


Power Laws in Socio-Economics

Jan Schulz and Jan David Weber

Working Paper No. 203

March 2025



Bamberg Economic Research Group
Bamberg University
Feldkirchenstraße 21
D-96052 Bamberg
Telefax: (0951) 863 5547
Telephone: (0951) 863 2687
felix.stuebben@uni-bamberg.de
<http://www.uni-bamberg.de/vwl/forschung/berg/>

ISBN 978-3-949224-24-9

Redaktion:

Dr. Felix Stübben*

* felix.stuebben@uni-bamberg.de

Power Laws in Socio-Economics

Jan Schulz and Jan David Weber

March 2025

Power laws are pervasive in economics and social sciences, particularly in the upper tails of distributions such as wealth, income, firm size, and city populations. Their scale-free property makes them a universal framework to understand phenomena spanning several orders of magnitude. This chapter explores their mathematical and often counterintuitive statistical properties, empirical evidence, and the stochastic processes that generate them. Emphasis is placed on their universal applicability, particularly to firm size, wealth, and income distributions and their potential to address the pressing issues of our time.

Keywords: Distribution, Growth Processes, Extreme Values, Concentration, Inequality

Jan Schulz: Economics Institute, University of Bamberg. Jan David Weber: Institute for Socio-Economics, University of Duisburg-Essen. We thank Jan Lorenz, Leonhard Ipsen and Christophe van Langenhove for their helpful comments and suggestions. This work was supported by the grant “e-lernen” by the German Ministry for Education and Science. We also thank Vera Petkova for her able research assistance. This chapter was prepared for the Encyclopedia of Evolutionary and Institutional Economics.

1. Introduction

“[I]t is the so-called ‘normal’ distribution that is abnormal in the sense that it has many unique properties not possessed by any other.”
– E. T. Jaynes (2003, p. 729)

Power laws or Pareto distributions are ubiquitous across various domains. Although power laws appear in natural phenomena e.g. in the intensity frequency of earthquakes, they also characterize the upper tails of many economic and social indicators of interest, such as income and wealth distributions or the distribution of firm and city sizes. While quantitative economists and statisticians are usually trained and equipped to deal with ‘normal’ or Gaussian distributions, the unique properties of this distributional regularity bear important implications for economic and social theory as well. The case of human height that is very well approximated by a Gaussian distribution illustrates the intuition here well: While even knowing the height of one human being is instructive for the typical height of human beings, one wealth level that is, for the largest values, characterised by a power law, does tell us little about the typical wealth level. To put it another way: If human height were distributed like wealth, it would be common to observe people over 10,000 meters high. While human height has a so-called characteristic scale, distributions characterised by power laws lack this scale. Statistical inference based on characteristic scales might become spurious, as discussed in this chapter.

As a direct corollary of the lack of a characteristic scale, power laws describe potentially extremely concentrated phenomena. This extreme concentration that a power law upper tail¹ implies carries important implications for policy and institutional design and might aid in identifying promising levers to influence aggregate outcomes. At the same time, their apparent universality across institutional contexts challenges some central tenets of institutional economics and might point to fundamental limits for the impact of institutional design. We will revisit these issues in turn.

The remainder of this chapter is organised as follows. We discuss the universality of power laws (Section 2), followed by the basic mathematical concepts and characteristics of the power law distribution (Section 3). We continue discussing the estimation procedures most commonly employed in the literature and elaborate on their potential pitfalls (Section 4). In the subsequent Section 5, we introduce several theoretical mechanisms that generate power laws and discuss, how they reflect real-world economic processes. Section 5). We finally situate the preceding mathematical and formal discussion in the context of socio-

¹The upper tail refers to the highest values of the distribution.

economics and highlight the implications of power laws in economic theory and analysis (Section 6).

2. Universality of Power Laws

This section will discuss empirical evidence on power laws in economics and the social sciences without claiming completeness. A power law or Pareto upper tail was first established by Pareto for income distributions in England, several Italian cities, several German states, Paris, and Peru (Persky 1992). Subsequent research has indeed confirmed the universality of an upper Pareto tail (Silva and Yakovenko 2004; Schneider 2199; Tao et al. 2019; Ludwig and Yakovenko 2022). Similarly, for wealth distributions, we also find an upper power law tail (cf. Appendix G of Schulz and Milaković 2023, for a review, and Disslbacher et al. (2020) for an extensive discussion). The tail exponents are consistently lower for wealth than for income distributions, reflecting the fact that wealth distributions as a whole are usually more unequal. More generally, consumption expenditures, labour income, capital income, and wealth distributions all exhibit power law tails, with tail exponents in decreasing order (Gaillard et al. 2023). Conventional neoclassical models typically have trouble replicating this finding (Stachurski and Toda 2019), which provides an interesting opportunity for further research to identify the determinants of this inequality ranking.

Apart from power laws situated broadly in the field of household finance, they also characterise several other economic and social phenomena. Firm size distributions, measured by sales or employment, and the distribution of city sizes typically follow a special case of a particularly heterogeneous power law called Zipf’s law with a tail exponent of unity (Gabaix 1999, 2011; Axtell 2001; Córdoba 2008a,b; Schulz and Mayerhoffer 2021). Zipf’s law in firm and city sizes imply an extreme concentration of economic activity and population that far exceeds the concentration of income and wealth among households. We will revisit a candidate explanation for this phenomenon based on scale-dependence in stochastically multiplicative growth in section 5.

