Business Groups as Knowledge-Based Hierarchies of Firms[*]

Carlo Altomonte (Bocconi University)

Gianmarco I.P. Ottaviano (Bocconi University and LSE)

Armando Rungi (IMT Lucca)

JUNE 2018 - PRELIMINARY

Abstract

We propose a new theory of business groups as knowledge-based hierarchies that arise when contractual incompleteness may lead to dissipation of firms' specific knowledge. The theory predicts that a parent firm choosing to organize its activities as a business group rather than as a single integrated entity is more likely to emerge in good institutional environments. When this happens, a 'hierarchical' business group with several layers of subsidiaries controlled by the parent is more likely to appear than a 'flat' one with fewer layers if the firm has better production possibilities (which require more challeging problem solving), faces lower communication costs between hierarchical layers, and incurs a lower skill premium in hiring good managers. We provide empirical support for these theoretical predictions exploiting the unique features of a dataset in which we observe the ownership structures (number of subsidiaries, countries and industries in which subsidiaries operate, and the subsidiaries' positions at different hierarchical layers) of 178,190 business groups incorporated in OECD countries and controlling more than 1,150,000 (domestic and foreign) subsidiaries worldwide in the year 2010.

JEL classification: D23; L23; F23; L25; G34

Keywords: business groups, knowledge hierarchies, property rights, contract theory, multinational enterprises, organization of production

[*]We are indebted with Luis Garicano for initial discussions about the topic of this paper. We are also grateful to Esteban Rossi-Hansberg for useful comments.

1 Introduction

A business groups (BG) is a non-standard organizational form consisting of a collection of at least two legally autonomous firms that function as a single economic entity through a common source of hierarchical control via equity stakes. BGs are crucial components of the global economy. The world's largest businesses by consolidated revenue (as classified in the Fortune 500 list) as well as the top 100 multinational enterprises (as listed by UNCTAD) are all organized as BGs. These multinational groups have on average around 300 affiliates/subsidiaries each, and up to 10 hierarchical layers of control (UNCTAD, 2016). According to BEA data, at least 75% of total US trade can be linked to firms operating in the US as parts of BGs (either as US headquarters or US subsidiaries of foreign groups).[1] Yet, despite the practical relevance of BGs, it is still broadly true what Baker, Gibbons and Murphy (2002) wrote several years ago: "The economics literature has not had much to say about non-standard organizational forms [...] now much discussed in the business and organizational literatures, including [. . .] business groups".[2]

The aim of the present paper is to contribute to filling this gap in the economics literature by developing and testing a knowledge-based theory of BGs, in which this organizational form emerges as the optimal solution to a problem of knowledge creation, transmission and potential dissipation in multiteam production when contracts are incomplete.

In principle, a BG can emerge as the result of four types of decisions by a common source of hierarchical control (henceforth called 'headquarter' or simply HQ). First, there is the 'portfolio decision' on *which* activities the HQ would like to be performed. Second, there is the 'integration decision' on *who* should be put in charge of performing any given activity, that is whether the activity should be performed within the boundaries of the HQ (through divisions or branches) or outside the boundaries of the HQ's organization (through controlled subsidiaries or outsourcing).[3] Third, there is the 'hierarchy decision' on *how* the activities have to be structured, in particular which type of hierarchical structure should serve the purpose of carrying out them in a systematic way. Fourth and last, there is the 'location decision' on *where* the various activities should be performed within and beyond national borders. While all four decisions together constitute the overall organizational decision of the HQ, it is the integration and hierarchy decisions that define the peculiar traits of BGs. These defining decisions are the focus of the present paper.

Both the integration and hierarchy decisions have received extensive attention in the literature, but never before jointly or even separately in the specific case of BGs. For istance, Belenzon

[1] In the case of France around 65% of aggregate imports or exports can be attributed to firms that belong to BGs (Altomonte et al., 2013).

[2] BGs have been studied by the business literature, although often confined to emerging countries (Colpan and Hikino, 2010; Khanna and Yafeh). Williamson (1975) hints at the view of BGs as organizational forms located between markets and hierarchies. A number of studies point out the difficulty of classifying these network-like organizational forms (Powell, 1990; Granovetter, 1995; Hennart, 1993).

[3] Divisions are organizational units defined within the premises of the HQ. Branches are production units physically separated by the HQ but not legally independent from it. Subsidiaries are legally independent companies controlled by the HQ. Outsourcing takes place through independent suppliers, thus outside the boundaries of the Business Group.

et al. (2013) highlight the presence of internal capital markets as a crucial advantage of BGs. Bertrand et al. (2002) relate the creation of BGs to the 'tunneling' of profits from subsidiaries to HQs. Lewellen and Robinson (2013) stress the tax arbitrage motive behind the emergence of BGs. Almeida and Wolfenzon (2006) study the pyramidal structures of BGs through the lenses of separation between cash flow and voting rights. Differently from but complementarily to all these approaches, we conjecture instead that an additional reason behind the creation of BGs is the efficient management of HQ-specific knowledge, created and transmitted within the boundaries of the group in order to protect it from possible dissipation when contractual incompleteness undermines intellectual property rights (IPR).

Within the conceptual framework we propose, specific knowledge is communicated to, enriched by and embedded in the human resources the HQ relies on for its activities. Specific knowledge can be used repeatedly as an input in production. However, in doing so, the HQ has to deal with two crucial issues. The first issue is that human resources face a time constraint that limits how often their embedded knowledge can be used. The time constraint can be relaxed when human resources are organized as a hierarchical collection of teams according to their knowledgeability. This arrangement allows less knowledgeable teams to perform activities with lower knowledge intensity under the supervision of more knowledgeable teams, which in parallel can specialize in more knowledge intensive activities. In this perspective, the *hierarchy decision* solves the problem of how to use knowledge efficiently and how to communicate it among human resources so as to minimize the cost of using it as a production input (Garicano and Rossi-Hansberg, 2015).

Specifically, we assume that each activity is performed by a team consisting of a manager and an endogenous number of workers. We abstract from the internal hierarchy of individuals within teams, an issue already studied by Garicano (2000) and Garicano and Rossi-Hansberg (2006) among others. We focus, instead, on the external hierarchy of teams considered as separate production units within the BG. We study a situation in which the HQ already knows the production possibilities (i.e. owns the 'blueprints') of a given set of products and acts as a monopolist for each of them in a perfectly integrated economy in partial equilibrium as in Garicano (2000). The fact that the portfolio and the location decisions are both immaterial in this setup allows us to emphasize the integration and hierarchy decisions.

Each team is involved in the supply of one and only one product. All products face identical CES demand functions as well as identical production possibilities. However, before a team can turn the production possibilities of its product into actual production, its manager has to solve a problem. The problem comes in versions of different level of difficulty. If solved, more difficult versions allow for more efficient production. Their solution, however, requires more knowledgeable managers, whose hiring is more expensive. Problem solving also requires the HQ's supervision, which can be direct or indirect through a managerial hierarchy. Indirect supervision arises because helping a manager absorbs time due to communication costs and the HQ has only a limited amount of time available. It is feasible because more knowledgeable managers not only can solve more difficult versions of the problem, but they can also help less

knowledgeable managers solve less difficult versions of the problem. Nonetheless, as the HQ, each manager has only a limited amount of time to devote to supervision. Given all these time constraints, the optimal way to deal with problem solving when direct supervision is not viable is through a knowledge-based hierarchy of teams such that more (less) knowledgeable managers solve harder (easier) versions of the problem in teams assigned to higher (lower) hierarchical layers.

The second crucial issue the HQ has to deal with is that its specific knowledge is only imperfectly contractible and thus dissipable were human resources to walk away. This issue of potential knowledge dissipation is what the *integration decision* has to deal. To model such decision we rely on the property rights approach to the theory of the firm (Antràs and Rossi-Hansberg, 2009), giving the HQ a choice between integrating a team inside its boundaries without any legal autonomy ('division') and having alternatively the team as a legally autonomous entity under its control ('subsidiary'). The latter arrangement defines the distinctive trait of a BG.[4]

Potential knowledge dissipation arises from the fact that, when the HQ assigns a team with supplying a product, it has to reveal the corresponding production possibilities, i.e. the problem that need to be tackled in order to start production. For simplicity, we assume that the contract between the HQ and a division is complete: both the quality of the product supplied and the knowledge revealed are contractible. Differently, contractual incompleteness mars the specific relationship between the HQ and a subsidiary as neither the quality of the product supplied nor the knowledge revealed are contractible in this case. This contractual situation therefore compounds the standard 'quality holdup' with a novel 'knowledge holdup'.

We also assume that, whereas a product of sub-par quality is worthless both inside and outside the relationship between the subsidiary and the HQ, this is not the case for the revealed production possibilities as these have positive value for the subsidiary even outside the specific relationship with the HQ. The underlying idea is that, due to different allocation of residual property rights, the value of the outside option and thus the risk of knowledge dissipation are lower (and, for simplicity, actually inexistent in our setup) when the team is a division with no legal autonomy rather than a legally autonomous subsidiary. Under these circumstances, a trade-off for the HQ arises from the fact that, thanks to its more valuable outside option, a subsidiary has a stronger incentive than a division to invest in the specific relationship with the HQ, but at the same time the more valuable outside option also implies that the HQ can extract less rent from the relationship. In other words, the surplus the HQ can generate through a subsidiary is larger than through a division, but the share of surplus the HQ can extract is smaller. To enrich the analysis, we further allow subsidiaries to be more productive than divisions for given investment in the relationship with the HQ once the production possibilities

[4]A third canonical option for the HQ would be to keep the team as an independent agent from which to source ('outsourcing'). However, the tradeoff between in-house and arm's length sourcing has been investigated in depth by the existing literature. For parsimony we thus prefer to abstract from the implied 'make or buy' decision by assuming that circumstances are such that the HQ has already ruled out outsourcing as a profit maximizing option. This allows us to shift the focus of our analysis from the standard 'make or buy' decision to the non-standard 'how to make' decision.

have been revealed to them. The assumption here is that the HQ has a comparative advantage in defining the production possibilities whereas outside teams have a comparative advantage in actual production thanks to specialization.

To summarize, in our model an *integration decision* determines whether teams are organized as divisions or subsidiaries: the decision strikes the optimal balance between rent extraction with knowledge protection on the one side, and team incentivization (plus comparative advantage) with knowledge dissipation on the other. A BG emerges when the integration decision leads to teams being run as subsidiaries, which the model predicts to be more likely to happen in institutional environments that better protect IPR. How subsidiaries are arranged in hierarchical order is then determined by a *hierarchy decision* that optimally solves the problem of how to use, communicate and enrich HQ-specific knowledge efficiently. In this respect, the model predicts that a BG is more likely to arrange its subsidiaries on a larger number of layers when the HQ is endowed with a larger amount of intangibles, problem solving is more challenging, communication costs between layers are lower, and the skill premium for more knowledgeable managers is smaller.

We test these predictions exploiting the unique features of an original dataset we constructed from Orbis ownership data provided by Bureau Van Dijk. The dataset covers 178,190 parent (HQ) companies incorporated in OECD countries and controlling more than 1,150,000 (domestic and foreign) subsidiaries worldwide in year 2010.[5] It includes information on the ownership structures of BGs in terms of number of subsidiaries, countries and industries in which subsidiaries operate, and the subsidiaries' positions at different hierarchical layers. It also covers around 4.5 million independent firms not part of BGs in 2010 and before, which we use as control group.

In line with our theory, the empirical test of the model confirms that firms in a given industry and of a given size/age are more likely to set up subsidiaries when operating in countries with better IPR protection. This finding is robust to the inclusion of other country-specific variables, such as the level of financial development, the past level of income and growth, and the general quality of institutions. We also find evidence consistent with the idea that BGs are more likely to be structured on a larger number of hierarchical layers when communication between parents and subsidiaries is easier, when subsidiaries face less standardized assignments, and when the skill premium for better managers is lower. These results are robust to the inclusion of additional controls at the group level, including locational characteristics of the countries in which the BG operates (such as the tax level, number of patent per inhabitants, financial development, quality of the business environment) that might affect the hierarchical structure of the BG in addition to (or in correlation with) the explanatory variables we target.

The rest of the paper is structured as follows. Section 2 characterizes the BGs in our data,

[5] Orbis data have already been used in the literature to study BGs in terms of innovation (Belenzon et al. 2010), the international transmission of shocks (Cravino and Levchenko, 2017), or the effect of managerial culture on firm boundaries (Gorodnichenko et al., 2017). Other studies, related to the global reach of international groups (e.g. Alfaro et al, 2009 and 2017), have relied instead on data sourced from Dun & Bradstreet (D&B), which is one of the different sources integrated in the Orbis Ownership database.

also highlighting some stylized facts on their integration and hierarchy patterns. Section 3 presents our knowledge-based theory of BGs. The main empirical implications of the theory are first discussed in Section 4 and then tested in Section 5. Section 6 concludes.

