

The distribution functions of vessel calls and port connectivity in the global cargo ship network

Michael T. Gastner, César Ducruet

▶ To cite this version:

Michael T. Gastner, César Ducruet. The distribution functions of vessel calls and port connectivity in the global cargo ship network. Maritime Networks. Spatial structures and time dynamics, pp.242-261, 2015, 9781138911253. halshs-01374944

HAL Id: halshs-01374944 https://shs.hal.science/halshs-01374944

Submitted on 12 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The distribution functions of vessel calls and port

connectivity in the global cargo ship network

Michael T. GASTNER, Institute of Technical Physics and Materials Science, Research Centre for Natural Sciences, Hungarian Academy of Sciences, Budapest, Hungary César DUCRUET, CNRS & UMR 7235 EconomiX, Nanterre, France

Pre-final version of the chapter published in Ducruet C. (Ed.) (2015) *Maritime Networks. Spatial Structures and Time Dynamics.* Routledge Studies in Transport Analysis, pp. 242-261.

Recent years have seen intense research activity in the modelling and analysis of complex networks, mainly driven by the availability of new large-scale databases for social, biological, and technological networks (see for example Newman, 2010 for a review). Maritime transport networks are one area where these new ideas and techniques have found fertile ground (Deng et al., 2009; Hu and Zhu, 2009; Kaluza et al., 2010; Ducruet, 2013). In this study, we analyse a database generated from *Lloyd's Shipping Index*, a weekly publication of cargo ship movements by *Lloyd's List*, over the period 1890 to 2008. For 20 selected years, an entire volume of the *Index*, each containing data for one week, was extracted and the data transformed into a network where the nodes are ports and links are nonstop ship voyages. Because cargo shipping is the dominant transport mode for world trade (UNCTAD, 2013), it is of great economic relevance for understanding the importance of the nodes. Here we measure importance in two ways:

- the number of vessel calls; and
- the degree, defined as the number of ports that the node is connected to by at least one arriving or departing ship.

The call and degree distributions are arguably the two most important summary statistics of the network (see Chapter 4 by Barthelemy for a more general discussion). They may not allow a complete reconstruction of traffic on the links. However, unweighted and weighted degree distributions are an important feature of a network's topology and have often been used as circumstantial evidence for mechanistic models of the network's evolution (Barabási and Albert, 1999; Tadić, 2002; Moore et al., 2006). The call distribution also plays a crucial role for predicting the full origin-destination matrix (i.e. the traffic between all pairs of ports) because it is an input in transport forecasting (e.g. in the gravity model or the intervening opportunities model (Wilson, 1967)).

In the early phase of complex network science, many degree distributions of realworld networks were characterised as power laws (Barabási and Albert, 1999). In the parlance of statistics, an integer-valued power law is a probabilistic model that assigns the probability Pr(k) to the event that an arbitrary node has degree or weight k so that

$$\Pr(k;\tau) = \frac{k^{-\tau}}{\zeta(\tau)}, \quad k = 1, 2, 3, \dots$$
 (14.1)

Here $\zeta(\tau)$ is the Riemann zeta function and $\tau > 1$ a fixed parameter that has to be fitted to the data. Because $Pr(ak, \tau) = a^{-\tau} Pr(k, \tau)$, the distribution of equation 14.1 is also called scale-free. The interest in power law distributions stems mainly from the fact that these are particularly heavy-tailed (i.e. the tail of the distribution decays more slowly than an exponential function). As a consequence, a power-law distribution has a large range in degrees: while most nodes have only a small degree, some nodes possess a much larger degree than the average. At first glance, such a heterogeneity in degrees is indeed found in many empirical networks (Albert and Barabási, 2002).

Most of the time, a power-law degree distribution was inferred from straight-line fits to the log-log diagram of the degree frequency, but this is now generally viewed as an unsatisfactory approach (Clauset et al., 2009). Identifying a region where the data appear more or less linear is largely arbitrary because most distributions are too noisy and substantially curved on double-logarithmic scales. Straight-line fits based on standard least-squares algorithms can also lead to a bias in the estimated exponents. Most importantly, however, there is no a priori reason why Pr(k) has to be a power law. Many other common probability distributions are also heavy-tailed and may fit the data better. Recent studies in fact doubt that power laws are as ubiquitous as once believed (Edwards et al., 2007; Willinger et al., 2009; Stumpf and Porter, 2012).

This study proposes that statistical methods should be applied to the maritime network to understand which type of call or degree distribution best explains the observed data. We apply information-based model selection (Burnham and Anderson, 1998) and statistical tests to compare different candidate distributions. Another goal of the paper is to assess whether the call and degree distributions of the cargo ship network have significantly changed over time. One hypothesis from the geographic literature is that the call and degree distributions may structurally change alongside major technological transformations of the shipping industry and their consequences on port operations and maritime network configurations (McCalla, 2004). The studied period goes across different dominant ship technologies, such as sail, steam, combustion, specialised vessels (e.g. container, tanker), and mega-carriers. Such technological evolutions are believed to have been selective, as some ports were dropped from the network and replaced or superseded by new ones better adapted to changing standards, sometimes resulting in an increasing concentration of port activity favouring fewer and larger ports. Containerisation is seen as a revolution in itself with profound impacts on network configuration and world trade (Cullinane and Khanna, 1999;

Bernhofen et al., 2013). In this study we find indeed evidence that the call distributions have evolved so that the fraction of small ports has decreased. At the same time the Gini coefficient, a common measure for inequality of a distribution, has slightly decreased over the study period.

Before proceeding with the statistical analysis, we emphasize one caveat. The voyages reported by Lloyd's certainly form only a subset of the entire global traffic distribution and may possibly be biased, for example if certain ports, ships, or routes are systematically underreported. The quality of reporting may also differ between different years. We currently have insufficient knowledge whether such biases are present and thus cannot apply any corrections. *Lloyd's Shipping Index*, however, is the most complete and consistent data source available to study the development of the cargo ship network over the investigated time period. Therefore, we are confident that the trends reported below are genuinely representative of the network's evolution.

Probabilistic models

We investigate eight different models that have frequently been used to fit empirical degree distributions in complex networks (Table 14.1). We restrict our study to discrete distributions

- (a) whose support are all positive integers; and
- (b) that depend on maximally two parameters.

Restriction (a) reflects that degree or port calls only have integer values. One might argue that k=0 should also be included and that the maximum degree should have an upper bound because the network is finite. However, from *Lloyd's Shipping Index* we cannot directly infer which ports were in principle open to traffic but remained unused. Consequently, distributions with infinite support but excluding k=0 are more appropriate in the present context. The number of ports with k>0 can differ slightly between the call and degree distributions: some ports have a positive number of calls but degree zero because in the raw data some vessels are reported to call at one, not two ports (namely origin and destination) in their latest known voyage. For the call distribution we kept these isolated ports but removed them from the degree distribution since zero is unlikely to be their true degree. Restriction (b) avoids overfitting of the data but still includes the "usual suspects" for degree distributions in socio-economic networks.

[Table 14.1 here]

We include four one-parameter models: besides the power law of equation 14.1 (also known as zeta distribution), we assess the likelihood of Poisson, geometric, and Yule-Simon distributions. The Poisson distribution describes the node degrees of large sparse Erdős-Rényi random graphs, a common null model in network studies. The tail of a Poisson distribution decays faster than exponentially so that degrees in Erdős-

Rényi graphs are effectively limited to values near the mean degree. The geometric distribution decays exponentially, whereas the Yule-Simon distribution has a power law tail and only differs mildly from a strict power law for small k. Because the Yule-Simon distribution is the exact solution of popular «preferential attachment» models (Simon, 1955) (i.e. models where nodes are constantly added to the network and linked preferentially to nodes of high degree), we have included it in our list.

