

DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY


Power Laws in Economics: An Introduction

Citation

Gabaix, Xavier. 2016. “Power Laws in Economics: An Introduction.” Journal of Economic Perspectives 30 (1) (February): 185–206. doi:10.1257/jep.30.1.185.

Published version

https://doi.org/10.1257/jep.30.1.185

Link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:34651705

Terms of use

This article was downloaded from Harvard University’s DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFlOWI2M2EwYTg

Accessibility

https://accessibility.huit.harvard.edu/digital-accessibility-policy

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story

Journal of Economic Perspectives—Volume 30, Number 1—Winter 2016—Pages 185–206

Power Laws in Economics: An Introduction

Xavier Gabaix

Paul Samuelson (1969) was once asked by a physicist for a law in economics that was both nontrivial and true. This is a difficult challenge, as many (roughly) true results are in the end rather trivial (for example, demand curves slope down), while many nontrivial results in economics in fact require too much sophistication and rationality on the part of the agents to actually hold true in practice.[1] Samuelson answered, "the law of comparative advantage." The story does not say whether the physicist was satisfied. The law of comparative advantage is a qualitative law, and not a quantitative one as is the rule in physics. Indeed, many of the insights of economics seem to be qualitative, with many fewer reliable quantitative laws.

This article will make the case that a modern answer to the question posed to Samuelson would be that a series of power laws count as actually nontrivial and true laws in economics—and that they are not only established empirically, but also understood theoretically.[2] I will start by providing several illustrations of empirical power laws having to do with patterns involving cities, firms, and the stock

[1] In Gabaix (2014), I explore this issue by proposing a behavioral version of basic microeconomics—specifically consumer and equilibrium theory. Many nontrivial predictions of microeconomics (for example, whether demand systems have the property of Slutsky symmetry) turn out to fail when agents are boundedly rational, though more trivial predictions (for example, demand curve slope down) typically hold.

[2] At the same time, in some highly incentivized and abstract environments, other "laws" (roughly) hold, like the Black–Scholes formula, or some successful mechanisms from implementation theory.

■ *Xavier Gabaix is Martin J. Gruber Professor of Finance, Stern School of Business, New York University, New York City, New York. His email address is xgabaix@stern.nyu.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
http://dx.doi.org/10.1257/jep.30.1.185

doi=10.1257/jep.30.1.185

market. I summarize some of the theoretical explanations that have been proposed. I suggest that power laws help us explain many economic phenomena, including aggregate economic fluctuations. I hope to clarify why power laws are so special, and to demonstrate their utility. In conclusion, I list some power-law-related economic enigmas that demand further exploration.

A formal definition may be useful. A power law, also called a scaling law, is a relation of the type $Y = aX^{\beta}$, where Y and X are variables of interest, β is called the power law exponent, and a is typically an unremarkable constant. For instance, if X is multiplied by a factor of 10, then Y is multiplied by 10^{β}—one says that Y "scales" as X to the power β.

Some Empirical Power Laws

City Sizes

Let us look at the data on US cities with populations of 250,000 or greater, plotted in Figure 1. We rank cities by size of population: #1 is New York, #2 Los Angeles, and so on, using data for Metropolitan Statistical Areas provided in the *Statistical Abstract of the United States* (2012). We regress log rank on log size and find the following:

$$\ln(Rank) = 7.88 - 1.03 \ln(Size).$$

The relationship in Figure 1 is close to a straight line ($R^2 = 0.98$), and the slope is very close to 1 (the standard deviation of the estimated slope is 0.01).[3] This means that the rank of a city is essentially proportional to the inverse of its size (indeed, exponentiating, we obtain $Rank = a\, Size^{-1.03}$) with $a = e^{7.88}$). A slope of approximately 1 has been found repeatedly using data spanning many cities and countries (at least after the Middle Ages, when progress in agriculture and transport could make large densities viable, see Dittmar 2011). There is no obvious reason to expect a power law relationship here, and even less for the slope to be 1.

