	-0.10
			(1.78)*	(-2.93)***
Individual Fixed Effects	Firm	Firm	Firm	Firm
Controls	YES	YES	YES	YES
Number of Obs	566,260	535,436	566,260	535,436
Number of Groups	110,453	103,470	110,453	103,470
R-squared	0.018	0.05	0.013	0.05

Note: Robust standard errors in parentheses.
$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Controls: year, 2-digit NACE industries, size category (indicator variables for Micro (left out), Small, Medium and Large), leverage, age and age^2

Table 5.2: Interaction with size and services industry with the dependency status change. I provide interactions with Micro, Small, Medium and Large categories in columns 1 and 2 and the service sector in columns 3 and 4, for LP and job growth. Micro category is left out in the size interactions, and other than service industries are left out in the service interaction. The services sector is defined by 2-digit NACE codes 49-96

In Table 5.2 above, there appears to be significant size dependent differences in the statistical relationship with LP and growth. The coefficient is largest for micro firm category in both LP (column 1) and job growth (column 2). Micro firms indeed seem to be the driver of the result. Looking at the services interaction, we obtain evidence that the relationship is driven by the service industries. The coefficient is twice the size of the main treatment (in Table 5.1). If we control for service interaction,

[7]The corresponding figure for EU countries is 0.9% and for the US is 1.13%, see the OECD's website at data.oecd.org.

the coefficient is no longer significant for the rest of the industries. Similarly for job creation (column 4), the coefficient about doubles from the one appearing in Table 5.1, and the coefficient is no longer significant for the rest of the firms.

In other words, the positive (negative) dynamics in terms of productivity (job growth) are driven by the small firms in the service sector, which forms a major part of the Finnish economy in terms of employment.

Based on this evidence, the implications of joining a business group are not homogeneous across sizes and industries, and the role of business groups in the economy might be most relevant in the samples which are usually not analyzed due to lack of data.

Next, we will explore the model assumptions and address some concerns about endogeneity.

5.3.1 Parallel trends

Within firm regressions suggest that there is a positive relationship with labor productivity and a negative relationship with growth rate. The standard assumption in DiD specifications is the condition on pre-treatment trends, requiring that there is no statistically different trend prior to the treatment. I obtain evidence of this point by inspecting the time periods before and after treatment using the following specification:

$$y_{i,t} = \alpha_i + \sum_{s=1}^{4} \beta_{t-s}\mu_{t-s}T_{i,t} + \sum_{s=1}^{4} \beta_{+s}\mu_{t+s}T_{i,t} + X^{'}_{i,t}\gamma + \mu_t + \varepsilon_{i,t}, \tag{5.4}$$

where $\mu_{t-s}T$ and $\mu_{t-s}T$ are interactions with the treatment. They take values 1 during the years $t-s$ before and $t+s$ after the treatment. α_i, μ_t are firm and year fixed-effects, and $X_{i,t}$ contains the firm specific time varying characteristics (as above). I plot the β coefficients obtained from these regressions in the following figures.

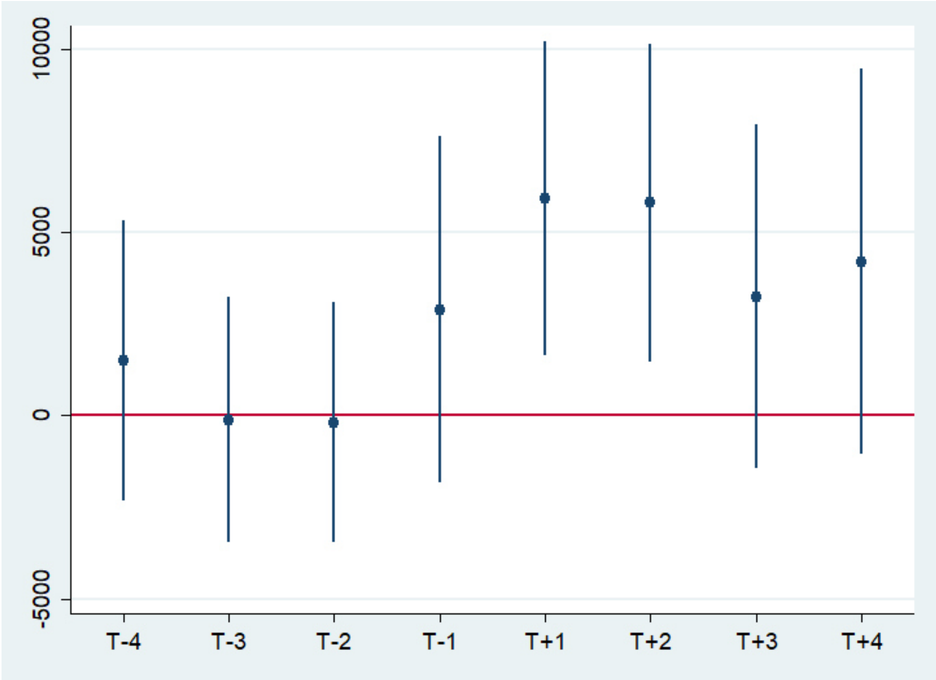

Figure 5.1: LP, firm joins a business group, 95% CI, T+1 indicates the first period of treatment

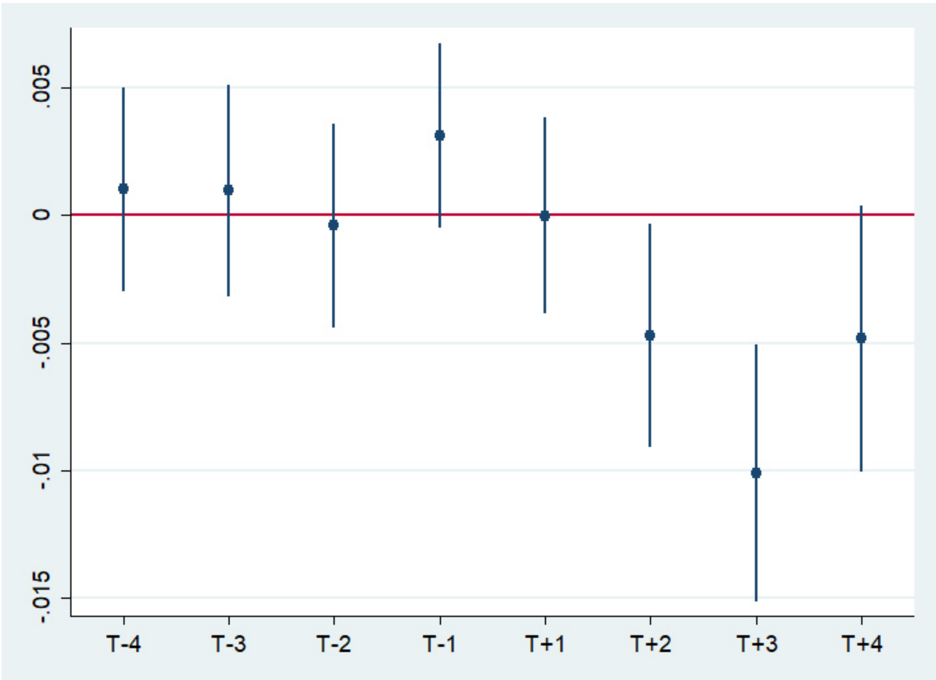

Figure 5.2: Growth rate, firm joins a business group, 95% CI, T+1 indicates the first period of treatment

The figures above do not display pre-treatment trends different from 0. For labor productivity (Figure 5.1), the treatment seems to have a significant impact already at t+1. Interestingly, the negative job growth is statistically different from 0 starting from t+2 onwards (Figure 5.2). This observation can have a natural explanation in fixed nature of employment contracts. Based on these graphs, we cannot conclude that the result is not causal due to violation of pre-treatment trends.

5.3.2 Further evidence and other explanations

The results based on the within firm regressions can be questioned at least on the grounds that a simultaneous event affects both the decision to join a business group and future productivity. Indeed there are several possibly related events that may occur at the same time as the dependency status changes, which might explain the results. For instance, the main person owner may change, mergers of physical resources may occur, or foreign direct investments can be the reason for dependency status change. I try to control for these alternative explanations below. Moreover, it is useful to provide more details on the change in dependency status. A firm may become dependent by joining a group as an affiliate, or it may become a mother by establishing a group either by obtaining ownership of other firms, or by reorganizing its existing activities. The data allows to distinguish these different roles. I use the specification (5.1), and look at different controlling variables (defined as dummies, 1 for the duration of the event) besides the dependency status change. First I fix the firm and the main owner, to see if the result holds for any given firm-owner pair, making the regression more neutral to the possible matching effect. I also add indicator variables for the following events: firm becomes foreign affiliate, firm changes the main owner (largest shareholder), firm is involved in a merger, firm becomes group head, or firm becomes an affiliate. Merger data is obtained from the tax office, and it is defined as an event where some resources are transferred between enterprises (including split-off and other legal restructuring). Firm can join an enterprise group as a result of a merger, but it is conceptually different to transfer voting rights than to transfer productive assets between firms. Even though the statistical office has adjusted the growth rates to represent organic growth rate, there can still be positive impacts within the firm trend on growth. For example, the firm may be on the decline before the merger, and during the subsequent period after the merger firms involved will be sharing the same growth pattern, leading to higher growth rates (the adopted methodology computes growth rates as if the merged firms were operating as one already the year before).

Dep Variable	LP	Jobs	LP	Jobs	LP	Jobs	LP	Jobs	LP	Jobs	LP	Jobs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Dependent	1,476 (2.24)**	-0.07 (-3.59)***	1,170 (2.02)**	-0.045 (-2.45)**	1,240 (2.75)**	-0.05 (-2.78)***	1,254 (2.2)**	-0.05 (-2.7)***				
Foreign affiliate			911 (0.15)	-0.04 (-1.01)**								
M&A					1,566 (1.29)	0.08 (2.99)***						
Owner change							-100 (-0.72)	0.02 (2.26)**				
Mother									1,533 (1.56)	-0.007 (-0.31)		
Affiliate											553 (0.89)	-0.05 (-2.74)***
Individual FE	Firm-Owner	Firm-Owner	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
Controls	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Number of Obs	566,260	535,436	566,260	535,436	566,260	535,436	566,260	535,436	566,260	535,436	566,260	535,436
Number of Groups	147,394	136,024	110,453	103,470	110,453	103,470	110,453	103,470	110,453	103,470	110,453	103,470
R-squared	0.018	0.05	0.018	0.05	0.018	0.05	0.018	0.05	0.018	0.05	0.017	0.05

Note: Robust standard errors in parentheses.
*p<0.1; **p<0.05; ***p<0.01
Controls: year, 2-digit NACE industries, size category (indicator variables for Micro, Small, Medium and Large), leverage, age and age^2

Table 5.3: Further evidence: the dependency relationship with LP and job growth by different related variables and alternative specifications. The first two regressions are fixing firm and owner - as opposed to fixing only the firm (columns 1-2). The other regressions analyze what happens when firm becomes a foreign affiliate (columns 3-4), is involved in M&A (columns 5-6), owner changes (columns 7-8), firm becomes a mother (columns 9-10), or firm becomes an affiliate (columns 11-12).

In Table 5.3 above, once I fix the firm and main owner, the relationship with dependency is positive with labor productivity and negative with job growth (columns 1 and 2), both quantitatively stronger than in the case where the main owner is not fixed. Fixing firm, foreign ownership has a negative coefficient in job growth regression, but does not have a statistically significant coefficient in LP regressions (columns 3 and 4). Merger has a positive relationship with job growth, and we do not observe a significant relationship with LP (columns 5 and 6). In columns 9-12 the negative correlation on growth rate is significant only if the firm becomes an affiliate, and we do not obtain statistically significant coefficients if mothers and affiliates are separated. This result indicates that the job creation dynamics are mostly driven by the instances where the firm becomes a part of a group as an affiliate.

5.3.3 More controls

I go a step further to clarify the impact on productivity and try to make sure the results are not plagued by omitted variables bias. I add multiple productivity related controls to the baseline set, and then try to abstract from the effects of higher quality managers (CEO) and the effect of ownership structure, by fixing firm-CEO, firm-main owner-CEO. By fixing CEO and owner we can also abstract from the endogenous matching effects.

The literature has discussed several aspects that can be related to productivity (see, e.g. Syverson, 2011 for a survey). I control for size of the business group by summing up the observed domestic employment[8], I also include CEO salaries, CEO tenure, CEO age, HHI ownership concentration index, average salaries, market share, and the interactions of year and 2-digit NACE codes. This set of controls should account for improved management quality, labor input quality, ownership structures and for the bias arising from unobserved prices. As we have seen, the improved productivity levels are mainly associated with small services firms (which are more labor intensive), and therefore wage levels and employee quality should be an important part of the explanation. Although salaries in particular suffer from the issue of reverse causality (see, e.g.Van Reenen (1996)), it will serve as a tight controlling variable in our case.