As a mixed case we introduce the exponentially truncated power law as one of our two-parameter models in Table 14.1. The negative binomial is another two-parameter example that decays more slowly than an exponential if its parameter r exceeds 1, but with much less weight in the tail than a power law. All the distributions mentioned so far are discrete: they are defined for integer numbers k. In the case of the Poisson and negative binomial distributions, k=0 is conventionally included in the distributions' support. In order to restore restriction (a) from above, we constrain these distributions to exclude k=0, which explains why the equations in Table 14.1 differ from the ordinary textbook form.

Among continuous distributions, there are two further canonical candidates whose decay is between an exponential and a power law: the lognormal distribution and the

Weibull distribution (also known as stretched exponential if $\beta < 1$ in the formula

stated in Table 14.1). We include these models in our study because previous studies have reported lognormal (Stumpf et al., 2005; Bhattacharya et al., 2008; Gómez et al., 2008; Todor et al., 2012) and Weibull distributions (Lahererre and Sornette, 1998; Broido and Claffy, 2001; He et al., 2007; Rocha et al., 2010) in real-world networks. To allow a direct comparison with the other models, we have to discretise the continuous lognormal and Weibull distribution, which can be accomplished in a variety of ways. Here we have chosen to integrate the continuous distributions between subsequent integers k and k+1. In terms of the cumulative distribution F, we can express the integral as the probability F(k+1)-F(k). For the lognormal distribution, we assign this probability to the integer at the upper boundary k+1 and call this the "discrete lognormal" distribution. In the case of the Weibull distribution preliminary tests showed that the likelihoods are in general slightly larger if F(k+1)-F(k) is assigned to k instead of k+1. Imposing the constraint (a), yields the expression for the "discrete Weibull" distribution in Table 14.1.

One subtlety to note is that in this study we do not judge the fit of the distributions by the tails alone, as is often done elsewhere (Barabási and Albert, 1999; Clauset et al., 2009). In the study of continuous phase transitions in physics it is justified to restrict attention to the tails because only these are important for determining "universal" power-law features. However, in the present context it is far-fetched to assume that the cargo ship network has anything to do with a physical phase transition. Instead we will assess the match of the distribution over the full set of positive integers

 $k=1, 2, 3, \rightleftharpoons$ with the same motivation as in the study of city size distributions by

Eeckhout (2004). Although it is in principle possible to restrict the analysis to the tails by introducing a lower cutoff k>1, this would introduce an additional parameter and ignore the bulk of the data, which consists of ports with only few calls and low degree. On the contrary, we regard it as valuable information for practitioners to model lowtraffic ports too, not only the small fraction of busy hubs that make up the distribution's tail.

Akaike information criterion

Assuming that all calls and degrees are independent, the likelihood function for any of the models in Table 14.1 has the general form

$$L(\mathbf{v}) = \prod_{i=1}^{n} \Pr(k_i; \mathbf{v}), \qquad (14.2)$$

where k_i is the number of calls (or the degree) at port *i*, *n* is the number of ports in the

sample, and v is the set of the parameters in the second column of the table. The calls (or degrees) may in reality depend on each other so that L(v) in equation 14.2 is more properly thought of as a composite likelihood. We can justify the use of a composite likelihood in our present context because the call (or degree) distribution Pr(k) that we would like to model is a marginal rather than the complete joint distribution of all calls (or degrees). Therefore, the full dependence structure is in statistical terms a "nuisance parameter" which neither matters to us nor is it clear how to specify the full likelihood. In such cases, composite likelihood methods have proved to be a well-behaved alternative (Cox and Reid, 2004; Varin et al., 2011). Another, more pragmatic, point of view is that there is no straightforward method to establish from our data how the k_i may depend on each other so that assuming independence is the most parsimonious choice.

For a specified model *j*, we determine the parameter $\hat{\mathbf{v}}_j$ that maximizes *L* and hence also the log-likelihood ln(*L*). A comparison between different models can then be performed by ranking their Akaike information criterion (AIC) (Akaike, 1974),

$$AIC_{j} = -2\ln(L(\hat{\mathbf{v}}_{j})) + 2K_{j}, \qquad (14.3)$$

where K_j is the number of parameters in the respective model. The AIC not only tells us which model is closest to the data in information content, properly taking into account that higher K_j generally allows better fits to the data, but weakens the explanatory power of the model. We can also make quantitative comparisons between different models based on the differences. If AIC_{min} is the minimum AIC over all models, then the difference

$$\Delta_j = \text{AIC}_j - \text{AIC}_{\min} \tag{14.4}$$

estimates the relative expected information gain between model *j* and the estimated best model. Because the likelihood of model *j* given the number of calls (or degrees) $k_1, k_2,...$ is proportional to $\exp(-\Delta_j/2)$ (Burnham and Anderson, 1998), the relative likelihood is the so-called Akaike weight

$$w_j = \frac{\exp(-\Delta_j/2)}{\sum_i \exp(-\Delta_i/2)},$$
(14.5)

where the summation in the denominator is over all models included in the comparison. Model selection by Akaike weights has become an increasingly popular tool to compare different hypothesized probability distributions (Edwards et al., 2007; Stumpf et al., 2005; Hamilton et al., 2008; Prieto and Sarabia, 2011). The Akaike weights for our data sets are summarized in Table 14.2 for the call distributions and in Table 14.3 for the degree distributions.

As a quick glance at the tables reveals, the maximum Akaike weights are achieved by the discrete lognormal and Weibull distributions and, in the case of the calls, in some years a truncated power law. These two-parameter models always perform better than even the best one-parameter model, which is in all cases the Yule-Simon distribution.

The added term $2K_i$ in equation 14.3 for introducing a second parameter is therefore

more than compensated by an increased likelihood for the best-performing models.

The effect of including a second parameter can be seen in Figure 14.1(a) and (b) where we compare the observed call distributions in 1910 and 2000 with the maximum-likelihood estimates for the Yule-Simon, truncated power law, lognormal, and Weibull distributions. Both observed distributions are substantially curved on a log-log scale and thus difficult to fit by an asymptotic power law such as the Yule-Simon distribution. All other plotted candidate distributions have the flexibility to

follow the curvature more accurately. Among these, the truncated power law decays in the limit $k \rightarrow \infty$ most rapidly and the lognormal distribution most slowly in the tail.¹

Comparing the observed call distributions in 1910 and 2000, the most obvious difference is that the initial decay on the left-hand side appears less curved in Figure 14.1(a) than in (b). For this reason, the truncated power law that fits well in 1910 is no longer a suitable candidate in 2000. In general, we observed by visual inspection that in the small to medium port range the complementary cumulative call distribution tends to be more curved in later years. This trend is detected by the Akaike weights in Table 14.2 that have shifted over the years from the truncated power law to the Weibull distribution. In practice, this change in the distribution implies that there is now a smaller fraction of ports listed with maximally 10 calls in one week. The inset in Figure 14.1(b) confirms this trend, showing a statistically significant decrease between the years 1890 and 2008 from 75.4 percent to 69.0 percent of ports having no more than 10 calls.²

[Figure 14.2 here]

The interpretation of the degree distribution is a little trickier. The Akaike weights in Table 14.3 seem to suggest a clear distinction: before the mid-1960s the most likely model is in all but one case a lognormal, but afterwards it is always a Weibull distribution. However, Figure 14.2(a) shows that in 1910 the lognormal and Weibull distributions are visually more or less equally good fits. Only in 2000 (Figure 14.2b) does the maximum-likelihood Weibull distribution clearly fit better in the tail than the lognormal. Unlike for the call distribution, we do not find a significant trend that the percentage of low-degree ports has decreased (inset in Figure 14.2b). In order to shed light on the significance of the apparent trend in the Akaike weights for the degree distribution, in the next section we will compare the performance of the maximum-likelihood models with another statistical technique.

