Cracks in the Pyramid: Business Group Structure and Natural Disasters *

Carler Altomonte† G.I.P. Ottaviano ‡ Tommaso Sonno§

Extended Abstract

Abstract

We study how business groups adjust their organizational complexity in response to natural disasters, using a novel panel covering over 200,000 group-year observations worldwide. By matching firm-level data on subsidiaries, ownership hierarchies, and geographic location with geolocated disaster shocks, we construct new exposure measures and examine their effects on internal group structure. We find that disaster exposure leads to short-run reductions in complexity, particularly in the number of subsidiaries per ownership layer. These effects are stronger when subsidiaries are directly hit and are confirmed by an event study around large, isolated disasters. We also document a decline in the skewness of the distribution of subsidiaries across layers, especially among initially pyramid-shaped groups. Results are robust to controls for disaster intensity and point to dynamic, structure-dependent responses to shocks.

*We thank ...

†Bocconi University, Baffi Carefin and Bruegel, carlo.altomonte@unibocconi.it
‡Bocconi University, Baffi Carefin, IGIER, CEP (LSE), CEPR, gianmarco.ottaviano@unibocconi.it
§University of Bologna, CEP (LSE), tommaso.sonno@unibo.it

1 Introduction and Literature Review

We study the evolution of business groups' (BGs) organizational structure worldwide, from 2007 to 2018, focusing on the distribution of subsidiaries across layers, i.e. a BG's organizational shape, and observe how the latter is affected by natural disasters.

We define a business group as an organizational form of economic activity in which at least two legally autonomous firms function as a single economic entity through hierarchical control: a parent company ('headquarter' - HQ) owns, directly or indirectly, the majority of the equity shares of at least one legally independent firm ('subsidiary'). Subsidiaries in the first layer are directly controlled by the HQ; subsidiaries in the second layer are controlled directly by subsidiaries in the first layer, and thus indirectly by the HQ; and so on. Under this definition, multinational enterprises (MNEs) are BGs with at least one subsidiary incorporated in a country different than the HQ's. In line with Altomonte et al. (2021), as long as a business group has at least two layers of subsidiaries, a sufficient statistic to characterize a BG's organizational shape is the skewness of the distribution of its subsidiaries across layers. When skewness is negative, subsidiaries are denser at lower layers leading to a pyramid-shaped hierarchy. When skewness is positive, subsidiaries are denser at higher layers leading to an inverted pyramid-shaped hierarchy. The closer skewness is to zero, the more symmetric the distribution is, resembling to diamond-shaped structures.

In the context of foreign direct investment, natural disasters effectively discourage multinational corporation to enter disaster-struck markets, especially when other multinationals from the same home-country are also absent from said market (Oh, Oetzel, et al., 2020). Natural disasters also increase the likelihood of exit for multinational joint ventures, particularly when firms hold highly specific assets (Bowman, Foulser-Piggott, and Beamish, 2023). Post-disaster divestment can be observed not only at the country-, but also at the regional-level; this is largely driven by intra-national relocations toward more developed, less disaster-prone regions, contributing to greater regional economic divergence (Friedt and Toner-Rodgers, 2022). However, the overall number of subsidiaries in BGs appears relatively unaffected by natural disasters if compared to terrorist attacks or technological disasters, suggesting that firms perceive these events as less critical (Oh and

Oetzel, 2011).

Innovation outcomes are similarly impaired, with significant reductions in both the quantity and quality of corporate innovation following natural disasters. Financial constraints emerge as the primary channel, alongside increased inventor mobility, lower inventor productivity, and a heightened aversion to risky innovation projects (Le et al., 2024).

The effects of natural disasters differ substantially across disaster types and severity. Disasters with even impact dispersion, high recurrence, and advance warning, such as hurricanes and floods, prompt more anticipatory firm responses (McKnight and Linnenluecke, 2019). Furthermore, when paired with a smaller scale of impact flooding and storms may even produce positive outcomes, such as increases in labor and capital growth (Zhou and Botzen, 2021). In addition they may even spur agricultural growth, as long as their scale and severity is contained. Disasters of greater size indeed seem to exclusively have negative impacts, which become more pronounced in developing economies (Panwar and Sen, 2019).