Dep Var	LP	LP	LP	LP	LP	LP	LP	LP
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent	823.3	597	1,206	1,284	1,339	1,174	1,553	1,666
	(1.45)	(0.96)	(2.15)**	(2.03)**	(2.03)**	(1.59)	(2.01)**	(1.93)*
Fixed	Firm	Firm	Firm	Firm	Firm-CEO	Firm-CEO	Firm-CEO-Owner	Firm-CEO-Owner
Controls	YES	YES	no wages	no wages	YES	YES	YES	YES
Obs	555,406	512,942	555,418	512,954	555,406	512,942	555,406	512,942
Groups	109,499	105,844	109,500	105,845	209,064	196,807	237,159	222,479
R-squared	0.06	0.08	0.04	0.04	0.06	0.08	0.065	0.08

Note: Robust standard errors in parentheses.
*p<0.1; **p<0.05; ***p<0.01
Controls: Year, leverage, firm age, firm age squared, size categories, NACE2, CEO salaries, CEO tenure, CEO age, HHI ownership concentration index, average salaries, market share, interactions of year*NACE2

Table 5.4: Further controls and fixed effects: Dependent variable is always LP. I fix first only the firm (1-4), then firm-CEO (5-6), and then firm-CEO-main owner (7-8). I add a large set of controls that can potentially explain productivity. In columns 3 and 4 I drop wages as controls. In columns 2, 4, 6 and 8 I exclude multiplant firms in order to conduct a very strict tests, controlling for the possibility of shocks spilling over from firm's other industries.

In the above Table 5.4, regressions in columns 2, 4, 6 and 8 exclude multiplant firms in order to make sure that shocks spilling over from other industries of the firm do not affect our results. Columns 1 and 2 describes the regression fixing the firm only, and using the full set of controls as explained above. We see that the treatment is not significant, and this can be explained by the inclusion of wage levels of employees, as seen from columns 3 and 4 which yield positive significant coefficients in a regression where firm is fixed and average wages are removed from the set of controls. The confounding effect of wage structures is not surprising given that large part of the positive relation of LP and dependencies comes from small (services) firms, which should be more labor intensive. It is fairly obvious that the reverse causality is a problem, and wage bill can simply increase due to higher margins that are distributed to employees. Nevertheless, wages are a good control for the quality of the workforce.

Business group may impose human resource policies that lead to changes in employees (including the CEO) and invest in new higher skilled employees. We can try to abstract from some of the changes in personnel that happen in the firm by fixing its key members, such as the CEO and main owner. Fixing firm-CEO, and including the full set of controls in columns 5 and 6, the coefficient is significant in 5,

[8]we don't observe the foreign parts

and its size is similar than in the main results (Table 5.1). Fixing firm-CEO-owner yields significant coefficients, and their size are even slightly larger than in Table 5.1. The test we perform in 8 is very demanding also in terms of controlling for factors outside the firm. If I estimate regressions where the key members of the firm are not allowed to change, I am unable to find productivity related observables that have the capacity to confound the impact of dependency on productivity increase[9].

The evidence seems to imply that in the Finnish data, where the business group membership have a positive association with small services firms, an important mechanism works through changes in human resource policies, but many simultaneous dynamics do occur, as we are going to see next.

Overall, using a large set of productivity related controls on employees, owners, firm, group, demand shocks, and ability to set prices, business group membership has a positive relationship with productivity for each firm-CEO pairs and each firm-CEO-owner in the sample.

5.4 Indicative evidence on mechanism

The different mechanisms for improving productivity at dependent firms are not well documented in the empirical literature. However, several hypotheses can be drawn, based on prior literature on productivity and size.

Various market selection models explain differential growth rates of small and large firms by the uncertainty they face (the seminal paper is Jovanovic (1982)), where the basic intuition is that new firms entering the market are not aware of their true "type", i.e. their competitiveness in relation to the incumbent firms in the industry. The intuition is as follows: After entering, firms start receiving information of their true ability from the realized earnings, and in the early post-entry years the updating effect is stronger and leads to higher growth rates among survivors. This can be interpret also in terms of *experimentation* processes (see e.g. Brynjolfsson et al., 2007, Hyytinen and Maliranta, 2013, and Gabler and Poschke, 2013), where firms deliberately expose themselves to productivity risk in order to achieve higher levels of productivity, and this effect is stronger among young firms that fight for survival. This is why I'm interested in a measure of risk after the firm becomes dependent.

Another important explanation of firm outcomes are financial constraints. In particular, because it is costly to fire and hire employees, firms that are financially constrained cannot fire enough people during bad times (see,e.g. Moscarini and Postel-Vinay, 2012). In terms of productivity, there is an intuitive link with financial constraints. Firms that are constrained cannot organize production optimally, because it is costly to invest in e.g. IT systems, machinery, or hire the most skilled individuals (for empirical evidence of the link between financial constraints and productivity, see e.g. Ferrel et al. (2016)).

It has been suggested that transfers of intangible resources, such as management practices and know-how are an important part of the dynamics. Atalay et al. (2014) show with U.S. manufacturing data, that firms do not transfer tangible resources particularly intensively, but rather intangible resources

[9] I have added cost of debt, sales volumes, share of high skilled staff, job turnover. I choose not to report those because adding cost of debt or job turnover removes many observations, sales volume is already incorporated in value added, and wage structure should incorporate the information on high skill share.

are transferred in a vertically integrated structure. Bloom et al. (2016) finds that European affiliates owned by American firms were able to benefit from more advanced IT systems and management practices, thus leading to improved productivity levels. Exploring the employment data, I assess whether dependency status change is associated with changes in CEO and other executives, which would be a straightforward way to implement new policies at the firm.

In the spirit of the property rights approach to firm boundaries (see, e.g. Hart and Holmström, 2010), another way that a firm may become more productive is that fewer executives are needed to run it, and that higher quality managers are now responsible for more productive assets. This is why I look at number of executives and their salaries. I posit that CEO salary should be related to CEO quality in a competitive labor market. The impact of CEO quality on firm outcomes is also a topic of large literature (e.g. Bertrand and Schoar (2003)).

Furthermore, if the explanation on financial constraints affecting the ability to fire and hire is correct, we should observe an increase in job turnover after the treatment. Job turnover may also play an important part in productivity, if the firm is able to hire new skills to replace old redundant ones. Related to this, I also look at investments in skills, such as the proportion of wages paid out to high-skilled staff, and changes in wage structures more generally. In a competitive labor market, higher salaries would imply more productive and more skilled workforce (Abowd et al., 2005 and Fox and Smeets, 2011).

We inspect these variables within firms by looking at post-treatment years. The within regression is specified as

$$y_{i,t+1} = \alpha_i + \beta T_{i,t} + X_{i,t}^{'}\gamma + \mu_t + \varepsilon_{i,t}, \tag{5.5}$$

where $y_{i,t+1}$ is the lead of the dependent variable, in order to make sure that the potential policy change occurs after the treatment. I regress 1) standard deviation of sales, measured from monthly turnovers and averaged over the year, 2) cost of debt, measured as interest payments over external capital, 3) Dummy indicating a change of CEO, 4) CEO salary, 5) number of executives working at the firm, 6) average salary of the executives, 7) share of wages paid out to skilled staff, 8) job turnover, defined as (number of job terminations / average workforce)*100, 9) job turnover is also inspected with an interaction term coinciding with a particularly interesting year, 2009, when the Finnish economy experienced a dramatic fall, making financial constraints more binding, 10) average salary in the firm.

Dep Var	$\sigma sales$	Cost debt	Change CEO	CEO salary	Executives	Salary exe	Skilled	Job Turnover	Job Turnover	Salary
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Dependent	-0.60	-0.0007	0.04	4,968	-0.25	15,952	0.01	2.22	2.17	1,256
	(-3.37)***	(-5.03)***	(5.63)***	(6.47)***	(-5.44)***	(1.92)*	(3.92)***	(6.28)***	(6.12)***	(4.95)***
Dependent *Crisis (2009)									0.66	
									(2.53)***	
Fixed	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
Controls	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Obs	424,006	443,715	443,715	443,715	443,715	177,563	443,715	384,985	384,985	443,712
Groups	86,675	90,761	90,761	90,761	90,761	44,757	90,761	85,119	85,119	90,761
R-squared	0.01	0.94	0.001	0.01	0.004	0.004	0.01	0.04	0.04	0.011

Note: Robust standard errors in parentheses.
*p<0.1; **p<0.05; ***p<0.01
Controls: year, 2-digit NACE industries, size category (indicator variables for Micro, Small, Medium and Large), leverage, age and age^2

Table 5.5: Indicative mechanisms: Dependent variables are leads of standard deviation of sales (col. 1), cost of debt (col. 2), dummy indicating a change of CEO (col. 3), CEO salary (col. 4), number of executives (col. 5), average salary of executive (col. 6), share of salary by skilled workers (col. 7), job turnover (col. 8-9), and average salary (col. 10). Treatment is 1 if firm is part of a business group, and I add interaction with crisis year (2009) in a regression 9 with job turnover as the dependent variable.

We observe interesting dynamics that take place after the treatment in Table 5.5. In column 1 we see that the impact on standard deviation of sales is negative, indicating that the risk levels of the firm goes down. This finding is consistent with the market selection or experimentation process explanations where the risk levels, or uncertainty is a key element in explaining growth. The results in column 2 indicates that the cost of debt goes down, pointing towards the fact that financial constraints should be less binding for the dependent firm. Column 3 shows, that the probability of CEO change increases and column 4 shows, that the salary of CEO goes up. These two observations together form an important channel by which the mother firm can influence the decision-making in the affiliated firm. The new CEO is appointed to implement new policies, and the new CEO might be of a better quality than the previous one, which is evidenced by the salary increase. Columns 5 and 6 show that the number of executive positions goes down, while the average executive salary increases. The firm may benefit from the better management practices and managers of the mother firm, thus allowing gains in productivity. Column 7 documents increased investments in skilled staff and this might help firm to become more productive. Columns 8 and 9 show that the job turnover rates increase, and that job turnover increases relatively more if dependency status was changed during the crisis year of 2009. This finding is consistent with the idea that financial constraints prevent especially small firms from adjusting labor input optimally during economic downturns. While the management literature usually has found that job turnover has a negative impact on productivity (see, e.g. Hancock et al., 2013), it might be an important element of the creative destruction process (e.g. Aghion and Howitt, 2009) and especially so if the firm is able to replace low productivity positions with higher productivity ones. This is evidenced in column 10, which shows that the average salary is positively affected by the dependency status change. Overall, we document changes in risk levels, cost of debt, changes in key management positions and changes in human resources policies of the firm that are related to productivity and job creation dynamics. The results are similar, if we consider these dependent variables contemporaneously.

The most obvious advantage of the group structure is the possibility to shift financial resources among members. In those instances, the firm should report these as a separate items in the balance sheet. The group can extend either loans, or provide direct transfers (called "concern aid"). Direct transfers can be useful in order to minimize taxes (receiver records these transfers as part of income, and donor gets to deduct these from the final result, thus reducing the overall taxable amount if the receiver is running losses). We obtain these two items for our treated firms, and compute concern aid/turnover, and enterprise group debt/ external capital, after the treatment has occurred. The following table provides a t-test to show that these amounts are not negligible for the treated firms after dependency status change.

	Concern aid/turnover	enterprise group debt/ external capital
Mean	0.0019	0.129
t-statistic	(2.6) $***$	(71.47) $***$

Note: H0=0.
*p<0.1; **p<0.05; ***p<0.01

Table 5.6: Financial transfers: t-test with null hypothesis: concern aid=0 or enterprise group debt over total debt = 0

In our sample of treated firms, direct transfers amount to 0.19% of turnover, and enterprise group debt is around 13% of the total debt, as seen from Table 5.6. It seems that the lending mechanism is quantitatively quite important. This points towards gains in terms of diversification, which reduces the overall risk levels, and it may allow the mother firm to negotiate better interests on behalf of the group members.

While interesting in terms of showing the micro level dynamics within firms, the results in this subsection are not enough to document the aggregate effects of dependencies. The aggregate effects are important in order to see if these dynamics have persistent and sizable economic effect.

5.5 Aggregate productivity contributions of dependent small firms

This section is based on Fornaro and Luomaranta (2017)

Fornaro and Luomaranta (2017) uses a productivity growth decomposition model to reveal evidence of dependent and independent small companies' contribution to the aggregate productivity growth dynamics, explaining how dependent small firms can positively contribute to the economic outcomes. The idea of the model is to decompose productivity growth into its micro-level components by firm categories. The components of interest are: 1) within (productivity growth of an average employee within the firm), 2) between (productivity growth due to reshuffling of labor inputs between firms in the category), 3) exit (contribution to productivity by the exiting firms, which is positive if exiting firms have lower than average productivity) and 4) entering (which is negative if the entering firms have lower productivity than the average incumbent). The cross terms are a residual terms meant as a corrective measure to force the different components to add up to the industry totals (see Maliranta, 2003 for a thorough discussion of the model). The following table describes the contribution of dependent small firms to the aggregate outcomes in the light of the adopted decomposition model. Notice that the right hand side of the table includes normalized components (taking into account the labor input share), and the left hand side includes the absolute components.