In social science more broadly, so-called scale-free networks for which their measures of centrality follow a power law distribution also are seemingly ubiquitous (Barabási and Albert 1999; Barabási 2009): Power laws are claimed for the degree distribution of production networks (Acemoglu et al. 2012), the follower counts in social networks (Aparicio, Villazón-Terrazas, and Álvarez 2015), citations in citation networks (Zhong and Liang 2024) and centrality in the world trade network (Serrano and Boguná 2003). Yet, there exists an active discussion of how ubiquitous scale-free networks really are (Broido and Clauset 2019; Artico et al. 2020). In any case, they often provide a good first-order approximation for the network phenomenon of interest, even though their precise nature and generating process might be more complicated than the canonical “preferential attachment” model

of Barabási and Albert (1999) suggests. We revisit generating mechanisms in section 5.

Power laws also appear in financial economics, particularly in the tails of empirical return distributions (Gabaix et al. 2003; Franke and Westerhoff 2016) which might hint at the importance of self-reinforcing “chartalist” trading rules and herding behaviour. Finally, they appear in firms’ productivity distributions (Yang et al. 2025) and show that there does not exist a “representative” or Gaussian shock to firms as assumed e.g. by Gabaix (2011) and criticised in Dosi et al. (2019).²

3. Basic Concepts of Power Laws

While the apparent universality of power laws is interesting in its own right, this empirical regularity also has important consequences for statistical inference and our understanding of concentration and economic inequality. This subsection will review some of the technical details of power law distributions based on their formal definition. For a non-technical introduction, please skip to the next subsection.

3.1. PDF and CCDF of a Power Law

Power laws are distributional regularities of a random variable w that can be expressed as the cumulative complementary distribution function (CCDF)

$$(1) \quad P(W > w) \sim w^{-\alpha}, \quad \text{where } \alpha > 0$$

or, equivalently, as a probability density function (PDF)

$$(2) \quad f(w) \sim w^{-(\alpha+1)}.$$

The so-called tail index α is an inverse measure of inequality - the lower α , the higher the degree of inequality in the system. By taking logs, these equations become linear:

$$(3) \quad \log P(W > w) \sim -\alpha \log w + c$$

$$(4) \quad \log f(w) \sim -(\alpha + 1) \log w + c'$$

Given the notoriously challenging formal verification of power laws, this linear relationship of the PDF, CCDF or rank-size relationship (a variant of the CCDF) for a double-logarithmic scale is often used for visual identification. The α th moment of the distribution

²In particular, the results of Gabaix (2011) rely on a winsorizing procedure that eliminates extreme productivity shocks. As Schulz-Gebhard (2023, chp. 2) demonstrates formally, this should decrease the explanatory power of his model rather than being a necessary condition to generate “granularity”.

does not exist, which implies for the empirically relevant case of $1 \leq \alpha < 2$ that the second moment (the variance or standard deviation) is undefined, highlighting the fat-tailed nature of the distribution (Mitzenmacher 2004).

3.2. Properties of Power Laws

Power laws describe distributions that are extremely concentrated for the empirically relevant cases. The inequality of the distribution is measured by the so-called tail exponent α . The higher α , the more equal the distribution. For a tail exponent of unity, we recover what is called Zipf's law which might aid in illustrating the extreme degree of concentration a power law implies: For Zipf's law, the size of an entity is directly (inversely) proportional to its rank. This implies that the second-largest observation is approximately half the size of the maximum observation, the third-largest observation is approximately a third of the size of the top observation, and so on. For Zipf's law, almost all mass is thus concentrated at the very top.

This concentration is perhaps best exemplified by the non-existence of the second moment of the distribution for the empirically relevant range of the tail exponent of $1 \leq \alpha < 2$. In practice, this implies that the value for a sample's standard deviation or variance does not converge if we continue to sample from a population characterised by such a power law. Figure 1 illustrates this issue by comparing the standard deviations of samples from a normal (left) and power law (right panel) distribution for an increasing sample size $N = 1, \dots, 1000000$.




FIGURE 1. Standard deviation of samples with increasing sample size $N = 1, \dots, 1000000$ samples from a standard normal distribution and a power law with $\alpha = 1.01$.

While the standard deviation for the normal distribution quickly and smoothly converges to the population value of 1, the calculated standard deviation for the power law never converges and exhibits discrete large “jumps” due to the inclusion of extremely large values that drive the standard deviation up (compare the scale on the vertical axes).

The second moment of a variable sampled from a power law distribution is thus simply not informative, as it strongly depends on a few extreme observations.




FIGURE 2. Violin plot for three measures of statistical association (the Pearson correlation coefficient r , Spearman's ρ and Kendall's τ) for the maximum and mean sampled from a power law with differing α for $1 \leq \alpha \leq 2$ for 100 samples each with size 10,000.

Even though the mean as the first moment exists for $\alpha > 1$, measures of central tendency are still heavily dependent on the maximum observation that is sampled. Figure 2 depicts three measures of statistical association (the Pearson correlation coefficient r , Spearman's ρ and Kendall's τ) for the maximum and mean sampled from a power law with differing α for $1 \leq \alpha \leq 2$ for 100 samples each with size 10,000. As is immediately apparent, especially for the usually observed tail exponents below 1.3 in social systems, the mean strongly depends on the maximum observation that is sampled, even though sample sizes are comparatively large. Thus, again, measures of central tendency and totals depend on the maximum observation that is sampled, and sample means are therefore unreliable or unrepresentative of the actual population mean for the empirically relevant sample sizes.