Vuong's likelihood ratio tests

For models with an equal number of parameters, Akaike weights compare the models purely by the differences in their log-likelihood. Proponents of AIC-based model selection have argued that the Akaike weights are sufficient to judge the significance of the best model (Burnham and Anderson, 1998). However, others argue that the loglikelihood alone does not in itself allow an assessment when we should reject the second-ranked model in favour of the model with the highest Akaike

¹ Depending on the parameters, k may have to be larger than the maximum port size for the lognormal to exceed the Weibull distribution. For example in Figure 14.1(a) we are not yet far enough in the asymptotic regime on the right-hand side for the Weibull distribution to fall below the lognormal.

 $^{^{2}}$ Because the number of ports, however, has more than doubled between 1890 and 2008 (see second column of Table 14.2), the *absolute* number of ports with less than or equal to 10 calls has of course still increased.

weight (Stephens et al., 2005). Likelihood ratio tests, on the other hand, can inform us how significant the difference in the log-likelihood is (Clauset et al., 2009).

In this section we investigate the three two-parameter models that achieved maximal Akaike weight in at least one year for either the call or degree distribution: truncated power law, discrete lognormal, and Weibull distributions. For these nonnested models we apply the likelihood ratio test devised by Vuong (1989). The test statistic for comparing models r and s is the ratio of their likelihoods from equation 14.2 or equivalently its logarithm

$$R = \ln\left(\prod_{i=1}^{n} \frac{p_{r}(k_{i})}{p_{s}(k_{i})}\right) = \sum_{i=1}^{n} (\ln p_{r}(k_{i}) - \ln p_{s}(k_{i})), \qquad (14.6)$$

where $p_r(k_i)$ is the probability $\Pr(k_i; \hat{\mathbf{v}}_r)$ assigned to observing degree k_i in model r with the maximum-likelihood parameters $\hat{\mathbf{v}}_r$. If we assume that all observed k_i are independent (as we discussed after equation 14.2), then all terms $\ln p_r(k_i) - \ln p_s(k_i)$ in the sum on the left-hand side of equation 14.6 are also independent. With the shorthand notation $l_i = \ln p_r(k_i) - \ln p_s(k_i)$, the variance of one term in the sum can be estimated as

$$\sigma^{2} = \frac{1}{n} \sum_{i=1}^{n} l_{i}^{2} - \left(\frac{1}{n} \sum_{i=1}^{n} l_{i}\right)^{2}.$$
(14.7)

For sufficiently large *n*, the random variable *R* thus becomes normally distributed with an estimated variance $n\sigma^2$. We can then apply a conventional *Z*-test to determine whether the observed value *R* is significantly different from zero given the observed variance. The *p*-value can be expressed as

$$p = \operatorname{erfc}\left(\frac{|R|}{\sqrt{2n\sigma^2}}\right),\tag{14.8}$$

where erfc is the complementary error function.

This handwaving derivation, which follows essentially that of Clauset et al. (2009), is admittedly oversimplified. For both models *r* and *s* in Eq. 14.6 we have fitted the parameters $\hat{\mathbf{v}}_r$ and $\hat{\mathbf{v}}_s$ to the same data so that there are nontrivial correlations between *R* and σ^2 . However, Vuong (1989) proved that Eq. 14.8 still remains true. One noteworthy point about this equation is the appearance of the variance σ^2 . The variance of the data is a crucial piece of evidence whether one of the two models in question is likely to be significantly better. The Akaike weights, by contrast, did not account for the variance.

[Table 14.4 here]

We list the *p*-values for all pairwise comparisons between truncated power laws, lognormal, and Weibull distributions in Table 14.4. We highlight in bold type all *p*-values less than 0.1 and, where the likelihood ratio test indicates a deviation from randomness at this 10 percent significance level, we list in parentheses the more likely model. As an overall pattern for the call distribution (left half of the table), the tests are for most years indifferent between the three candidate models. However, for the degrees (right half of the table) the test strongly rejects for most years the truncated power law, consistent with its small Akaike weights in Table 14.3. Also in agreement with our earlier findings, the tests favour the Weibull distribution in some of the more recent years for both calls and degrees. There are examples where the Akaike weights suggest a high likelihood for one particular model, yet the likelihood-ratio test does not lend strong support to it. For example, in 2008 the Weibull distribution has an Akaike weight 0.957 for the calls, but, after factoring in the variance in the data, the likelihood ratio test does not reject the possibility that the data could be from a lognormal distribution (Akaike weight 0.043) or even a truncated power law despite

its much lower Akaike weight (5.93×10^{-5}) .

There is thus for most years no simple answer if the call distribution is better described by a Weibull or lognormal distribution. Over the range of observed calls (between 1 and approximately ; the precise upper bound, of course, depends on the year in question) the maximum-likelihood distributions from the lognormal and Weibull family do in fact not differ very much as can be seen in Figure 14.1(a). Likewise, for the degree distribution there is no clear support in favour of the lognormal hypothesis prior to 1960 (except in 1940) despite generally having the highest Akaike weight. Afterwards, there is increasing evidence in favour of the Weibull distribution, which might have to do with an increasing number of ports in the sample that allows us to distinguish more clearly between the models.

One has to bear in mind that neither the Akaike weight nor the likelihood ratio test can tell us that a model is good in an absolute sense, only that it is more plausible than its competitors. In other words, if all our candidate models are bad, then "in the country of the blind, the one-eyed man is king." We will now apply two classic goodness-of-fit tests that show that some of our candidate models are indeed a good match for the observations.

Kolmogorov-Smirnov and Anderson-Darling tests

The key idea behind both the Kolmogorov-Smirnov (KS) and Anderson-Darling (AD) tests is to compare the difference between the observed and hypothesized cumulative distribution functions. If the observed degree (or number of calls) are $k_1, k_2, \rightleftharpoons, k_n$, then the observed cumulative distribution $F_{obs}(x)$ is the number of ports with $k_j \ge x$ divided by *n*. The KS statistic is defined by

$$D_{\text{ks}} = \max_{\mathbf{x}=1,2,3,\pm} \left| \mathcal{F}_{\text{obs}}(\mathbf{x}) - \mathcal{F}_{\text{model}}(\mathbf{x}) \right|, \tag{14.9}$$

where F_{model} is the cumulative distribution function of the model to be tested (Kolmogorov, 1933). In words, D_{KS} is the maximum absolute difference between observed and hypothesized cumulative distribution function for any possible value x. While D_{KS} has a very intuitive interpretation, it also has one shortcoming when applied to heavy-tailed distributions: $|F_{\text{obs}}(x) - F_{\text{model}}(x)|$ is typically maximized where $F_{\text{model}} \approx 0.5$ and therefore D_{KS} does not effectively sample the tail where $F_{\text{model}}(x)$ is close to 1. This phenomenon can be understood as follows. If the model were correct and the difference between F_{obs} and F_{model} only the consequence of random chance, the jackknife estimate of the standard deviation in the difference is $\sqrt{F_{\text{model}}(x)[1-F_{\text{model}}(x)]/(n-1)}$, which has a maximum at $F_{\text{model}}(x)=1/2$.