	Absolute Components						*Normalized Components*					
	Prod.	Within	Between	Entry	Exit	Cross terms	Prod.	Within	Between	Entry	Exit	Cross Terms
Micro												
Dependent	0.13	0.047	0.003	0.06	−0.034	0.04	10.9	4.18	0.30	5.01	−2.73	4.12
Independent	0.34	0.48	0.056	−1.00	0.94	−0.14	2.59	3.61	0.42	−6.52	6.11	−1.02
Small												
Dependent	0.62	0.21	0.08	0.15	−0.01	0.17	7.62	2.71	1.06	1.85	−0.22	2.22
Independent	0.63	1.04	0.01	−1.54	1.46	−0.35	2.28	3.73	0.07	−5.09	4.84	−1.26
SME												
Dependent	1.39	0.64	0.14	0.15	0.08	0.35	5.81	2.72	0.61	0.63	0.37	1.45
Independent	0.56	1.14	−0.03	−1.75	1.61	−0.42	1.79	3.56	−0.10	−5.05	4.68	−1.30

Table 5.7: Adopted from Fornaro and Luomaranta (2017): Contributions of dependent and independent SMEs to the productivity of the business economy (2002-2014). Results are in real terms and are reported in percentage points.

The findings in Table 5.7 indicate that in absolute terms, dependent micro and small firms have lower productivity contributions than the independent firms in those categories, but dependent firms have larger contribution if also the medium category is included. Dependent firms in all the size categories have larger productivity contributions if we look at normalized components, where the labor input shares are taken into account. Among dependent SMEs, it is the micro firms that create the largest normalized productivity contribution. This pattern emerges because there are much more independent small and micro businesses. However, in the SME category as a whole, it is indeed the firms in business groups that provide most of the productivity contribution. The dependent SMEs are responsible for almost 3 times more of the aggregate productivity growth. Another interesting pattern emerges: the between component of dependent firms is much higher than the one of independent firms. This means that the more productive firms receive more labor inputs, and less productive firms receive less labor inputs. This can be interpret as the creative destruction component. Also, the entering and exiting dependent firms have an opposite impacts as one would normally expect. The dependent entries have immediate positive effect on productivity (entering firms are more productive than incumbents). The opposite is usually true for independent businesses. Overall, there is a strong presence of dependent SMEs, and their contribution to aggregate SME productivity is substantial. This is why more detailed (industries, and individual groups) studies are still needed for assessing how business groups contribute to the business renewal and competitiveness.

5.6 Conclusions

In this work, I contribute to literature on the role of enterprise groups in the economy by analyzing job creation and productivity dynamics of firms joining/leaving business groups. I find that business group dependency has heterogenous implications, depending on size and industry.

Small firms that become dependent of a business group, experience a notable increase in productivity and a decrease in job creation. I provide suggestive evidence of mechanisms that can explain these results. First of all, the risk levels go down. In the light of experimentation process or market selection mechanisms, this might have a negative relationship with growth. Second, the cost of borrowing goes down, giving support to the explanation that credit constraints are alleviated. This has been proposed as one mechanism by which small firms have different growth patterns (e.g. Moscarini and Postel-Vinay, 2012). Finally, and importantly, we observe a number of changes in firm employees. Notably, the CEO is likely to change, number of managers is reduced, job turnover increased, and there is an increased investments in high-skilled staff. If firm is fixed, adding average wages as control renders the statistical relationship with productivity insignificant. Abstracting from changes in CEO and owner by fixing Firm, CEO and owner, and adding average wages as control, yields significant and large coefficients explaining the relationship between labor productivity and dependency. These results would indicate that the human resource policies (by changing the key members of the firm) are an important reason behind the obtained results.

If we keep the key members of the firm fixed, dependency status change has significant positive (unexplained) impact on productivity. This result is obtained by using a large set of productivity related controls on industries, employee quality, ownership structure, firm characteristics, group size, possible demand shocks, and ability to set prices.

The findings can be taken together with the observations in Fornaro and Luomaranta (2017) where the small dependent firms drive productivity of SMEs. It seems that business groups have a significant role in improving aggregate productivity through ownership relationships. In addition, group structures is an important channel by which the small and large firms are interlinked, besides the input-output networks.

One may draw conflicting policy implications. The support of new start-ups as providers of new jobs might be justified, but if the policy goal is to improve productivity, then the ability to promote activities and ownership of business groups can be an advantage.

The analysis conducted in this paper can be extended in many ways. First of all, one can examine different aspects of dependent and independent SMEs. For example, the dependency status as a binary variable hides potentially interesting dynamics in relation to the entire ownership or production network, such as the channels by which the headquarters direct investments inside the firm. Moreover, an interesting direction of research is to analyze the aspects of allocative efficiency of business groups' internal markets, versus that of external markets. Finally, more industry specific investigation is warranted in order to see why service firms become more productive in a group.

5.7 Appendix A: Data description

Variable	Firm-years	Mean	Std. Dev.
LP	566,266	53,695	41,844
log LP	550,336	10.75	0.539
CEO changes	566,266	0.36	0.47
Owner changes	566,266	0.08	0.28
Mergers	566,266	0.005	0.07
Dependent	566,266	0.12	0.33
Foreign affiliate	566,266	0.026	0.159
Mother	566,266	0.041	0.199
Affiliate	566,266	0.085	0.279
Employees	566,266	16.957	150.756
Firm age	566,266	14.589	12.606
Firm age^2	566,266	371.7	780.3
HHI ownership	566,266	5,484	3,656
Job growth	535,442	0.39	1.58
Leverage (D/E)	566,260	77,016	471,247
STD. sales	515,306	33.266	15.837
Cost capital	566,266	0.0217	0.025
Change CEO	566,266	0.238	0.426
Average salary	566,253	32,002	19,554
Average Salary Exec	210,219	55,371	37,578
CEO salary	566,266	48,334	58,307
CEO tenure	555,431	7.04	6.74
CEO age	566,266	44.68	10.33
Firm executives	566,266	1.012	11.009
Salary executives	210,219	199,097	1,870,860
High-skill share	566,266	0.120	0.261
Employment turnover	406,079	5.400	22.351
Micro	566,266	0.78	0.42
Small	566,266	0.18	0.39
Medium	566,266	0.03	0.18
Large	566,266	0.008	0.09
Services	566,266	0.64	0.47
Mkt Share (%)	563,303	0.1	1.26
Group Debt	565,461	0.03	0.13
Group transfers	566,253	0.007	2.2

Table 5.8: Descriptive statistics of all the variables that appear in the analysis

Our sample consists of mostly small firms, as can be seen from Table 5.9 below:

	Firm-year	Employees	LP	job growth rate
Large	4,691	850.9	70,889	0.03
Medium	19,066	99.6	63,051	0.08
Small	103,364	20.3	56,612	0.20
Micro	439,145	3.7	51,289	0.45

Table 5.9: Average FTE, LP and job growth rates by firm-years divided into size categories. Size classes are defined, in terms of FTEs, as: micro (<10), small (10-49), medium (50-249) and large (>250).

Most of the variation in the dependency status (-1,0,+1) is from small firms in our within regressions.

	-1	0	+1
Large	6	4,022	21
Medium	108	15,965	445
Small	344	87,634	1,663
Micro	366	343,496	1,743

Table 5.10: Dependency status change in firm-years divided into size categories. Size classes are defined, in terms of FTEs, as: micro (<10), small (10-49), medium (50-249) and large (>250).

5.8 Appendix B: Adjustment for legal restructuring.

In this appendix, we discuss the details the procedure adopted by Statistics Finland to control for merger and split-offs in a set of enterprises. Assume that firm 1 is examined after an event (merger or split-off) where N firms are involved. Then the estimated employment of firm 1 one year ago is calculated by:

$$emp(firm_{1,t-12}) = \frac{emp(firm_{1,t}) * emp(firm_{1,t-12}, firm_{2,t-12} ... firm_{N,t-12})}{emp(firm_{1,t}, firm_{2,t} ... firm_{N,t})}$$

where t is the time periods in which the adjustment is computed, and N is the number of firms involved in a merger or split-off. The sum of the previous year employment levels in all the firms involved in the event is divided for each continuing firm weighted by their relative size at present time t. Let us go through some simple numerical examples to see how this works:

1. Assume a firm A with 2 employees in period t, that had 1 employee in t-12. Firm A acquires firm B with 1 employee at time t, m and 1 employee one year ago. Firm A, which continues existing, will be assigned a new estimated number of employees for the comparison year, in order to make the growth rates comparable year-on-year. The comparison values of firm A is estimated as $\frac{2(1+1)}{(2+1)} = 4/3$, and the rate of change for A becomes $(2+1)/(4/3) = 2.25$ (as opposed to 3 if no correction is done)

2. Consider the situation where firm A is split into smaller units, say B and C. A has 3 employees at time $t-12$, B has 3 employees at t and C has 2 workers at t. B and C did not exist at $t-12$, so their comparison values become: $(3/3)3 = 3$ and $(2/3)3 = 2$, resulting in the rate of change for B and C to be 3/3 and 2/2 (equal to 1 for both firms). The growth rate is forced to be the same among the continuing firms after a split-off.

5.9 Appendix C: Identifying the CEO

We are interested in identifying the CEO in each firm, interpreted as the person who has control on the firm's operations. We employ the following sequential procedure, as e.g. in Queiró (2016). First, we identify a person as the CEO if he or she is explicitly defined as such among the list of employees. This is the case for 7% of the firms. For the remaining firms, we consider those employees identified as having managerial responsibilities, and say that the CEO is the manager with the highest salary. This identifies an additional 30% of the CEOs. For the remaining firms, we look at whether an active entrepreneur (as classified by the tax administration) appears in the list of employees, in which case the person is identified as the CEO.[10] This is the case for 23% of our CEOs. The remaining 41% of the CEOs are defined as the highest paid worker in the firm. As a validation test, we notice that 86% of the CEOs explicitly defined as such (our first criterion) also have the highest salary in the firm.

[10]The tax administration identifies an active entrepreneur in a firm if a person owns at least 30% of the shares and receives a significant income from the firm (at least 9,663 euros in 2006).

Chapter 6

Job Creation and the Role of Dependencies

Paolo Fornaro*, Henri Luomaranta**

*Research Institute of the Finnish Economy, Finland

**Statistics Finland and TSM, University of Toulouse Capitole, Toulouse, France

Abstract

We contribute to the extensive literature on the relationship between firm size and job creation, by examining dependencies between enterprises. Using Finnish monthly data encompassing the population of Finnish private businesses, we calculate the gross job creation and destruction, together with the net job creation, for different size classes and industries. Importantly, we divide firms into a dependent (i.e. owned, at least partially, by a large company) and independent category. Due to the quality of the data, we are able to isolate the 'organic' growth of firms, disregarding the effects of mergers, split-offs and other legal restructuring. We find that independent companies have shown a considerably higher net job creation, regardless of their size class. However, dependent firms do not show particularly different behaviors with respect to the sensitivity to aggregate conditions, compared to their independent counterparts. Once we control for age, we find that independent firms generate more (net) jobs during the early years of their existence but destroy more jobs once they become older.

JEL Classifcation Code: D22, E24, E32, L25

Keywords: dependencies; firm size; firm age; employment creation

6.1 Introduction

The relationship between employment generation and firm's size has been the focus of extensive research. Since the seminal article of Birch (1981), there has been a lot of discussion about whether small firms are the main force underlying employment growth. This view has been the center of political debate, where public support to small businesses has been advocated in the light of their large growth enhancing capabilities. However, the original insights by Birch have been contested in multiple empirical works, which have pointed out issues underlying the data and the methodology adopted. For example, Davis et al. (1996) argue that the procedure that Birch (1981) uses to classify a firm as small or large (i.e. using the base year on which the growth rate is computed) leads to an overestimation of the job creation stemmed from smaller businesses. Subsequent works studying the effects of firms' size and job creation are, among many others, Davis et al. (1996) and Neumark et al. (2011). In these papers it has been found that, after adjusting for the statistical biases of Birch (1981), small firms do not create more net jobs compared to large ones, or at least not in such a dramatic way as found in Birch's work. For the Finnish economy, there has been a number of studies where the relation between firm sizes and net job flows is examined. Some examples of these analyses comprise Hohti (2000), Ilmakunnas and Maliranta (2003) and, more recently, Wit and Kok (2014) and Anyadike-Danes et al. (2015).

The enterprise size has not been the only firm's characteristic analyzed in regards to employment creation. Another important feature that has been considered as a contributing factor to net job growth is firm's age. A key study in this respect is Haltiwanger et al. (2013), where the authors show that once we control for firm's age, small and large firms do not show discrepancies in net job creation. Other studies which are interested in the effect of the firms' age on job creation are Criscuolo et al. (2014), Distante et al. (2014) and Anyadike-Danes et al. (2015). The common finding of these studies is that young firms are the main drivers of job creation, with start-ups being especially important.