For these sample sizes, the infinite variance of the population distribution also leads to a downward bias in common inequality measures, namely top shares and the Gini coefficient (Taleb and Douady 2015; Fontanari, Taleb, and Cirillo 2018). This is primarily because the asymptotic distribution of their (non-parametric) estimators is non-normal and non-symmetric for the case of a finite mean but infinite variance, i.e. a tail exponent α between 1 and 2. Additionally, top shares are affected by sample and population size and totals and thus might give the impression of changes in inequality but merely reflect changes in size or aggregate levels. Thus, they violate two commonly agreed-upon axioms, namely scale independence and the population principle, for inequality metrics, as, e.g. discussed in Cowell (2000).

This sensitivity to extreme observations extends to other commonly employed measures such as the Theil index (Theil 1967).³ As (Cowell and Flachaire 2007) demonstrate for several standard inequality indices, semi-parametric methods that rely on estimating the parameters of the power law tail before applying the metric should be preferred for finite

³Depending on the application, the Theil index might, of course, have other advantages such as straightforward decomposability to analyse the within-group and between-group inequality and highlight structural differences (Conceicao and Ferreira 2000).

samples. The top shares in a power law distribution also exhibit super-additivity, i.e., in the presence of fat tails, a weighted average of samples shows a lower top share than these samples evaluated together. Thus, for example, an average measure across European countries would show lower inequality than a top share for Europe as a whole (Taleb and Douady 2015). In line with Cowell and Flachaire (2007), both Taleb and Douady (2015) and Fontanari, Taleb, and Cirillo (2018) provide some guidance on mitigating these problems and recommend using parametric methods for estimating these inequality metrics that practitioners should use. Yet, as we discuss in section 6, inference of tail exponents from survey data is often itself biased due to differential reporting and sampling biases.

4. Estimation and Measurement

Identifying power laws and estimating their parameters is fraught with challenges. To identify a power law in empirical data, practitioners often exploit the linearity of a power law on a double-logarithmic scale. Yet, this method often leads to false positives, as many distributions exhibit an approximately linear upper tail when plotted on a double-logarithmic scale. This issue was eloquently summarized in Mar’s law (Akin 2003): “Everything is linear if plotted log-log with a fat magic marker.”




FIGURE 3. Rank-size plots for a standard normal distribution, a log-normal distribution with $\sigma = 2$ and a power law distribution with $\alpha = 1.01$ on a double-logarithmic scale. The minimum is normalised to unity for all three cases.

Figure 3 illustrates this issue. Depicted are the rank-size relationships of the upper tails (10,000 observations) of samples from standard normal distribution, a high-variance log-normal distribution ($\sigma = 2$) and a power law with tail exponent close to unity with $\alpha = 1.01$. While the normal distribution can be clearly distinguished from the power law, the high-variance log-normal displays approximately linear behaviour with values spanning more than two orders of magnitude, which is usually considered to be a minimum requirement

for identifying a power law (Stumpf and Porter 2012). Thus, establishing if a power law is indeed present is a challenging task that naturally also hinges on the pre-analytic vision of the researcher and what they consider plausible as a generating mechanism. There also exist several more statistically grounded approaches to distinguish power laws from other distributions, yet the literature is far from a consensus as to which method is most appropriate (Broido and Clauset 2019; Artico et al. 2020).

Even if the existence of a power law can be established, estimating its parameters is similarly challenging. For the tail exponent, two methods should yield unbiased and asymptotically normal estimates: Estimating it using the corresponding Maximum Likelihood estimator, the so-called “Hill estimator” (Haeusler and Teugels 1985) or exploiting the linearity of the power law on the double-logarithmic scale (cf. eq. (??) and eq. (??)) and estimating the exponent using an Ordinary Least Squares regression (Gabaix and Ibragimov 2011). Yet, estimating the tail exponent also requires determining the minimum observation from which onwards the power law applies – and the estimate of the tail exponent is very sensitive to misspecification of this minimum level caused by noisy estimates (Eckerstorfer et al. 2016). Methods to determine the threshold observation typically rely on estimating the tail exponent for increasing observations going from high to low and examining the goodness of fit. This is either determined by visually inspecting the stabilisation of the estimated tail exponent for different thresholds (Castaldi and Milaković 2007) or by using a statistical test such as the Kolmogorov-Smirnov test (Clauset, Shalizi, and Newman 2009) or the Cramér–von Mises test (Eckerstorfer et al. 2016). Yet, similarly to distinguishing a power law from other distributional types, there is no consensus in the literature as to determining the threshold observation yet, and practitioners should thus use different methods in agreement.

Finally, the estimation of the tail exponent is also affected by differential reporting and sampling as an issue of virtually all types of survey data. If all observations are reduced proportionally, the Hill estimator remains unbiased. Yet, if specific groups are not represented at all (“undersampling”) or their reported values are downwards biased (“underreporting”), the estimator becomes biased (Schulz and Milaković 2023). In the empirically relevant case of underreporting or undersampling of the (super) rich in wealth and income surveys, the estimator is upward biased and underestimates the actual degree of inequality in the population. Thus, estimating power laws and gauging the reliability of estimates requires a high degree of domain knowledge, arguably to a higher extent than for other distributional types, as estimates crucially depend on a few top observations.

5. Generating Mechanisms

Modelling underlying generative processes is a fundamental task in economics. A model that precisely matches empirical distributions but lacks a plausible mechanism may be

overfitted and fail to generalise beyond observed data (Arthur 2023). In contrast, models that treat outcomes probabilistically often yield more robust insights, particularly when dealing with power laws due to their distinctive statistical properties. Rather than imposing power laws directly through an overly rigid structure, it is crucial to identify the minimal set of assumptions that naturally produces the same statistical patterns as observed in empirical data. Distinguishing between essential and incidental features of observed distributions helps both theoretical and empirical research focus on mechanisms that genuinely reflect the underlying generative processes.