There is one obvious cure to this problem: we divide the difference to be maximized in equation 14.9 by the expected standard deviation,

$$D_{AD} = \max_{\mathbf{x}=1,2,3,\rightleftharpoons} \frac{\left| F_{obs}(\mathbf{x}) - F_{model}(\mathbf{x}) \right|}{\sqrt{F_{model}(\mathbf{x})[1 - F_{model}(\mathbf{x})]}},$$
(14.10)

where we dropped the term $\sqrt{n-1}$ because it is independent of x. D_{AD} is called the Anderson-Darling statistic (Anderson and Darling, 1952). We have decided to carry out tests for both D_{KS} and D_{AD} because these two statistics measure different features

of the distribution. A good model should be able to have small values of D_{KS} as well as D_{AD} .

We test the significance of the same three two-parameter models as in the likelihood ratio test (truncated power law, lognormal, and Weibull distribution), but also include for comparison the Yule-Simon distribution which the Akaike weights identified as the best one-parameter model. We calculate *p*-values with Monte Carlo simulations based on the following algorithm.

First, we determine for a given model the maximum-likelihood parameters $\hat{\mathbf{v}}_{obs}$ that fit the *Lloyd's Shipping Index* data best. For the model distribution with parameters $\hat{\mathbf{v}}_{obs}$ we calculate the observed test statistics $D_{\text{KS, obs}}$ and $D_{\text{AD, obs}}$. Next we generate *n* random numbers drawn from the model distribution with parameters $\hat{\mathbf{v}}_{obs}$. We then pretend that we do not know $\hat{\mathbf{v}}_{obs}$ and determine the maximum-likelihood parameters $\hat{\mathbf{v}}_{md}$ that fit the random numbers best. In general, $\hat{\mathbf{v}}_{md}$ will differ slightly from $\hat{\mathbf{v}}_{obs}$. From the difference between the random numbers (now treated as surrogate observation) and the model with $\hat{\mathbf{v}}_{md}$ we calculate $D_{\text{KS, md}}$ and $D_{\text{AD, md}}$. We repeat drawing *n* random numbers 10^5 times and estimate the *p*-value p_{KS} for the KS test by the fraction of runs with $D_{\text{KS, md}} \ge D_{\text{KS, obs}}$. The same calculations are also carried out for the AD statistic. The repeated calculation of $\hat{\mathbf{v}}_{md}$ slows down the simulation, but is necessary to mimic the steps in the calculation of $D_{\text{KS, obs}}$ and $D_{\text{AD, obs}}$. Otherwise we obtain *p*-values with a strong downward bias that would lead us to accept the null hypothesis (i.e. that the real data follows the model distribution) more often than truly justified (Clauset et al., 2009).

[Table 14.5]

The *p*-values are listed in Table 14.5. The highlighted entries in bold type are those cases where there is no reason to suspect at the 10 percent significance level that the model is wrong, either in terms of the KS or the AD statistic. It is striking that the Yule-Simon distribution fails as a null hypothesis for call and degree distributions in all years. The truncated power law is accepted only for the call distribution (left half

of the table) and mostly in the early years of our data base. By contrast, the lognormal distribution is a suitable null hypothesis for the call distribution in all except one year (1995), and even then the null hypothesis would not be rejected at a 5 percent significance level. For the degree distribution (right half of the table), a lognormal null hypothesis is accepted in most, but not all years. Especially in the later years, the Weibull distribution shows better performance than the lognormal, confirming the trends we observed in the Akaike weights and the likelihood ratio tests. However, the KS and AD tests in earlier years only sporadically support a Weibull distribution.

Discussion

The overall picture that emerges from the KS and AD tests is a surprisingly consistent performance of the lognormal model for the call distribution. The only rejection, namely by the KS test for the data of 1995, could plausibly be by random chance. After all, at the chosen 10 percent significance level it is likely that at least one false positive exists among the 20 years which we have tested. The lognormal hypothesis also gains support from a recent analysis of world container port throughput (Ding and Teo, 2010) that reported a good fit between the number of containers handled at 300 top ports and a lognormal distribution. There is also a simple mechanistic model that could explain how a lognormal distribution might come about: Gibrat's law of proportionate growth (Gibrat, 1931). It is, in principle, possible to carry out further tests whether *Lloyd's Shipping Index* supports the key principle behind Gibrat's law, namely that the growth rates in calls are independent of the number of calls. Such a test will make more stringent demands on the data quality than what we currently have available. Right now such an effort would be hampered, for example, by the irregular time intervals between the samples. As more data becomes available, an analysis of port growth rates is clearly an intriguing research direction.

For the time being, we can instead view the call distribution from yet another angle. As an alternative to plotting the complementary distribution function directly (Figure 14.1), economists frequently employ Lorenz curves (Lorenz, 1905) to visualise inequality in distributions. Translated to our application, the Lorenz curve y(x) shows the percentage of ship calls that were made at the x percent of lowest ranked ports (ranked by the number of calls). We plot the Lorenz curves for three representative years in Figure 14.3(a). If all ports had an equal number of calls, the Lorenz curve would be the dashed diagonal line. One measure of inequality is the area between this diagonal and the actually observed Lorenz curve: the more unequal the distribution, the larger this area. Multiplied by two, this measure is known as the Gini coefficient (Gini, 1912). The coefficient itself, as well as a jackknife estimate of its standard error, can be conveniently calculated with ordinary least-squares regression (Giles, 2004). The results for all years in our data base (Figure 14.3b) reveal that the Gini coefficient for the calls has decreased from 0.80 in 1890 to 0.74 in 2008. Although this is a subtle decline, it is statistically significant: the values are more than five standard errors apart.

[Figure 14.3 here]

Why is the inequality declining? The maximum number of calls has increased (from 822 in 1890 to 2,422 in 2008), which at first glance suggests increasing inequality. The resolution to this apparent paradox lies in the inset of Figure 14.1(b): the fraction of small ports with ≤ 10 calls has decreased. As the total number of ports has grown over the years, an overproportional number of new medium-sized ports were added to the network. Together with a flattening global hierarchy this has reduced the gap between core and periphery, thereby making the network more polycentric. This trend more than compensates the growth in maximum port size and has led to an overall decrease in the Gini coefficient.

There is more than just the Gini coefficient that we can infer from the Lorenz curve. One complementary measure is the Lorenz asymmetry coefficient (LAC) (Damgaard and Weiner, 2000). A Lorenz curve is defined to be symmetric if it has the same slope as the diagonal "line of equality" (i.e. a slope of 1) at the point where the curve and the antidiagonal line y=1-x (i.e. the dotted line in Figure 14.3a) intersect. One can show that for a continuous cumulative distribution F with mean μ the slope equals 1 at

$$x = F(\mu)$$
. At this point $y = \int_{0}^{\infty} (x/\mu) dF(x)$, so a criterion for symmetry is

LAC =
$$F(\mu) + \int_{0}^{\mu} (x/\mu) dF(x) = 1.3$$
 If LAC < 1, the Lorenz curve is skewed such that it

has slope 1 below the dotted antidiagonal. Conversely, if LAC > 1, the Lorenz curve is parallel with the line of equality above the antidiagonal symmetry axis.

The LAC is of interest because curves with the same Gini coefficient can have different asymmetries. If LAC <1, the inequality in the distribution is caused by a large gap between a roughly equal number of small and large ports. By contrast, if LAC >1, the inequality is due to a small number of very busy ports, whereas the majority of ports experiences approximately equally low traffic. The base case is a lognormal distribution where LAC=1 regardless of the parameters μ and σ (Damgaard and Weiner, 2000). For the call distribution, we see in Figure 14.3(b) that LAC=1 is always included in the error bar which represents a jackknife estimate of the standard deviation. This observation gives additional credence to the lognormal distribution as a working hypothesis for ship calls.