In this paper, we investigate another possible source of heterogeneity among firms which might affect job growth, i.e. external ownership and dependence. In particular, we look at how firms belonging to an enterprise group contribute to the employment generation (both gross and net), within different size classes. Large corporations are a key player in modern economies, accounting for a large share of aggregate output and potentially have substantial effects on the business cycle (see, e.g., Gabaix, 2011). However, as pointed out in the previous literature, large firms are usually associated with lower job creation compared to small enterprises. The fact that previous analyses do not separate dependent and independent enterprises might be a decisive factor behind these results. In a recent Eurostat report (Airaksinen et al., 2015), the share of dependent firms' employment within the small and medium enterprises (SME) category is documented to be substantial in several European countries, including Finland. This consideration casts doubt on many previous conclusions in the small versus large literature, where the SME status is systematically defined by the number of employees only, regardless of the ownership structure. For example, the statistical result that small firms tend to create more jobs, on average, could stem from large firms investing through affiliates. Even in the case of looser control, it is arguable the employment generation of small dependent enterprises could be impacted by the decisions of the mother company. If dependent, small firms are behind the large job creation rates of SMEs, then the narrative of small businesses being the driver of employment generation should actually be interpreted in the light of large corporations creating jobs through subsidiaries.

The contribution to job creation by dependent and independent enterprises has not been studied extensively in the literature. A notable exception is Boccara (1997), where the author examines the job growth stemmed from small and medium firms belonging to enterprise groups in France during the 1984-1992 period, finding that the small firms belonging to large business groups exhibit higher job creation. Another work which touches the issue of dependencies and employment growth is the OECD report Schreyer (2000), in which the author discusses possible economic channels behind the relationship. Small firms might have multiple benefits from belonging to a large corporation. Subsidiaries owned (even partly) by a large company might have a better access to financing (both internal and external), together with more informal advantages such as access to a wider human capital and information related to market conditions and technology. However, there are possible channels leading to a negative impact of dependencies onto job creation. Large firms could consider their small subsidiaries as a small part of the production chain, which must perform a well defined and limited amount of tasks, without the need to grow in size.

We use monthly employment data of Finnish firms to study how the dependence to large companies affect the job creation (both gross and net) of enterprises, controlling for size and age. The data, extracted from the Statistics Finland database, allows us to verify if an enterprise belongs to a business group and how large is the share of the firm owned by the mother company, giving us the possibility to disentangle control from more informal dependencies and networks. The employment figures are adjusted to represent the "organic" growth of the firms, disregarding the effects of merger, split-offs and other legal restructuring. In addition, we examine the possible heterogeneity between the different industries of the economy (e.g. manufacturing and services), which might have an impact on how belonging to an enterprise group affects the job creation of a company. For example, it is likely that in the service industry, where human capital plays a larger role, firms benefits more from dependencies and connections than in, e.g., constructions. Finally, we analyze how dependent and independent enterprises respond to different aggregate economic conditions. In particular, we examine the job flows of firms with different ownership structures during periods of economic expansion (which we identify as periods in which monthly output is above trend) and economic downturns (output below trend).

We find that small, medium and micro independent firms have experienced consistently higher growth rates compared to their dependent counterparts, regardless of the size classification methodology and size class considered. Once we control for age, we find that young independent firms have generated more jobs compared to their dependent counterparts, but this relationship is reversed for older companies. This pattern can be explained by the fact that young independent firms are more uncertain about their productivity potential, causing them to create more jobs during the early stages of their life and subsequently destroy more jobs when they get older. We also find that the effect of dependencies is not unique across industries. In particular, while dependent firms exhibit lower job creation rates inside the trade, services and construction industries, the negative effect of dependencies disappears or reverts in the manufacturing and financial sectors. Finally, we do not find a clear effect of dependencies onto the sensitivity to the business cycle for small and medium firms.

The remainder of the paper is structured as follows: In Section 6.2 we introduce the main methodological issues underlying the analysis, in Section 6.3 we briefly describe the data and in Section 6.4 we present the results. Section 6.5 concludes.

6.2 Methodological issues

The analysis of job creation and its relation with the firm size is highly sensitive to the data source and the methodology adopted. For example, the criterion to determine if a given enterprise should be included in the small or large size class is not uniform over the literature and using different selection procedures can yield very different results. In the work of Birch (1981), firms are included in the small class if the number of employees during the base year of the job growth calculation is below a threshold. This criterion, as argued by Davis et al. (1996) among others, can lead to a serious overestimation of the job creation stemmed from smaller businesses. In particular, using the base year to classify a firm will cause the inclusion of many enterprises affected by temporary negative shocks in the small class (this phenomenon is addressed in the literature as the regression to the mean bias). Neumark et al. (2011) find that, using the base year classification of Birch (1981), small firms are generating a substantially larger share of employment compared to big enterprises. However, when they use the firms' average size to classify them, the gap between the job creation of small and large businesses shrinks substantially.

In this analysis, we use two size classification methodologies. The first one is the *dynamic size classification* method: enterprises are classified each year, using the average size between the two years on which the growth is computed. The number of full time equivalents (FTE) obtained is then compared to the cutoff points used by Statistics Finland to determine the size class of a company. As discussed in papers such as Davis et al. (1996) and Haltiwanger et al. (2013), this type of classification is robust to the regression to the mean bias. However, allowing companies to change size class over time tends to exacerbate the sensitivity of small enterprises to the business cycle. As discussed in Moscarini and Postel-Vinay (2012), during times of economic hardship we can expect firms to move to the small category and vice versa during expansions.

The second size classification criterion we use is called *average size classification* and it is based on the average number of employees (full time equivalents in our case) computed over the existence of the firm. As in the case of the *dynamic classification*, this methodology is robust to the regression to the mean bias. However, contrary to the *dynamic classification*, this methodology does not suffer from procyclicality issues. One problem with the *average size classification* is that it relies on the assumption that firms reach a long-term scale of operations during their lifespan, implying that the process underlying a firm's size is stationary.

The key measures of the analysis reported in Section 6.4 are the gross job creation, gross job destruction and net job growth. The gross job creation is defined as the sum of positive changes in the number of FTEs within a certain firm category, i.e. we have $jc_t = \sum_i^N dE_{it}^+$ where dE_{it}^+ are the positive changes in employment between time t and $t-12$ and which are then summed over the N firms belonging to a certain class. Job destruction is defined as $jd_t = \sum_i^N |dE_{it}^-|$, with dE_{it}^- being the negative change in the number of FTEs for company i. Importantly, we use the adjusted values for the FTEs in the base year, to control for mergers and acquisitions (details on the methodology are provided in the Appendix) and to obtain a measure of the organic growth of a firm. The net job creation is defined as the difference between gross job creation and job destruction. Finally, we compute two measures of net job creation rate. The first one is used to compute the contribution to the overall creation of jobs in the economy due to a category of companies. Denoting the net job creation at time t for category C as $NJC_{t,C}$, we compute

$$NJCR_{t,C}^1 = \frac{NJC_{t,C}}{(1/2E_t + 1/2E_{t-1}),} \tag{6.1}$$

where E_t is total employment. The denominator in (6.1) is suggested throughout the literature (e.g. Moscarini and Postel-Vinay, 2012) because it is more robust the regression to the mean bias. Another interesting measure is

$$NJCR_{t,C}^2 = \frac{NJC_{t,C}}{(1/2E_{t,C} + 1/2E_{t-1,C}),} \tag{6.2}$$

where $E_{t,C}$ indicates the total number of employees in category C, making (6.2) an indicator of how a certain category is growing over time.

A final issue worth discussing in this section is the role of firms' age. As pointed out in Haltiwanger et al. (2013), the age of a company is a key determinant in explaining its job creation. In particular they show that, after controlling for age, there is no clear difference in the net job creation rate of small and larger companies. To make sure that our results are not driven by the longevity of the firms we examine, we consider a subset of companies which are present throughout our sample period. Moreover, we analyze the impact of dependencies onto job creation while separating SMEs into different age categories.

6.3 Data Description

The data is extracted and anonymized at the premises of Statistics Finland, the Finnish national statistical agency, from the short term business statistics dataset (which is used internally to produce aggregate indexes). The data contains monthly observations of persons employed (as full time equivalents, FTEs) for the entire business sector, excluding public sector and primary producers. Thus, we analyze the employment generation patterns of enterprises that are active in the business economy. The analyzed enterprises are classified by Statistics Finland into broad activity categories based on the classification of economic activities system in the EU (NACE Rev. 2). In order to control for heterogeneities arising from the different activity categories, we group the enterprises into manufacturing, construction, trade, services and finance industries.

The Finnish Business register contains information on ownership links between the enterprises that belong to a group. Furthermore, the register holds information on the nationality of the enterprise group, and thus the Statistics agency is able to distinguish between foreign and domestically owned enterprise groups. By linking these data sources at micro-level, we are able to pinpoint whether at any given time an enterprise is "independent" (no enterprise group links), "dependent" (the enterprise is at least partly owned by a mother company, or the enterprise is a mother company itself), "controlled" (the mother company owns over 50% of the enterprise), or "foreign controlled" (the enterprise group head is foreign, and its ownership exceeds 50%). After applying these classifications to the enterprises, we use two sets of data. The first sample includes monthly observations of employment destruction and creation for all the enterprises that are active at any given month between January 1998 and September 2014, and the second sample includes employment creation and destruction of only those enterprises that are present for the full sample period, thus excluding entries and exits. Net job creation computation are based on adjusted FTEs, where the effects of mergers and split-offs are eliminated by the methodology of Statistics Finland. For the foreign controlled enterprises, the data is available only from January 2007 onward and hence is analyzed in a separate subsection.

The sample including entries and exits contains 253,685 enterprises in September 2014 and 234,257 enterprises in January 1998. The sample where only long lasting enterprises are included contains 70,356 enterprises. The following tables provide the number of enterprises in each of the analyzed categories by industry (Table 6.1) and size category (Table 6.2) in 09/2014 for both samples, in order to characterize the data and the Finnish business economy.

	Manufacturing	Construction	Trade	Services	Finance
Full sample					
Independent	20,716	37,565	41,813	124,439	2,021
Dependent	2,541	804	2,543	6,216	631
Controlled	2,340	758	2,438	5,532	597
Foreign controlled	543	87	1,023	937	141
Long lasting enterprises					
Independent	7,952	8,104	12,902	29,910	338
Dependent	1,307	339	1,299	2,035	195
Controlled	1,230	320	1,270	1,855	186
Foreign controlled	246	22	505	307	42

Table 6.1: Number of enterprises in September 2014, divided by industry and dependency status.

	Micro	Small	Medium	Large
Full sample				
Independent	216,093	9,634	775	52
Dependent	6,643	3,816	1,727	549
Controlled	5,840	3,611	1,672	542
Foreign controlled	1,110	918	500	203
Long lasting enterprises				
Independent	54,041	4,728	400	37
Dependent	2,163	1,862	855	287
Controlled	1,963	1,783	832	283
Foreign controlled	313	448	261	100

Table 6.2: Number of enterprises in September 2014, divided by size class and dependency status.

While the figures reported in Table 6.1 point toward dependent firms being a small share of the overall population of enterprises, Table 6.2 provides key information to motivate this analysis. The number of dependent medium-sized and small enterprises represents a large share of the total, highlighting the fact that disregarding the possible links between larger companies and subsidiaries might bias the results for two important size class of firms such as the small and medium enterprises.

6.4 Results

We start our empirical analysis by studying the relationship between firm size and the measures of interest reported in Section 6.2. In this fashion, we can compare the Finnish setting with the findings obtained in studies as, e.g. Davis et al. (1996) and Haltiwanger et al. (2013).

In particular, in Table 6.3, we report the total number of employees, the gross job creation and destruction, together with net job creation, for large and SMEs (i.e. the category encompassing small, medium and micro firms) companies. Moreover, we compare enterprises with different dependencies status, even though we do not separate firms of different size class within the same dependency class. We report results for both dynamic and average size classification and the results are expressed in terms of FTEs.

	Total Number of Employees	Gross Creation	Destruction	Net Job Creation
Average Size Classification				
Large	495, 383	28, 465	43, 924	−15, 458
Medium	235, 594	21, 627	26, 558	−4, 930
Small	256, 658	34, 670	33, 554	1, 115
Micro	317,340	74, 912	69, 480	5, 431
Dependent	695,932	50,470	60,589	−10, 119
Control	678,087	48,344	58,826	−10, 482
Independent	609,046	109,205	112, 929	−3, 723
Dynamic Size Classification				
Large	513,171	29, 137	42, 663	−13,526
Medium	230, 965	21, 051	27, 377	−6,325
Small	254, 768	33, 466	33, 596	−130
Micro	306, 072	76, 020	69,881	6, 139
Dependent	695,932	50,470	60,589	−10, 119
Control	678,087	48,344	58,826	−10, 482
Independent	609,046	109,205	112, 929	−3, 723

Table 6.3: Average number of total number of employees, gross creation, destruction and net job creation. Enterprises are divided by size class and dependency status. All values reported are FTEs.

The figures reported in Table 6.3 are somewhat similar to what has been found in the literature. Firms of smaller size exhibit large gross job creation and destruction, especially the enterprises in the micro category. Independently from the size classification methodology, large firms are the most important employer of the Finnish economy, considering the average number of FTEs between 1998 and 2014. At the same time, they have experienced the lowest net job creation, shredding on average more than 10,000 jobs on a year-on-year basis. Micro enterprises, on the other hand, seem to be the ones contributing the most to net job growth. This result holds regardless of the size classification method, even though the net job creation of these enterprises is slightly smaller if we use the average classification methodology. Interestingly, by using the dynamic size classification, micro firms are the only ones generating positive net job creation.