The universality of power laws suggests a common mechanism explaining its emergence in various domains. In the literature, two main types of processes are discussed to generate power laws: Firstly, stochastically multiplicative growth and secondly exponential growth of an exponentially distributed variable. We discuss both explanatory approaches in turn. The exposition is based on income and wealth distributions, yet, the discussed mechanisms are sufficiently general to also apply to other domains.⁴

5.1. Multiplicative Growth

Stochastically multiplicative growth is capturing a rich-get-richer phenomenon that is well known across various domains with differing labels such as the “Matthew effect”, “cumulative causation”, “preferential attachment” or “path-dependent growth” (Weber and Schulz 2024).⁵ The basic idea is that wealth, income or another variable of interest grows multiplicatively, i.e., the absolute wealth gain between period $t + 1$ and t depends on the wealth in t . We discuss this stochastically multiplicative growth using a pedagogical example below. Champernowne (1953) considers the evolution of normalised wealth or wealth shares, W_t^i , governed by:

$$(5) \quad W_{t+1}^i = \gamma_{t+1}^i W_t^i$$

Here, γ is an independent and identically distributed random variable drawn from a

⁴For example, as Franke and Westerhoff (2016) show, the fat-tail property of empirical return distribution can be explained in a chartist-fundamentalist framework that can be interpreted as a self-reinforcing mechanism with a stabilising force. For the second explanatory attempt, consider Coad (2010) which derives Zipf’s law in the upper tail of empirical firm size distribution from their exponential firm age distribution and multiplicatively growing firm sizes.

⁵The emergence of a steady-state power law requires a stabilising force such as entry and exit as in Gabaix et al. (2016) or a reflecting boundary like (small) fixed entry costs preventing very small entities from operating Gabaix (1999, 2009) for the distribution to not degenerate into a log-normal with variance increasing in time without bounds (Gabaix 2009). The precise definition of these fixed costs or entry and exit can strongly affect the aggregate outcome, though. As e.g. Schulz and Mayerhoffer (2021) show, if the entry of firms in a stochastically multiplicative growth model is localised and reflects the average level only of a specific market segment, this might affect the tail exponent of the resulting firm size distribution. They also demonstrate that Zipf’s law is only consistent with a situation where market entry conditions are not localised.

distribution with finite variance. He shows that the steady-state distribution of the upper tail converges to a power law with a tail exponent α satisfying:

$$(6) \quad \mathbb{E}[\gamma^\alpha] = 1$$

Assuming that γ follows a Geometric Brownian Motion (GBM) by

$$(7) \quad \gamma_t^i = \exp\left(\mu - \frac{\sigma^2}{2} + \sigma \epsilon_t^i\right), \quad \epsilon_t^i \sim \mathcal{N}(0, 1)$$

implies that the tail exponent α is given by:

$$(8) \quad \alpha = 1 - \frac{2\mu}{\sigma^2}$$

This allows for a straightforward interpretation of the tail exponent concerning the two parameters of the GBM μ and σ (Gabaix 2009). The volatility in the growth process defined by σ unambiguously and perhaps unsurprisingly increases tail inequality. The case of μ is more interesting. Recall that W is normalized such that $\sum_i W^i = 1$. This implies that for scale-independent growth, all expected growth rates are the same and thus, shares are expected to stay constant, implying $\mu = 0$. This condition of growth rates being independent of size is called Gibrat's law or the law of proportionate effect. It directly implies Zipf's law, i.e., $\mu = 0$ implies (for our setup) that $\alpha = 1$ (Gabaix 1999). We observe Zipf's law for firm size and city size distributions but typically not for income and wealth distributions. This suggests that firm and city growth processes are independent of size, while for income and wealth distributions, income and wealth levels located in the upper tail of the distribution grow slower than the average, leading to $\mu < 0$ to explain $\alpha > 1$. One plausible explanation for this phenomenon might be the role of entrants: For cities, there is no new entry and firms will need some kind of minimum level of assets to enter the market. By contrast, entrants for income and wealth distributions will tend to grow from almost nothing with thus huge initial growth rates. This increases the average growth rate, leaving the incumbents in the tail with lower-than-average growth for their portfolios. Thus, scale dependence might shape the differences in the distributional outcomes we observe for different economic variables of interest.

A more recent literature both in mainstream and heterodox outlets has emphasised the role of scale- and type-dependence in the recent surge in income and wealth inequality using more involved models Gabaix et al. (2016). Here, type-dependence refers to specific investor types gaining higher returns, e.g. due to more sophisticated investment strategies, while scale-dependence refers to higher rates of returns for higher levels of wealth, e.g., due to minimum investment inlays. Empirical findings generally support the relevance and existence of differential rates of return according to type or size of the wealth portfolio

(Bach, Calvet, and Sodini 2020; Fagereng et al. 2020; Ederer, Mayerhofer, and Rehm 2021; Lagemann and Rehm 2025). Yet, empirically, it is very hard to distinguish if an investor is rich with high rates of return because they are sophisticated or they enjoy high rates of return due to them being rich (Schulz and Milaković 2023). Similarly, estimation biases for power laws themselves (cf. section 6) also confound estimates of scale- and type-dependence (Schulz and Milaković 2023). Also, there is no agreement about the causes of scale- and type-dependence yet. Given the relevance of this topic to explain the historical evolution of inequality, this is an exciting area for further research.