³ Strictly speaking, this is the definition only for a continuous distribution F. For discrete distributions the Lorenz curve is a polygon instead of a smooth curve so that there is typically no point where the slope is exactly 1. However, one can generalize the definition so that it still works for the discrete distributions obtained from finite samples, see Damgaard and Weiner (2000) for details.

While the lognormal model is hence a generally promising candidate for the call distribution, it is not equally good for the degree distribution because LAC <1 in the later years (Figure 14.3b). These numbers confirm the results from Tables 14.3-5 showing that the lognormal distribution is in those years not the best model for the degrees. There is, however, no immediate contradiction between this finding and a lognormal call distribution. The relationship between the call and degree distributions is complicated: regardless of whether one or multiple voyages are made between the same two ports, it always only adds 1 to the ports' degrees. In other words, the call distribution is a measure of the weighted multigraph of voyages, whereas the degree distribution is derived from the unweighted network that forms a so-called simple graph. The collapse of multiple voyages into one unweighted link can conceivably change the distribution so that it appears to come from a completely different probabilistic law. In future analysis, we will explore with suitable null models (e.g. random graphs with a fixed weighted degree sequence; see Molloy and Reed, 1995) how the degree distribution changes when a heavy-tailed (in particular lognormal) multigraph is mapped to a simple graph and if this may explain our observations for the cargo ship data.

Conclusion

We have statistically analysed call and degree distributions of the cargo ship network extracted from snapshots of *Lloyd's Shipping Index* in 20 different years between 1890 and 2008. We have applied information-based model selection and statistical hypothesis tests to quantify the empirical distribution in a mathematically principled manner. For the call distribution a lognormal null model passes the Anderson-Darling goodness-of-fit test in all years and the Kolmogorov-Smirnov test in 19 out of 20 years at a 10 percent significance level. In some early years the Akaike weight is higher for truncated power laws than a lognormal distribution; in later years a Weibull distribution is preferred. However, Vuong's likelihood ratio test does not reject the lognormal null hypothesis at a 5 percent significance level for any tested year, neither compared with the maximum-likelihood truncated power law nor Weibull distribution. When the empirical call distribution is replaced by the degree distribution, the lognormal model is plausible in early years, but in later years Weibull distributions fit the data better.

As in all model selection problems, one should bear in mind that reality is, of course, more complex than any of the candidate models. In our case, it might be possible to reduce the AIC further by allowing more than two parameters, but we feel that two parameters are a good compromise between simplicity of the model and goodness-of-fit. With additional data it might become possible to analyse the dynamics of port calls in more detail, especially to test if Gibrat's law applies in our case, which could explain a lognormal call distribution. It may also become feasible to test the assumption of independence between ports and apply full likelihood methods to the dynamics of the network (Wiuf et al., 2006).

Acknowledgment

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. [313847] "World Seastems". M. T. G. thanks the European Commission for financial support (project number FP7-PEOPLE-2012-IEF 6-4564/2013).

References

Akaike H. (1974) A new look at the statistical model identification. *IEEE T. Automat. Contr.*, 19: 716-723.

Albert R., Barabási A.L. (1999) Emergence of scaling in random networks. Science, 286: 509-512.

Albert R., Barabási A.L. (2002) Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74: 47-97.

Anderson T.W., Darling D.A. (1952) Asymptotic theory of certain 'goodness of fit' criteria based on stochastic processes. *Ann. Math. Stat.*, 23: 193-212.

Bernhofen D.M., El-Sahli Z., Kneller R. (2013) *Estimating the Effects of the Container Revolution on World Trade*. Lund University Working Paper 2013:4, Department of Economics, School of Economics and Management.

Bhattacharya K., Mukherjee G., Saramäki J., Kaski K., Manna S.S. (2008) The international trade network: weighted network analysis and modelling. *J. Stat. Mech.*, 2008: P02002.

Broido A., Claffy K.C. (2001) Internet topology: connectivity of IP graphs. *Proc. SPIE*, 4526: 172-187.

Burnham K.P., Anderson D.R. (1998) Model Selection and Multimodel Inference. New York: Springer.

Clauset A., Shalizi C.R., Newman M.E.J. (2009) Power-law distributions in empirical data. *SIAM Rev.*, 51: 661-703.

Cox D.R., Reid N. (2004) A note on pseudolikelihood constructed from marginal densities. *Biometrika*, 91: 729-737.

Cullinane K.P.B., Khanna M. (2000) Economies of scale in large containerships: Optimal size and geographical implications. *Journal of Transport Geography*, 8(3): 181-195.

Damgaard C., Weiner J. (2000) Describing inequality in plant size or fecundity. *Ecology*, 91: 1139-1142.

Deng W.B., Long G., Wei L., Xu C. (2009) Worldwide marine transportation network: Efficiency and container throughput. *Chinese Physics Letters*, 26(11): 118901.

Ding D., Teo C.P. (2010) World container port throughput follows lognormal distribution. *Maritime Policy and Management*, 37: 401-426.

Ducruet C. (2013) Network diversity and maritime flows. *Journal of Transport Geography*, 30: 77-88.

Edwards A.M., Phillips R.A., Watkins N.W., Freeman M.P., Murphy E.J., Afanasyev V., Buldyrev S.V., da Luz M.G.E., Raposo E.P., Stanley H.E., Viswanathan G.M. (2007) Revisiting Lévy flight search patterns of wandering albatrosses, bumblebees and deer. *Nature*, 449: 1044-1047.

Eeckhout J. (2004) Gibrat's law for (all) cities. Am. Econ. Rev., 94: 1429-1451.

Gibrat R. (1931) Les Inégalités Economiques. Paris: Librairie du Recueil Sirey.

Giles D.E.A. (2004) Calculating a standard error for the Gini coefficient: some further results. *Oxford B. Econ. Stat.*, 66: 425-433.

Gini C. (1912) Variabilità e mutabilità. Bologna: P. Cuppini.

Gómez V., Kaltenbrunner A., López V. (2008) Statistical analysis of the social network and discussion threads in Slashdot. *Proc. 17th Int. Conf. World Wide Web*, pp. 645-654.

Hamilton D.T., Handcock M.S., Morris M. (2008) Degree distributions in sexual networks: a framework for evaluating evidence. *Sex. Transm. Dis.*, 35: 30-40.

He Y., Siganos G., Faloutsos M., Krishnamurthy S. (2007) A systematic framework for unearthing the missing links: measurements and impact. *Proc. 4th USENIX Conf. Networked Systems Design & Implementation*, p. 14.

Hu Y., Zhu D. (2009) Empirical analysis of the worldwide maritime transportation network. *Physica A*, 388(10): 2061-2071.

Kaluza P., Kölzsch A., Gastner M.T., Blasius, B. (2010) The complex network of global cargo ship movements. *Journal of the Royal Society Interface*, 7(48): 1093-1103.

Kolmogorov A. (1933) Sulla determinazione empirica di una legge di distribuzione. G. Ist. Ital. Attuari, 4: 83-91.

Lahererre J., Sornette J. (1998) Stretched exponential distributions in nature and economy: fat tails with characteristic scales. *European Physical Journal B*, 2: 525-539.

Lorenz M.O. (1905) Methods of measuring the concentration of wealth. *Publ. Am. Stat. Assoc.*, 9: 209-219.

McCalla R.J. (2004) From 'Anyport' to 'Superterminal'. In: Pinder D., Slack B. (Eds), *Shipping and Ports in the Twenty-first Century*, London: Routledge, pp. 123-142.