2 Dataset construction and sources

BGs structure. In order to identify BGs' hierarchies, we exploit the Orbis Ownership Database of Bureau van Dijk, which provides ownership control links among more than 200 million entities, and reports, for each company, information on all shareholders. Starting from these data, Sonno (2017) elaborates an algorithm that obtains the network of ownership for each business group. In particular, the final result stems from the definition of direct or indirect majority ($\geqslant 50.01\%$) of the voting rights provided by Bureau Van Djik, and from the binary links joining each firm with each of its shareholders, and with its global ultimate owner (GUO). This definition of control is consistent with the international standards for multinational corporations (OECD, 2005; EUROSTAT, 2007; UNCTAD, 2009). Using the information on the GUO, the algorithm delimits the perimeter of business groups and identifies headquarters, while following ownership links backwards, it is able to disentangle the exact position of each subsidiary in the control hierarchy. The latter information is summarized employing the concept of hierarchical layers: headquarters are always at layer 0, subsidiaries they control directly by the HQ are at layer 1 while subsidiaries controlled by layer n subsidiaries are posed at layer n+1.

This provides the hierarchical structure of each BG in terms of number of layers and number of subsidaries on each layer with their location (by geolocating them using zipcodes), industry and year of incorporation. Thus, this algorithm constructs the network of business groups for more than 6.3 million business groups, with 12.8 million affiliates in more than 200 countries, from 2007 to 2018. This is the first global, firm-level dataset documenting multinationals' hierarchies and activities in a panel setting. In order to analyze the hierarchical structures of BG, we compute, separately for each group in each year, the skewness of the distribution of subsidiaries among layers. In this setting, right skewed distributions (positive skewness) correspond to groups organized in the form of inverted pyramids with a large number of firms in the top layers (close to the headquarter) and less subsidiaries in further layers. On the other hand, left skewed distributions (negative skewness) correspond to groups that are organized as Pyramids (few firms at the top and many at the bottom of the hierarchy).

We obtain geographical coordinates for all subsidiaries in the dataset through a Google API procedure.

Natural Disasters. We obtain information on natural disasters from the Emergency Events Database (EM-DAT) and the Geocoded Disasters database (GDIS). EM-DAT contains official information at the country and regional level on the occurrence of natural and technological disasters between 1980 and 2020, associated victims, affected people (defined as requiring food, water, shelter, or immediate medical assistance) and economic damage. An event is recorded in EM-DAT if it meets at least one of the following criteria: at least ten victims recorded; at least one hundred affected people recorded; international relief aid sought; or a state of emergency declared. GDIS contains disasters from 1960 to 2018 and it is a geocoded extension of a selection of EM-DAT.
We consider all disasters except Technological and Biological. The EM-DAT dataset provides information for 22126 disasters in the years 1980-2018. Of these, only around 2200 are provided with coordinates. We obtain coordinates for the rest of the disasters in two steps. First, we merge EM-DAT with the GDIS dataset. There are however observations from EM-DAT that do not match GDIS disasters; for these observations we obtain coordinates through a geocoding procedure with

Google API.
The final result is a coordinates-augmented EM-DAT dataset.

A major step for the purposes of this research is the identification of disaster shocks at the subsidiary level. We assign each disaster to the subsidiaries it hit according to the following procedure. We spatially identify disasters at GADM 2 level, meaning that a disaster with coordinates inside a GADM 2 region will hit all subsidiaries with coordinates inside that region. Using a Geographic Information System, we join cities and disasters with GADM 2 polygons.
We finally obtain a panel with all subsidiaries through time[1] and information on disasters hitting those subsidiaries.

3 Descriptive Statistics

3.1 Dataset Evolution

In Table 1 we present how many parent and affiliate companies we observe in 2008 and in 2018, i.e. at the beginning and end period of analysis. Parent companies have increased by roughly 65%, and affiliates by more than 42%.

Table 1: Parent and affiliate companies: beginning-end period

Companies	2008	2018
Parent	63,002	104,025
Affiliates	890,642	1,271,702

Notes: The table describes how many parent and affiliate companies we observe at the beginning and at the end of our period of analysis.

In Figure 1 we provide descriptive statistics on the evolution of BGs over time by their size. From the left panel, using 2008 as base year, we can observe the percentage variation in the number of BGs by four size clusters. The right panel provides the beginning and end period composition of the sample by BG's size.