5.2. Age and Hierarchy

For the second type of explanatory attempt, we need exponential growth for an exponentially distributed variable. The literature has invoked two variants of this process: (Jones 2015) assumes exponential growth of wealth due to a constant multiplicative rate of return for the portfolio. He also assumes that agents face a constant probability of death in each period, which directly implies that age is exponentially distributed. The result is exponential growth over an exponentially distributed amount of time, letting wealth levels be distributed according to a power law. The second, Simon-Lydall approach is based on hierarchy within firms. If each person situated in the hierarchy has a specific number of subordinates or “span of control” (except for the bottom tier), the number of direct and indirect subordinates grows exponentially in hierarchical level (Simon 1957; Lydall 1959). If income grows at a given exponential rate with the number of subordinates, the income distribution exhibits a power law tail (Fix 2018a, 2019).

Which type of explanatory attempt is more appropriate is naturally strongly domain-dependent. While the Champernowne approach is undoubtedly more general and applies to a large host of domains, the Simon-Lydall hierarchy argument and the approach based on exponential age distributions arguably have more testable implications concerning their underlying assumptions and, consequently, provide more potential to affect the overall level of inequality.

6. Implications of Power Laws

As we have illustrated for the case of the cross-sectional standard deviation, aggregate measures of central tendency (especially the mean), ratios, or measures of dispersion are heavily dependent on the maximum observation(s). This bears important implications both for the measurement of inequality and the aggregate behavior of the system.

Power laws also pose problems for the measurement of inequality. Especially for wealth (and, to a lesser extent, income) distributions, the problem of the “missing rich” exists. The super-rich tend not to respond to survey requests, underestimating overall

inequality and total wealth levels (Eckerstorfer et al. 2016). This issue is compounded by the extreme concentration a power law distribution implies: Schulz and Milaković (2023) demonstrate that accounting for these missing rich can increase estimated total wealth by almost one order of magnitude, i.e., a factor of 10. Even if sampling is completely unbiased and participation probability is homogeneous in the population, the estimated totals and means based on the sample crucially depend on the maximum observation that is included in the sample (see Figure 2). Thus, if researchers have reason to suspect a power law distribution with a low tail exponent describing the quantity of interest, the usual feasible sample sizes typically do not allow inference about measures of central tendency or dispersion.

Apart from these statistical and epistemological issues, the existence of a power law also determines how micro-level entities affect the behaviour of the aggregate system. If either the size of firms or the centrality within the production network of a producer follows a power law with a tail exponent close to one, idiosyncratic shocks affecting only one entity at a time do not necessarily cancel out in the aggregate anymore (Gabaix 2011; Acemoglu et al. 2012), in contrast to the famous “second Lucas critique” claiming that idiosyncratic shocks should not matter for aggregate fluctuations of GDP (Lucas 1977). This is due to the fact that the assumptions of the central limit theorem are violated for a power law with infinite variance (Gabaix 2011). As he demonstrates, the idiosyncratic destinies of the 100 largest firms in the US appear to statistically account for almost a third of GDP fluctuations. In the jargon of this literature, these firms are thus “granular”. This finding has been corroborated for many other countries and regions (Blanco-Arroyo et al. 2018; Carvalho and Grassi 2019; Miranda-Pinto and Shen 2019; Silva and Da Silva 2020; Shin, Li, and Lee 2024) and for context other than firm-level shocks and GDP (Esquierro and Da Silva 2024; Freund and Pierola 2015; Ipsen, Aminian, and Schulz 2023; Kovalenko, Schnabel, and Stüber 2022). Yet, there also exist criticisms of granularity, Wagner and Weche (2020) do not find evidence for granularity to exist in Germany, and (Dosi et al. 2019) criticise the fundamentally supply-side approach based on productivity shocks employed in (Gabaix 2011), but finding support for demand-side granularity with demand shocks. Nevertheless, the vast majority of studies point to the existence and relevance of the granular channel - with important implications for macroprudential and industrial policy: Granularity suggests that this usually high-dimensional problem can be addressed by monitoring and targeting only very few large and strongly interconnected entities with potentially large spill-over effects for the aggregate economy.

Power laws also challenge conventional wisdom in another domain as well, namely marginal productivity theory. As Tao et al. (2019) and Ludwig and Yakovenko (2022) demonstrate, empirical income distributions are seemingly universally characterised by a two-class structure with an exponential bulk characterising the vast majority of 95 - 99 % lower

incomes, while the upper tail follows a power law tail. This remarkably robust finding points to a general societal class division, where the upper tail is dominated by capital income that grows multiplicatively due to financial markets' constant revaluation of assets. By contrast, the exponential bulk consists primarily of wage income that is caused by an additive growth process that results from a form of zero-sum competition, in which either the total bargaining power of workers is conserved (Dos Santos 2017) or workers value their wage relative to some societal average (Ragab 2014). The empirical two-class structure thus suggests that income inequality is driven by stock valuation and thus the share of income subject to multiplicative returns, as Silva and Yakovenko (2004) show for the Gini coefficient of capital income. It is worth emphasizing that differences in productivity do not primarily drive income inequality but that the whole income distribution is shaped by differences in assets people hold and their valuation. This finding is thus a direct challenge to marginal productivity theory and its development in human capital theory (Fix 2018b).

While the power law upper tail of firm size and centrality distributions helps identify levers with impact for policy, the apparent universality of this distributional regularity in income and wealth distributions also points to the limitations of policy regarding distributional outcomes. Indeed, claiming universality has been frequently criticized as being “naturalizing inequality” (Cook 2020). If this two-class structure is indeed as universal across welfare state regimes, varieties of capitalism or bargaining institutions, it appears that policy and institutional changes cannot affect income inequality. We argue that this argument is partially misguided: While the universal two-class structure appears to be a fundamental property of a class society (Shaikh, Jacobo et al. 2020), policy and institutional structure can nevertheless affect the parameters of this two-class distribution. Indeed, the share of capital income dominating the upper tail appears to be at least partially shaped by workers' bargaining power that, in turn, depends on policy (Stansbury and Summers 2020). The two-class structure might thus be best interpreted as a basic constraint of a capitalist economy in which policy can operate in – not a constraint on policy per se.