From this very simple analysis, we can already draw some interesting conclusions regarding the dependency effect on job creation. On average, dependent firms represent the majority of the population, employing almost 100,000 employees more than the independent enterprises (but this is most likely due to the presence of large mother companies). Moreover, the vast majority of employees within the dependent firms class work in controlled enterprises. In other words, most dependent enterprises are tightly controlled by their mother company (in terms of ownership). Independent firms, on the other hand, have a much higher gross creation and destruction, together with the highest net job growth. However, in Table 6.1 we are not separating the size effect and the dependency effect. For example, it might be that very low net job creation of dependent firms is due to the fact that larger companies are more likely to belong to this category and hence distort their actual contribution to job creation. Below, we report similar figures for SMEs and considering different type of dependency.

	Total Number of Employees	Gross Creation	Destruction	Net Job Creation
Average Size Classification				
Medium Dependent	161,656	13,288	16,513	−3,224
Small Dependent	72,757	8,665	9,608	−942
Micro Dependent	16,137	3,341	3,378	−36
SMEs Dependent	250,551	25,296	29,500	−4,203
Medium Controlled	155,015	12,594	15,953	−3,359
Small Controlled	67,371	7,937	9,028	−1,090
Micro Controlled	14,223	2,945	3,028	−82
SMEs Controlled	236,609	23,477	28,010	−4,533
Medium Independent	73,938	8,339	10,045	−1,706
Small Independent	183,901	26,004	23,946	2,058
Micro Independent	301,202	71,570	66,102	5,468
SMEs Independent	559,042	105,913	100,093	5,819
Dynamic Size Classification				
Medium Dependent	152,675	12,970	16,278	−3,307
Small Dependent	67,119	8,308	9,299	−991
Micro Dependent	14,736	3,253	3,562	−309
SMEs Dependent	234,531	24,532	2,9141	−4,608
Medium Control	146,267	12,332	15,746	−3,414
Small Control	61,856	7,585	8,749	−1,164
Micro Control	12,883	2,816	3,218	−401
SMEs Control	221,007	22,735	27,715	−4,979
Medium Independent	78,289	8,080	11,098	−3,017
Small Independent	187,649	25,158	24,297	861
Micro Independent	291,335	72,767	66,318	6,449
SMEs Independent	557,274	106,006	101,714	4,292

Table 6.4: Average number of total number of employees, gross creation, destruction and net job creation for small, medium and micro enterprises, divided by dependency status.

The results reported in Table 6.4 underline some substantial differences between dependent and independent firms, with respect to employment creation and destruction. Within the small and medium enterprises, independent firms represent the largest category, with more than the double the FTEs of dependent companies. Moreover, independent firms have experienced a much larger gross job creation and destruction, during our sample. Finally, companies which belong to the independent category seem to be the main source of the positive net job creation observed for small and micro enterprises.

The channels underlying the effect of dependency on firms' job creation does not have a clear *a priori* positive or negative impact. On the one hand, we expect that small firms belonging to a corporation benefits to the access of a large stock of human capital and knowledge which is likely to be available to the mother company. Moreover, the subsidiary can benefit from participation to the formal and informal networks of a large corporation, e.g. the ability to reach new clients and suppliers. These benefits can lead to a better performance of the small company, which in turn can lead to an increase in its size and hence to a larger job creation. On the other hand, a mother company can consider its subsidiaries as small parts of its production process, which are highly specialized. For example, a large mother company might be in charge of the administrative side of multiple subsidiaries, which would not require separate staff to handle managerial duties. In this way, the small enterprises belonging to a large corporation would be organized in a way to achieve maximum productivity and hence they might actually reduce the number of employees, leading to a lower job creation of dependent companies.

The findings outlined in this subsection point toward a negative impact of dependency onto job growth, with small companies belonging to a corporation showing negative job creation. Small dependent firms seem to be restricted to a specialized task and do not increase in size. The fact that they have been shredding jobs can be interpreted as an attempt of their mother companies to achieve high levels of productivity. Another possible explanation is that small dependent enterprises have been dragged down by the poor performance of their large mother companies, which have been declining in terms of job creation.

6.4.1 Dependencies and the role of age

Even though the results of Table 6.4 are extremely interesting in the light of showing the dependency effect against the size effect in job creation, we should examine another factor that has been regarded in the literature (see, e.g., Haltiwanger et al., 2013) as key in explaining the net job creation of different types of enterprises, i.e. firm age. To address this issue, we use two different datasets containing dependent and independent SMEs. The first dataset is the same adopted to obtain the results in Table 6.3 and 6.4 and considers entries and exits of firms, while the second one includes only continuous firms, i.e. present throughout our sample. In this way, we compare companies which have been long lasting, at least toward the end of the sample, and hence the effect of age should be milder. For example, Haltiwanger et al. (2013) show that the effect of age on the job creation of firms of different size is especially strong on start-up companies, while it reduces substantially for older enterprises.

In Table 6.5, we report the net job creation rates for dependent and independent medium, small and micro firms, computed using (6.1) and (6.2). To keep the analysis contained, we consider the results for the average size classification methodology only.

	$NJCR^1$%	$NJCR^2$%	$NJCR^1$% Continuous	$NJCR^2$% Continuous
Medium Dependent	−0.23	−1.71	−0.03	0.04
Small Dependent	−0.07	−1.04	−0.006	0.11
Micro Dependent	−0.002	-0.45	−0.007	−0.74
SMEs Dependent	−0.31	−1.39	−0.04	0.07
Medium Controlled	−0.25	−1.85	−0.04	−0.05
Small Controlled	−0.08	−1.52	−0.01	−0.12
Micro Controlled	−0.006	−0.92	−0.006	−0.21
SMEs Controlled	−0.33	−1.62	−0.06	−0.05
Medium Independent	−0.12	−3.21	0.15	0.80
Small Independent	0.17	0.87	0.28	0.79
Micro Independent	0.43	2.07	0.19	0.46
SMEs Independent	0.47	0.97	0.63	0.64

Table 6.5: Net job growth rates for micro, small and medium sized enterprises, divided by dependency status. Both the dataset with entries and exits and the one with long-lasting firms only are considered and results are obtained using the average size classification.

The results included in Table 6.5 confirm the strong effect of dependencies on the net job creation and the rate of growth of firms of different size class. Enterprises which depend or are controlled by a mother company have lower job creation rates and seem to grow less. The effect is especially pronounced for small and micro enterprises, while medium independent enterprises seem to have a lower growth rate, with respect to their initial size (i.e. looking at $NJCR^2$), compared to their dependent counterparts. However, they have a larger net job creation with respect the overall number of employees.

These considerations are not affected by shifting our focus to continuous enterprises. When we consider more stable companies, the net job creation rates and the growth rates of dependent firms become less negative or even turn positive. However, independent firms are still the ones that have contributed the most to employment generation.

As additional robustness check, we look at the effect of dependencies on the net job creation rate of SMEs of different age groups. We divide firms in "new" (0-1 year), "young" (2-5 years)," middle-age" (6-10 years) and "old" (10 or more years) and compute their net job creation rates using both formula (6.1) and (6.2). Notice that category "new" includes the very important entrant group, for which net job creation corresponds to their gross job creation.

The age of a firm is based on the procedure adopted by Statistics Finland, i.e. by looking at the age of the legal unit. This method is not flawless because a legal unit can be considered new if it is the result of legal restructuring. As pointed out in Hyytinen and Maliranta (2013), using the administrative age tends to make the firms look younger. Notice that we are interested in comparing dependent and independent enterprises, so the problem is centered on how the administrative age of dependent and independent firms is sensitive to this issue. It is arguable that young subsidiaries tend to include older enterprises which are formed after restructuring. However, if there is no new legal unit formed after the dependency status change, the age of the enterprise does not change (in other words, the age of the firm does not reset after becoming dependent or independent).

There are however some adjustments that milden this issue: first of all we are not considering large firms, which are the most sensible to this problem. Moreover, in case a firm is considered new because of a restructuring, we have access to its adjusted previous year value. In case that value is present for a given entrant, we omit that firm because it is not a real new entrant (greenfield entry). Finally, we want to stress that we are looking at organic changes of FTEs, so we are already filtering out the effects of mergers and split-offs when calculating the net job creation of the different groups. To see the effect of this adjustment, we also report the results where we consider all new firms based on the age of the legal unit (i.e. without making a greenfield entry adjustment). Results are reported in Table 6.6.

	$NJCR^1\%$	$NJCR^2\%$
Greenfield Dependent	0.13	
New Dependent	0.14	36.50
Young Dependent	−0.014	−0.28
Middle-age Dependent	−0.06	−2.14
Old Dependent	−0.40	−2.91
Greenfield Controlled	0.11	
New Controlled	0.12	34.46
Young Controlled	−0.02	−0.79
Middle Age Controlled	−0.06	−2.46
Old Controlled	−0.40	−3.09
Greenfield Independent	1.13	
New Independent	1.16	95.36
Young Independent	0.35	4.75
Middle-age Independent	−0.10	−2.19
Older Independent	−1.01	−3.75

Table 6.6: Net job growth rates for SMEs, divided by dependency status and age. Results are reported in percentage points.

First of all a word on the difference between greenfield entries and entrants based on their administrative age. As it can be seen, for all categories, the job creation of greenfield entrants is very similar to the net job creation of formal entrants. We have checked the average number of firms that are not real entrants and find that the proportion of non-greenfield entrants in the dependent category is 37% while it is 22% for independent firms. Moreover, looking at the results, it seems that most of the net job creation of new firms is due to greenfield entrants. Finally, and most importantly, using greenfield entrants does not remove the effect of dependencies, with the job creation of new independent firms 1% higher than the one of dependent firms. Notice that we do not report $NJCR^2$ for the greenfield entrants because the number of workers in that group corresponds to the job creation for that group.

Looking at the rest of the results in Table 6.6, we see that after we control for different age groups, we find an interesting pattern in the effect of dependency status. In particular, it seems that there is an inverse relationship between the effect of dependency on net job creation and the age of the enterprise. For new and young firms, we clearly see that being dependent has a negative effect on the net job creation, especially for new firms. The employment generated by new enterprises which are independent is almost 10 times higher than their dependent counterparts, while for young firms we find that dependent firms have a negative job creation rate against the positive one of independent companies. These considerations are even more clear when we look at the $NJCR^2$, i.e. how the group we are examining grew. We find that independent young firms have experienced an average yearly growth of 4.75% while dependent firms have a mildly negative job creation. Moreover, we can see that the group of independent new firms has experienced a growth rate which is almost three times the one of dependent new companies (notice that the very large values for $NJCR^2$ can be explained by the strong effect of entrants).

Things become radically different when we look at older enterprises. For middle-age firms (6 to 10 years old), we find that the dependency status does not have a large effect, especially when looking at the $NJCR^2$. However, when looking at older firms, we find that independent companies have had a substantially lower net job creation rate, both with respect to the overall economy (i.e. $NJCR^1$) and to their own size (even though to a smaller extent).

This interaction between firm age and the dependency status can be explain by the *experimentation* process (see Brynjolfsson et al., 2007, Hyytinen and Maliranta, 2013, and Gabler and Poschke, 2013) that new firms face when entering the market. It is plausible that a newly formed or young dependent enterprise has a better idea of its productivity potential (for example because it needs to perform a specific task for its mother company), compared to an independent one. This can be reflected in young independent firms creating more employment because they are too optimistic of their production possibilities. In time, when they achieve their long-run level of productivity, independent companies need to shred excessive jobs which they have created during the learning phase. On the other hand, a dependent or controlled company hires less during the initial stage of its life and hence does not need to decrease its labor input as much as its independent counterpart, when it gets older.

To sum up the results of this subsection, we find that controlling for age does not render the dependency status uninfluential in explaining the heterogeneity in the net job creation of Finnish enterprises. However, we find that the impact of dependencies changes as firms get older.

6.4.2 Cyclical Analysis

The results discussed in the previous subsections evidence the strong impact of ownership structure onto the average gross and net job creation. It is also interesting to analyze how dependency from a mother company affects the sensitivity of a firm to the business cycle. To do this we compute the euclidean distance between the mean net job creation of a certain category of firms during periods of low and high economic growth. A contractionary period is defined as month in which the indicator of real economic activity[1] is below its trend and vice versa for an expansionary period.

In other words, our measure of sensitivity to aggregate economic conditions is given by:

$$\Gamma_C = \frac{\sqrt{\overline{NJC}^2_{Rec,C} + \overline{NJC}^2_{Exp,C}}}{\overline{E}_C}, \tag{6.3}$$

where $\overline{NJC}_{Rec,C}$ is the average net job creation for category C during periods of slow economic growth and $\overline{NJC}_{Exp,C}$ is the same measure taken during period of good aggregate economic conditions. Finally, $\overline{E}_C$ is the average number of FTEs for category of firms C, which is used to make the figure comparable across companies of different class sizes and dependency status. Intuitively, a low value of Γ_C indicates that the employment generation of a certain type of enterprises does not vary substantially during different macroeconomic conditions. On the other hand, a large value of this indicator points toward a remarkable sensitivity of certain classes of firms to the business cycle.

