Supplementary material for 'The Corporate Elite Community Structure of Global Capitalism' by Eelke Heemskerk & Frank Takes

1: Filtering ties in the network

We considered and tested a number of approaches for reducing 'administrative' ties. Ideally we do this with a simple, clear and sensible intervention. But all of approaches to filter out the administrative ties resulted had shortcomings. In particular, all filter approaches removed social ties as well. Therefore, we decided not to apply any additional filters before running the algorithm. The filter approaches we examined included the following:

For each firm we have and indication of its independence in terms of ownership. In theory, filtering out all firms with a majority shareholder (over 380,000 firms) would decrease administrative ties through holding pyramids. However, there are many large companies that we do want to include but that have a majority owner. These include certain state-owned firms, firms with partial state ownership such as Volkswagen AG and also that issue share certificates without voting rights on the market (for instance the Dutch financial multinational Aegon).

A similar approach is based on information on the ultimate owner of each firm. A seemingly sensible filter would be one that excludes those firms that have majority owners, and whose ultimate owner is part of the set of firms that we have selected. This would narrow the selection to exclude subsidiary firms (over 160,000 firms). However, this would also lead to the exclusion of firms that are part of larger international corporations, but that do play an important role in the network of interlocking directorates in particular business communities. These include for instance the local branches of Siemens or Royal Dutch Shell.

We also considered removing firms in the financial services sector with relatively high levels of assets but with relatively few employees because these are probably shell companies.

Although this would remove some of the connections through administrative ties between shell companies, it is only a partial fix.

Another approach is to filter out edges within holding pyramids by comparing the names of the companies. If a director serves on two firms with very similar names, this is probably a position within a holding structure. This approach was not effective for two reasons. First, it only deleted a certain proportion of intra-holding interlocks. Second, there are plenty of holding structures wherein a short holding name is followed by a longer ‘surname’, for instance EDP Financial Services and EDP Production Facilities. The overlap in names here is too small to be detected by a filter.

Alternatively, we could search for people in the database with a very large number of positions and remove these from the database as outliers. This will reduce the number of administrative ties. However, the drawback is that directors often create sets of both social and administrative ties. Excluding these directors would thus also delete a potentially significant number of social ties.

There was one other option to filter our administrative ties for which we lacked the data necessary for execution. Administrative board interlocks between firms often go together with significant equity ties as well. Just as there is a network of interlocking directorates, there is also a network of ownership (Vitali, Glattfelder & Battiston 2011). Removing all board interlocks that are accompanied by significant directional ownership relationships may lead to a significant reduction in administrative ties. Because we were not able to execute this filter it was also not possible to see if there are any serious drawbacks to this method

2: How representative is the data?

An important question regarding the quality of the data is how representative it is. We know that we have data on board composition of corporations all over the globe and well beyond the western-industrialized world, such as South-America, Africa, and Asia. But it may still be the case that there is an over- or underrepresentation in the number of firms from particular countries. In order to get an impression of the quality of the data we looked at how the presence of companies from particular countries in our sample resembles that country's global economic position.

As baseline for the comparison we take GDP of a country as percentage of the sum of world GDP. This gives us a relative number for a countries global economic position. We than calculated the number of firms in a country in our sample as percentage of all firms in our sample. And third we calculated the sum of revenues for all firms included in the selection in a country as percentage of the sum of revenue for all firms in the sample. In figure S1,the blue vertical bars show the difference between the percentage share of GDP and the percentage share of firms per country. The red bars show the difference between GDP and revenue. This gives a graphical representation of under- and overrepresentation. We only include countries that are under- or overrepresented by at least 0.1 percent point on one of the two measures; most countries fall under this threshold.

For instance, compared to its contribution to world GDP Brazil is underrepresented in our selection by both measures. South Africa is overrepresented in terms of relative share of revenue. The stark underrepresentation of China in terms of total corporate revenue in the sample compared to worldwide GDP contribution is because of a large share of incomplete firm data. Because we do not use revenue data for our analysis, this is not a problem. For the USA the underrepresentation on the other hand is a result from its very large contribution to global GDP.