[1]Notice that the panel is unbalanced as subsidiaries have different life spans.

Figure 1: Business Groups demography, 2008-2018

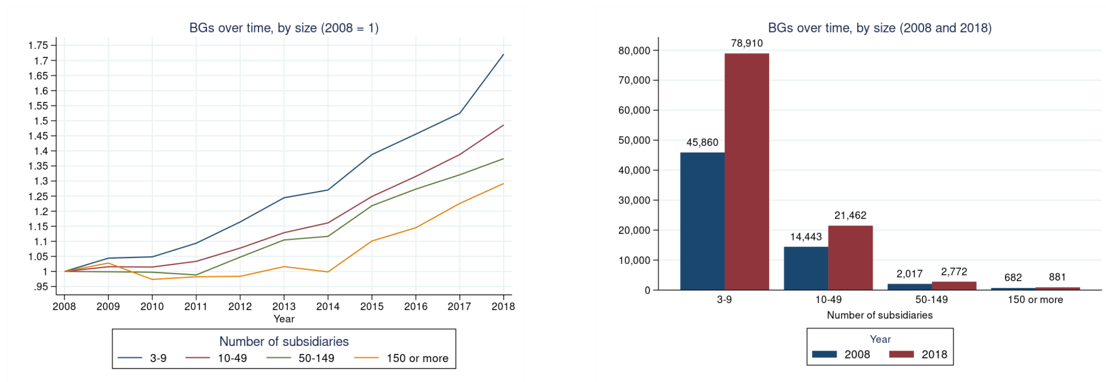

3.2 Evolution of skewness over time

In Figure 2, in the left panel we plot the average skewness for each year. Though always positive, on average the level of skewness across our entire sample has decreased, i.e. BGs became more pyramidal. Also, though small in nominal terms, the downward shift is statistically significant and corresponds to a more than 20% decrease. To further check the overall trend, on the right hand side of Figure 2 we plot the weighted average of BG's skewness for each year, weighted by firm size. The results show a similar, if not more accentuated, decreasing trend in the yearly mean level of skewness. This suggests that for bigger BGs there has been a more evident decreasing trend in skewness.

Figure 2: Overall variation of skewness over time

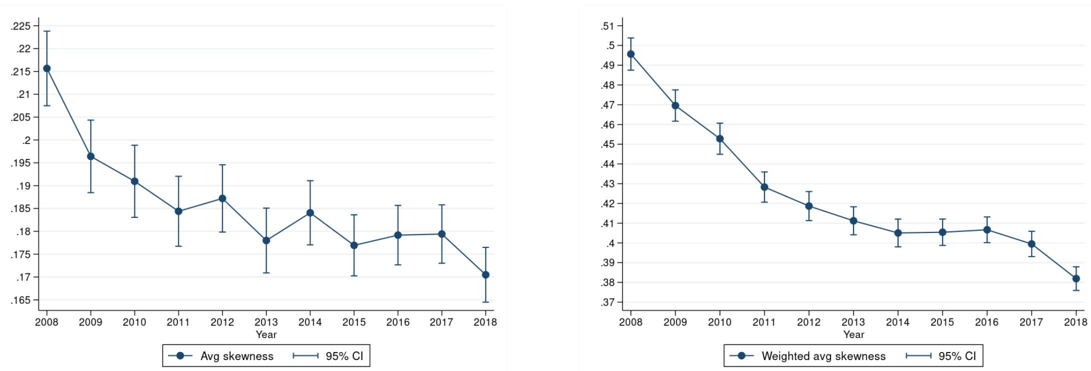

Notes: The figure on the left shows the average skewness of the overall sample, for each year. The right-hand side graph shows the weighted average skewness of the firms, weighted by the firm size.

We decompose this aggregate result in order to understand its driving factors.

To do so, we look at the intensive (i.e., the internal re-organization of "surviving" BGs over time) and extensive margin (i.e. the organization of entering and exiting BGs) variations, respectively. In each graph we plot the mean skewness for each year of the corresponding sample.