7. Conclusion

Dealing with power laws is likely challenging for applied economists and social scientists used to work in a Gaussian and statistically well-behaved world. Yet, as Jaynes (2003) reminds us in the introductory quote, empirical reality is often not well-behaved in this sense. Given the ubiquity and apparent universality of this distributional regularity, one might thus respond to Jaynes that it is power laws, not Gaussian distributions that should be characterized as “normal”. Our suggestion with this chapter is to take up the challenge – and leverage the opportunity this remarkably robust and ubiquitous distributional regularity provides. As we have discussed, power laws might aid in understanding the

fundamental drivers of income and wealth inequality and can help to identify systemically significant entities for aggregate behaviour. Many open questions still need to be tackled by (young) researchers: How can we gauge the reliability of statistical inference in the presence of fat tails? Which explanatory attempt best explains the emergence of power laws in different domains? What is the role of scale- and type-dependencies in stochastic growth? How can industrial and macroprudential policy best exploit power law distributions in firm size, network centrality and many other economic variables of interest? It is our sincere hope that this chapter provides some guidance in how to address these exciting open questions.

References

- Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The network origins of aggregate fluctuations." *Econometrica* 80 (5): 1977–2016.
- Akin, David. 2003. "Akin's laws of spacecraft design." *UMD Space Systems Laboratory*: 25.
- Aparicio, Sofía, Javier Villazón-Terrazas, and Gonzalo Álvarez. 2015. "A model for scale-free networks: application to Twitter." *Entropy* 17 (8): 5848–5867.
- Arthur, W Brian. 2023. "Economics in nouns and verbs." *Journal of Economic Behavior & Organization* 205: 638–647.
- Artico, Igor, I Smolyarenko, Veronica Vinciotti, and Ernst C Wit. 2020. "How rare are power-law networks really?" *Proceedings of the Royal Society A* 476 (2241): 20190742.
- Axtell, Robert L. 2001. "US firm sizes are Zipf distributed." *Science* 93: 1818–1820.
- Bach, Laurent, Laurent E Calvet, and Paolo Sodini. 2020. "Rich pickings? Risk, return, and skill in household wealth." *American Economic Review* 110 (9): 2703–2747.
- Barabási, Albert-László. 2009. "Scale-free networks: a decade and beyond." *Science* 325 (5939): 412–413.
- Barabási, Albert-László, and Réka Albert. 1999. "Emergence of scaling in random networks." *science* 286 (5439): 509–512.
- Blanco-Arroyo, Omar, Alba Ruiz-Bufo, David Vidal-Tomás, and Simone Alfarano. 2018. "On the determination of the granular size of the economy." *Economics Letters* 173: 35–38.
- Broido, Anna D, and Aaron Clauset. 2019. "Scale-free networks are rare." *Nature Communications* 10 (1): 1017.
- Carvalho, Vasco M, and Basile Grassi. 2019. "Large firm dynamics and the business cycle." *American Economic Review* 109 (4): 1375–1425.
- Castaldi, Carolina, and Mishaël Milaković. 2007. "Turnover activity in wealth portfolios." *Journal of Economic Behavior & Organization* 63 (3): 537–552.
- Champernowne, David G. 1953. "A model of income distribution." *The Economic Journal* 63 (250): 318–351.
- Clauset, Aaron, Cosma Rohilla Shalizi, and Mark EJ Newman. 2009. "Power-law distributions in empirical data." *SIAM review* 51 (4): 661–703.
- Coad, Alex. 2010. "The exponential age distribution and the Pareto firm size distribution." *Journal of Industry, Competition and Trade* 10: 389–395.