2 Characterizing Business Groups

2.1 Business Groups as Hierarchies of Firms

We define a Business Group (BG) as a collection of at least two legally autonomous firms that function as a single economic entity through a common source of hierarchical control. Control is exerted by a headquarter/parent company on (one or more) subsidiaries via equity stakes. Under this definition, multinational enterprises (MNEs) can be considered as a subset of BGs, since they have at least one legally autonomous subsidiary located abroad, ultimately controlled by a parent located in the origin country. In the case of economic entities having more than one productive plant (multi-plant firms), or organized internally through multiple divisions, if all plants/divisions are commanded by the same firm under a single legal status we consider them as branches of that firm, i.e. they are not a BG according to our definition.[6]

Figure 1: A business group as a hierarchy of firms

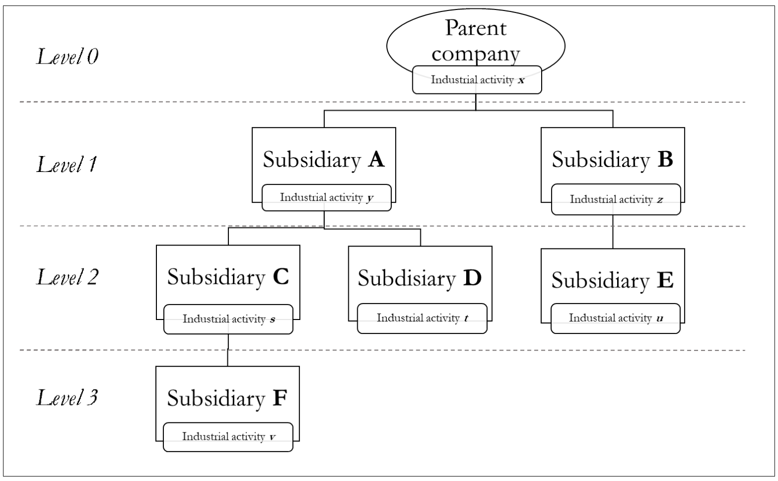

In Figure 1, we represent the structure of a BG represented as a hierarchical graph, with the parent company/HQ located at the top layer 0 and its legally autonomous subsidiaries arranged at different lower layers in a hierarchy of control. As it can be seen, subsidiaries can be directly or indirectly controlled by the parent, and each of them can perform a different activity for the

[6] The notions of branches/divisions and subsidiaries/affiliates tend to overlap in some contexts. In this paper, in accordance with international standards (for example UNCTAD, 2009) we define a branch as a new productive location, division, department or office set up by a corporation and positioned within the original legal boundary of the company. As a result, our definition of BG rules out strategic business alliances, but it includes in principle joint ventures, since in this case corporate assets are owned (and controlled) by more than one proprietary firm.

parent's organization.[7]

To identify the boundaries of a BG, we rely on a notion of corporate control established in international accounting standards (OECD 2005; UNCTAD, 2009; Eurostat, 2007) when a parent has a command of direct or indirect majority (> 50%) of voting rights.[8] Such a notion of control is not exhaustive, as it leaves outside cases of affiliates that are *de facto* controlled through minority ownership (< 50%), or peculiar forms of control derived from some form of market advantage (e.g. a monopsony), as well as particular forms of government regulations (e.g. 'golden shares'). Yet, it has some clear advantages. First, the majority of voting rights is a unique standard for both domestic and multinational Business Groups. Second, it allows to rule out cases of double (or triple) accounting of affiliates among different groups, thus generating a definition of the boundaries of a BG which is univocal, i.e. in our data a Business Group is a closed set of firms. Third, such a definition of corporate control allows for a straightforward comparison with official statistics, as the majority of voting rights is the criterion commonly used in international standards on foreign subsidiaries (Eurostat or OECD FATS) and for international tax purposes (IAS, IFRS).[9] Appendix A provides further details on the relationship between ownership and corporate control.

To build our dataset of BGs, we have sourced worldwide proprietary linkages and firm-level financial accounts from the Orbis ownership database by Bureau van Dijk for the year 2010. Based on these linkages, we have identified those companies for which a parent company has a command of direct or indirect majority (> 50%) of voting rights: these are our subsidiaries (or majority-owned affiliates), and together with the parent they form a business group. Firm-level data of subsidiaries are then stratified according to their position on hierarchical levels in each BG. For each parent company and each subsidiary along the control chain, we also collect industry affiliations following the 6-digit NAICS rev. 2007 classification. We end up with a global dataset of 270,374 parent companies controlling a total of 1,519,588 (domestic or foreign) subsidiaries in 207 countries in the year 2010.

Table 1 describes the sample. Two thirds of our BGs' parent companies are headquartered in OECD economies, controlling around 75% of affiliates worldwide. Around 20% of the groups incorporated in OECD countries are multinational companies, i.e. they control at least one affiliate in a host country different from the parent country. The proportion of MNEs is only 14% when originating from developing countries. This is in line with previous findings (Khanna and Yafeh, 2007), because emerging countries have a relatively larger proportion of domestic firms organized as BGs. The vast majority of parent companies reports a primary activity in

[7] Technically, affiliates are defined as legally independent companies whose shares are partially owned by another corporation. If the share of control exceeds 50 per cent, the affiliate becomes a subsidiary. In other words, a majority-owned affiliate is a subsidiary.

[8] Corporate control can be derived by a direct, indirect or consolidated concentration of voting rights (Faccio and Lang, 2002; Chapelle and Szafarz, 2007; Del Prete and Rungi, 2017). Company H can control 60% of shares of company A, which controls 70% of shares of company B. Although company H does not formally control company B directly, it does indirectly, via company A. The latter is known as the principle of the Ultimate Controlling Institution in the OECD FATS Statistics (or Ultimate Beneficial Owner in UNCTAD data).

[9] We are not the first to use such a notion of control to identify the boundaries of a Business Group, e.g. Belenzon (2010; 2013).

services industries, especially in OECD countries. The share of primary industries for parents is slightly higher in developing economies. Clearly, a parent can be active in a service industry, e.g., it can be a holding company classified in the financial industry, but it can control subsidiaries that operate in manufacturing, primary or services industries.

Table 1: Sample distribution of Business Groups

	Total sample	OECD sample
N. of parents	270,374	178,190
of which:		
- Multinational HQ	49,897	36,314
HQ main activity:		
- Agric. & Mining	6,840	3,467
- Manufacturing	25,718	14,634
- Services	237,816	160,089
HQ controlling:		
Subsidiaries	1,519,588	1,154,138
Independent firms (control group)	-	4,160,047

To validate our dataset, we rely on UNCTAD (2011), where details on the numbers of parents and subsidiaries of MNEs are reported by country for the same year of our sample. The correlations of UNCTAD figures with our relevant sample of BGs are .94 and .93 when measured by country at the parent or subsidiary level, respectively (see Appendix A for more details).

At the bottom of Table 1 we also include information on independent firms that are neither controlled by parents nor they control subsidiaries in year 2010, which we use as a control group. The latter set of firms might suffer from a potential sample selection bias driven by the heterogeneity in the quality of data across countries. Indeed, we may assume that in a given subset of countries, especially developing countries, information is available only for the largest firms. In this case, our control group would under-represent the actual population of firms and, to the extent that parent companies are relatively larger, it might bias some of our results. To deal with this possible sample selection problem, we employ for the empirical analyses of the paper a subset of BGs originated in OECD countries, in which the representativeness of firm-level data is overall better. The latter implies working with 178,190 business groups in 32 countries (vs. the original 270,374 groups in 196 countries) and a control group of 4,167,873 independent firms, again retrieved from the ORBIS database in 2010.[10] For the parent firms located in OECD countries we retain all the information on their worldwide network of affiliates (around 1.15 million affiliates, as reported in Table 1). The issue of potential sample selection of the control group together with a number of robustness checks is addressed in the empirical section of the paper.

In our total sample a BG is composed on average of a parent controlling 5.6 subsidiaries, with a highly skewed distribution. The left panel of Figure 2 shows in fact that 57% of BGs

[10] We do not have enough information on BGs operating in 2010 out of Estonia, Latvia and Slovenia, three small-open economies now members of the OECD.

in our dataset consist of one parent and one subsidiary, while about 13% of BGs have more than five subsidiaries and only 0.7% of parents control more than 100 subsidiaries.[11] When we look at the hierarchy of control, BGs on average are organized on 1.3 layers of corporate control, with larger groups being in general more complex, as shown in the right panel of Figure 2: groups characterized by a higher number of hierarchical layers (measured on the horizontal axis) tend to have a larger number of subsidiaries.[12] There is however substantial heterogeneity also in the hierarchy of BGs, with some groups organized around no more than one or two hierarchical layers of control reporting a large number of subsidiaries (the dots in the left part of the figure), and some other groups organized in several layers of control reporting however only few subsidiaries (as shown by the bottom bars of the boxplots).

Figure 2: N. of affiliates and N. of layers across Business Groups

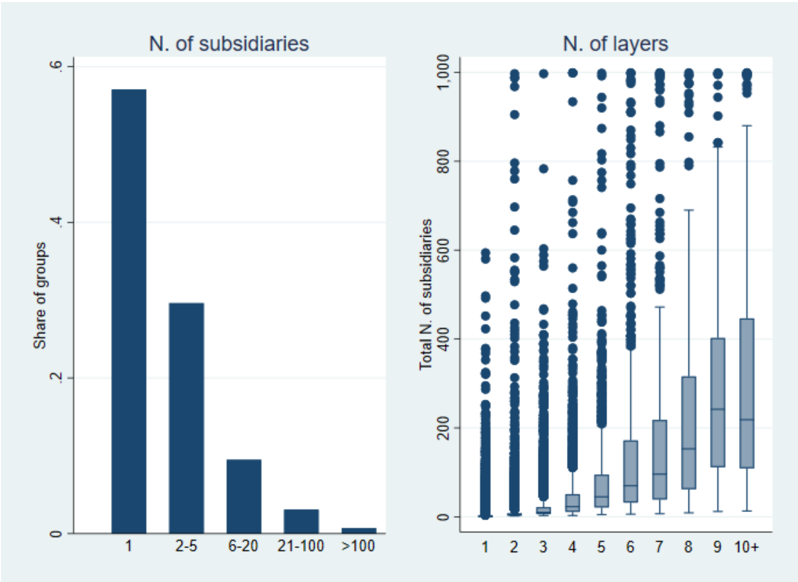

Table 1 presents another feature of the organization of our BGs, and namely the fact that the average number of affiliates per layer decreases with the number of layers. In other words, Business Groups seems to be organized as 'inverted' pyramids, with a first hierarchical layer of control relatively more dense in the number of firms, while lower layers are progressively less populated.

Restricting the analysis to BGs whose headquarter is located in OECD economies does not alter the key underlying distributions of our original sample: 55% of business groups whose parent is active in an OECD country have only one affiliate (vs. 57% in the overall sample reported in Figure 2), while 80% of parents organize affiliates on one layer of control (vs. 82% in the overall sample). In the total sample, 165 business groups have 10 (or more) layers of

[11]Looking at the 208,181 parents for which we could retrieve a complete set of balance sheet information, we have that 0.7% of BGs with more than 100 affiliates are responsible for more than 70% of value added recorded in our data.

[12]The horizontal bar within the boxplots represents the median of the distribution of subsidiaries for groups characterized by that given number of hierarchical layers. The boundaries of the boxplot are the first and third quartile of the same distribution. The lower and upper bars visualize respectively the 1st and 99th percentiles.

Table 2: Number of affiliates per layer across Business Groups

	BG with: 10 layers	7 layers	4 layers	3 layers	2 layers
1	62.6	64.8	19.5	11.1	5.8
2	51.8	41.6	14.0	7.4	2.5
3	42.7	34.0	8.5	2.8	
4	40.9	24.2	3.2		
5	30.8	15.0			
6	29.5	7.8			
7	23.9	3.0			
8	21.6				
9	15.7				
10	12.6				
N. of BGs	165	347	3,068	8,697	32,823

Note: Average number of affiliates per layer, for different sized BGs

hierarchical control, while in the OECD sample 146 groups display the same characteristic.