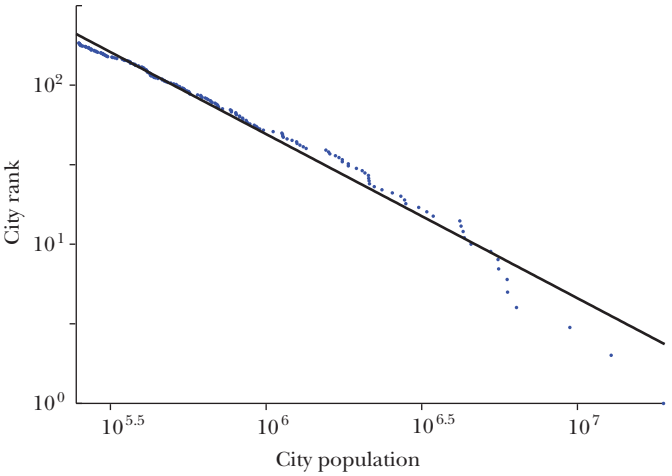
To think about this type of regularity, it is useful to be a bit more abstract and see the cities as coming from an underlying distribution: the probability that the population size of a randomly drawn city is greater than x is proportional to $1/x^{\zeta}$ with $\zeta \approx 1$. More generally,

$$\mathrm{P}(Size > x) = a/x^{\zeta},$$

at least for x above a cutoff (here, the 250,000-inhabitants cutoff used by the *Statistical Abstract of the United States*). An empirical regularity of this type is a power law. In

[3] Actually, the standard error returned by an ordinary least squares calculation is incorrect. The correct standard error is $|\text{slope}| \times \sqrt{2/N} = 0.11$, where $N = 184$ is the number of cities in the sample (Gabaix and Ibragimov 2011).

Figure 1
A Plot of City Rank versus Size for all US Cities with Population over 250,000 in 2010



Source: Author, using data from the S*tatistical Abstract of the United States* (2012).
Notes: The dots plot the empirical data. The line is a power law fit ($R^2 = 0.98$), regressing ln *Rank* on ln *Size*. The slope is −1.03, close to the ideal Zipf's law, which would have a slope of −1.

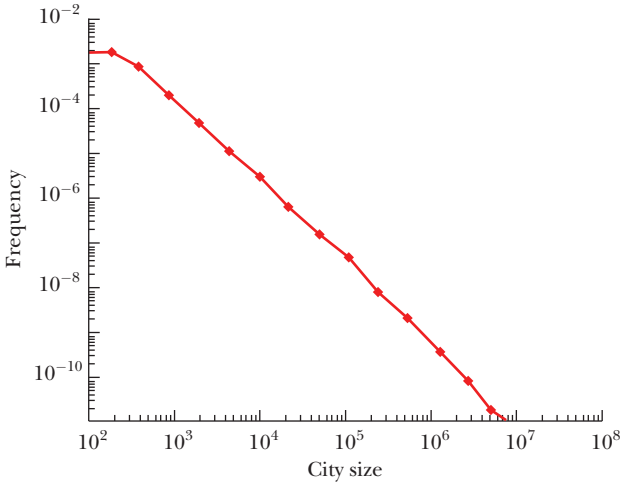
a given finitely sized sample, it generates an approximate relation of type shown in Figure 1 and in the accompanying regression equation.

The interesting part is the coefficient ζ, which is called the power law exponent of the distribution. This exponent is also sometimes called the "Pareto exponent," because Vilfredo Pareto discovered power laws in the distribution of income (as discussed in Persky 1992). A "Zipf's law" is a power law with an exponent of 1. George Kingsley Zipf was a Harvard linguist who amassed significant evidence for power laws and popularized them (Zipf 1949).

A lower ζ means a higher degree of inequality in the distribution: it means a greater probability of finding very large cities or (in another context) very high incomes.[4] In addition, the exponent is independent of the units (inhabitants or thousands of inhabitants, say). This makes it at least conceivable, a priori, that we might find a constant value in various datasets. What if we look at cities with size less than 250,000? Does Zipf's law still hold? When measuring the size of cities, it is better to look at agglomerations rather than the fairly arbitrary legal entities, but this is tricky. Rozenfeld et al. (2011) address the problem using a new algorithm that constructs the population of small cities from fine-grained geographical data. Figure 2 shows the resulting distribution of city sizes for the United Kingdom,

[4] Indeed, the expected value of S^α is mathematically infinite if α is greater than the power law exponent ζ, and finite if α is less than the power law exponent ζ. For example, if $\zeta = 1.03$, the expected size is finite, but the variance is formally infinite.