We report, in Table 6.7, this measure of sensitivity to the business cycle for SMEs of various

[1]We use the Trend Indicator of Output (TIO), produced by Statistics Finland, as monthly measure of real economic activity.

ownership structure, considering both the dataset which includes entry and exit and the one with only continuous companies.

	Γ_C%	Γ_C % Continuous
Medium Dependent	3.70	2.32
Small Dependent	3.10	2.10
Micro Dependent	1.86	1.82
SMEs Dependent	3.37	2.22
Medium Control	3.89	2.37
Small Control	3.45	2.20
Micro Control	2.10	1.94
SMEs Control	3.63	2.29
Medium Independent	4.25	3.37
Small Independent	2.86	2.95
Micro Independent	2.89	1.60
SMEs Independent	2.36	2.04

Table 6.7: Sensitivity of micro, small and medium sized enterprises to aggregate economic conditions. Higher numbers indicate more sensitivity to the business cycle. Both the dataset with entries and exits and the one with long-lasting firms only are considered.

Looking at Table 6.7, we see that the sensitivity to business cycles varies widely across different types of firms. First of all, it seems that micro firms tend to be more stable with respect to different aggregate economic conditions. While this can be surprising in the light of works such as Gertler and Gilchrist (1994), where smaller enterprises are seen as especially sensitive to economic downturn, it resembles the conclusions obtained in Moscarini and Postel-Vinay (2012). In their analysis, the authors found that larger firms employment behavior exhibits stronger correlation to the business cycle.

Firms' age also plays an influential role in terms of the sensitivity to the macroeconomic cycle. This comes as no surprise, given that we expect older firms to fluctuate less and because we omit entries and exits, which are heavily affected by different economic conditions. The dependency status, however, does not seem to have a clear effect on the cyclicality of job creation. For example, independent medium and micro enterprises seem to be more sensitive to the aggregate economic environment compared to their dependent counterpart, while the opposite holds for small firms and SME category as a whole.

Overall, while dependencies have a strong effect on the average job creation, it does not seem to have a substantial impact on their cyclical behavior.

6.4.3 Sectoral Analysis

So far, we have analyzed firm-level data without distinguishing the industry to which a certain enterprise belongs to. We can expect the effect of dependencies to vary across different industries. For example, the sharing of know-how between the mother company and its subsidiaries might be more relevant in firms working in the service sector compared to the ones working in the construction or manufacturing sector.

In Table 6.8, we examine the net job creation rate defined following (6.1), where we use the total number of employees belonging to an industry as the denominator. We do this for dependent and independent SMEs belonging to various industries. For the sake of brevity, we limit our analysis to the dataset including entries and exits, and to the average size classification method.

	Dependent	Independent
Medium Construction	−0.003	0.03
Small Construction	−0.01	0.76
Micro Construction	−0.01	1.44
SMEs Construction	−0.036	2.23
Medium Finance	0.11	−0.16
Small Finance	0.023	0.07
Micro Finance	0.04	0.15
SMEs Finance	0.18	0.06
Medium Trade	−0.07	-0.04
Small Trade	−0.07	0.14
Micro Trade	−0.03	0.11
SMEs Trade	−0.17	0.21
Medium Services	−0.10	0.03
Small Services	−0.01	0.39
Micro Services	−0.01	0.83
SMEs Services	−0.12	1.25
Medium Manufacturing	−0.65	−0.40
Small Manufacturing	−0.17	−0.16
Micro Manufacturing	0.03	−0.06
SMEs Manufacturing	−0.80	−0.62

Table 6.8: Net job creation for micro, small and medium sized enterprises, divided by industry and dependency status. We consider only the average classification methodology and the dataset with entries and exits.

Table 6.8 highlights some interesting industry specific features to the relationship between dependency and job creation. Importantly, there is no a unique effect of dependency across industry. While we see that independent companies belonging to the service, trade and construction industries show substantially larger net job creation with respect to their dependent counterparts, the same cannot be said for the finance and manufacturing industries. In particular, the net job creation of enterprises in manufacturing do not seem to be affected greatly by the dependency status. Moreover, we find that independent firms in the finance industry have experienced lower growth compared to the dependent ones.

The results reported in this subsection shed some more light onto the possible interpretation of the general finding of the negative impact of dependency on job creation. One can argue that in the service and trade industries the mother company can intervene strongly in the administration side of its subsidiaries, which are then limited to some specialized tasks. On the other hand, in the manufacturing industry it is likely that the mother company cannot centralize some activities in the same fashion.

6.4.4 The Role of Foreign Ownership

As mentioned before, our data on foreign controlled enterprises start in January 2007. Given that this period is of particular importance, in the light of the Great Recession, and given the possible idiosyncrasies characterizing firms belonging to a foreign corporation, we decided to analyze them separately. In tables 6.9 and 6.10, we report both net job creation rate measures defined in Section 6.2 and the business cycle sensitivity indicator, respectively. We compare foreign controlled firms with the behavior of independent companies during the same sample period and consider both the data including entry and exit and the one with continuous firms only.

	$NJCR^1$%	$NJCR^2$%	$NJCR^1$% Continuous	$NJCR^2$% Continuous
Medium Foreign	−0.13	−1.61	−0.08	−1.89
Small Foreign	−0.03	−1.62	−0.01	−0.76
Micro Foreign	0.01	0.83	−0.0003	−0.07
SMEs Foreign	−0.15	−1.42	−0.10	−1.52
Medium Independent	−0.47	−5.40	−0.016	−0.35
Small Independent	−0.20	−0.73	−0.050	−0.35
Micro Independent	0.71	1.53	−0.08	−0.44
SMEs Independent	0.03	0.08	−0.14	−0.40

Table 6.9: Net job growth rates for micro, small and medium sized enterprises, divided by dependency status. Both the dataset with entries and exits and the one with long-lasting firms only are considered and results are obtained using the average size classification.

	Γ_C%	Γ_C % Continuous
Medium Foreign	5.67	4.22
Small Foreign	4.67	3.21
Micro Foreign	4.73	1.75
SMEs Foreign	4.75	3.82
Medium Independent	3.27	2.07
Small Independent	3.34	2.34
Micro Independent	2.30	1.80
SMEs Independent	2.81	2.03

Table 6.10: Sensitivity of micro, small and medium sized enterprises to aggregate economic conditions. Higher numbers indicate more sensitivity to the business cycle. Both the dataset with entries and exits and the one with long-lasting firms only are considered.

Tables 6.9 and 6.10 highlight some surprising results which go in a different direction compared to what we have found so far. Medium and small foreign-controlled firms show a higher (albeit still negative) net job creation rate from 2007 to 2014, compared to their independent counterparts. On the other hand, micro independent enterprises had a much better performance, in terms of job creation. Looking at the overall SMEs group, we find that both net job creation rate measures indicate a superiority of independent firms in generating employment. For long-lasting enterprises we find that the for all SMEs the dependence status has a positive effect on the net job creation rate, but they have grown less (as evidenced by lower $NJCR^2$).

The results contained in Table 6.10 evidence a clear characteristic of foreign-controlled firms, i.e. their high sensitivity to the business cycle. Enterprises that are controlled by a foreign corporation adjust better to different economic conditions and adjust their employment level accordingly. This holds true for both the data including entry and exits and the one with only continuous firms. An explanation for this result can be found in the fact that a foreign corporation can adjust production across different countries and reallocate resources based on the business cycles of the various economies in which it operates more easily.

6.4.5 A possible explanation: cross correlations of net job creations

So far, we have described the data agnostically, i.e. we did not seek a possible explanations to why dependent and independent SMEs show different patterns of employment behavior. In particular, we have found that independent firms have been more successful in creating new jobs, even though we have not determined a possible cause.

One of the most simple explanations is that dependent small enterprises are heavily affected by the performance of their mother company. This kind of relationship would be reflected in substantially higher correlations between the net job creation of large and dependent companies against the one between small independent enterprises and big firms. We report these correlations in Table 6.11, where we use the average size classification methodology and examine both data with entry and exit and continuous firms.

	Medium	Small	Micro
Dependent	0.773	0.740	0.275
Controlled	0.776	0.743	0.244
Foreign	0.398	0.429	−0.056
Independent	0.664	0.662	0.601
	Continuous	Firms	
Dependent	0.814	0.766	0.641
Controlled	0.813	0.743	0.634
Foreign	0.761	0.808	0.648
Independent	0.747	0.811	0.675

Table 6.11: Correlations of net job creation rates between SMEs and large firms.

In Table 6.11, we can see that the net job creations of independent firms are substantially less correlated with the one of large firms, compared to the controlled and dependent enterprises. This can point out to a "dragging down" explanation for the lower net job growth of dependent companies. However, there is a caveat: the performance of large companies, especially in a small economy as the Finnish one, can be an indicator of aggregate economic conditions and a lower correlation to the net job creation of large companies can simply indicate a lower sensitivity to the business cycle of independent firms. We touched on this point in section 6.4.2, but we could not identify a clear relationship between dependence status and business cycle sensitivity. For the micro enterprises, however, it seems that independent firms are more correlated to large enterprises than the dependent ones. This result might be driven by the very small values of net job creation of micro dependent enterprises.

Another interesting fact is the low correlation between foreign controlled firms and big Finnish companies. We have found, in Section 6.4.4, that foreign controlled enterprises are especially sensitive to the aggregate conditions of the Finnish economy. In the light of these results, the low correlations of Table 6.11 can be explained by the effect of large Finnish companies performance and not by a business cycle explanation, giving support to a possible dragging down effect underlying the lower net job growth of dependent firms.

Finally, in the case of continuing firms, we see slightly higher correlations coefficients with the net job growth of large companies (especially for the foreign controlled SMEs). This can be explained by the intrinsically higher stability of this kind of enterprises which show a common correlation with the overall trend underlying the economy.

6.5 Conclusions

We contribute to literature on the relationship between firms' size and job creation by investigating an additional source of heterogeneity within the SMEs, i.e. their dependency status. In particular, we separate the small and medium enterprises population using different degrees of control and examine their gross job creation and destruction, together with their net job growth.

We find that independent SMEs have experienced, on average, higher net job creation compared to firms which depend on a mother company. This result holds for all the size classes and different degrees of control. Moreover, we find that the negative effects of dependency onto job creation is present also when we examine only long-lasting enterprises. However, controlling for firm age introduces an inverse relationship between the effect of dependencies on job generation and the age of the company. In particular, young independent firms generate much more employment than their dependent counterparts, but older independent companies have slightly lower net job creation rates compared to subsidiaries. We also find that SMEs in different industries exhibit different patterns. Importantly, dependency status does not seem to play a large role in the job creation for the manufacturing industry, or at least not to the same extent as in the services and trade sector. Finally, we do not find a specific impact of dependencies onto the sensitivity of SMEs to the business cycle.

There are multiple channels that can explain the negative effect of being a subsidiary on the job creation of small firms. First of all, dependent enterprises are more than likely influenced by their mother company in their hiring decisions. If the mother company is shredding jobs, as it can be seen in the very negative net job creation of large companies in Table 6.3, it will probably impact its subsidiary, by blocking the creation of new jobs or even imposing job cuts to its small affiliates. This explanation is partially supported by the findings reported in Table 6.11, where dependent firms exhibit higher correlations with large companies. Another explanation can be found in the attempt to achieve higher productivity. It is possible for the mother company to centralize some tasks which were previously conducted within the subsidiary. In this view, the mother company sees the small subsidiary as a small part of the production process and does not have particular incentives in increasing the scale of its controlled firms. This reasoning can explain the results of Table 6.6, where dependent firms hire less during the early years but destroy fewer jobs once they get older.

The analysis conducted in this paper can be extended in multiple ways. First of all, we can examine different aspects of dependent and independent SMEs, other than employment. For example, we could look at labor productivity or the value added produced in different types of small enterprises, based on their dependency status. This productivity study could indicate if mother companies focus on keeping their subsidiaries small and efficient, explaining their lower job creation. Moreover, it could be interesting to analyze the share of firms contributing to the negative and positive job creation inside a given category. In this way, we could see if the negative job creation is generated by the largest companies within a size class or if the contribution to the job creation is evenly distributed.

6.6 Appendix: Adjustment for legal restructuring

In this appendix, we discuss the details the procedure adopted by Statistics Finland to control for merger and split-offs in a set of enterprises. Assume that firm 1 is examined after an event (merger or split-off) where N firms are involved. Then the estimated employment of firm 1 one year ago is calculated by:

$$emp(firm_{1,t-12}) = \frac{emp(firm_{1,t}) * emp(firm_{1,t-12}, firm_{2,t-12}...firm_{N,t-12})}{emp(firm_{1,t}, firm_{2,t}...firm_{N,t})}$$

where t is the time periods in which the adjustment is computed, and N is the number of firms involved in a merger or split-off. The sum of the previous year employment levels in all the firms involved in the event is divided for each continuing firm weighted by their relative size at present time t. Let us go through some simple numerical examples to see how this works:

1. Assume a firm A with 2 employees in period t, that had 1 employee in t-12. Firm A acquires firm B with 1 employee at time t, m and 1 employee one year ago. Firm A, which continues existing, will be assigned a new estimated number of employees for the comparison year, in order to make the growth rates comparable year-on-year. The comparison values of firm A is estimated as $\frac{2(1+1)}{(2+1)} = 4/3$, and the rate of change for A becomes $(2+1)/(4/3) = 2.25$ (as opposed to 3 if no correction is done)

2. Consider the situation where firm A is split into smaller units, say B and C. A has 3 employees at time $t-12$, B has 3 employees at t and C has 2 workers at t. B and C did not exist at $t-12$, so their comparison values become: $(3/3)3 = 3$ and $(2/3)3 = 2$, resulting in the rate of change for B and C to be 3/3 and 2/2 (equal to 1 for both firms). The growth rate is forced to be the same among the continuing firms after a split-off.