2.2 Stylized facts on Business Groups

In order to start making sense of the heterogeneous set of organizational choices of our BGs, namely their integration and hierarchy patterns, some recent results in the literature on the property rights approach to the theory of the firm, and on knowledge-based hierarchies internal to the firm's organization, can be linked to a number of stylized features of our data on Business Groups.

Looking at the integration decision of a BG, Biancini and Bombarda (2017) exploit the U.S. Related-Party Trade database and find empirical evidence supportive of a positive link between the level of Intellectual Property Right (IPR) protection and the relative share of imports from the foreign affiliates of US multinational groups. Their key intuition is that the risk of technology expropriation can occur also when technology transfers happen within firm's boundaries, e.g. due to former licensees and employees infringing the related trademarks and patents. As a result, an increase in IPR restores the advantages of vertical integration in the bargaining problem, reinforcing the HQ's control over both physical capital and intangible assets. Eppinger and Kukharskyy (2017) look at the ownership share of some half a million HQ-affiliate pairs from more than one hundred countries, retrieved again from the Orbis ownership database by Bureau van Dijk. They study the traditional trade-off of the property rights theory of the firm in which, by integrating the producer, the headquarter obtains residual control rights over non-contractible inputs but undermines the producer's incentives to invest into these inputs. They show that a well-functioning contract enforcement may effectively substitute for the need to incentivize the producer, and thus derive that the optimal degree of integration (the share of equity control of the affiliate by the parent) should be higher in countries with better contracting institutions.

Figure 3 below provides evidence suggestive of a similar pattern in our BG data. The

left hand panel reports the partial scatterplot relationship between the likelihood of observing a business group at the home country level, on the vertical axis, and a normalized index of IPR protection measured through the standard Park (2008) index at the country level.[13] The right hand panel correlates the same likelihood of observing a BG with the normalized average quality of institutions (the 'rule of law' index from the Worldwide Governance Indicators, as in Kaufmann et al., 2010) in the countries in which the group operates.[14] Consistently with the findings of the previously quoted literature, also in our data the integration decision is more likely to be observed for business groups that operate across jurisdictions characterized by a better protection of IPR or higher institutional quality.

Figure 3: Stylized facts - Integration decision

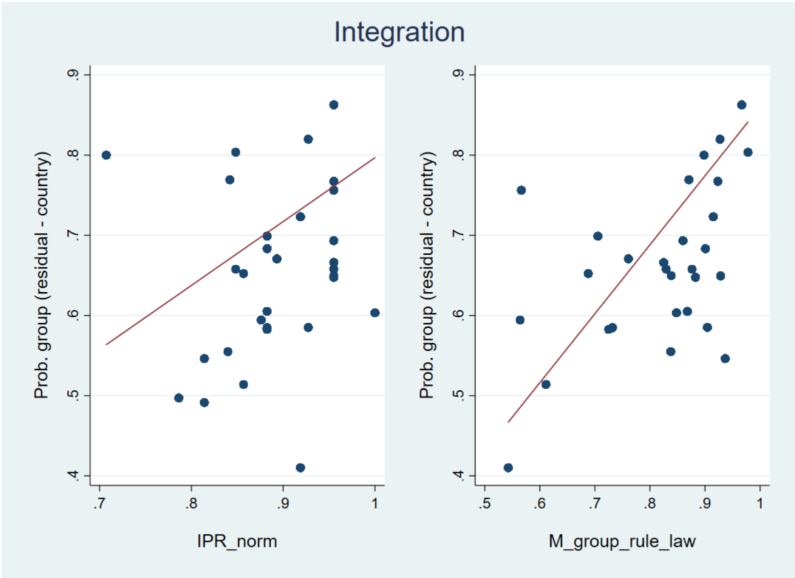

Turning to the hierarchy decision, the literature on knowledge-based hierarchies (Garicano and Rossi-Hansberg, 2004; Caliendo and Rossi-Hansberg, 2012; Caliendo et al., 2015) has shown both theoretically and empirically that, within the (internal) hierarchy of a firm's employees organized in teams of production, easier communication allows for a higher number of layers, with lower hierarchical layers characterized by relatively more standardized, lower value-added activities. We bring these insights to our data, in which the (external) hierarchy of teams within the BG is implicitely derived by observing the position of each subsidiary in the control chain set by the parent, as shown in Figure 1. If the logic of organizing the control hierarchy is similar to the one driving the internal organization of the firm, we should observe that business

[13] The vertical axis in the left panel is the (normalized) residual of an unweighted OLS regression of the (log) number of parent companies in each home country regressed on each country's (log) GDP, population and industry FE. Poschke (2017) shows how features of the firm size distribution tend to be correlated with countries' income, with richer countries having fewer entrepreneurs and fewer small firms.

The IPR protection index is calculated in periods of 5 years for 122 countries and updated to 2005 by Park (2008). Its original version ranges between 0 and 5, with 5 being the highest level of IPR protection.

[14] Correlating the likelihood of observing a BG with the quality of institutions measured only in the same home country in which the group parent is incorporated (vs. the average of the home and host countries in which the group operates) yields the same result, i.e. a positive and significant correlation of 0.35 in both cases.

Figure 4: Stylized facts - Hierarchy decision

groups characterized by easier communication should have a larger number of hierarchies, and that subsidiaries placed on lower hierarchical layers of control should be active in industries characterized by a relatively higher degree of standardization.

In this respect, the left panel of Figure 4 reports the partial scatterplot relationship between the ease of communication across subsidiaries and parent within each group and the average number of layers characterizing each BG. Ease of communication is proxied through the similarity of the activities pertaining to the subsidiary and the parent (the share of equal NAICS-2 digit industries within each group), the prior being that if both subsidiaries and parent operate in exactly the same industries, communication costs should be at a minimum. On the vertical axis of the left panel we report equally spaced class intervals (bins) of the communication variable. On the horizontal axis we report the average number of hierarchical layers of the business groups operating within each communication bin. The right panel of Figure 4 plots the average level of standardization of subsidiaries operating in a given hierarchical layer across our business groups, on the vertical axis, vs. the position of the same affiliates in the hierarchical layer in which they are placed (horizontal axis). Standardization is measured through a (normalized) industry-level index capturing the relative standardization of tasks of the industries in which the HQ is operating: the index is originally sourced by the Current Population Survey in the US and discussed by Blinder and Krueger (2013). The preliminary evidence in our data hints at a relation between the span of the hierarchy of subsidiaries, communication costs, and the standardization of tasks, a pattern similar to the drivers of the internal organization of teams witin the firm identified by the literature.

In the next section we will build on these stylized facts to propose a unified theoretical framework in which BGs can emerge as the profit-maximizing organizational form when the use and the transmission of firm-specific knowledge absorbs scarce resources and faces incomplete contracts.

3 A Knowledge-Based Theory of Business Groups

A parent company (henceforth, 'headquarter' or simply HQ) owns the 'blueprints' of a large number of M differentiated products and knows their production possibilities. To turn the production possibilities of a product into actual production, a problem has to be solved. The HQ knows how to solve the problem but this requires time which the HQ has only in limited amount. The problem comes in different versions, indexed $\wp = 1, ..., P$ in decreasing order of difficulty, and the HQ decides which version to tackle. Solving a more difficult version of the problem is harder but allows for production at lower marginal cost than solving an easier version.[15] The HQ has to decide how many products to make and how to organize their production.

Each product faces isoelastic demand $y = Ap^{-\sigma}$ where $\sigma > 1$ is the demand elasticity and $A > 0$ is a demand shifter, and is produced by a team consisting of a 'manager' and an endogenous number of 'workers' hired by the manager. Workers wage is taken as the unit of value ('numeraire'). The manager receives the blueprint from the HQ and is in charge of solving the problem associated with the supply of the assigned product. If he solves the problem, workers produce. If he does not solve the problem, the unit's productivity is zero. If the manager solves a version $\wp$ of the problem, her unit can produce at productivity $\theta_\wp = e^{-\theta\wp}$ with $\theta \geq 0$, i.e. solving more difficult versions of the problem (lower $\wp$) entails higher productivity (lower marginal costs).

Solving more difficult versions requires, however, higher managerial skills that not all managers have. In particular, there are different types of managers that the HQ can hire, indexed $\wp = 1, ..., P$ in decreasing order of ability, so that $\wp$ refers indifferently to the difficulty of the problem's version and to the ability of the managers who can solve it. Each manager has only a limited amount of time in which she can solve at most one problem. The ability differential of managers is reflected in different hiring costs, with $w\theta_\wp$ denoting the fixed hiring costs of a manager whose problem-solving ability allows her team to achieve productivity $\theta_\wp$. At $w\theta_\wp$ the supply of managers of ability $\wp$ is infinitely elastic. Contracts signed by the HQ with managers and workers are incomplete.

On top of adequate skills, solving a given version $\wp$ of a problem by the manager also requires the supervision of the HQ in implementing the blueprint. Supervision by the HQ can be direct (one degree of hierarchical separation between the HQ and the production unit in which the manager operates), or indirect (more than one degree of hierarchical separation between the HQ and the manager's unit). In the latter case, supervision can be performed through managers of higher ability, i.e. at least equal to $\wp - 1$.[16]. Supervision entails a communication time cost for the supervisor, captured by a parameter $\varphi_\wp$, with $\varphi_\wp = e^{\varphi\wp}$ and $\varphi > 0$, i.e. the higher the skill of the supervised manager (lower $\wp$), the lower the communication costs associated with its supervision.

[15] One could think of the different versions of the problem as characteristics of the production process. More complex production processes are harder to design, but allow for production at lower marginal cost.

[16] In other words, a manager of ability $\wp$ cannot be supervised by her peers or by lower ability managers.

All managers as well as the HQ have only a limited amount of time they can devote to supervision. This amount is the same for the HQ and the managers, and equals $\tau_{\wp} = e^{\tau}$ with $\tau > 0$ for all $\wp$. On the other hand, in order to solve a problem version characterized by productivity $\theta_{\wp}$, a manager of ability $\wp$ requires $\varphi_{\wp}\theta_{\wp}$ units of supervision time, directly from the HQ or indirectly through managers of higher ability $\wp-1$. Hence there is a trade-off between supervising several lower ability managers (with high $\varphi_{\wp}$) solving easier problem versions (with low $\theta_{\wp}$) and few higher ability managers (with low $\varphi_{\wp}$) solving more difficult problem versions (with high $\theta_{\wp}$). The trade-off is captured by the following recursive time constraint requiring that, for each problem-solving ability type $\wp$, the amount of available supervision time matches its required amount:

$$\tau_{\wp-1}n_{\wp-1} = \varphi_{\wp}\theta_{\wp}n_{\wp} \tag{1}$$

where $n_{\wp}$ is the number of managers of ability $\wp$ hired by the HQ, with $\wp = 0$ referring to the HQ, in which case $n_0 = 1$ holds. Solving the recursion under the assumed functional forms for $\tau_{\wp}, \varphi_{\wp}$ and $\theta_{\wp}$ yields

$$n_{\wp} = \prod_{s=1}^{\wp} \frac{\tau_{s-1}}{\varphi_s \theta_s} = e^{\tau + \sum_{s=1}^{\wp}(\theta-\varphi)s} = e^{\tau + \frac{1}{2}\wp(\wp+1)(\theta-\varphi)}. \tag{2}$$

According to (2), $n_{\wp}$ is increasing (decreasing) in $\wp$ for $\theta > (<)\varphi$. This constraint on supervision time will crucially affects the HQ's decision on how to use knowledge efficiently and how to communicate it between teams so as to optimize its use as a production input.

3.1 Integration

The *organizational choice* of the HQ consists of two choices: one on 'integration' and the other on 'hierarchy'. The *integration decision* is on whether the supply of each product should be assigned to a production unit under full control by the HQ (a 'division', d), or to a production unit with operative and legal autonomy (a generic 'affiliate', f) still operating within the boundaries of the HQ's organization, that is, under the HQ's supervision.[17]

This choice is driven by a trade-off between gains from specialization and losses from holdup. Specifically, the supply of each product requires an intermediate input, which is produced by workers after their manager has solved the relevant assigned problem. While a unit of intermediate input produces a unit of final output within the HQ's organization, the number of units of labor needed to produce a unit of that input depend on whether it is supplied by a division or an affiliate. In particular, a division of the HQ is less efficient in production than an affiliate. We use $\delta/\theta_{\wp}$ and $\kappa/\theta_{\wp}$ with $\kappa < \delta$ to denote the marginal costs associated with the solution of problem version $\wp$ of a problem in the cases of division and affiliate respectively. Affiliates thus have a comparative advantage in production with respect to divisions.