Figure 2
Density Function of City Sizes (Agglomerations) for the United Kingdom



Source: Rozenfeld et al. (2011).
Notes: We see a pretty good power law fit starting at about 500 inhabitants. The Pareto exponent is actually statistically non-different from 1 for size $S > 12{,}000$ inhabitants.

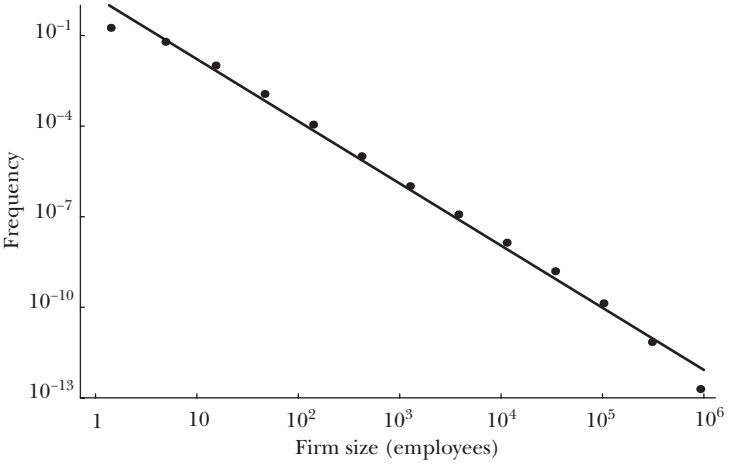
where the data is particularly good. Here we see the appearance of a straight line for cities of about size 500 and above. Zipf's law holds pretty well in this case, too.

Why might social scientists care about this relationship? As Krugman (1996) wrote 20 years ago, referring to Zipf's law, which remained unexplained by his work of economic geography: "The failure of existing models to explain a striking empirical regularity (one of the most overwhelming empirical regularities in economics!) indicates that despite considerable recent progress in the modeling of urban systems, we are still missing something extremely important. Suggestions are welcome." We shall see that since Krugman's call for suggestions, we have much improved our understanding of the origin of the Zipf's law, which has forced a great rethinking about the origins of cities—and firms, too.

Firm Sizes

We now look at the firm size distribution. Using US Census data, Axtell (2001) puts firms in "bins'" according to their size, as measured by number of employees, and plots the log of the number of firms within a bin. The result in Figure 3 shows a straight line: again, this is a power law. Here we can even run the regression in "density"—that is, plot the number of firms of size approximately equal to x. If a power law relationship holds, then the density of the firm size distribution is $f(x) = b/x^{\zeta+1}$, so the slope in a log-log plot should be $-(\zeta + 1)$ (because $\ln f(x) = -(\zeta + 1) \ln x$ plus a constant). Impressively, Axtell finds that the exponent $\zeta = 1.059$. This demonstrates a "Zipf's law" for firms.

Figure 3
Log *Frequency* versus log *Size* of US firms (by Number of Employees) for 1997



Source: Axtell (2001).
Notes: Ordinary least squares (OLS) fit gives a slope of 2.06 (s.e. = 0.054; $R^2 = 0.99$). This corresponds to a frequency $f(S) \sim S^{-2.059}$, which is a power law distribution with exponent 1.059. This is very close to an ideal Zipf's law, which would have an exponent $\zeta = 1$.