Intensive Margin. The "surviving" BGs, i.e. groups which we first observe in 2007 and are still present in 2018, are almost 27 thousands. As depicted in Figure 3, the intensive margin evolution presents, on average, an opposite direction with respect to the overall variation in skewness. This positive trend is statistically significant. This suggests that groups observed throughout the whole period shifted towards a more inverted-pyramid setting. However, by carefully observing the actual value of the average skewness, we can see how the range of the latter corresponds to higher values than the range of Figure 2, indicating that surviving BGs have on average a higher skewness than the whole sample. Thus, surviving BGs have added affiliates relatively more in upper layers. This is in line with previous evidence that incumbent groups have more subsidiaries in their first layer.

Figure 3: Intensive margin variation of skewness across time

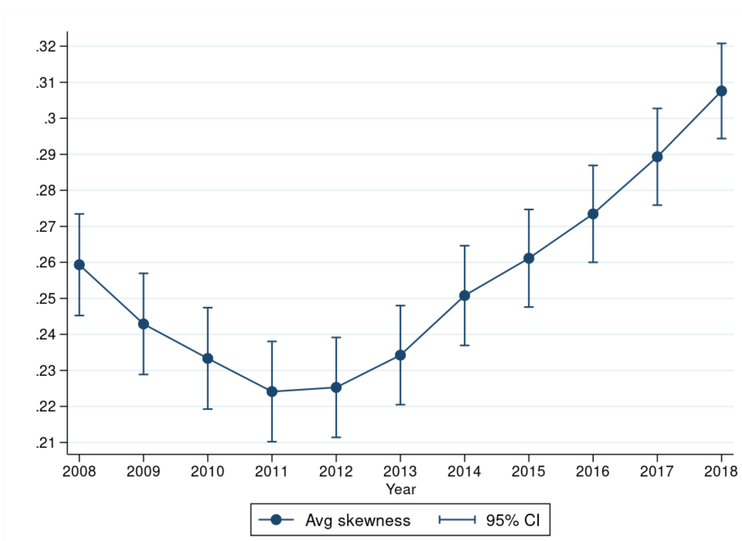

Notes: The graph shows the yearly mean level of skewness of BG which we observe from 2008 to 2018.

Extensive Margin. Here we provide evidence for the evolution of the skewness at the extensive margin. To do so, in Figure 4, left panel, we compute the average skewness of entering and exiting BGs, for each year. Each dot in the graph corresponds to the average skewness of the groups entering/exiting in each year. This implies, however, that the samples used for the average skewness for the two groups change every year. Therefore, in the right panel of Figure 4 we represent the difference in the average skewness of entering (exiting) firms rescaled by their numerosity. In this way, we can interpret the blue line as the absolute increase or decrease in the average skewness due to the churning of our dataset. More formally:

$$\bar{X}_1 = \frac{N_0}{N_1}\left(\frac{\sum_{N_0} skewness}{N_0}\right) + \left\{\frac{E}{N_1}\left(\frac{\sum_E skewness}{E}\right) - \frac{X}{N_1}\left(\frac{\sum_X skewness}{X}\right)\right\}$$

where $\bar{X}_1$ is the average skewness in year 1; N_0 is the number of firms in year 0; E and X are the number of firms entering and exiting in year 1, respectively.