[17] As the choice between in-house and arm's length production (outsourcing) has been investigated at length in the existing literature, we abstract from this choice by assuming that the HQ has already decided in favor of in-house production, hence shifting the focus from the 'make or buy' decision to the 'how to make' decision.

Alternative ownership arragements, however, also differ in terms of their exposure to *holdup* problems. In particular, when the HQ assigns a production unit with supplying a product, it has to reveal the corresponding blueprint to the manager, together with the versions of a problem that need to be solved before the product can be actually made. The holdup arises for two reasons. The first is related to the fact that the *quality* of the intermediate input is not contractible and an intermediate input of the wrong quality is useless for final production ('quality holdup'). The second reason for holdup is that also the description of the *production possibilities* and the associated problem are not contractible, and it may have a value for the manager outside the specific relation with the HQ ('knowledge holdup'). Specifically, having access to the details on the production possibilities is more valuable for an affiliate manager operating outside the specific relation with the HQ, due to her farther reaching legal liability with respect to a division manager. For simplicity, we thus assume no holdup in the case of a division, as this is fully under the legal and functional control of the HQ. If the affiliate starts independent production, it is however less efficient than under the (direct or indirect) supervision of the HQ, which we capture by assuming that $\rho > 1$ units of intermediate input are required to produce a unit of final output when the affiliate manager operates outside the specific relation with the HQ. We interpret this parameter as a measure of the protection of intellectual property rights on 'blueprints'.[18]

The timing of events is as follows. First, the HQ meets the manager and reveals the production possibilities of the product together with the problem to be solved before production can take place; the manager solves the problem's version that matches her ability. Second, the manager decides whether to walk away or not with the revealed pieces of information. Third, if the manager does not walk away, she hires the workers needed to produce the intermediate input, which is later transformed into final output in the affiliate under the (direct or indirect) supervision of the HQ. Fourth, after the intermediate input has been produced, the HQ and the manager bargain on the joint surplus from final production. We assume Nash bargaining over this surplus with the HQ's and the manager's bargaining weights equal to $1 - \omega$ and ω respectively with $\omega \in (0, 1)$. Importantly, as we identify BGs in our data through a notion of control that considers (directly or indirectly) majority-owned affiliates (i.e. subsidiaries) of parent companies, bargaining weights can be assumed exogenous to the specific ownership share of the parent, as long as the latter is >50%.

After the intermediate input has been produced the two parties have zero outside options: the input has no value for the manager outside his relation with the HQ, and it is too late for the HQ to find another supplier for its product. However, once the HQ has revealed its information and the manager has solved the problem, the latter may decide to walk away and exploit the acquired knowledge in order to start independent production, hiring workers for intermediate supply and

[18] Arguably both types of holdup characterizing an affiliate would be even more salient in the case of arm's length suppliers. At the same time, arm's length suppliers may have a stronger comparative advantage in production. In this paper, we concentrate on the integration choice of the HQ between division and affiliate, i.e. those cases in which the comparative advantage of arm's length suppliers is not strong enough to compensate their bigger holdup problem. The latter thus defines the boundaries of the HQ.

transforming the intermediate input into final output without the assistance of the HQ. Note that, once the HQ has revealed its information, it has no option outside its relation with the manager. Differently, while the manager has no outside option after the intermediate has been produced by the affiliate, she does have an outside option after being instructed by the HQ and before the intermediate has been supplied. This contractual situation therefore compounds the standard 'quality holdup' with a novel 'knowledge holdup': in the quality holdup a product of sub-par quality is worthless both inside and outside the relationship between the subsidiary and the HQ, but this is not the case for the knowledge holdup, as the revealed production possibilities have positive value for the subsidiary even outside the specific relationship with the HQ.

Consider now the different alternatives. When the supply of a product is assigned to an internal *division* of the HQ, there is no holdup. The division maximizes profit $\pi_d = p_d y_d - (\delta/\theta_\wp)\, x_d$, where x_d is the amount of the intermediate input produced, $y_d = x_d$ is the final output and p_d is its price. Maximized operating profit under our isoelastic demand is then given by

$$\pi_d(\theta_\wp) = \frac{A}{\sigma}\left(\frac{\sigma}{\sigma-1}\frac{\delta}{\theta_\wp}\right)^{1-\sigma} = \theta_\wp^{\sigma-1}\,\overline{\pi} \tag{3}$$

with $\overline{\pi} \equiv \frac{A}{\sigma}\left(\frac{\sigma}{\sigma-1}\delta\right)^{1-\sigma}$.

Differently, when the supply of a product is assigned to a (legally independent and majority-owned) *affiliate* of the HQ, there is a double holdup. In this case the Nash bargaining solution requires the maximization of the joint surplus of the HQ and the affiliate, taking into account their respective bargaining weights $(1-\omega)$ and ω as a solution to the standard quality holdup, as well as the positive value of the affiliate's outside options O_f associated with independent production (knowledge holdup). This requires the maximization of

$$(p_f y_f)^{1-\omega}(p_f y_f - O_f)^{\omega} \tag{4}$$

subject to the splitting constraint of joint surplus $\pi_f + s_f = p_f y_f$ where π_f and s_f are the surplus shares accruing to the HQ and the affiliate respectively. The FOC for the maximization of (4) entails $s_f = \omega p_f y_f + (1-\omega)\, O_f$. To determine O_f, we assume that independent production will happen through a division.[19] Hence, the outside option of the affiliate's manager equals the revenue from independent production $p_r y_r = A^{\frac{1}{\sigma}}\,(x_f/\rho)^{\frac{\sigma-1}{\sigma}}$ where $y_r = x_f/\rho$ is final output, p_r is its price, x_f is the amount of the intermediate input produced and ρ is the penalization incurred by the protection of the HQ property rights in case of autonomous production by the affiliate. Then the affiliate chooses x_f so as to maximize

$$v_f = s_f - \kappa x_f = \omega\,(p_f y_f - p_r y_r) + p_r y_r - (\kappa/\theta_\wp)\, x_f \tag{5}$$

with $p_f y_f = A^{\frac{1}{\sigma}}\,(x_f)^{\frac{\sigma-1}{\sigma}}$. That the marginal cost equals $\kappa/\theta_\wp$ is due to the fact that the version of the problem corresponds to where the manager is located in the hierarchy.

[19] We make this simple assumption in order to abstract from the integration decision of the runaway manager.

The corresponding profit maximizing amount of intermediate input is

$$x_f(\theta_\wp) = A\left\{\frac{\sigma-1}{\sigma}\frac{\theta_\wp}{\kappa}\left[\omega+(1-\omega)\rho^{\frac{1-\sigma}{\sigma}}\right]\right\}^\sigma \tag{6}$$

This is the optimal quantity of input chosen by the affiliate taking into account the quality holdup as well as the outside option associated with the knowledge holdup. With respect to quality, while producing with a comparative advantage in production reduces the marginal cost (due to $\kappa/\theta_\wp < \delta/\theta_\wp$), the holdup problem materializes in a reduction of the amount of intermediate input supplied for given marginal cost (due to $\omega < 1$). As for knowledge, when ρ goes to infinity, the outside option of the affiliate goes to zero and in the limit we have $x_f(n) = A\left(\frac{\sigma-1}{\sigma}\frac{\theta_\wp}{\kappa}\omega\right)^\sigma$, which is the amount of intermediate input supplied when there is quality holdup but no knowledge holdup. When ρ goes to one (lowest protection of HQ's property rights), independent production by the affiliate's manager faces no penalty. The HQ's knowledge of production possibilities is completely expropriated by the manager, whose team's intermediate input and final output supplied equal $A\left(\frac{\sigma-1}{\sigma}\frac{\theta_\wp}{\kappa}\right)^\sigma$, which is the quantity produced by a division with marginal cost $\kappa/\theta_\wp$. Expression (6) then shows that the smaller is ρ, the larger is the quantity of intermediate input supplied. In this respect, the knowledge holdup mitigates the underprovision of the intermediate input caused by the quality holdup. The implied operating profit for the affiliate then evaluates to

$$v_f(\theta_\wp) = \left[\omega+(1-\omega)\rho^{\frac{1-\sigma}{\sigma}}\right]^\sigma \left(\frac{\kappa}{\delta}\right)^{1-\sigma}\pi_d(\theta_\wp) \tag{7}$$

which converges to the profit of a putative division with marginal cost $\kappa/\theta_\wp$ as ρ goes to one.

Turning to the problem of the HQ, the FOC for the maximization of (4) entails that the HQ receives operating profits $\pi_f = (1-\omega)(p_f y_f - p_r y_r)$, which evaluates to

$$\pi_f(\theta_\wp) = \Omega\left(\frac{\kappa}{\delta}\right)^{1-\sigma}\theta_\wp^{\sigma-1}\,\overline{\pi} \text{ with } \Omega \equiv \sigma(1-\omega)\left(1-\rho^{\frac{1-\sigma}{\sigma}}\right)\left[\omega+(1-\omega)\rho^{\frac{1-\sigma}{\sigma}}\right]^{\sigma-1} \tag{8}$$

Then (8) shows that $\pi_f(\theta_\wp)$ deviates from $\pi_d(\theta_\wp)$, the operating profits of the HQ if she produces through a division, due to two factors linked to holdup (Ω) and comparative advantage $(\kappa/\delta)^{1-\sigma}$. Through the first factor, the HQ's operating profit is a hump-shaped function of ρ. As ρ goes to one, $\pi_f(\theta_\wp)$ tends to zero as the HQ's knowledge is fully dissipated. As ρ goes to infinity, the affiliate's outside option vanishes and $\pi_f(\theta_\wp)$ tends to the standard case in the literature with quality holdup only.[20] Between these two extremes two opposite effects are at work. On the one hand, larger ρ makes the affiliate supply a larger quantity of the intermediate input as its outside option gains strength. This is good for total surplus and thus the HQ's profit. On the other hand, by strengthening the affiliate's outside option, larger ρ reduces the HQ's bargaining power over total surplus, which is bad for the HQ's profit.

[20] This could be interpreted as the situation of a 'branch', i.e. a production unit physically separated by the HQ (and thus characterized by the standard quality holdup), but not legally independent, i.e. without a knowledge holdup.

We are now ready to determine when the HQ prefers to run teams as affiliates rather than divisions, thus giving rise to a BG. This is the case whenever $\pi_f(\theta_\wp) \geq \pi_d(\theta_\wp)$. Hence, given (3) and (8), a necessary and sufficient condition for a BG to be the profit maximizing integration decision by the HQ is

$$\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} \geq 1 \tag{9}$$

i.e. the efficiency gains from comparative advantage dominate the efficiency losses from holdup.

3.2 Hierarchy

The *hierarchy decision* relates to the HQ choice on how many products to produce and how to structure problem solving across the corresponding production units, given the supervision time constraint (2) and a fixed cost $F > 0$ of activating a hierarchical layer. The outcome of this choice can be characterized recursively going layer by layer from the top. For parsimony, we assume that (9) holds so as to focus on the case of affiliates. The case of divisions, arising when (9) is violated, can be studied analogously.[21] We also focus on a 'contiguous' hierarchy, such that managers of ability $\wp$ supervise managers of ability $\wp + 1$ and are supervised by managers of ability $\wp - 1$. We assume such contiguity is an equilibrium outcome and then characterize the conditions under which that is indeed the case.

First, at layer $\ell = 0$ there is only the HQ and no operating profit is generated at that layer. Second, as the HQ cannot produce without opening at least one affiliate, the minimum number of layers of an active hierarchy is two ($\ell = 0$ and $\ell = 1$). Third, given (8), the HQ profits generated by affiliates placed at layer $\ell = 1$ are an increasing function of the ability of managers (i.e. a decreasing function of $\wp$), which implies that the HQ has an incentive to appoint the managers with the highest ability (i.e. the lowest $\wp = 1$) at that layer. We thus have $\ell = \wp = 1$, with the HQ receiving profits from each affiliate equal to

$$\Pi_f(\theta_1) = \left[\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\theta(\sigma-2)}\, \overline{\pi} - w\right] e^{-\theta}$$

Fourth, due to the time constraint (2), the number of affilates that can be opened at layer $\ell = 1$ equals

$$n_1 = \frac{\tau_0 n_0}{\varphi_1 \theta_1} = e^{\tau + (\theta - \varphi)}.$$

Hence, the total profits received by the HQ from layer $\ell = 1$ evaluate to

$$\Pi_f(\theta_1) n_1 - F = \left[\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\theta(\sigma-2)}\, \overline{\pi} - w\right] e^{\tau - \varphi} - F$$

where $F > 0$ is the fixed costs of activating a layer. It then follows that layer $\ell = 1$ will be activated at all if and only if

$$\Pi_f(\theta_1) n_1 - F \geq 0$$

[21] In our model the ownership and hierarchy decisions are separable. A simple way to break separability would be to assume that communication costs are different for divisions and affiliates.