- Conceicao, Pedro, and Pedro M. Ferreira. 2000. "The Young Person's Guide to the Theil Index: Suggesting Intuitive Interpretations and Exploring Analytical Applications." *UTIP Working Paper* 14.
- Cook, Eli. 2020. "Naturalizing inequality: the problem of economic fatalism in the age of Piketty." *Capitalism: A Journal of History and Economics* 1 (2): 338–378.
- Córdoba, Juan Carlos. 2008a. "A generalized Gibrat's law." *International Economic Review* 49 (4): 1463–1468.
- Córdoba, Juan-Carlos. 2008b. "On the distribution of city sizes." *Journal of urban Economics* 63 (1): 177–197.
- Cowell, F.A. 2000. "Chapter 2 Measurement of inequality." vol. 1 of *Handbook of Income Distribution*, 87–166: Elsevier.
- Cowell, Frank A, and Emmanuel Flachaire. 2007. "Income distribution and inequality measurement: The problem of extreme values." *Journal of Econometrics* 141 (2): 1044–1072.
- Disslbacher, Franziska, Michael Ertl, Emanuel List, Patrick Mokre, and Matthias Schnetzer. 2020. "On Top of the Top: A Generalized Approach to the Estimation of Wealth Distributions." *INEQ Working Paper Series* 20.
- Dos Santos, Paulo L. 2017. "The principle of social scaling." *Complexity* 2017 (1): 8358909.
- Dosi, Giovanni, Mauro Napoletano, Andrea Roventini, and Tania Treibich. 2019. "Debunking the granular origins of aggregate fluctuations: from real business cycles back to Keynes." *Journal of Evolutionary Economics* 29: 67–90.
- Eckerstorfer, Paul, Johannes Halak, Jakob Kapeller, Bernhard Schütz, Florian Springholz, and Rafael Wildauer. 2016. "Correcting for the missing rich: An application to wealth survey data." *Review of Income and Wealth* 62 (4): 605–627.
- Ederer, Stefan, Maximilian Mayerhofer, and Miriam Rehm. 2021. "Rich and ever richer? Differential returns across socioeconomic groups." *Journal of Post Keynesian Economics* 44 (2): 283–301.
- Esquierro, Leon, and Sergio Da Silva. 2024. "Granular Cities." *Economies* 12 (7): 179.
- Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and persistence in returns to wealth." *Econometrica* 88 (1): 115–170.
- Fix, Blair. 2018a. "Hierarchy and the power-law income distribution tail." *Journal of Computational Social Science* 1 (2): 471–491.
- Fix, Blair. 2018b. "The Trouble With Human Capital Theory." *Real-World Economics Review* (86): 15–32.
- Fix, Blair. 2019. "Personal income and hierarchical power." *Journal of Economic Issues* 53 (4): 928–945.
- Fontanari, Andrea, Nassim Nicholas Taleb, and Pasquale Cirillo. 2018. "Gini estimation under infinite variance." *Physica A: Statistical Mechanics and its Applications* 502: 256–269.
- Franke, Reiner, and Frank Westerhoff. 2016. "Why a simple herding model may generate the stylized facts of daily returns: explanation and estimation." *Journal of Economic Interaction and Coordination* 11: 1–34.
- Freund, Caroline, and Martha Denisse Pierola. 2015. "Export superstars." *Review of Economics and Statistics* 97 (5): 1023–1032.
- Gabaix, Xavier. 1999. "Zipf's Law and the Growth of Cities." *American Economic Review* 89 (2): 129–132.
- Gabaix, Xavier. 2009. "Power laws in economics and finance." *Annual Review of Economics* 1 (1):

255–294.

- Gabaix, Xavier. 2011. “The granular origins of aggregate fluctuations.” *Econometrica* 79 (3): 733–772.
- Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H Eugene Stanley. 2003. “A theory of power-law distributions in financial market fluctuations.” *Nature* 423 (6937): 267–270.
- Gabaix, Xavier, and Rustam Ibragimov. 2011. “Rank- 1/2: a simple way to improve the OLS estimation of tail exponents.” *Journal of Business & Economic Statistics* 29 (1): 24–39.
- Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. “The dynamics of inequality.” *Econometrica* 84 (6): 2071–2111.
- Gaillard, Alexandre, Christian Hellwig, Philipp Wangner, and Nicolas Werquin. 2023. “Consumption, Wealth, and Income Inequality: A Tale of Tails.” *Federal Reserve Bank of Chicago Working Paper* 2023 - 43.
- Haeusler, Erich, and Jozef L Teugels. 1985. “On asymptotic normality of Hill’s estimator for the exponent of regular variation.” *The Annals of Statistics*: 743–756.
- Ipsen, Leonhard, Armin Aminian, and Jan Schulz. 2023. “Stress-testing inflation exposure: Systemically significant prices and asymmetric shock propagation in the EU28.” *BERG Working Paper Series* 188.
- Jaynes, Edwin T. 2003. *Probability theory: The logic of science.*: Cambridge University Press.
- Jones, Charles I. 2015. “Pareto and Piketty: The macroeconomics of top income and wealth inequality.” *Journal of Economic Perspectives* 29 (1): 29–46.
- Kovalenko, Tim, Claus Schnabel, and Heiko Stüber. 2022. “Is the German labour market granular?.” *Applied Economics Letters* 29 (1): 41–48.
- Lagemann, Theresa, and Miriam Rehm. 2025. “Differential returns in Germany.” *ifso working paper* 45.
- Lucas, Robert E. 1977. “Understanding business cycles.” *Carnegie-Rochester Conference Series on Public Policy* 5 (1): 7–29.
- Ludwig, Danial, and Victor M Yakovenko. 2022. “Physics-inspired analysis of the two-class income distribution in the USA in 1983–2018.” *Philosophical Transactions of the Royal Society A* 380 (2224): 20210162.
- Lydall, Harold F. 1959. “The distribution of employment incomes.” *Econometrica*: 110–115.
- Miranda-Pinto, Jorge, and Yuanting Shen. 2019. “A granular view of the Australian business cycle.” *Economic record* 95 (311): 407–424.
- Mitzenmacher, Michael. 2004. “A brief history of generative models for power law and lognormal distributions.” *Internet Mathematics* 1 (2): 226–251.
- Persky, Joseph. 1992. “Retrospectives: Pareto’s law.” *Journal of Economic Perspectives* 6 (2): 181–192.
- Ragab, Amr. 2014. “Three essays on the incomes of the vast majority.” PhD dissertation, The New School.
- Schneider, Markus P. A. “Revisiting the Thermal and Superthermal Two-Class Distribution of Incomes.” 88 (1): 5.
- Schulz-Gebhard, Jan. 2023. “Essays on Aggregation With Heterogeneous and Interacting Agents.” PhD dissertation, Otto-Friedrich-Universität Bamberg, Fakultät Sozial-und Wirtschaftswissenschaften.
- Schulz, Jan, and Daniel M Mayerhoffer. 2021. “Equal chances, unequal outcomes? Network-based evolutionary learning and the industrial dynamics of superstar firms.” *Journal of Business*