Molloy M., Reed B. (1995) A critical point for random graphs with a given degree sequence. *Random Struct. Algor.*, 6: 161-180.

Moore C., Ghoshal G., Newman M.E.J. (2006) Exact solutions for models of evolving networks with addition and deletion of nodes. *Physical Review E*, 74: 036121.

Newman M.E.J. (2010) Networks: An introduction. Oxford: Oxford University Press.

Prieto F., Sarabia J.M. (2011) Fitting the degree distribution of real-world networks. *Int. J. Complex Systems in Science*, 1: 129-133.

Rocha L.E.C., Liljeros F., Holme P. (2010) Information dynamics shape the sexual networks of Internet-mediated prostitution. *Proc. Nat. Acad. Sci.*, 107: 5706-5711.

Simon H.A. (1955) On a class of skew distribution functions. *Biometrika*, 42: 425-440.

Stephens P.A., Buskirk S.W., Hayward G.D., Martínez del Rio C. (2005) Information theory and hypothesis testing: a call for pluralism. *J. Appl. Ecol.*, 42: 4-12.

Stumpf M.P.H., Ingram P.J., Nouvel I., Wiuf C. (2005) Statistical model selection methods applied to biological networks. In: Priami C., Merelli E., Gonzalez P., Omicini A. (Eds.), Transactions on computational systems biology III, *Lecture notes in computer science*, 3737, Berlin: Springer, pp. 65-77.

Stumpf M.P.H., Porter M.A. (2012) Critical truths about power laws. Science, 335: 665-666.

Tadić B. (2002) Temporal fractal structures: Origin of power laws in the world-wide web. *Physica A*, 314: 278-283.

Todor A., Dobra A., Kahveci T. (2012) Uncertain interactions affect degree distribution of biological

networks. IEEE Int. Conf. Bioinformatics and Biomedicine, pp. 457-461.

UNCTAD (2013) Review of Maritime Transport. Geneva: United Nations.

Varin C., Reid N., Firth D. (2011) An overview of composite likelihood methods. Stat. Sinica, 21: 5-42.

Vuong Q.H. (1989) Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57: 307-333.

Willinger W., Alderson D., Doyle J.C. (2009) Mathematics and the Internet: a source of enormous confusion and great potential. *Not. Am. Math. Soc.*, 56: 586-599.

Wilson A.G. (1967) A statistical theory of spatial distribution models. Transport. Res., 1: 253-269.

Wiuf C., Brameier M., Hagberg O., Stumpf M.P.H. (2006) A likelihood approach to analysis of network data. *Proc. Nat. Acad. Sci*, 103: 7566-7570.

Distribution	Parameters	$\Pr(k), \ k=1, 2, 3, \rightleftharpoons$
Poisson (POIS)	$\lambda > 0$	$\frac{\lambda^k}{(e^{\lambda}-1)k!}$
geometric (GEOM)	$p \in (0,1)$	$p(1-p)^{k-1}$
power law (ZETA)	$\tau > 1$	$\left[\zeta(\tau)k^{\tau} ight]^{-1}$
Yule-Simon (YULE)	ho > 0	$\rho B(k, \rho+1)$
negative binomial (NEGB)	$p \in (0,1), r > 0$	$\frac{\Gamma(k+r)\boldsymbol{p}^{k}(1-\boldsymbol{p})^{r}}{k!\Gamma(r)[1-(1-\boldsymbol{p})^{r}]}$
truncated power law (TPOW)	$q \in (0,1), \ \tau > 1$	$\frac{q^k}{\operatorname{Li}_{\tau}(q)k^{\tau}}$
discrete lognormal (DLGN)	$\mu \in \mathbf{R}, \sigma > 0$	$\begin{cases} \frac{1}{2} + \frac{1}{2} \operatorname{erf}\left(-\frac{\mu}{\sqrt{2}\sigma}\right) & \text{if } k = 1\\ \frac{1}{2} \left[\operatorname{erf}\left(\frac{\ln k - \mu}{\sqrt{2}\sigma}\right) - \operatorname{erf}\left(\frac{\ln(k-1) - \mu}{\sqrt{2}\sigma}\right) \right] & \text{if } k = 2, 3, \dots \end{cases}$
discrete Weibull (DWEI)	$\boldsymbol{q} \in (0,1), \boldsymbol{\beta} > 0$	$\frac{1}{q} \left[q^{(k^{\beta})} - q^{((k+1)^{\beta})} \right]$

Table 14.1: The investigated probabilistic models

Year	Ports	POIS	GEOM	ZETA	YULE	NEGB	TPOW	DLGN	DWEI
1890	904	0	0	2.61×10 ⁻¹⁶	1.60×10 ⁻¹⁰	4.61×10-30	3.64×10 ⁻¹	6.10×10 ⁻¹	2.58×10-2
1910	1200	0	0	2.90×10 ⁻²⁵	4.56×10 ⁻¹⁷	4.76×10-31	9.37×10 ⁻¹	6.13×10 ⁻²	2.12×10-3
1915	992	0	0	1.51×10 ⁻²³	1.23×10 ⁻¹⁶	3.55×10 ⁻²³	7.70×10 ⁻¹	2.26×10 ⁻¹	3.94×10 ⁻³
1920	994	0	0	9.36×10 ⁻²⁰	2.07×10 ⁻¹³	2.42×10-30	8.83×10 ⁻¹	1.10×10 ⁻¹	6.22×10 ⁻³
1925	1205	0	0	2.51×10 ⁻²⁹	4.17×10 ⁻²⁰	6.09×10 ⁻³⁰	6.97×10 ⁻²	8.27×10 ⁻¹	1.03×10 ⁻¹
1930	1254	0	0	4.78×10-33	3.16×10 ⁻²³	5.61×10 ⁻²⁹	9.64×10-3	9.64×10 ⁻¹	2.66×10-2
1935	1282	0	0	4.29×10 ⁻³³	3.15×10 ⁻²²	5.01×10 ⁻²⁹	1.03×10 ⁻²	5.24×10 ⁻¹	4.66×10 ⁻¹
1940	1309	0	0	1.42×10 ⁻²³	2.62×10 ⁻¹⁵	1.50×10 ⁻⁴¹	9.79×10 ⁻¹	1.96×10 ⁻²	1.20×10-3
1946	1281	0	0	6.05×10 ⁻²⁶	7.30×10 ⁻¹⁷	9.37×10 ⁻³⁷	8.47×10 ⁻¹	8.47×10 ⁻²	6.80×10 ⁻²
1951	1321	0	0	1.91×10 ⁻³⁴	1.06×10 ⁻²³	1.42×10 ⁻²⁵	3.21×10 ⁻¹	5.67×10 ⁻¹	1.13×10 ⁻¹
1960	1541	0	0	1.49×10 ⁻⁴⁷	1.51×10 ⁻³⁴	2.66×10 ⁻²³	8.99×10 ⁻²	8.78×10 ⁻¹	3.22×10 ⁻²
1965	1554	0	0	2.58×10-59	1.07×10 ⁻⁴⁴	1.42×10 ⁻¹⁷	3.52×10-3	6.40×10 ⁻¹	3.56×10-1
1970	1512	0	0	5.70×10-57	2.29×10 ⁻⁴³	4.73×10 ⁻¹⁴	5.83×10-1	3.44×10 ⁻¹	7.25×10 ⁻²
1975	1610	0	0	6.08×10-60	1.08×10 ⁻⁴⁴	2.98×10 ⁻²²	1.15×10-4	4.91×10 ⁻¹	5.09×10 ⁻¹
1980	1637	0	0	1.40×10-76	1.40×10 ⁻⁵⁹	3.60×10 ⁻¹⁵	2.23×10-6	2.17×10 ⁻¹	7.83×10 ⁻¹
1985	1925	0	0	3.21×10 ⁻¹⁰⁷	1.22×10 ⁻⁸⁴	2.54×10 ⁻¹⁹	8.01×10 ⁻¹⁴	9.46×10 ⁻²	9.05×10 ⁻¹
1990	1903	0	0	1.14×10 ⁻¹⁰⁷	2.37×10 ⁻⁸⁵	1.12×10 ⁻¹⁸	6.86×10 ⁻¹⁴	4.29×10 ⁻²	9.57×10 ⁻¹
1995	1953	0	0	2.93×10 ⁻¹⁰⁷	6.45×10 ⁻⁸⁵	1.71×10 ⁻¹⁵	8.34×10 ⁻¹¹	2.09×10-2	9.79×10 ⁻¹
2000	2050	0	0	1.41×10 ⁻¹¹³	1.36×10 ⁻⁸⁹	2.15×10 ⁻¹⁵	4.87×10 ⁻¹¹	1.13×10 ⁻²	9.89×10 ⁻¹
2008	2157	0	0	1.82×10-96	3.97×10 ⁻⁷⁴	8.03×10 ⁻¹⁵	5.93×10 ⁻⁵	4.31×10 ⁻²	9.57×10 ⁻¹