Consider now layer $\ell = 2$. Given (8), also profits generated by affiliates at layer $\ell = 2$ are an increasing function of managers' ability, which implies that the HQ appoints the managers with the highest feasible ability. This is $\wp = 2$ as managers of ability $\wp = 1$ potentially assigned to level $\ell = 2$ cannot be supervised by the managers of the same ability assigned to level $\ell = 1$. We thus have $\ell = \wp = 2$ with the profit of each affiliate equal to

$$\Pi_f(\theta_2) = \left[\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-2\theta(\sigma-2)}\, \overline{\pi} - w\right] e^{-2\theta}$$

Due to the time constraint (2), the number of affilates that will be opened at layer $\ell = 2$ equals

$$n_2 = e^{\tau + 3(\theta - \varphi)}$$

with total profit of the layer

$$\Pi_f(\theta_2) n_2 - F = \left[\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-2\theta(\sigma-2)}\, \overline{\pi} - w\right] e^{\tau + \theta - 3\varphi} - F$$

Layer $\ell = 2$ will thus be activated at all if and only if

$$\Pi_f(\theta_2) n_2 - F \geq 0$$

This constraint is more stringent than $\Pi_f(\theta_1)n_1 - F \geq 0$ as long as $\Pi_f(\theta_1)n_1 > \Pi_f(\theta_2)n_2$ holds. The latter condition is always verified if $\theta < \varphi$, as in this case $n_2 < n_1$ ('inverted pyramid') while $\Pi_f(\theta_\wp)$ is always decreasing in $\wp$. If instead $\theta > \varphi$, and thus $n_2 > n_1$ ('pyramid'), the condition holds for φ large enough.[22] As long as this restriction holds, a necessary condition for $\ell = 2$ to be worth activating is that $\ell = 1$ is itself worth activating. Vice versa, a sufficient condition for $\ell = 1$ to be worth activating is that $\ell = 2$ is itself worth activating. In other words, when the restriction holds, the hierarchy is contiguous.

These results obtained for $\ell = 1$ and $\ell = 2$ can be generalized by induction to the generic layer. For this generic layer we will have $\ell = \wp$, with affiliate profit

$$\Pi_f(\theta_\wp) = \left[\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\wp\theta(\sigma-2)}\, \overline{\pi} - w\right] e^{-\wp\theta}$$

as long as the hierarchy is contiguous, which is the case if and only if

$$\varphi > \varphi_c \equiv \frac{\wp}{\wp+1}\theta + \ln \frac{\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-(\wp+1)\theta(\sigma-2)}\, \overline{\pi} - w}{\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\wp\theta(\sigma-2)}\, \overline{\pi} - w} \tag{10}$$

[22] Specifically, in the case of pyramidal hierarchies (i.e. when $\theta > \varphi$) the condition always holds if $\varphi > \frac{\theta}{2} + \frac{1}{2} \ln \frac{\Omega\left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-2\theta(\sigma-2)}\, \overline{\pi} - w}{\Omega\left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\theta(\sigma-2)}\, \overline{\pi} - w}$.

Given (2), the corresponding number of affilates will be

$$n_{\wp} = e^{\tau + \frac{1}{2}\wp(\wp+1)(\theta-\varphi)}$$

with total profit

$$\Pi_f(\theta_{\wp})n_{\wp} - F = \left[\Omega\left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\wp\theta(\sigma-2)}\,\overline{\pi} - w\right] e^{\tau+\frac{1}{2}\wp(\wp+1)(\theta-\varphi)-\wp\theta} - F \tag{11}$$

The activation of the layer will happen at all if and only if

$$\Pi_f(\theta_{\wp})n_{\wp} - F \geq 0$$

Hence, the hierarchy stops at layer $\ell = \wp^*$ where $\wp^*$ is the largest integer $\wp$ compatible with

$$\left[\Omega\left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\wp\theta(\sigma-2)}\,\overline{\pi} - w\right] e^{\tau+\frac{1}{2}\wp(\wp+1)(\theta-\varphi)-\wp\theta} - F \geq 0$$

At that layer there are

$$n_{\wp^*} = e^{\tau+\frac{1}{2}\wp^*(\wp^*+1)(\theta-\varphi)}$$

affiliates. Note that we have a pyramid (inverted pyramid) hierarchical structure for $\theta > (<)\varphi$, as $n_{\wp}$ is increasing (decreasing) in $\wp$. Moreover an inverted pyramidal structure, a structure of hierarchy consistent with our stylized fact (see Table 1) is always contiguous: $\theta < \varphi$ implies $\varphi > \varphi_c$ as in (10) we have $0 < \varphi_c < \theta\wp/(\wp+1) < \theta$.

3.3 Organization

We can now combine the findings on the integration and hierarchy decisions to obtain the following:

Proposition 1 *A BG arises in equilibrium iff*

$$\Omega\left(\frac{\kappa}{\delta}\right)^{1-\sigma} > 1$$

When this condition holds together with (10): (i) the BG is organized as a hierarchy of $\wp^*$ layers of affiliates, where $\wp^*$ is the largest integer $\wp$ such that the profit of the BG at layer $\wp$ is larger than the fixed costs of activating the layer, i.e.

$$\left[\Omega\left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\wp\theta(\sigma-2)}\,\overline{\pi} - w\right] e^{\tau+\frac{1}{2}\wp(\wp+1)(\theta-\varphi)-\wp\theta} - F \geq 0$$

(ii) the total number of affiliates of the BG is

$$M^* = \sum_{\wp=1}^{\wp^*} n_{\wp} = \sum_{\wp=1}^{\wp^*} e^{\tau+\frac{1}{2}\wp(\wp+1)(\theta-\varphi)}$$

(iii) the number of affiliates assigned to layer $\wp$ of the BG is

$$n_\wp = e^{\tau + \frac{1}{2}\wp(\wp+1)(\theta-\varphi)}$$

(iv) the total profit of the BG is

$$\sum_{\wp=1}^{\wp^*} \left[\Omega \left(\frac{\kappa}{\delta}\right)^{1-\sigma} e^{-\wp\theta(\sigma-2)}\ \overline{\pi} - w \right] e^{\tau + \frac{1}{2}\wp(\wp+1)(\theta-\varphi)-\wp\theta} - F\wp^*$$

In terms of comparative statics, we thus have that, controlling for the parameters that drive the integration condition $(\Omega\,(\kappa/\delta)^{1-\sigma}\,\overline{\pi})$, higher communication costs (higher φ) are associated on average with a smaller number of group affiliates M^* and lower profits per layer, and thus a lower likelihood of more hierarchical layers to be activated (as higher $\wp$ are associated with relatively lower profits). Groups characterized on average by simpler production processes that, *ceteris paribus*, yield a lower productivity $\theta_\wp$ should in turn display a lower number of subsidiaries, and shorter hierarchies. Analogously, an increase in the cost of managers (higher w), higher fixed costs per layer (higher F) or a reduction of the time available for supervision (lower τ) will all be associated to shorter hierarchies. We now turn to empirically test some of these predictions.

4 From Theory to Empirics

We first test the integration condition (9) leading to the choice of a parent company to create majority-owned affiliates ('subsidiaries'), thus becoming the HQ of a BG. To that extent, we exploit our dataset on BG together with the control group of independent firms. Once this is checked in the data, we then test on the restricted sample of BGs the drivers of the hierarchy choice as predicted by our theoretical Proposition 1, controlling for the covariates that affect the integration decision. The general setup of our empirical specification will thus be of the form:

$$group_{isj} = \alpha + \beta\ \rho_j + \Gamma\ X_{i,s,j} + \varepsilon_{isj} \tag{12}$$

$$H_{g(i)sj} = \alpha + \Lambda\ Z_g + \Gamma\ X_{i,s,j} + \Phi\ Y_{g(j)} + \varepsilon_{gsj} \tag{13}$$

Specifically, in the *integration equation* (12) we evaluate the probability that a firm i operating in an industry s in a country j has to establish a group by creating subsidiaries, a decision based on the set of covariates dictated by our integration condition. The dependent variable in regression (12) is thus a dummy $group_{isj} = 1$, indicating whether a firm formed a business group or not before 2010. We employ in this specification the sample of 178,190 business groups operating in 32 OECD countries vs. the control group of 4,167,873 independent firms for which $group_{isj} = 0$, also retrieved from the ORBIS database in 2010.

In terms of covariates, Equation (9) in our model postulates that BGs are more likely to arise when the efficiency gains in organizing production through subsidiaries vs. internal divisions

dominate the efficiency losses from holdup. The latter are captured in our theoretical model by our parameter $\rho > 1$, which measurs the units of intermediate inputs required to produce a unit of final output when the team manager operates outside the specific relation with the HQ. As ρ goes to one, the HQ's operating profit tends to zero as the HQ's knowledge is fully dissipated. As ρ goes to infinity, the subsidiary's outside option vanishes and profits tend to the standard case in the literature with quality holdup only (see equation (8)). As a result, consistenlty with our theoretical model and the stylized facts presented in Figure 3, a parent choosing to organize itself through subsidiaries is *ceteris paribus* more likely to emerge in environments in which the protection of intellectual property rights (IPR) held by the HQ is stronger.[23] We measure IPR protection at the country level ρ_j using the normalized measure of the standard Park (2008) index employed in our stylized facts.

We remain agnostic about the nature of the efficiency gains driving the integration choice, and assume that a firm chooses to integrate its activities through subsidiaries vs. internal divisions based on some specific characteristics of the industry in which she produces her main output. We thus incorporate in the estimation a full set of industry fixed effects X_s measured at the three-digit level of aggregation (NAICS), as well as a country-industry varying measure of capital intensity X_{sj} retrieved from the WIOD database (the standard measure of gross fixed assets to sales, as e.g. in Acemoglu et al., 2009).[24] These covariates also help controlling for the overarching choice of outsourcing vs. integration.[25] We incorporate in the integration regression additional country-specific controls X_j, namely market size (population) and market potential (average GDP growth between 2005 and 2010), as well as the business environment (total tax rate, cost of opening a new business, number of patent per inhabitant), retrieved from World Bank Development Indicators. We also add to the estimating equation two controls X_i at the individual firm-level, namely firm age and firm size (small, medium, large, very large), an information reported in the Orbis database.

In the *hierarchy equation* (13) we restrict the sample to our set of business groups $g(i)$ having firm i as a parent, and test some of the group characteristics Z_g that might shape a BG's organization as identified by our theoretical model. To that extent we use as dependent variable $H_{g(i)sj}$ either the total number of subsidiaries in the group, or the number of hierarchical layers in which the business group is organized, i.e. a discrete variable ranging from 1 to 10 (the maximum number of layers recorded by the ownership database in ORBIS). In particular the available firm-level data allow us to test two main implications of our theoretical Proposition 1.

[23] As already discussed, this result is also in line with evidence in the literature, e.g. Biancini and Bombarda (2017) or Eppinger and Kukharskyy (2017).

[24] Out of the 32 OECD members included in our dataset, 26 are available in WIOD (Chile, Switzerland, Iceland, Israel, Norway and New Zealand are not included). For every country, data is available for 35 industries aggregated at the 2-digit level or higher. We employ country-specific World Input-Output tables for the year 2007 (i.e. before the credit crisis) to retrieve our variable at the country-industry level.

[25] Antràs and Chor (2013) or Antràs et al. (2017) model the integration vs. outsourcing strategy of parent companies based on the interaction of a series of industry-specific characteristics (demand elasticity, downstreamness, contractibility) linked to the I/O structure of the production process of a given firm. Antràs (2003) also presents evidence that the share of intra-firm U.S. imports on total U.S. imports is positively related to the capital intensity (and R&D intensity) of the industry.

The first implication links the total number of subsidiaries in a group to the communication costs between firms (our parameter φ in the model), and to the complexity of the problems solved, associated to a higher productivity ($\theta_{\wp}$). The prior is that lower communication costs and a higher complexity (lower standardization) of the problems solved should be associated to a larger number of subsidiaries. The second testable implication of the model relates to the number of hierarchical layers. By looking at the total profit expression (11) of the BG, we have that higher communication costs (higher φ) are associated on average with lower profits per layer, and thus a lower likelihood of more hierarchical layers to be activated, i.e. less vertical groups. The same negative effect on the number of layers is found for skill premia of good managers (higher w). The theory also posits that more standardized tasks should be performed by subsidiaries placed at lower hierarchical levels, as teams located there are predicted to solve easier versions of the production problem.