Figure 4: Extensive margin variation of skewness over time

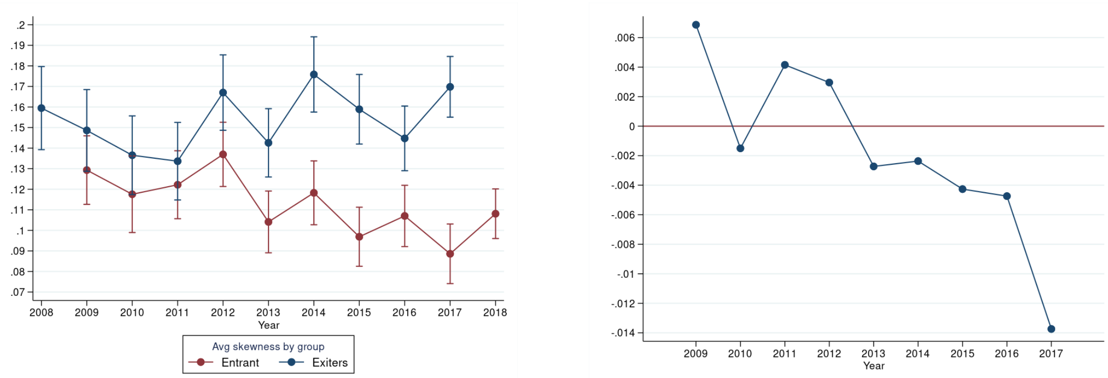

((a)) Yearly mean skewness

((b)) Weighted skewness

From the left panel of Figure 4, we can observe that entering BGs have on average a significantly lower skewness than exiting BGs. This is in line with the overall decreasing average skewness over time. Moreover, the right panel confirms the results. In fact, we can observe how the difference in the average skewness of entering and exiting firms rescaled by their numerosity is, on average, below zero.

3.3 Business Group Life-cycle

In this section we study the evolution over time of business groups. It seems that in general as the groups get older they tend to be organized as inverse pyramids see Figure 5. Figures 5 and 6, select 10 cohorts of groups that are observed over 10 years (from 2008 to 2017). The first cohort (C0) only contains groups that are 1 year old in 2008 and 10 years old in 2017; the second cohort (C10) only contains groups that are 11 years old in 2008 and 20 years old in 2017. The others cohorts are built analogously. This helps visualizing both the changes in averge skewness within each cohort over time (due to macroeconomic trends or other factors that affect simultaneously all cohorts) and the changes in average skewness across cohorts (due to the natural life cycle of business groups).

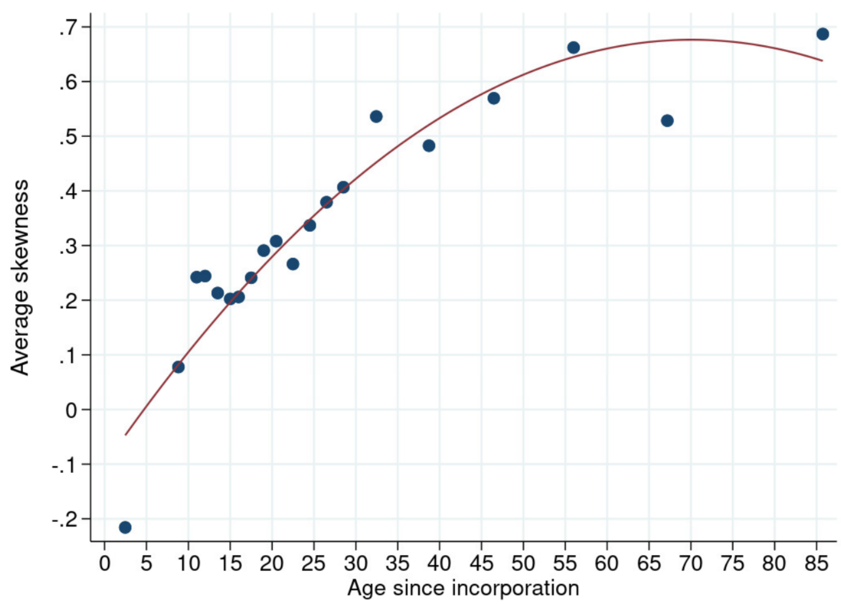

Figure 5: Group Structure over Age

Notes: Here we plot a restricted number of 1,808 GUOs that have been observed over the entire 10 years time-span (2008 to 2017). Among the group of GUOs that are observed over the 10 years, these are selected as members of non-overlapping age cohorts. In particular, cohort C0 consists of 165 groups (born in 2008 and 10 y.o. in 2018), cohort C10 consists of 689 groups (11 y.o. in 2008 and 20 y.o. in 2018), cohort C20 consists of 425 groups, cohort C30 of 147 groups, cohort C4 of 106 groups, cohort C50 of 83 groups, cohort C60 of 82 groups, cohort C70 of 45 groups, cohort C80 of 41 groups, cohort C90 of 25 groups. Other cohorts have fewer observations and are therefore excluded from the graph. This binscatter groups obervations in 20 points representing a total of 17,087 observations.

Within this aggregate picture, we can imagine there are different dynamics for Pyramids and Inverse Pyramids (see Figure 6).

We can study the observed trends analytically by recourse to a linear regression model specification (see Table 2).