Table 14.2: Akaike weights for the distribution of vessel calls

Values below 10-200 are rounded to zero. The largest Akaike weight in each year is highlighted in bold type.

Year	Ports	POIS	GEOM	ZETA	YULE	NEGB	TPOW	DLGN	DWEI
1890	895	0	3.05×10 ⁻¹³²	4.96×10-37	1.22×10 ⁻²⁵	8.99×10 ⁻⁸	3.62×10 ⁻⁴	6.21×10 ⁻¹	3.79×10 ⁻¹
1910	1186	0	3.66×10 ⁻¹⁸⁶	1.97×10 ⁻⁶⁹	5.02×10 ⁻⁵³	1.20×10-9	1.54×10 ⁻⁸	9.82×10 ⁻¹	1.80×10 ⁻²
1915	970	0	9.77×10 ⁻¹⁴⁰	4.07×10 ⁻⁶²	8.23×10 ⁻⁴⁹	1.48×10 ⁻⁴	1.87×10 ⁻⁴	6.04×10 ⁻¹	3.96×10 ⁻¹
1920	955	0	8.66×10 ⁻¹⁴⁷	4.53×10-60	1.01×10 ⁻⁴⁶	5.13×10-7	1.18×10-6	9.65×10-1	3.49×10 ⁻²
1925	1170	0	1.79×10 ⁻¹⁸⁸	3.85×10-67	3.23×10-52	5.22×10-5	3.74×10 ⁻⁴	1.05×10-1	8.94×10 ⁻¹
1930	1231	0	1.65×10 ⁻¹⁷⁹	6.11×10 ⁻⁸⁹	1.59×10-71	9.58×10-9	8.04×10-9	7.94×10 ⁻¹	2.06×10 ⁻¹
1935	1259	0	3.08×10 ⁻¹⁸¹	1.07×10-90	2.22×10-72	5.79×10 ⁻¹¹	5.18×10 ⁻¹¹	9.92×10 ⁻¹	8.34×10-3
1940	1273	0	3.41×10 ⁻¹⁹⁶	1.25×10 ⁻⁸⁵	6.69×10 ⁻⁶⁷	2.44×10 ⁻¹³	5.43×10 ⁻¹³	>9.99×10 ⁻¹	4.92×10 ⁻⁴
1946	1220	0	3.38×10 ⁻¹⁹⁷	5.90×10-57	1.13×10 ⁻⁴¹	1.13×10-6	1.51×10-3	5.11×10 ⁻¹	4.87×10 ⁻¹
1951	1294	0	3.85×10 ⁻¹⁸⁰	3.94×10 ⁻⁸⁸	2.04×10-69	1.19×10-9	1.29×10-9	9.94×10 ⁻¹	6.48×10 ⁻³
1960	1506	0	0	8.78×10 ⁻¹⁰⁵	1.09×10 ⁻⁸³	3.47×10-9	3.26×10-9	9.57×10-1	4.29×10 ⁻²
1965	1534	0	0	6.26×10 ⁻¹¹⁹	3.03×10-97	2.53×10-7	1.12×10-7	1.18×10 ⁻¹	8.82×10 ⁻¹
1970	1487	0	0	2.72×10 ⁻¹⁰⁹	2.20×10 ⁻⁸⁹	2.71×10-4	1.74×10 ⁻⁴	6.07×10 ⁻³	9.93×10-1
1975	1579	0	0	1.92×10 ⁻¹⁰³	3.00×10 ⁻⁸²	1.57×10-6	2.00×10-6	2.75×10-1	7.25×10 ⁻¹
1980	1591	0	9.20×10 ⁻¹⁹⁸	3.34×10 ⁻¹²⁴	2.44×10 ⁻¹⁰²	4.49×10-5	2.19×10-5	1.25×10-4	>9.99×10 ⁻¹
1985	1872	0	0	2.18×10 ⁻¹⁵⁰	1.57×10 ⁻¹²³	1.64×10 ⁻⁸	5.10×10-9	6.41×10 ⁻⁴	9.99×10-1
1990	1875	0	0	4.46×10 ⁻¹⁶²	1.44×10 ⁻¹³⁶	1.96×10-5	7.40×10 ⁻⁶	1.15×10 ⁻⁸	>9.99×10 ⁻¹
1995	1897	0	4.51×10 ⁻¹⁹¹	5.31×10 ⁻¹⁷⁶	3.15×10 ⁻¹⁴⁹	3.07×10 ⁻⁴	8.99×10 ⁻⁵	6.78×10 ⁻¹⁰	>9.99×10 ⁻¹
2000	1969	0	0	2.60×10-148	6.90×10 ⁻¹²¹	1.84×10-3	9.28×10 ⁻⁴	1.17×10-7	9.97×10 ⁻¹
2008	2007	0	8.85×10 ⁻¹⁸³	2.55×10 ⁻¹²⁷	1.78×10 ⁻⁹⁹	2.01×10 ⁻¹	1.61×10 ⁻¹	9.33×10 ⁻⁸	6.38×10 ⁻¹

Table 14.3: Akaike weights for the degree distribution

Values below 10-200 are rounded to zero. The largest Akaike weight in each year is highlighted in bold type.