Figure S1: Share of firms, revenues and GDP compared

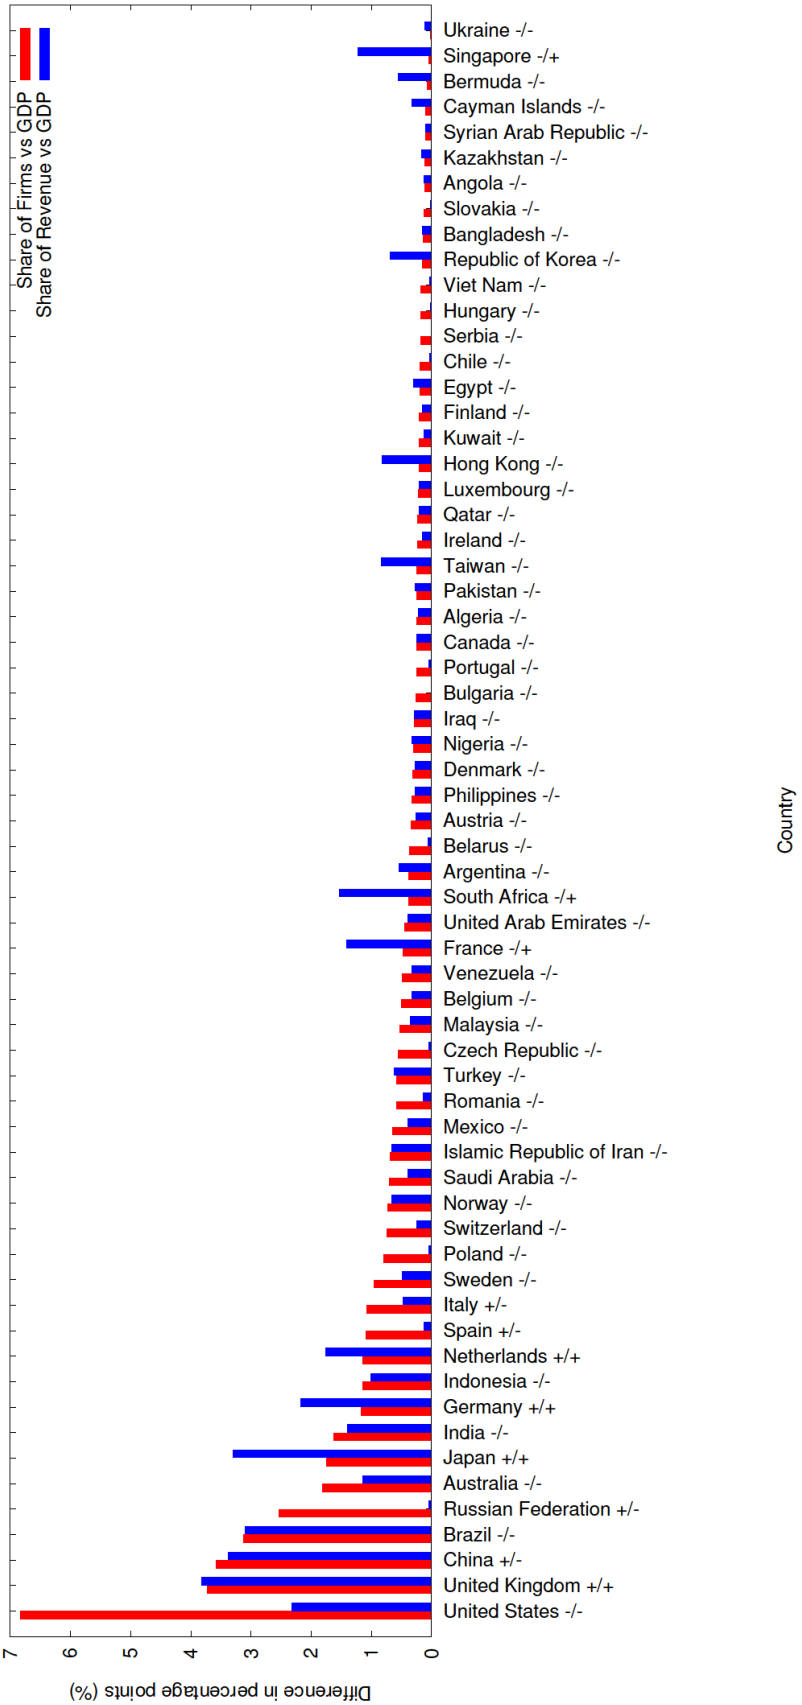

3. Geographic distance of board interlocks

The global network of corporate boards is connected through 1,712,060 interlocking directorates. For most firms in the selection, we have information on headquarter location. Using the country and city, we determined the geographic location in longitude and latitude for these firms. For a number of firms, we were unable to determine the geographic location. In most cases, this is because the name of a town or city is not unique in a particular country. In the end, we have been able to determine the geographic location of both endpoints (firms) of 1,079,150 edges (or board interlocks). The lion share of these edges connect firms in the same country. And of this set, near 600,000 edges occur within the same city.

In the community analysis, we exclusively looked at transnational ties. For a total of 177,946 transnational interlocks we have determined the geographic location of both firms. The average geographic distance of these transnational interlocks is 3022 kilometres. However, this is not normally distributed, as figure S2 illustrates. An intriguing pattern emerges. Most of the transnational board interlocks do not bridge very long distances. One fifth of all transnational interlocks connect firms of that are no more than 425 km apart. Interestingly, there is another ‘ bump’ from 500 kilometres onward. From 1000 kilometres onward, the interlocks become much more rare. This suggest that traveling up to 1000 kilometre is still acceptable for board members, while longer distances become less acceptable. The maximum distance of a board interlocks is 20,000 km; half of the diameter of the world.

Figure S2: Geographic distance of board interlocks