Table 2: Intensive and Extensive variations in BG's structures

	Intensive Change (*Dep Var:* Skewness)		Extensive Change (*Dep Var:* Change of Type)	
Interactions:	None	Type	None	Type
	(1)	(2)	(3)	(4)
Δ years	-0.003***	-0.022***	-0.000	0.001***
	(0.001)	(0.001)	(0.000)	(0.000)
log(Age)	0.107***	0.013	-0.009***	-0.013***
	(0.005)	(0.008)	(0.002)	(0.003)
log(Size)	-0.077***	-0.059***	-0.015***	-0.017***
	(0.007)	(0.008)	(0.002)	(0.002)
Δ years - Pyr		0.066***		-0.003***
		(0.002)		(0.000)
log(Age) - Pyr		0.047***		0.013***
		(0.011)		(0.004)
log(Size) - Pyr		-0.019		0.005
		(0.013)		(0.005)
FE	GUO	GUO	GUO	GUO
Observations	573,439	549,588	362,980	362,980
R^2	0.866	0.870	0.251	0.251

Notes: In columns 1-2 the dependent variable is the skewness of the group. In columns 3-4, the dependent variable is a dummy taking on value 1 if the group classification changes from pyramid to inverse pyramid or viceversa. In every specification, our controls of interest are: the number of years since the GFC (Δ years); the log of the group's age (log Age); the log of the number of group's subsidiaries (log Size). Every specification uses group fixed effects for time-constant latent group heterogeneity. For each dependent variable, we propose two types of model specifications: a baseline (labelled "None"), where controls are not interacted and a second one (labelled "Type") where the regressors are interacted with a dummy indicating groups that were Pyramids when first observed. We only keep groups that change of group type (extensive margin) maximum once (89.31% of observations).

4 Preliminary Evidence

As a first step, we label as "major disasters" those natural disasters characterized by an above median mortality. We then try to see whether BGs hit by at least one major disaster in one of their affiliates experience a significant change in their ownership structure in the following year. To assess structural shifts we however only look at those affiliates in the BG that are untouched by disasters. More succinctly, we ask whether a disaster event affecting one subsidiary of the group prompts the HQ to adopt different structural decision with regards to the other subsidiaries.