As in our stylized facts, communication costs are proxied through a variable measuring the similarity (the share of equal NAICS-2 digit industries) of the activities performed by subsidiaries and parent within each group, the prior being that, if both subsidiaries and parent operate in exactly the same industries (share equal to one), communication costs should be at a minimum. The ability to solve more difficult problems is measured, as already discussed, through the standardization index (Blinder and Krueger, 2013) of the main industry in which each firm is operating within the group. Managers'skill premia are proxied through the average share of population with tertiary education in the countries in which the group is operating, the prior being that the higher the share of tertiary education in a country, the higher the wage of managers.

In terms of model setup, the hierarchy equation identifies the effect of the theoretical covariates through the variation across groups. For two of our variables this is straightforward: groups that, on average, have higher communication costs given (a lower share of similar industries they are active in), and operate in countries with higher skill premia (higher share of tertiary education), should be relatively less vertically organized. We thus use in the regression the simple group-specific measures of these two covariates and identify through the across group variation.

The identification of the effect of standardization on BG's hierarchies is less straightforward. The theoretical model predicts a monotonic relationship between standardization of tasks and position in the hierarchy of layers, with more standardized tasks performed by subsidiaries placed at lower hierarchical levels. Within groups, the latter is indeed confirmed by one of our stylized facts in Figure 4, in which we plot the average standardization index of subsidiaries conditional on their position on a hierarchical layer. Our hierarchy equation (13) identifies instead the variation across groups, i.e. considering the average standardization of all the activities undertaken by a BG, irrespective of their position on the layers. To reduce the underlying heterogeneity, standardization will enter in our estimates as a binary variable, equal to one if the group has an average level of standardization higher than the sample mean, and zero otherwise. In other words, our prior is that BGs characterized by a relatively higher level of standardization (vs.

the sample mean) are more likely to be organized into relatively flatter hierarchies.

In our theoretical model the hierarchical organization of the BG is also conditional on the HQ-related characteristics $X_{i,j,s}$ that drive the integration choice, leading to possibly inconsistent estimates if these characteristics are unaccounted for. In addition to our main covariates, we thus always control in equation (13) for the age and size (small, medium, large, very large) of each parent i, the level of IPR protection in the country j in which the HQ is incorporated, as well as a full set of fixed effects for the industries s in which the parent is active.

Another critical element for the correct identification of the hierarchy choice is the possibility that, especially for multinational groups, the organization is partly driven by some local characteristics (such as financial opportunities or tax savings). For example, an European affiliate of a US group that, according to our covariates, should be placed relatively close to the parent in the hierarchy (e.g. because of higher communication costs, or low standardization of activities), might be placed under the European HQ of the same group for some regulatory reasons (e.g. taxes), thus ending up in a relatively lower layer of the organization with respect to our theoretical prior. To the extent that these local characteristics are correlated with our theory-based covariates, this could induce a spurious correlation between our postulated drivers and the hierarchical organization of the BG, invalidating our identification. We thus augment our hierarchy regression (13) with a set of controls $Y_{g(j)}$ related to the locational characteristics of the j countries across which the group operates. In particular, we include the average tax level faced by the BG, the average number of patent per inhabitants, the average level of financial development, and the average quality of the business environment (average number of days required to enforce a contract), an information retrieved across the host countries of each BG from the World Bank Development Indicators.[26]

Finally, when we test the hierarchy regression (13) using the number of layers as dependent variable, we also control for the total number of subsidiaries of the group, as the latter might mechanically influence our measure of hierarchy. In other words, we condition our equation to explain variation in the hierarchy choice within groups constituted by the same number of subsidiaries.

5 Empirical Results and Robustness

Table 3 presents the results of a linear probability estimation of our integration regression (12). As our dependent variable varies at the individual firm level, while our main covariate of interest varies at the country level, all the estimates employ standard errors clustered at the home country level. As predicted by our theoretical model, and consistently with our stylized facts, a better protection of IPR in the home country is associated to a higher probability of firms turning into business groups. Specifically, in our baseline estimate (Column 1) a ten percent higher (normalized) IPR index is associated to a 3 percent higher probability of a firm to create

[26] For domestic groups, these averages obviously coincide with the values of the home country in which the HQ is incorporated.

subsidiaries. The result is robust to (NAICS 3-digit) industry fixed effects potentially driving the integration decision, as well as the additional country and HQ-specific covariates previously discussed. In Column 2 we add country-industry covariates on capital intensity (gross fixed assets to sales, as retrieved from WIOD), a standard control in the literature on the integration decision of the firm (e.g. Antràs and Chor, 2013); in Column 3 we replicate our results restricting the sample only to parent companies originally operating in the manufacturing industry, as the latter is the typical sector analyzed in this literature. The IPR index always remain positive and significant.

In estimating our integration regression we have to consider the possibility that our initial sample of firms, even if restricted to OECD economies in which the coverage of firm-level data is generally good, might still be affected by some residual form of sample selection, with smaller firms under-represented in the data. The latter is a well-known coverage problem for other users of the same Orbis data (e.g. Kalemli-Ozcan et al. 2015), and it could be problematic for our purposes to the extent that smaller firms might be less likely to start a BG. In fact, if the sample selection is correlated to the quality of the home country institutions, the coefficient of our IPR variable in the integration equation will be biased.[27] For this reason, we have introduced a robustness check of the integration regression in the form of a two-step Heckman selection model. In the first stage we control for the probability that a smaller firm is found in our sample in relationship with GDP per capita, population and the country-level quality of institutions. We proxy the latter with the standard (normalized) 'rule of law' index from the Worldwide Governance Indicators (Kaufmann et al., 2010). In the second stage we augment the integration regression with the inverse mills ratios obtained from the first stage. In Column 4 we replicate the specification of Column 2 controlling for the selection equation in the first stage: results point to a higher and more significant IPR coefficient, indicating that sample selection might be stronger in countries with better IPR protection. Hence, if anything, our baseline specification is a conservative estimate of the actual effect of IPR protection on the probability of a parent creating a subsidiary. In Column 5 we replicate the selection model adding as a control the level of financial development in the home country, a variable identified in the literature as influencing both the rate of birth of business groups (Belenzon, 2013) and the vertical integration choices of firms (Acemoglu et al., 2009). Not surprisingly the coefficient of the IPR variable decreases by around 20 per cent when financial development is included, but remains positive and significant, and in any case higher than in our specifications without controls for the sample selection.[28]

All this evidence points to a significant and robust role of IPR protection in driving the integration choice of parents, thus confirming the intuition behind our theoretical condition (9) pointing at potential knowledge dissipation as a key driver for the emergence of a business group.

[27] The direction of the bias can go either way. If our control group would under-represent smaller firms in countries where the quality of the institutions is poorer, the integration equation may overestimate the impact of institutions. If however stricter IPR rules prevent commercial databases as Orbis to access information on smaller firms, the bias would go in the opposite direction.

[28] Adding also the home country rule of law index to the specification in the second stage does not change the magnitude and significance of the IPR coefficient.

Table 3: Integration decision

Dep Var	(1) Group	(2) Group	(3) Group	(4) Group	(5) Group
IPR	0.280**	0.309**	0.285**	0.441***	0.335**
	[0.143]	[0.140]	[0.115]	[0.157]	[0.140]
Capital Intensity		0.031**	-0.145	0.029*	0.027*
		[0.015]	[0.103]	[0.015]	[0.014]
Cost of startup	-0.000	-0.000	-0.001	0.001	0.002*
	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]
Total tax rate	-0.001***	-0.002***	-0.001**	-0.001*	-0.001*
	[0.000]	[0.000]	[0.001]	[0.001]	[0.000]
GDP Growth	-0.003	-0.003	-0.001	-0.001	-0.002
	[0.004]	[0.005]	[0.005]	[0.005]	[0.005]
Population	-0.007	-0.007	-0.001	0.007	0.006
	[0.006]	[0.007]	[0.006]	[0.009]	[0.008]
Patent / inhabitant	-0.006	-0.006	-0.011	-0.007	-0.005
	[0.007]	[0.008]	[0.008]	[0.008]	[0.007]
Age	0.003	0.004*	0.004**	0.004*	0.003*
	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]
Financial development					0.021***
					[0.008]
Inverse Mills Ratio				-0.143**	-0.163***
				[0.054]	[0.053]
Constant	-0.027	-0.042	-0.101	-0.393**	-0.335*
	[0.103]	[0.151]	[0.154]	[0.178]	[0.192]
Size HQ (dummy)	YES	YES	YES	YES	YES
Industry FE (3-digit)	YES	YES	YES	YES	YES
First-stage selection eq.	NO	NO	NO	YES	YES
Observations	4,273,058	3,737,847	458,154	3,737,847	3,736,256
R-squared	0.273	0.285	0.408	0.288	0.282

Standard errors clustered at the country level in brackets
*** p<0.01, ** p<0.05, * p<0.1

Looking at the hierarchy regression (13), Table 4 presents the baseline estimates of our two testable implications, on the total number of subsidiaries (Columns 1 and 2) and on the number of hierarchical layers in a BG (Columns 3 and 4), respectively. Column 1 shows that, in line with our priors, BGs characterized by lower communication costs and more complex (less standardized) activities tend to have a larger number of subsidiaries. The result holds controlling for the HQ-specific variables affecting the parent integration choice, and namely size and age of the parent, the level of IPR protection in the home country in which the HQ is incorporated, as well as a full set of NAICS 3-digit FE. In Column 2 we replicate the analysis on the number of subsidiaries including some locational characteristics of the countries in which the BG operates (tax level, number of patent per inhabitants, financial development, quality of the business environment), since these might affect the hierarchical organization of the BG on top of (or in correlation with) our covariates. Our main results on the number of subsidiaries do not change. In Columns 3 and 4 we repeat the exercise using an ordered probit specification, where the dependent variable is the number of hierarchical layers in each BG on our covariates. We find that BGs tend to be characterized by a significantly higher number of hierarchical layers when communication costs between parents and affiliates are lower, the skill premia for good managers are not too high, and more complex production processes (that is, a relatively lower

standardization of the industries across which the group operates) are present. As before, results are robust to the inclusion of HQ-specific variables as well as the set of locational characteristics at the group level, and they are entirely consistent with our theoretical model.[29]

Table 4: Hierarchy decision

	(1)	(2)	(3)	(4)
Dep Var	Ssubsidiaries	Subsidiaries	Layers	Layers
Estimation method	OLS	OLS	O. Probit	O. Probit
Ease of communication	0.084***	0.083***	0.151***	0.153***
	[0.009]	[0.009]	[0.018]	[0.018]
Standardization	-0.152***	-0.151***	-0.029**	-0.029**
	[0.009]	[0.008]	[0.013]	[0.013]
Skill premium			-0.097***	-0.088***
			[0.016]	[0.018]
N. of subsidiaries			1.143***	1.148***
			[0.010]	[0.010]
Fin. Development (group)		-0.001		0.022
		[0.008]		[0.015]
Tax rate (group)		0.011***		-0.004***
		[0.001]		[0.001]
Patent (group)		-0.089***		0.020*
		[0.008]		[0.011]
Contract enforcem.t (group)		-0.288***		0.112***
		[0.026]		[0.038]
IPR index (HQ)	2.505***	3.238***	0.738***	0.416**
	[0.111]	[0.143]	[0.180]	[0.205]
Age (HQ)	0.055***	0.049***	-0.123***	-0.120***
	[0.004]	[0.004]	[0.006]	[0.006]
Constant	-2.130***	-0.980***	-	-
	[0.113]	[0.201]	-	-
Size HQ (dummy)	YES	YES	YES	YES
Industry FE (3-digit)	YES	YES	YES	YES
Observations	59,269	59,269	59,269	59,269
R-squared	0.441	0.448	-	-

Robust standard errors in brackets
*** p<0.01, ** p<0.05, * p<0.1

Table 5 presents some robustness checks of our findings on the number of layers. In Column 1 we augment the regression with a variable measuring the vertical integration of activities within the group, which is a proxy for the revealed preferences of the parent in terms of integration choices. In particular, in the absence of actual data on internal shipments of intermediate goods and services across firms, Acemoglu et al. (2009) proposed to proxy the vertical integration of a firm exploiting the information on the set of industries in which a firm is engaged, combined with the input coefficient requirements that link those industries as retrieved from input-output tables.[30] We thus employ in our estimates a refined measure of their index, adapted to the

[29] When testing for the number of layers, we always control for the total number of subdisidiaries within a group to avoid picking up a mechanical correlation, as larger groups in terms of subsidiaries are more likely to be organized over a larger number of hierarchical layers. In other words, we explain the variation in the hierarchy choice of layers across groups constituted by the same number of subsidiaries.