- Economics* 91 (9): 1357–1385.
- Schulz, Jan, and Mishaël Milaković. 2023. “How wealthy are the rich?” *Review of Income and Wealth* 69 (1): 100–123.
- Serrano, Ma Angeles, and Marián Boguná. 2003. “Topology of the world trade web.” *Physical Review E* 68 (1): 015101.
- Shaikh, Anwar, Juan Esteban Jacobo et al. 2020. “Economic arbitrage and the econophysics of income inequality.” *Review of Behavioral Economics* 7 (4): 299–315.
- Shin, Kotbee, Tian Li, and Jaimin Lee. 2024. “Granular origins of the Chinese business cycle.” *Applied Economics Letters*: 1–5.
- Silva, A Christian, and Victor M Yakovenko. 2004. “Temporal evolution of the “thermal” and “superthermal” income classes in the USA during 1983–2001.” *Europhysics Letters* 69 (2): 304.
- Silva, Murilo, and Sergio Da Silva. 2020. “The Brazilian granular business cycle.” *Economics Bulletin* 40 (1): 463–472.
- Simon, Herbert A. 1957. “The compensation of executives.” *Sociometry* 20 (1): 32–35.
- Stachurski, John, and Alexis Akira Toda. 2019. “An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity.” *Journal of Economic Theory* 182: 1–24.
- Stansbury, Anna, and Lawrence H Summers. 2020. “Declining worker power and American economic performance.” *Brookings Papers on Economic Activity* 156.
- Stumpf, Michael PH, and Mason A Porter. 2012. “Critical truths about power laws.” *Science* 335 (6069): 665–666.
- Taleb, Nassim Nicholas, and Raphael Douady. 2015. “On the super-additivity and estimation biases of quantile contributions.” *Physica A: Statistical Mechanics and its Applications* 429: 252–260.
- Tao, Yong, Xiangjun Wu, Tao Zhou, Weibo Yan, Yanyuxiang Huang, Han Yu, Benedict Mondal, and Victor M Yakovenko. 2019. “Exponential structure of income inequality: evidence from 67 countries.” *Journal of Economic Interaction and Coordination* 14 (2): 345–376.
- Theil, Henri. 1967. *Economics and Information Theory*.: North-Holland Publishing Company.
- Wagner, Joachim, and John P Weche. 2020. “On the granularity of the German economy—first evidence from the top 100 companies panel database.” *Applied Economics Letters* 27 (21): 1768–1771.
- Weber, Jan D., and Jan Schulz. 2024. “Growing Differently: European Integration and Regional Cohesion.” *Journal of Economic Issues* forthcoming.
- Yang, Jangho, Torsten Heinrich, Julian Winkler, François Lafond, Pantelis Koutroumpis, and J Doyné Farmer. 2025. “Measuring productivity dispersion: a parametric approach using the Lévy alpha-stable distribution.” *Industrial and Corporate Change* 34 (1): 79 – 117.
- Zhong, Xiaoshi, and Huizhi Liang. 2024. “On the Scale-Free Property of Citation Networks: An Empirical Study.” In *Companion Proceedings of the ACM on Web Conference 2024*,: 541–544.

BERG Working Paper Series (most recent publications)

- 187 Sarah **Mignot**, Fabio **Tramontana** and Frank **Westerhoff**, *Complex dynamics in a non-linear duopoly model with heuristic expectation formation and learning behavior*
- 188 Leonhard **Ipsen**, Armin **Aminian** and Jan **Schulz**, *Stress-testing Inflation Exposure: Systemically Significant Prices and Asymmetric Shock Propagation in the EU28*
- 189 Sarah **Mignot** and Frank **Westerhoff**, *Explaining the stylized facts of foreign exchange markets with a simple agent-based version of Paul de Grauwe's chaotic exchange rate model*
- 190 Roberto **Rozzi** and Stefanie Y. **Schmitt**, *Vertical product differentiation, prominence, and costly search*
- 191 Florian **Herold** and Christoph **Kuzmics**, *Farkas' Lemma and Complete Indifference*
- 192 Sarah **Mignot**, Paolo **Pellizzari**, and Frank **Westerhoff**, *Fake News and Asset Price Dynamics*
- 193 Fabian **Dietz** and Marco **Sahm**, *Fairness in Round-Robin Tournaments with Four Players and Endogenous Sequences*
- 194 Jan **Schulz**, Caleb **Agoha**, Anna **Gebhard**, Bettina **Gregg** and Daniel M. **Mayerhoffer**, *Excessive White Male Privilege Biases the Measurement of Intersectional Wage Discrimination*
- 195 Leonhard **Ipsen** and Jan **Schulz**, *The (Dis-)Equalizing Effects of Production Networks*
- 196 Ivan **Savin**, Philipp **Mundt** and Margherita **Bellanca**, *The paradox of climate policy diffusion*
- 197 Dominik **Bruckner** and Marco **Sahm**, *Party Politics: A Contest Perspective*
- 198 Lasare **Samartzidis**, Philipp **Mundt** and Jan **Schulz**, *Input specificity and labor's bargaining power: A production tree approach to functional income distribution*
- 199 Maybrit **Wächter**, Christian R. **Proaño**, and Juan Carlos **Peña**, *How Fitting is "one-size-fits-all"? Revisiting the Dynamic Effects of ECB's Interest Policy on Euro Area Countries*
- 200 Sarah **Mignot**, *Coevolution of stock prices and their perceived fundamental value*
- 201 Arne **Lauber**, Christoph **March** and Marco **Sahm**, *Round-Robin Tournaments in the Lab: Lottery Contests vs. All-Pay Auctions*
- 202 Stefanie Y. **Schmitt**, *Investments in environmental quality under limited attention*
- 203 Jan **Schulz** and Jan David **Weber**, *Power Laws in Socio-Economics*