References

Knut A. Aastveit and Torres G. Trovik. Estimating the output gap in real time: A factor model approach. *The Quarterly Review Of Economics and Finance*, 54(2008/23), February 2014.

John M. Abowd, John Haltiwanger, Ron Jarmin, Julia Lane, Paul Lengermann, Kristin McCue, Kevin McKinney, and Kristin Sandusky. The relation among human capital, productivity, and market value: Building up from micro evidence. In Carol Corrado, John Haltiwanger, and Daniel Sichel, editors, *Measuring Capital in the New Economy*, pages 153–198. University of Chicago Press, 2005.

Daron Acemoglu, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016, 2012.

Philippe Aghion and Peter W. Howitt. *The Economics of Growth*. MIT Press, 2009.

Aarno Airaksinen, Henri Luomaranta, Pekka Alajääskö, and Anton Roodhuijzen. Dependent and independent SMEs and large enterprises. Statistics explained, Eurostat, 2015.

Jyrki Ali-Yrkkö, Timo Seppälä, Tuomo Nikulainen, and Mika Pajarinen. Nokia and finland in a sea of change. Technical Report B 244, ETLA, 2010.

Filippo Altissimo, Riccardo Cristadoro, Mario Forni, Marco Lippi, and Giovanni Veronese. New Eurocoin: Tracking Economic Growth in Real Time. *The Review of Economics and Statistics*, 92(4): 1024–1034, November 2010.

Ronald C Anderson and David M Reeb. Founding-family ownership and firm performance: evidence from the s&p 500. *Journal of Finance*, 58(3):1301–1328, 2003.

James S Ang, Rebel A Cole, and James Wuh Lin. Agency costs and ownership structure. *Journal of Finance*, 55(1):81–106, 2000.

Michael Anyadike-Danes, Carl-Magnus Bjuggren, Sandra Gottschalk, Werner Hölzl, Dan Johansson, Mika Maliranta, and Anja Myrann. An international cohort comparison of size effects on job growth. *Small Business Economics*, 44(4):821–844, April 2015. doi: 10.1007/s11187-014-9622-0.

S. Boragan Aruoba, Francis X. Diebold, and Chiara Scotti. Real-Time Measurement of Business Conditions. *Journal of Business & Economic Statistics*, 27(4):417–427, 2009.

Spyros Arvanitis and Tobias Stucki. Do mergers and acquisitions among small and medium-sized enterprises affect the performance of acquiring firms? *International Small Business Journal*, 33(7): 752–773, 2015. doi: 10.1177/0266242614522280.

Enghin Atalay, Ali Hortaçsu, and Chad Syverson. Vertical integration and input flows. *American Economic Review*, 104(4):1120–48, April 2014. doi: 10.1257/aer.104.4.1120.

Jushan Bai and Serena Ng. Determining the number of factors in approximate factor models. *Econometrica*, 70(1):191–221, January 2002a.

Jushan Bai and Serena Ng. Boosting diffusion indices. *Journal of Applied Econometrics*, 24(4):607–629, 2009. doi: 10.1002/jae.1063.

Jushuan Bai and Serena Ng. Determining the Number of Factors in Approximate Factor Models. *Econometrica*, 70(1):191–221, 2002b.

Emanuele Baldacci, Dario Buono, George Kapetanios, Stephan Krische, Massimiliano Marcellino, Gian Luigi Mazzi, and Fotis Papailias. Big Data and Macroeconomic Nowcasting: from data access to modelling . Technical report, Eurostat, December 2016.

Vassiliki Bamiatzi, Salih Tamer Cavusgil, Liza Jabbour, and Rudolf R. Sinkovics. Does business group affiliation help firms achieve superior performance during industrial downturns? an empirical examination. *International Business Review*, 23:195–211, 2014.

Marta Banbura, Domenico Giannone, and Lucrezia Reichlin. Nowcasting. In *Oxford Handbook of Economic Forecasting*. Oxford University Press, 2011.

Oriana Bandiera Renata Lemos Andrea Prat Raffaella Sadun Oriana Bandiera, Renata Lemos, Andrea Prat, and Raffaella Sadun. Managing the family firm: Evidence from ceos at work. *The Review of Financial Studies*, 31(5):1605–1653, 2017.

Morten Bennedsen, Kasper Meisner Nielsen, Francisco Pérez-González, and Daniel Wolfenzon. Inside the family firm: The role of families in succession decisions and performance. *Quarterly Journal of Economics*, 122(2):647–691, 2007.

Morten Bennedsen, Francisco Pérez-González, and Daniel Wolfenzon. Evaluating the impact of the boss: Evidence from ceo hospitalization events. *Unpublished manuscript*, 2012.

Adolf A. Berle and Gardiner C. Means. *The modern corporation and private property*. Transaction Publishers, 1932.

Marianne Bertrand and Sendhil Mullainathan. Enjoying the quiet life? corporate governance and managerial preferences. *Journal of Political Economy*, 111(5):1043–1075, 2003.

Marianne Bertrand and Antoinette Schoar. Managing with style: The effect of managers on firm policies. *Quarterly Journal of Economics*, 118(4):1169–1208, 2003.

Ilke Van Beveren. Total factor productivity estimation: A practical review. *Journal of Economic Surveys*, 26(1):98–128, 2012.

Sumon Kumar Bhaumik, Saul Estrin, and Tomasz Mickiewicz. Ownership identity, strategy and performance: Business group affiliates versus independent firms in india. *Asia Pacific Journal of Management*, 34(2):281–311, 2017.

David L. Birch. Who creates jobs? *Public Interest*, 65(3):3–14, 1981.

Nicholas Bloom, Raffaella Sadun, and John Van Reenen. Management as a technology. Technical report, Harvard Business School, 2016.

Bloomberg. Nokia's Junk Debt Rating Dropped Lower by SP on Losses. `http://www.bloomberg.com/news/articles/2012-08-15/nokia-s-junk-debt-rating-dropped-lower-by-s-p-on-losses`, August 2012.

Frédéric Boccara. Mythes et réalités sur l'emploi et la croissance des PME : le rôle des groupes (1984-1992). Document de travail, INSEE, 1997.

Erik Brynjolfsson, Andrew McAfee, Michael Sorell, and Feng Zhu. Scale without mass: business process replication and industry dynamics. *Proceedings*, 2007.

Michael Carney, Marc Van Essen, Saul Estrin, and Daniel Shapiro. Business group prevalence and impact across countries and over time: What can we learn from the literature? *Multinational Business Review*, 25(1):52–76, 2017.

Vasco Carvalho and Xavier Gabaix. The Great Diversification and Its Undoing. *American Economic Review*, 103(5):1697–1727, August 2013.

CNN.com. MobilCom: The fallout begins. `http://edition.cnn.com/2002/BUSINESS/09/13/ericsson.mobilcom/index.html`, September 2002.

Timothy G. Conley and Bill Dupor. A Spatial Analysis of Sectoral Complementarity. *Journal of Political Economy*, 111(2):311–352, 2003.

Chiara Criscuolo. Foreign affiliates in oecd economies: Presence, peerformance and contribution to host countries' growth. Book, OECD, 2005.

Chiara Criscuolo, Peter N. Gal, and Carlo Menon. The Dynamics of Employment Growth: New Evidence from 18 Countries. CEP Discussion Papers dp1274, Centre for Economic Performance, LSE, June 2014.

Steven J. Davis, John Haltiwanger, and Scott Schuh. Small business and job creation: Dissecting the myth and reassessing the facts. *Small Business Economics*, 8(4):297–315, August 1996.

Jan De Loecker and Pinelopi Koujianou Goldberg. Firm performance in a global market. *Annu. Rev. Econ.*, 6(1):201–227, 2014.

Christine De Mol, Domenico Giannone, and Lucrezia Reichlin. Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components? *Journal of Econometrics*, 146(2):318–328, October 2008.

David J. Denis and Diane K. Denis. Performance changes following top management dismissals. *Journal of Finance*, 50(4):1029–1057, 1995.

Julian di Giovanni, Andrei A. Levchenko, and Isabelle Mejean. Firms, Destinations, and Aggregate Fluctuations. *Econometrica*, 82(4):1303–1340, 07 2014.

Roberta Distante, Ivan Petrella, and Emiliano Santoro. Size, Age and the Growth of Firms: New Evidence from Quantile Regressions. Working Papers 2014.69, Fondazione Eni Enrico Mattei, July 2014. URL `https://ideas.repec.org/p/fem/femwpa/2014.69.html`.

Catherine Doz, Domenico Giannone, and Lucrezia Reichlin. A two-step estimator for large approximate dynamic factor models based on kalman filtering. *Journal of Econometrics*, 164(1):188–205, September 2011.

Catherine Doz, Domenico Giannone, and Lucrezia Reichlin. A quasi–maximum likelihood approach for large, approximate dynamic factor models. *The Review of Economics and Statistics*, 94(4):1014–1024, November 2012.

Alex Edmans and Clifford G Holderness. Blockholders: A survey of theory and evidence. In *The handbook of the economics of corporate governance*, volume 1, pages 541–636. Elsevier, 2017.

Eurostat. Glossary:enterprise group, February 2019. URL `https://ec.europa.eu/eurostat/statistics-explained/index.php/Glossary:Enterprise`$_{g}roup$.

Martin D. D. Evans. Where Are We Now? Real-Time Estimates of the Macroeconomy. *International Journal of Central Banking*, 1(2), September 2005.

Kornelia Fabisik, Rüdiger Fahlenbrach, René M Stulz, and Jérôme P Taillard. Why are firms with more managerial ownership worth less? Technical report, National Bureau of Economic Research, 2018.

Fernando Fernandez-Rodriguez, Simon Sosvilla-Rivero, and Julian Andrada-Felix. Exchange-rate forecasts with simultaneous nearest-neighbour methods: evidence from the EMS. *International Journal of Forecasting*, 15(4):383–392, October 1999.

Allen Ferrel, Hao Liang, and Luc Rennebook. Socially responsible firms. *Journal of Financial Economics*, 122(3):585–606, 2016.

Andrew T. Foerster, Pierre-Daniel G. Sarte, and Mark W. Watson. Sectoral versus aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy*, 119(1):1 – 38, 2011.

Paolo Fornaro. Predicting Finnish economic activity using firm-level data. *International Journal of Forecasting*, 32(1):10–19, 2016.

Paolo Fornaro and Henri Luomaranta. Job creation and the role of dependencies. Technical Report 44, The Research Institute of the Finnish Economy, 2016.

Paolo Fornaro and Henri Luomaranta. Small and medium firms, aggregate productivity and the role of dependencies. Technical Report 47, The Research Institute of the Finnish Economy, 2017.

Paolo Fornaro and Henri Luomaranta. Aggregate fluctuations and the effect of large corporations: Evidence from finnish monthly data. *Economic Modelling*, 70:245–258, 2018.

Mario Forni, Marc Hallin, Marco Lippi, and Lucrezia Reichlin. The Generalized Dynamic-Factor Model: Identification And Estimation. *The Review of Economics and Statistics*, 82(4):540–554, November 2000.

Jeremy T. Fox and Valérie Smeets. Does input quality drive measured differences in firm productivity? Working Paper 16853, National Bureau of Economic Research, 2011.

Paul Frijters and Nemanja Antić. Can collapsing business networks explain economic downturns? *Economic Modelling*, 54:289 – 308, 2016.

Xavier Gabaix. The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772, 05 2011.

Alain Gabler and Markus Poschke. Experimentation by Firms, Distortions, and Aggregate Productivity. *Review of Economic Dynamics*, 16(1):26–38, January 2013. doi: 10.1016/j.red.2012.10.010.

Andrew Gelman, Aleks Jakulin, Maria Grazia Pittau, and Yu-Sung Su. A weakly informative default prior distribution for logistic and other regression models. *Ann. Appl. Stat.*, 2(4):1360–1383, 12 2008.

Mark Gertler and Simon Gilchrist. Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms. *The Quarterly Journal of Economics*, 109(2):309–340, 1994. doi: 10.2307/2118465.

Xavier Giroud and Holger M. Mueller. Capital and Labor Reallocation within Firms. *The Journal of Finance*, 70(4):1767–1803, 2015.

Yuriy Gorodnichenko, Enrique G. Mendoza, and Linda L. Tesar. The Finnish Great Depression: From Russia with Love. *American Economic Review*, 102(4):1619–1644, June 2012.