		calls		degrees				
Year	TPOW-DLGN	TPOW-DWEI	DLGN-DWEI	TPOW-DLGN	TPOW-DWEI	DLGN-DWEI		
1890	0.83	0.49	0.20	0.09 (DLGN)	0.04 (DWEI)	0.67		
1910	0.35	0.18	0.18	0.01 (DLGN)	0.01 (DWEI)	0.11		
1915	0.66	0.22	0.09 (DLGN)	0.16	0.05 (DWEI)	0.84		
1920	0.47	0.26	0.28	0.03 (DLGN)	0.01 (DWEI)	0.16		
1925	0.52	0.94	0.34	0.35	0.09 (DWEI)	0.19		
1930	0.25	0.85	0.10	0.02 (DLGN)	0.00 (DWEI)	0.64		
1935	0.37	0.48	0.94	0.00 (DLGN)	0.00 (DWEI)	0.13		
1940	0.19	0.15	0.35	0.00 (DLGN)	0.00 (DWEI)	0.02 (DLGN)		
1946	0.59	0.66	0.93	0.30	0.18	0.97		
1951	0.88	0.84	0.37	0.01 (DLGN)	0.00 (DWEI)	0.10		
1960	0.64	0.86	0.10	0.02 (DLGN)	0.00 (DWEI)	0.33		
1965	0.40	0.49	0.68	0.11	0.00 (DWEI)	0.55		
1970	0.92	0.73	0.39	0.64	0.10	0.07 (DWEI)		
1975	0.19	0.22	0.98	0.14	0.02 (DWEI)	0.72		
1980	0.13	0.08 (DWEI)	0.30	0.84	0.06 (DWEI)	0.01 (DWEI)		
1985	0.01 (DLGN)	0.00 (DWEI)	0.22	0.26	0.00 (DWEI)	0.08 (DWEI)		
1990	0.03 (DLGN)	0.01 (DWEI)	0.18	0.56	0.10	0.00 (DWEI)		
1995	0.12	0.03 (DWEI)	0.09 (DWEI)	0.29	0.17	0.00 (DWEI)		
2000	0.16	0.05 (DWEI)	0.06 (DWEI)	0.39	0.32	0.00 (DWEI)		
2008	0.70	0.54	0.23	0.19	0.86	0.00 (DWEI)		

Table 14.4: The p-values for Vuong's likelihood ratio test of the maximum-likelihood distributions

We show all pairwise comparisons between truncated power law, discrete lognormal and discrete Weibull distributions. We highlight p-values smaller than 0.1 in bold type and add in parentheses the more likely distribution. Values below 5×10-3 are rounded to zero.

	calls						degrees									
Year	YU	JLE	TP	OW	DL	GN	DV	VEI	YU	JLE	TPO	ЭW	DLGN		DV	VEI
	p _{KS}	$p_{_{\!\mathrm{AD}}}$	p _{KS}	$\pmb{p}_{\!\scriptscriptstyle{\mathrm{AD}}}$	p _{KS}	$\pmb{p}_{\!\scriptscriptstyle{\mathrm{AD}}}$	p _{KS}	$\pmb{p}_{\!\scriptscriptstyle{\mathrm{AD}}}$								
1890	0.00	0.06	0.22	0.37	0.98	0.78	0.04	0.38	0.00	0.02	0.00	0.11	0.93	0.63	0.34	0.47
1910	0.00	0.04	0.22	0.27	0.72	0.46	0.01	0.18	0.00	0.00	0.00	0.03	0.07	0.17	0.00	0.10
1915	0.00	0.04	0.24	0.21	0.81	0.69	0.01	0.20	0.00	0.00	0.00	0.06	0.50	0.43	0.03	0.31
1920	0.00	0.05	0.51	0.58	0.95	0.80	0.04	0.37	0.00	0.00	0.00	0.03	0.18	0.37	0.01	0.11
1925	0.00	0.03	0.04	0.06	0.96	0.82	0.07	0.40	0.00	0.00	0.00	0.07	0.78	0.44	0.29	0.55
1930	0.00	0.03	0.02	0.03	0.84	0.70	0.01	0.26	0.00	0.00	0.00	0.02	0.69	0.55	0.02	0.28
1935	0.00	0.03	0.00	0.04	0.91	0.77	0.42	0.49	0.00	0.00	0.00	0.03	0.35	0.51	0.00	0.10
1940	0.00	0.04	0.12	0.40	0.61	0.64	0.03	0.17	0.00	0.00	0.00	0.02	0.11	0.20	0.00	0.06
1946	0.00	0.03	0.02	0.02	0.31	0.22	0.03	0.10	0.00	0.00	0.00	0.06	0.30	0.30	0.01	0.24
1951	0.00	0.03	0.03	0.06	0.78	0.68	0.08	0.33	0.00	0.00	0.00	0.03	0.11	0.45	0.00	0.08
1960	0.00	0.01	0.11	0.04	0.55	0.65	0.01	0.25	0.00	0.00	0.00	0.02	0.03	0.22	0.00	0.09
1965	0.00	0.00	0.02	0.02	0.88	0.49	0.08	0.48	0.00	0.00	0.00	0.03	0.49	0.29	0.06	0.42
1970	0.00	0.00	0.07	0.03	0.84	0.54	0.00	0.23	0.00	0.00	0.00	0.07	0.23	0.25	0.09	0.53
1975	0.00	0.00	0.00	0.02	0.87	0.63	0.37	0.47	0.00	0.00	0.00	0.04	0.22	0.34	0.02	0.31
1980	0.00	0.00	0.00	0.01	0.60	0.52	0.20	0.75	0.00	0.00	0.00	0.04	0.23	0.14	0.21	0.75
1985	0.00	0.00	0.00	0.00	0.79	0.88	0.94	0.54	0.00	0.00	0.00	0.01	0.41	0.18	0.85	0.64
1990	0.00	0.00	0.00	0.00	0.26	0.46	0.26	0.24	0.00	0.00	0.00	0.00	0.02	0.09	0.34	0.58
1995	0.00	0.00	0.00	0.00	0.06	0.25	0.02	0.40	0.00	0.00	0.01	0.01	0.01	0.06	0.10	0.42
2000	0.00	0.00	0.00	0.00	0.60	0.25	0.56	0.16	0.00	0.00	0.00	0.00	0.09	0.07	0.16	0.34
2008	0.00	0.00	0.04	0.00	0.31	0.11	0.04	0.10	0.00	0.00	0.05	0.00	0.03	0.06	0.00	0.18

Table 14.5: The p-values for the Kolmogorov-Smirnov test and the Anderson-Darling test

Values below 5×10-3 are rounded to zero. We highlight those distributions in bold type where the null hypothesis (i.e. that the data is generated by the model) is not rejected at a 10% significance level.



Figure 14.1: Observed complementary cumulative call distribution function in (a) 1910, (b) 2000 together with the maximum-likelihood Yule-Simon, truncated power law, discrete lognormal and discrete Weibull distributions
The one-parameter Yule-Simon distribution fits the data far worse than any of the two-parameter alternatives. In 1910 the truncated power law has the highest
likelihood among all models. In 2000 the Weibull distribution fits best (Table 14.2). Inset in (b): The fraction of ports with no more than ten ports has dropped from around 75% to approximately 68% during the 1960s. This subtle, but statistically significant decrease is responsible for more curvature in later years on the left-hand side of the observed data. Error bars are jackknife estimates.



Figure 14.2: Observed complementary cumulative degree distribution function in (a) 1910, (b) 2000 together with the maximum-likelihood Yule-Simon, truncated power law, discrete lognormal and discrete Weibull distributions
The highest Akaike weight is achieved by the (a) lognormal, (b) Weibull distribution. Inset in (b): The fraction of ports with degree ≤10 does not show a clear trend.



Figure 14.3: (a) Lorenz curves for the call distributions in three exemplary years (black: 1890, dark grey: 1960, light grey: 2008)

The Gini coefficient measures inequality as twice the area between the Lorenz curve and the line of equality (dashed diagonal). The Lorenz asymmetry coefficient (LAC) measures whether curves have a slope of 1 at the intersection with the dotted antidiagonal line (LAC<1 if the slope equals 1 below the antidiagonal, LAC>1 if the slope equals 1 above it). For the three curves in the figure the Gini coefficient has decreased over time, but they all have a slope close to 1 at the symmetry axis. (b) Gini coefficients and LACs for all twenty years for which we have data. Error bars represent jackknife estimates of the standard deviation. The Gini coefficient for the call and degree distribution both show a slightly decreasing tendency that is statistically significant, similarly the LAC for the degrees. The LAC for the call distribution, however, shows no significant deviation from 1 in any of the investigated years