Figure 6: Group Structure over Age by Group Type

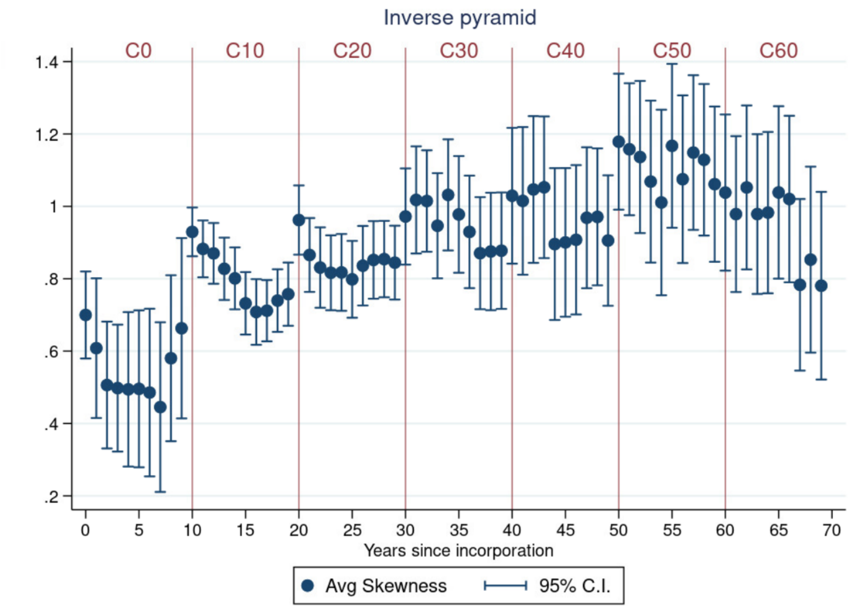

((a)) Inverse Pyramids

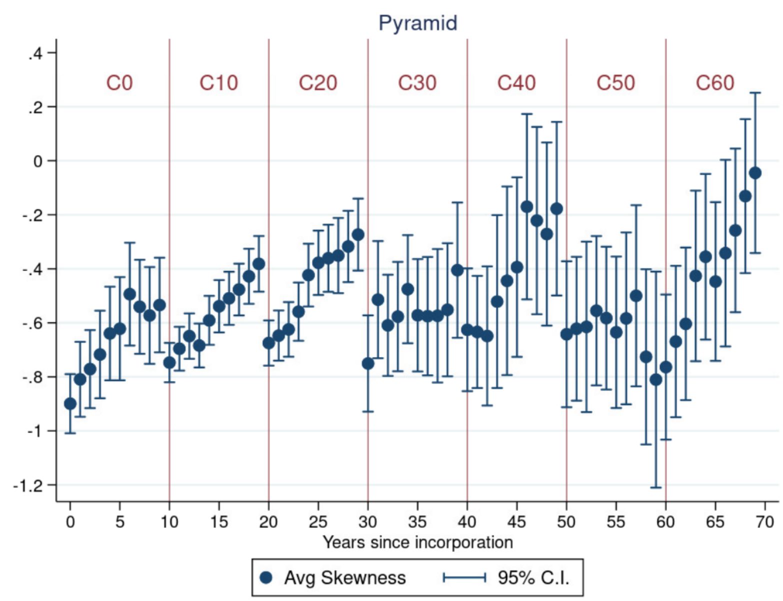

((b)) Pyramids

Notes: Here we split the sample of figure 5 in two subsamples, based on the observed level of skewness when the group is first observed in 2008. Pyramids are BGs with negative skewness, Inverse Pyramids are BGs with positive skewness. The plots represent a restricted number of 1,885 GUOs that have been observed over the entire 10 years time-span (2008 to 2017). Among the group of GUOs that are observed over the 10 years, these are selected as members of non-overlapping age cohorts. In particular, cohort C0 consists of 165 groups (born in 2008 and 10 y.o. in 2018), cohort C1 consists of 689 groups (11 y.o. in 2008 and 20 y.o. in 2018), cohort C2 consists of 425 groups, cohort C3 of 147 groups, cohort C4 of 106 groups, cohort C5 of 83 groups, cohort C6 of 82 groups, cohort C7 of 45 groups, cohort C8 of 41 groups, cohort C90 of 25 groups. Other cohorts have fewer observations and are therefore excluded from the graph. This binscatter groups obervations in 20 points representing a total of 17,840 observations.

Table 3: Cleaned outcome vars on major disasters lagged

	$\Delta Subs$_clean	$\Delta Layers$_clean
L.major	-0.0550***	-0.0545***
	(0.00660)	(0.0167)
FE	GUO year	GUO year
Obs	30766	39258
R^2	0.15	0.09

Standard errors in parentheses

$*p < 0.1, **p < 0.05, ***p < 0.01$

Notes: Main regressor is the lag at $t-1$ of a dummy taking value 1 when the BG is hit by a disaster with above-median mortality. Outcome variables represent the percentage change in the number of affiliates and layers of BGs, excluding from the computation any affiliates ever hit by a disaster.

By doing so, we ignore the mechanical effect of disasters, namely the fact that a natural disaster is bound to affect a BG structure insofar as it physically hits one of the firms that makes it up. We therefore look mainly at two outcome variables: the percentage change in the number of subsidiaries owned by the group ($\Delta Subs\_clean$) and the percentage change in the number of overall layers in the BG ($\Delta Layers\_clean$). The suffix "clean" indicates that all affiliates struck were ignored in the computation of such values, as specified prior. We include GUO and year Fixed effects and report results in Table 3.

As we can observe, BGs hit by major disasters tend to lose 5% of their affiliates in the following year. Such divestment response within the group then translates to an observable shrinkage of the ownership structure of the BG, which loses an analogous percentage of ownership layers. In turn the vertical hierarchy of the BG is cut shorter, with less firms and layers overall.