Jonathan Guryan. Desegregation and black dropout rates. *American Economic Review*, 94(4):919–943, 2004.

John Haltiwanger, Ron S. Jarmin, and Javier Miranda. Who creates jobs? small versus large versus young. *The Review of Economics and Statistics*, 95(2):347–361, May 2013.

Julie I. Hancock, David G. Allen, Frank A. Bosco, Karen R. McDaniel, and Charles A. Pierce. Meta-analytic review of employee turnover as a predictor of firm performance. *Journal of Management*, 39 (3):573–603, 2013.

Oliver Hart and Bengt Holmström. A theory of firm scope. *The Quarterly Journal of Economics*, 125 (2):483–513, 2010.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning: data mining, inference and prediction.* Springer, 2 edition, 2009.

Helsingin Sanomat. Nokia posts record quarterly results: share price plummets. `http://www2.hs.fi/english/archive/news.asp?id=20000728xx1`, July 2000.

Helsingin Sanomat. This is how a Helsingin Sanomat journalist tried to save Nokia. `http://www.hs.fi/kuukausiliite/a1381288411564?ref=hs-art-new-1`, October 2013.

Charles P Himmelberg, R Glenn Hubbard, and Darius Palia. Understanding the determinants of managerial ownership and the link between ownership and performance. *Journal of Financial Economics*, 53(3):353–384, 1999.

Satu Hohti. Job flows and job quality by establishment size in the finnish manufacturing sector 1980-94., 2000.

Michael Horvath. Sectoral shocks and aggregate fluctuations. *Journal of Monetary Economics*, 45(1): 69–106, 2000.

Rob Hyndman and Yeasmin Khandakar. Automatic time series forecasting: The forecast package for r. *Journal of Statistical Software, Articles*, 27(3):1–22, 2008.

Rob J. Hyndman and George Athanasopoulos. *Forecasting: principles and practice*. OTexts, 2018. ISBN 9780987507112. URL `https://books.google.fi/books?id=_bBhDwAAQBAJ`.

Aapo Hyvärinen. and Erkki Oja. Independent component analysis: Algorithms and applications. *Neural Netw.*, 13(4-5):411–430, May 2000. ISSN 0893-6080.

Ari Hyytinen and Mika Maliranta. Firm lifecycles and evolution of industry productivity. *Research Policy*, 42(5):1080–1098, 2013. doi: 10.1016/j.respol.2013.01.

Pekka Ilmakunnas and Mika Maliranta. The turnover of jobs and workers in a deep recession: evidence from the finnish business sector. *International Journal of Manpower*, 24(3):216–246, 2003.

Michael C. Jensen. Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review*, 76(2):323–329, 1986.

Michael C. Jensen and William H. Mecling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, 1976.

W. Bruce Johnson, Robert P. Magee, Nandu J. Nagarajan, and Harry A. Newman. An analysis of the stock price reaction to sudden executive deaths: Implications for the managerial labor market. *Journal of Accounting and Economics*, 7(1-3):151–174, 1985.

Julie Josse and François Husson. missmda: A package for handling missing values in multivariate data analysis. *Journal of Statistical Software, Articles*, 70(1):1–31, 2016. ISSN 1548-7660.

Boryan Jovanovic. Selection and the evolution of industry. *Econometrica*, 50(3):649–670, May 1982.

Daniel Kaufmann. Corruption, governance and security: Challenges for the rich countries and the world. *Available at SSRN 605801*, 2004.

Tarun Khanna and K. Palepu. The future of business groups in emerging markets: Long-run evidence from chile. *Academy of Management Journal*, 43:268–285, 2000.

Tarun Khanna and Yishay Yafeh. Business groups in emerging markets: Paragons or parasites? *Journal of Economic Literature*, 45(2):331–372, 2007.

Juha Kilponen, Seppo Orjasniemi, Antti Ripatti, and Fabio Verona. The Aino 2.0 model. Research Discussion Papers 16/2016, Bank of Finland, May 2016.

Mika Kuismanen and Ville Kämppi. The effects of fiscal policy on economic activity in Finland. *Economic Modelling*, 27(5):1315 – 1323, 2010.

Finn E. Kydland and Edward C. Prescott. Time to Build and Aggregate Fluctuations. *Econometrica*, 50(6):1345–70, 1982.

Borja Larrain, Giorgo Sertsios, and Francisco Urzúa I. The effects of losing business group affiliation. *Review Financial Studies*, hhy 120, 2018.

Markku Lehmus. Empirical macroeconomic model of the Finnish economy (EMMA). *Economic Modelling*, 26(5):926 – 933, 2009.

Ulf von Lilienfeld-Toal and Stefan Ruenzi. Ceo ownership, stock market performance, and managerial discretion. *Journal of Finance*, 69(3):1013–1050, 2014.

John B. Long and Charles I. Plosser. Real Business Cycles. *Journal of Political Economy*, 91(1):39–69, 1983.

Jacques Mairesse and Elizabeth Kremp. A look at productivity at the firm level in eight french service industries. *Journal of Productivity Analysis*, 4(1-2):211 – 234, 1993.

Vojislav Maksimovic and Gordon Phillips. Do conglomerate firms allocate resources inefficiently across industries? theory and evidence. *The Journal of Finance*, 57(2):721–767, April 2002. doi: 10.1111/1540-6261.00440.

Vojislav Maksimovic and Gordon M. Phillips. Conglomerate firms, internal capital markets, and the theory of the firm. *Annual Review of Financial Economics*, 5:225–244, November 2013.

Mika Maliranta. *Micro Level Dynamics of Productivity Growth. An Empirical Analysis of the Great Leap in Finnish Manufacturing Productivity in 1975-2000.* Number 38 in ETLA A. The Research Institute of the Finnish Economy, June 2003.

Mika Maliranta and Satu Nurmi. Business owners, employees, and firm performance. *Small Business Economics*, 52(1):111–129, 2019.

Troy D. Matheson, James Mitchell, and Brian Silverstone. Nowcasting and predicting data revisions using panel survey data. *Journal of Forecasting*, 29(3):313–330, 2010.

John J. McConnell and Henri Servaes. Additional evidence on equity ownership and corporate value. *Journal of Financial Economics*, 27(2):595–612, 1990.

Michele Modugno. Now-casting inflation using high frequency data. *International Journal of Forecasting*, 29(4):664–675, 2013.

Randall Morck, Andrei Shleifer, and Robert W. Vishny. Management ownership and market valuation: An empirical analysis. *Journal of Financial Economics*, 20:293–315, 1988.

Giuseppe Moscarini and Fabien Postel-Vinay. The contribution of large and small employers to job creation in times of high and low unemployment. *American Economic Review*, 102(6):2509–39, October 2012.

David Neumark, Brandon Wall, and Junfu Zhang. Do Small Businesses Create More Jobs? New Evidence for the United States from the National Establishment Time Series. *The Review of Economics and Statistics*, 93(1):16–29, August 2011. doi: 10.1162/REST-a-00060.

Whitney Newey and Kenneth West. A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–08, 1987.

Nokia press release. India emerges as the second largest market for Nokia. `http://company.nokia.com/en/news/press-releases/2007/08/23/india-emerges-as-the-second-largest-market-for-nokia`, August 2007.

Rickard Nyman and Paul Ormerod. Predicting Economic Recessions Using Machine Learning Algorithms. Papers 1701.01428, arXiv.org, January 2017.

Francisco Pérez-González. Inherited control and firm performance. *American Economic Review*, 96(5): 1559–1588, 2006.

Vasilios Plakandaras, Theophilos Papadimitriou, and Periklis Gogas. Forecasting Daily and Monthly Exchange Rates with Machine Learning Techniques. *Journal of Forecasting*, 34(7):560–573, November 2015.

Francisco Queiró. The effect of manager education on firm growth. *Quarterly Journal of Economics*, 118(4):1169–1208, 2016.

Reuters. Electronic component shortages may last through 2011. `http://uk.reuters.com/article/2010/05/24/tech-summit-components-idUKLNE64N01X20100524`, May 2010.

Reuters. Nokia to cut 3,500 jobs, close Romania plant. `http://www.reuters.com/article/2011/09/29/us-nokia-idUSTRE78S0DX20110929`, September 2011.

Antoinette Schoar. Effects of corporate diversification on productivity. *The Journal of Finance*, 57(6): 2379–2403, December 2002.

Paul Schreyer. High-growth firms and employment. Oecd science, technology and industry working papers, OECD, 2000.

Andrei Shleifer and Robert W. Vishny. A survey of corporate governance. *Journal of Finance*, 52(2): 737–783, 1997.

Jeremy C. Stein. Agency, information and corporate investment. In *Handbook of the Economics of Finance*, volume 1, pages 111–165. Elsevier, 2003.

Andrea Stella. Firm dynamics and the origins of aggregate fluctuations. *Journal of Economic Dynamics and Control*, 55:71–88, June 2015.

James H. Stock and Mark W. Watson. Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics*, 20(2):147–62, April 2002a.

James H. Stock and Mark W. Watson. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97:1167–1179, December 2002b.

James H. Stock and Mark W. Watson. Combination forecasts of output growth in a seven-country data set. *Journal of Forecasting*, 23(6):405–430, 2004.

Chad Syverson. What determines productivity? *Journal of Economic Literature*, 49(2):326–65, June 2011.

Juha Tarkka. Monetary policy in the BOF3 quarterly model of the Finnish economy. *Economic Modelling*, 2(4):298 – 306, 1985.

The Economist. More, more, more. `http://www.economist.com/node/6975882`, May 2006.

The Guardian. Nokia makes first loss for more than a decade in face of fierce competition. `http://www.theguardian.com/business/2009/oct /15/nokia-has-first-loss-for-a-decade`, October 2009.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.

John Van Reenen. The creation and capture of rents: Wages and innovation in a panel of u.k. companies. *Quarterly Journal of Economics*, 111(1):195–226, 1996.

Marno Verbeek and Theo Nijman. Testing for selectivity bias in panel data models. *International Economic Review*, pages 681–703, 1992.

Belen Villalonga and Raphael Amit. How do family ownership, control and management affect firm value? *Journal of Financial Economics*, 80(2):385–417, 2006.

Michael S. Weisbach. Ceo turnover and the firm's investment decisions. *Journal of Financial Economics*, 37(2):159–188, 1995.

Christopher K. I. Williams and David Barber. Bayesian classification with gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(12):1342–1351, Dec 1998.

Gerrit Wit and Jan Kok. Do small businesses create more jobs? new evidence for europe. *Small Business Economics*, 42(2):283–295, February 2014.

Klaus Wohlrabe and Teresa Buchen. Assessing the Macroeconomic Forecasting Performance of Boosting: Evidence for the United States, the Euro Area and Germany. *Journal of Forecasting*, 33(4):231–242, 07 2014.

Jefrey M. Wooldridge. *Econometric Analysis of Cross Section and Panel Data*. MIT Press, 2010.

Jing Xiao. The effects of acquisition on the growth of new technology-based firms: Do different types of acquirers matter? *Small Business Economics*, 45(3):487–504, October 2015. doi: 10.1007/s11187-015-9656-y.

Xianming Zhou. Understanding the determinants of managerial ownership and the link between ownership and performance: comment. *Journal of Financial Economics*, 62(3):559–571, 2001.

Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 67(2):301–320, 2005. ISSN 1369-7412.

Administrative registers maintained by statistical offices on vastly heterogeneous firms have much untapped potential to reveal details on sources of productivity of firms and economies alike.
It has been proposed that firm-level shocks can go a long way in explaining aggregate fluctuations. Based on novel monthly frequency data, idiosyncratic shocks are able to explain a sizable share of the Finnish economic fluctuations, providing support to the granular hypothesis.
The global financial crisis of 2007-2008 has challenged the field of economic forecasting, and nowcasting has become an active field. This thesis shows that the information content of firm-level sales and truck traffic can be used for nowcasting GDP figures, by using a specific mixture of machine learning algorithms.
The agency problem lies at the heart of much of economic theory. Based on a unique dataset linking owners, CEOs and firms, and exploiting plausibly exogenous variations in the separation of ownership and control, agency costs seem to be an important determinant of firm productivity. Furthermore, the effect appear strongest in medium-sized firms.
Enterprise group structures might have important implications on the voluminous literature on firm size, as large share of SME employment can be attributed to affiliates of large business groups. Within firm variation suggests that enterprise group affiliation has heterogeneous impacts depending on size, having strong positive impact on productivity of small firms, and negative impact on their growth. In terms of aggregate job creation, it is found that the independent small firms have contributed the most.
The results in this thesis underline the benefits of paying attention to samples encompassing the total population of firms in order shape more comprehensive policies. Researchers should continue to explore the potential of rich administrative data sources at statistical offices and strive to strengthen the ties with the official data producers.
Keywords: Granular effect, nowcasting, agency costs, CEO ownership, firm productivity, job creation, business group affiliation, firm size
Mots clés: choc granulaire, Prévisions du PIB en temps-réel , dilemme de l’agence, propriétaires-dirigeants, productivité de l’entreprise, création d’emplois , groupe de sociétés, taille des enterprises