[30] Alfaro et al. (2016) also use this index at the individual firm level on a global dataset.

input/output structure of Business Groups (see Appendix B for details). We also include a dummy capturing the multinational status of our BGs. Our main results do not change although, consistently with our story, the home country IPR protection index (the main covariate controlling in our regressions for the parent's integration choice) is not significant anymore when we include in the estimation the vertical integration measure at the group level. We also find, not surprisingly, that our multinational BGs tend to be structured *ceteris paribus* across more hierarchical layers than domestic groups. In Column 2, in line with one of our stylized facts, we add as a covariate the average rule of law index of the countries in which the group is operating, in order to capture the idea that multinational groups might have their controversies judged in jurisdictions different than the home country of the parent, and thus might not necessarily respond to the home country IPR protection index as a driver of their integration choices. In Column 3 we restrict our model to groups characterized by at least two subsidiaries, in order not to constrain the hierarchy choice to only one layer. In Column 4 we replicate the results dropping the proxy of group standardization, as this variable is available only on a subset of (manufacturing) industries, and thus induces a severe selection of our sample. All our main results do not change.

Table 5: Hierarchy decision - robustness checks

	(1)	(2)	(3)	(4)	(5)
Dep Var	Group	Group	Group	Group	Group
IPR	0.280**	0.309**	0.285**	0.441***	0.335**
	[0.143]	[0.140]	[0.115]	[0.157]	[0.140]
Capital Intensity		0.031**	-0.145	0.029*	0.027*
		[0.015]	[0.103]	[0.015]	[0.014]
Cost of startup	-0.000	-0.000	-0.001	0.001	0.002*
	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]
Total tax rate	-0.001***	-0.002***	-0.001**	-0.001*	-0.001*
	[0.000]	[0.000]	[0.001]	[0.001]	[0.000]
GDP Growth	-0.003	-0.003	-0.001	-0.001	-0.002
	[0.004]	[0.005]	[0.005]	[0.005]	[0.005]
Population	-0.007	-0.007	-0.001	0.007	0.006
	[0.006]	[0.007]	[0.006]	[0.009]	[0.008]
Patent / inhabitant	-0.006	-0.006	-0.011	-0.007	-0.005
	[0.007]	[0.008]	[0.008]	[0.008]	[0.007]
Age	0.003	0.004*	0.004**	0.004*	0.003*
	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]
Financial development					0.021***
					[0.008]
Inverse Mills Ratio				-0.143**	-0.163***
				[0.054]	[0.053]
Constant	-0.027	-0.042	-0.101	-0.393**	-0.335*
	[0.103]	[0.151]	[0.154]	[0.178]	[0.192]
Size HQ (dummy)	YES	YES	YES	YES	YES
Industry FE (3-digit)	YES	YES	YES	YES	YES
First-stage selection eq.	NO	NO	NO	YES	YES
Observations	4,273,058	3,737,847	458,154	3,737,847	3,736,256
R-squared	0.273	0.285	0.408	0.288	0.282

Standard errors clustered at the country level in brackets
*** p<0.01, ** p<0.05, * p<0.1

6 Conclusions

We have proposed and confronted with data a theory of business groups (BGs) as 'knowledge-based hierarchies' of firms (Garicano and Rossi-Hansberg, 2015) in a business environment characterized by 'incomplete contracts' (Antràs and Rossi-Hansberg, 2009). In our model the emergence of a BG and its hierarchical structure are the outcomes of two parallel decisions. The *integration decision* solves the tradeoff a headquarter (HQ) faces between better knowledge protection and easier rent extraction from integrated divisions on the one hand, and stronger team incentivization (plus comparative advantage) and more knowledge dissipation through legally independent subsidiaries on the other hand. The *hierarchy decision* solves the problem of how to use, communicate and enrich scarce HQ-specific knowledge efficiently.

The model predicts that a HQ is more likely to select the BG as the organizational form of its activities in the presence of better contractual institutions, and it is more likely to arrange its subsidiaries along a larger number of hierarchical layers when its intangibles are better, the problems it has to solve are more challenging, communication costs between layers are lower, and the skill premium for more knowledgeable managers is smaller.

We have tested these predictions exploiting the unique features of a dataset we have constructed from ORBIS. The dataset covers 178,190 business groups operating in 32 OECD countries and controlling more than 1,150,000 (domestic and foreign) affiliates in all countries worldwide in the year 2010. The dataset allows to observe the worldwide organization of BGs, with special emphasis on their ownership structure in terms of both the number of equity-controlled affiliates and their positions along the parent company's hierarchy of control.

In line with our theory, we have found significant and robust evidence that firms (in a given industry and of a given size) are more likely to set up subsidiaries when operating in countries characterized by better institutions. Conditional on this result, we have also found that BGs are more likely to be structured across several layers of hierarchy when communication costs between parents and affiliates are easier, the skill premia for good managers are smaller, and problem solving is more challenging (as proxied by lower standardization of the industries in which the BG operates). These results are robust to the inclusion of additional controls, as well as to different partitions of our sample in terms of industries or number of layers.

References

Acemoğlu, Daron, Simon Johnson, and Todd Mitton, "Determinants of Vertical Integration: Financial Development and Contracting Costs". The Journal of Finance, 63(2009), 1251-1290.

Alfaro, Laura, and Andrew Charlton, "Intra-industry Foreign Direct Investment", American Economic Review, 99(2009), 2096-2119.

Alfaro, Laura, Paola Conconi, Harald Fadinger, and Andrew Newman, "Do Prices Determine Vertical Integration ?", Review of Economic Studies 83 (2016), 1-35.

Almeida, Heitor V., and Daniel Wolfenzon, "A Theory of Pyramidal Ownership and Family Business Groups", The Journal of Finance, 61(2006), 2637-2680.

Altomonte, Carlo, Filippo di Mauro, Gianmarco I. P. Ottaviano, Armando Rungi, and Vincent Vicard, "Global Value Chains during the Great Trade Collapse: a Bullwhip Effect?", in S. Beugelsdijk, S. Brakman, H. van Ees, H. Garretsen (eds.) Firms in the International Economy, MIT Press, Cambridge MA., 2013.

Antràs, Pol, and Davin Chor. "Global Sourcing," Journal of Political Economy, 112(2004), pp. 552–580.

Antràs, Pol and Esteban Rossi-Hansberg, "Organizations and Trade," Annual Review of Economics, 1(2009), 43-64.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson. "Why Do Firms Own Production Chains?", NBER Working Paper N. 18020, 2012.

Baker, George, Robert Gibbons, and Kevin J. Murphy. "Bringing the Market Inside the Firm?", American Economic Review Papers and Proceedings, 91 (2001), 212-218.

– –, "Relational Contracts and the Theory of the Firm", The Quarterly Journal of Economics, 117 (2002), pp. 39-84

Belenzon, Sharon, Tomer Berkovitz and Luis Rios. 2013. "Capital markets and firm organization: How financial development shapes European corporate groups", Management Science, 59: 1326-1343

Belenzon, Sharon and Tomer Berkovitz. 2010. "Innovation in Business Groups", Management Science, 56(3): 519-535.

Blinder, A. S., and Krueger A. B. (2013) "Alternative Measures of Offshorability: A Survey Approach," Journal of Labor Economics 31, 97-128.

Bertrand, Marianne, Paras Mehta and Sendhil Mullainathan, "Ferreting Out Tunneling: An Application to Indian Business Groups," The Quarterly Journal of Economics Vol. 117, No. 1 (Feb., 2002), pp. 121-148

Caliendo, Lorenzo, and Esteban Rossi-Hansberg (2012). "The Impact of Trade on Organization and Productivity", The Quarterly Journal of Economics, 127(2012), 1393-1467.

Chapelle A. and Szafarz A. (2007). "Control consolidation with a threshold: an algorithm," IMA Journal of Management Mathematics vol. 18 pp. 235–243.

Colpan, Asli, and Takashi Hikino, "Foundations of Business Groups: Towards an Integrated Framework" in The Oxford Handbook of Business Groups, Colpan et al. (eds). Oxford University Press, 2010.

Cravino, J. and Levchenko, A. A. (2017). Multinational firms and international business cycle transmission. Quarterly Journal of Economics

Del Prete, Davide & Rungi, Armando, 2017. "Organizing the global value chain: A firm-level test," Journal of International Economics, Elsevier, vol. 109(C), pages 16-30.

Eurostat (2007). Recommendations Manual on the Production of Foreign Affiliates Statistics (FATS). European Commission.

Garicano, Luis, "Hierarchies and the Organization of Knowledge in Production", Journal of Political Economy, 108(2000), 874-904.

Garicano, Luis, and Thomas N. Hubbard, “The Return to Knowledge Hierarchies.”, NBER Working Paper N. 12815, 2007.

Garicano, Luis, and Esteban Rossi-Hansberg, “Inequality and the Organization of Knowledge,” American Economic Review, 94(2004), 197–202.

– –, “Organization and Inequality in a Knowledge Economy.”, The Quarterly Journal of Economics, 121(2006), 1383-1435.

– –, “Organizing Growth”, Journal of Economic Theory, 147(2012), 593–629.

– –, “Knowledge-based hierarchies: Using organizations to understand the economy. Annual Review of Economics, 7(2015), 1-30.

Gorodnichenko, Y., Kukharskyy, B., and Roland, G. (2017). Culture and Global Sourcing. University of Tübingen, mimeo.

Granovetter, Mark, "Coase revisited: Business Groups in the Modern Economy", Industrial and Corporate Change 4(1995), 93-130.

Grossman, Sanford, and Oliver Hart, "The costs and benefits of ownership: A theory of vertical and lateral Integration", Journal of Political Economy 94(1986), 691–719.

Grossman, Gene M., and Elhanan Helpman, "Integration versus Outsourcing in Industry Equilibrium", The Quarterly Journal of Economics 117(2002), 85-120.

– –, "Outsourcing versus FDI in Industry Equilibrium", Journal of the European Economic Association, 1(2003), 317-327.

– –, "Managerial incentives and the international organization of production", Journal of International Economics, 63(2004), pp. 237-262.

– –, "Outsourcing in a Global Economy," Review of Economic Studies, 72(2005), 135-159.

Hart, Oliver, and John Moore, "Property rights and the nature of the firm", Journal of Political Economy, 98(1990), 1119–1158.

Heckman, James (1976). "The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models", Annals of Economic and Social Measurement 5: 475–492.

Heckman, James (1978). "Dummy endogenous variables in a simultaneous equation system", Econometrica 46: 931–959.

Helpman, Elhanan, "Trade, FDI, and the Organization of Firms", Journal of Economic Literature, 44(2008), 589-630.

Hennart, Jean-François, "Explaining the Swollen Middle: Why Most Transactions are a Mix of 'Market' and 'Hierarchy'", Organization Science, 4(1993), 529-547.

IRI (2011). The 2010 EU Industrial R&D Investment SCOREBOARD. European Commission, Bruxelles (available at http://iri.jrc.ec.europa.eu/docs/scoreboard/2010/SB2010.pdf).

Khanna, Tarun, and Yishai Yafeh, "Business Groups in Emerging Markets: Paragons or Parasites?", Journal of Economic Literature, 45(2007), 331-372.

Lewellen, Katharina and Robinson, Leslie A., Internal Ownership Structures of U.S. Multinational Firms (2013). Available at SSRN: https://ssrn.com/abstract=2273553

OECD (2005). Guidelines for Multinational Enterprises. OECD, Paris.

Powell, Walter W., "Neither Market nor Hierarchy: Network Forms of Organization", Research in Organizational Behavior, 12(1990), 295-336.

UNCTAD (2016). World Investment Report 2016. United Nations, Geneva.

UNCTAD (2009). Training Manual on Statistics for FDI and the Operations of TNCs Vol. II. United Nations.

US BEA (2012). U.S. Multinational Companies. Operations of U.S. Parents and Their Foreign Affiliates in 2010. US Bureau of Economic Analysis, November 2012 (available at: http://www.bea.gov/scb/pdf/2012/11%20November/1112MN

Williamson, Oliver, "The Vertical Integration of Production: Market Failure Considerations", American Economic Review, 61(1971), 112-123.