We extend our analysis over a broader temporal horizon by conducting an event study, examining the impact of a major disaster on business group outcomes over a three-year window before and after the event. To do so, we consider "isolated major disasters", namely those major natural events that were neither preceded nor succeeded by another major disaster in the considered time frame. We then observe and register the impact of an isolated major disaster on the previously mentioned outcome variables. Figures 7 and 8 clearly show how major disasters reduce the count of subsidiaries and hierarchical layers in BGs both one and two years after the impact, while at the third year the BGs appear to revert to the pre-disaster levels. In this regard, the pre-period years do not seem to differ sig-

Figure 7: Event Study - Δ Subs Clean on Isolated Major Disasters

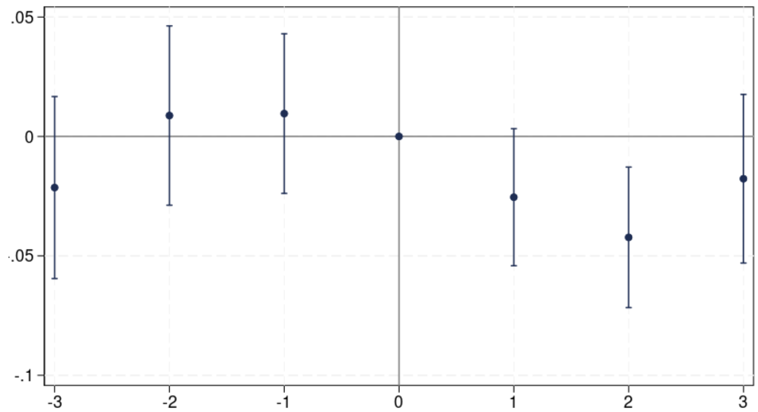

Figure 8: Event Study - Δ Layers Clean on Isolated Major Disasters

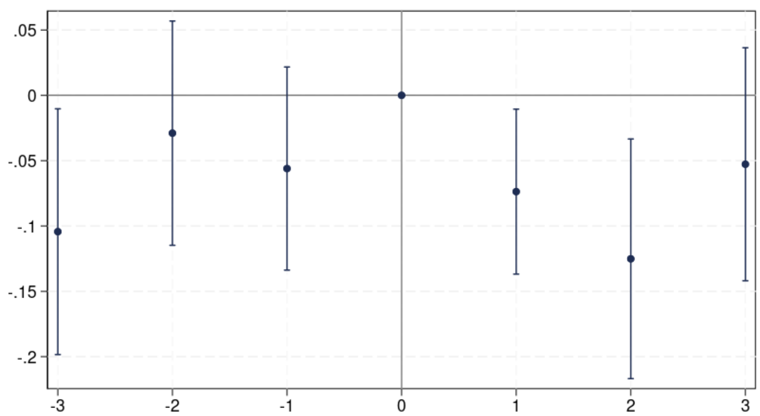

nificantly amongst each other nor, mostly, with respect to the year of the disaster. This suggests the absence of any pre-trends intrinsic to the BGs themselves, and leads to attribute the dips observable at $t + 1$ and $t + 2$ to the disaster impact itself.

References

Bowman, Gary, Roxane Foulser-Piggott, and Paul W Beamish (2023). "Natural disasters and MNE internalization: Reoptimizing subsidiary governance". In: *Journal of World Business* 58.2, p. 101387.

Friedt, Felix L and Aidan Toner-Rodgers (2022). "Natural disasters, intra-national FDI spillovers, and economic divergence: Evidence from India". In: *Journal of Development Economics* 157, p. 102872.

Le, Huong et al. (2024). "Natural disasters and corporate innovation". In: *The European Journal of Finance* 30.2, pp. 144–172.

McKnight, Brent and Martina K Linnenluecke (2019). "Patterns of firm responses to different types of natural disasters". In: *Business & Society* 58.4, pp. 813–840.

Oh, Chang Hoon and Jennifer Oetzel (2011). "Multinationals' response to major disasters: how does subsidiary investment vary in response to the type of disaster and the quality of country governance?" In: *Strategic Management Journal* 32.6, pp. 658–681.

Oh, Chang Hoon, Jennifer Oetzel, et al. (2020). "Natural disasters and MNC subnational investments in China". In: *Multinational Business Review* 28.2, pp. 245–274.

Panwar, Vikrant and Subir Sen (2019). "Economic impact of natural disasters: An empirical re-examination". In: *Margin: The Journal of Applied Economic Research* 13.1, pp. 109–139.

Zhou, Fujin and Wouter Botzen (2021). "Firm level evidence of disaster impacts on growth in Vietnam". In: *Environmental and Resource Economics* 79.2, pp. 277–322.