– –, Markets and Hierarchies: Analysis and Antitrust Implications. NewYork: Free Press, 1975

– –, The Economic Institutions of Capitalism. New York: Free Press, 1985.

Appendix A: Business Groups and ORBIS Ownership Database

Our two main sources of data are both compiled by Bureau Van Djik (BvD), a Belgian consulting firm, and comprise the Ownership Database, from which we derive information on intra-group control linkages, and the Orbis database, from which we retrieve companies' balance sheet information. We exploit the 2010 version of both databases in this paper. The Ownership Database, in particular, includes information on over 30 million shareholder/subsidiary links for companies worldwide. Information on proprietary linkages is collected directly from single companies, from official bodies when in charge, or from some national and international providers. In Table A.1 we include a list of the information providers, with the indication of the countries/areas they cover, as reported by the Ownership Database. In case of conflicting information among providers covering the same country/area, the Ownership Database is updated according to the latest available report. Among the international providers, Bureau van Djik enlists also Dun & Bradstreet, a data source that has already been exploited in other academic works mentioned in this paper (Acemoglu, Johnson and Mitton, 2009; Alfaro et al., 2009 and 2016).

Table A.1: Original sources of ownership linkages collected by Bureau Van Djik

CIBI Information, Inc. (Philippines), Creditreform (Bulgaria, Ukraine & Rep. of Macedonia) , Chamber of Commerce & Industry of Romania (Romania), CMIE (India), CFI Online (Ireland), Creditreform-Interinfo (Hungary), Infocredit Group Ltd, (Cyprus & Middle East), CreditInform (Norway), Creditreform Latvia (Latvia), Creditreform (Rep. of Macedonia), Informa Colombia SA (Colombia), Contact database, Credinform (Russia & Kazakstan), Creditreform Austria (Austria), Coface Slovenia (Slovenia), Dun & Bradstreet (USA, Canada, Latin America & Africa), DGIL Consult (Nigeria), MarketLine, (previously Datamonitor), PT. Dataindo Inti Swakarsa (Indonesia), DP Information Group (Singapore), Finar Enformasyon derecelendirme ve danismanlik hizmetleri A.S (Turkey), Suomen Asiakastieto (Finland), Factset, Worldbox (Switzerland), Honyvem (Italy), Creditreform Croatia (Croatia), Huaxia (China), Inforcredit Group (Cyprus), Informa del Peru (Peru), ICAP (Greece), Informa (Spain), InfoCredit (Poland), Ibisworld (Australia), Jordans (UK, Ireland), Patikimo Verslo Sistema (Lithuania), Krediidiinfo (Estonia), Købmandstandens Oplysningsbureau (Denmark), KIS (Korea), LexisNexis (Netherlands), Bureau van Dijk (Luxemburg), Creditreform Belgrade (Bosnia-Herzegovina, Serbia & Montenegro), Coface MOPE (Portugal), National Bank of Belgium (Belgium), Novcredit (Italy), Qatar Chamber of Commerce and Industry (Qatar), Annual return (UK), Coface SCRL (France), Creditinfo Schufa GmbH (Czech Republic, Slovakia, Iceland, Malta), SeeNews (Moldova, Albania, Georgia & Uzbekistan), Chinese source, Statistics Canada (Canada), China Credit Information Service Ltd (Taiwan), Taiwan Economic Journal (Taiwan), Teikoku Databank (Japan), Transunion (South Africa), UC (Sweden), Verband der Vereine Creditreform (Germany), Worldbox (New Zealand, Hong Kong, Switzerland, Monaco, Liechtenstein, Pakistan, Sri Lanka & Cuba)

The observation unit collected by the Ownership Database is the single link between a company and each of its shareholders, with additional information on the total (direct and indirect) equity participation when relevant. For the year 2010 there are 7,707,728 companies with information on shareholding structures in the original database. An algorithm provided by Bureau van Dijk allows to identify in principle the ultimate owners (UOs) of a single company. However, since our purpose is to track the whole network of firms developed by each Business Group and model it as a hierarchical graph (see Figure 1), we have in principle to depart from the complete shareholding structure of each company, in order to identify one ultimate parent company, its set of affiliates and their relative distance within the hierarchy. To that extent, we have slightly modified the original BvD algorithm in two ways: we reconcile conflicting information that can arise from a mismatch between controlling and controlled subjects, and we differentiate between corporate and individual ultimate owners.

Conflicting information deriving from controlling and controlled subjects can arise in presence of cross-participations. In accordance with international standards we apply a threshold criterion (>50.01%) for the definition of control on the basis of (direct and indirect) participation. The latter is the methodology currently used across international institutions (OECD 2005; UNCTAD, 2009; Eurostat, 2007), although it can lead to an overestimation of control in some bigger networks of affiliates. That is, even after adopting a majority threshold as a criterion of control, it is still possible to end up with one affiliate controlled by more than one ultimate parent company. To solve that problem we rely on information officially provided by companies' consolidated financial accounts, when available. In particular, if we find that an affiliate is enlisted in more than one Business Group, we give priority to the ultimate parent company that enlists that affiliate in its consolidated accounts. In case no consolidated accounts are available, we include the affiliate in the group where it is located at shorter control distance from the parent.

The other correction that we apply to the standard BvD algorithm relates to the fact that the latter reports every property linkage between a company and each of its shareholders, thus including as members of potential business groups (as previously defined) also affiliates that are directly controlled by individual (non-corporate) shareholders, and that are not controlling subjects of any other company. As we want to characterize the drivers of BGs starting from the maximization of a firm problem, we have excluded these non-corporate UOs from our sample, although we include in our analysis those corporate networks that involve at least one intermediate property linkage of a corporate nature.[31] Specifically, our modified algorithm partitions all firms for which information on ownership is available in two groups:

a) a set of independent companies, that have as controlling shareholder individuals or a family or no specific corporate entity, and that are not themselves controlling shareholders of any other company;

b) all the other companies for which information on property linkages is available; these companies are either owned by a corporate controlling (immediate) shareholder or are themselves independent, but act as controlling shareholders of other companies.

The set a) of independent firms is used as a control group. The algorithm then screens every firm belonging to group b) for the highest total (direct and indirect) participation in the equity of each company, as provided by the Ownership Database. Once it finds a corporate controlling entity A that sums up to more than 50.01% of control in a given company B, company B is classified as an affiliate, while the same algorithm checks the shareholding structure of company A. If the latter is in turn ultimately owned by another corporate entity C, the process is repeated until a controlling company that has no corporate controlling shareholder is found. The latter is considered as the ultimate parent company of affiliate companies A, B and C. In the case of quoted companies, we consider as ultimate parent the highest company in the path of proprietary linkages we can identify. The procedure run for the year 2010 has recovered a total of 270,374 parents and 1,519,588 majority-owned affiliates (or subsidiaries) according to our definition.

Having identified the set of affiliates and their parent, the algorithm then assigns a hierarchical level within each Business Group, counting from the parent how many steps of intermediate property are required for ultimate

[31] If for example an individual X directly controls affiliates A and B, we do not consider the X-A-B network as a business group. Whereas, in the case of an individual X that indirectly controls affiliates C and D through a third company E, we consider the E-C-D network as a Business Group, in which company E is the ultimate (corporate) owner.

control. In case the same affiliate is encountered more than once in the same path (due to cross-participations), we consider it as located on the closest level where we have finally encountered it. A limit of the Ownership Database concerns the maximum number of control levels that can be obtained after considering cross-participations: the algorithm allows to reach a maximum of 10 levels for a maximum of 1,000 affiliates. However, in our data less than 0.1% of BGs exceed such limits.

In the figure below we report the correlation between the number of headquarters controlling foreign affiliates abroad (left panel) and the number of foreign affiliates (right panel) located in each country, as retrieved from our sample and matched against the corresponding figures provided by UNCTAD (2011). The original source for data on affiliates in UNCTAD (2011) is Dun &Bradstreet, that is one of the sources of ownership data on which the ORBIS database also relies. The survey of UNCTAD (2011) refers to data in 2009, while our data are updated to 2010. We have excluded from the validation reported in the figure the datapoint on China, since the country does not adopt the international standard definition of control (>50.01%) in reporting the number of affiliates, preferring a less committal criterion of 'foreign-funded enterprises', leading to non comparable figures.

Table A.2: Sample validation

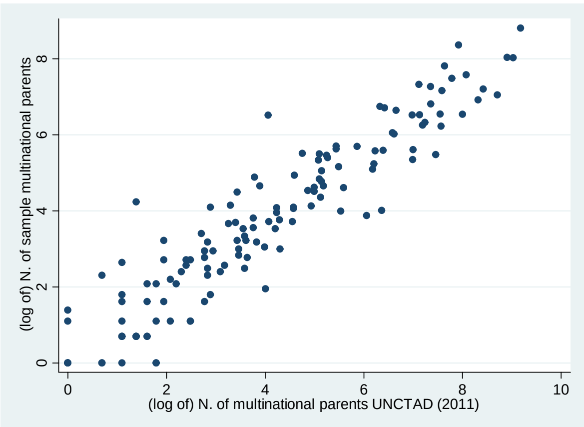

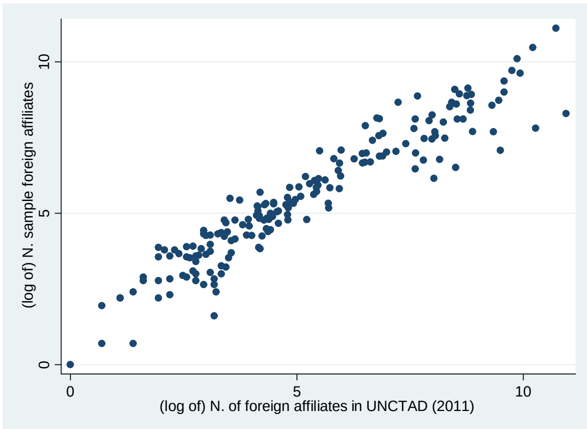

Appendix B: Vertical Integration in Business Groups

In absence of actual data on internal shipments of intermediate goods and services across firms, Acemoglu et al. (2009) proposed to proxy vertical integration exploiting the information on the set of industries in which a firm is engaged, combined with the input coefficient requirements that link those industries as retrieved from input-output tables. A firm-level index was therefore calculated summing up all input-output coefficients that linked each firm's primary activity to the secondary activities in which it was involved. The assumption is thus that a firm engaged in more industries, where backward and forward linkages in production are important, is supposed to have a higher capacity to source internally more inputs for its final output.

In order to take into account the BG dimension, we have refined the original index. In particular, we assume that within a group two sets of activities can be identified: a set of output activities $j \in N_H$, and a set of intermediate activities $i \in N_A$. The set of output activities coincides with the primary and secondary activities of the headquarter (N_H), whereas the range of intermediate activities at the group-level is represented by the set of primary and secondary activities in which controlled affiliates (N_A) are involved. With these assumptions, we can build a group-specific input-output table, where we report outputs in columns and inputs by row and where each combination VI_{ij} is the ith coefficient requirement to produce the jth output. As in Acemoglu et al. (2009) or Alfaro et al. (2016), we assume that industrial backward and forward linkages for all firms in our sample can be proxied by US input-output tables and adopt the industrial classification provided by the US Bureau of Economic Analysis, with 61 main industries mainly at a 3-digit level of disaggregation of the NAICS rev. 2002 classification. By summing up input coefficient requirements by column we obtain the vertical integration for each line of business in which the Business Group is involved.[32] To retrieve the vertical integration index for the whole group, we average the total of all input coefficient requirements (VI_{ij}) by the number of output activities ($|N_H|$), thus correcting for the potential conglomerate nature of the group. The result is the following group-specific (g) vertical integration index:

$$v_g = \sum_{\substack{i \in N_A \\ j \in N_H}} \frac{1}{|N_H|} VI_{ij} \tag{14}$$

where $VI_{ij}$are the input coefficient requirements for any output activity $j \in N_H$sourcing from all input activities $j \in N_A$. The group-specific vertical integration index can range from $0$to 1, where 1corresponds to complete vertical integration.

In our dataset the average vertical integration across groups (v_g) is .062 (that is, on average 6 cents worth of inputs are sourced within groups for a one dollar unit of output). For comparison, the figure obtained by Acemoglu et al. (2009) on their (unconstrained) sample is of .0487. Alfaro et al. (2016) also calculated in a similar way a vertical integration index for manufacturing firms with more than 20 employees, obtaining an average vertical integration of .063.

[32] In absence of actual data on internal shipments of intermediates, we can interpret this number as a mere propensity to be vertically integrated, where the sum of industry-level requirements gives us only the maximum possible integration of production processes.