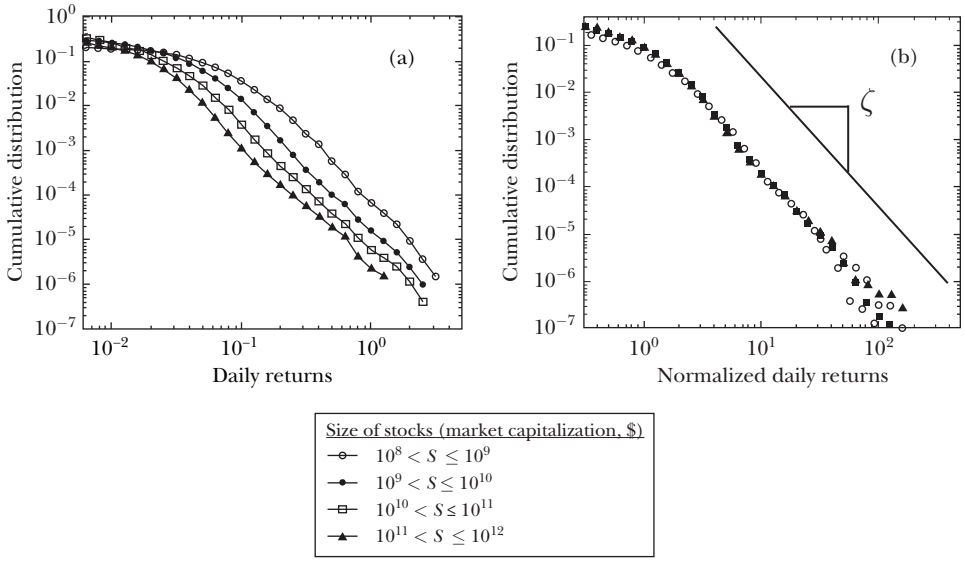
This finding has forced a rethinking of the underpinnings of firms: Most static theories of why firms exist—for example, theories based on economies of scope, fixed costs, elasticity of demand, and the like—would not predict a Zipf's law. Some other type of theory is needed, as we shall soon discuss.

Stock Market Movements

It is well-known that stock market returns are fat-tailed—that is, the probability of finding extreme values is larger than for a Gaussian distribution of the same mean and standard deviation. An energetic movement of physicists, the "econophysicists" (a term coined after the emergence of "geophysicists" and "biophysicists"), has quantified a host of power laws in the stock market. For instance, the size of daily stock market movements are represented in Figure 4. They are consistent with: $\mathrm{P}(|\, r_t \,| > x) = a/x^{\zeta}$ with $\zeta = 3$, the so-called "cubic" law of stock market returns. The left panel of Figure 4 plots the distribution for four different sizes of stocks. The right panel plots the distribution of normalized stock returns, which is calculated as the stock returns divided by their standard deviation: after this normalization, the four different distributions "collapse" onto the same curve. This is a type of "universality"—a term much used in the power law literature (and in physics) which means that different systems behave in the same way, after some rescaling. This cubic law appears to hold for a variety of other international stock markets too (Gopikrishnan et al. 1999).

Likewise, lots of other stock market quantities are distributed according to a power law (Plerou, Gopikrishnan, and Stanley 2005; Kyle and Obizhaeva 2014;

Figure 4
Cumulative Distribution of Daily Stock Market Returns for Different Sizes of Stocks



Source: Plerou et al. (1999).
Notes: The left panel shows the distributions for four different sizes (in terms of market capitalization) of stocks. The right panel shows the returns, normalized by volatility. The slopes $-\zeta$ are close to -3, reflecting the "cubic law" of stock market fluctuations: $P(|r| > x) \sim kx^{-3}$. The horizontal axis displays returns as high as 100 standard deviations.

Bouchaud, Farmer, and Lillo 2009). For instance, the number of trades per day is a power law distributed with exponent of 3, while the number of shares traded per time interval has an exponent of 1.5, and the price impact (that is, the size of the price movement when a large volume of shares is bought or sold) is proportional to the volume to the power of 0.5. Later, I discuss theories explaining those facts.

Again, why do social scientists care? One implication of the cubic law is that there are many more extreme events than would occur if the distribution were Gaussian; for example, if the distribution were the result of a series of events of fixed probability, like flipping coins or rolling dice. More precisely, under the cubic law, the chances of a 10 standard deviation event and a 20 standard deviation event are, respectively, $5^3 = 125$ and $10^3 = 1{,}000$ times less likely than a two standard deviation event, whereas if the distribution of returns was Gaussian, the chances of a 10 and 20 standard deviation event would be 10^{22} and 10^{87} times less likely than a two standard deviation event (much, much less likely). Indeed, in a stock market comprising about 1,000 stocks, a 10 standard deviation event happens in practice just about every day.

Seeking an explanation for these kinds of regularities presses us to rethink the functioning of stock markets—later we shall see theories that exactly explain these exponents in a unified way.

Other Examples of Power Laws

Income and wealth also follow roughly power law distributions, as we have known at least since Pareto (1896), who documented power laws relating to the distribution of income and who was the first to document power laws in economics or in any other area of social science (to the best of my knowledge). The distribution of wealth is more unequal than the distribution of income: this makes sense, because differences in growth rate of wealth across individuals (due to differences in returns or frugality) pile up and add an extra source of inequality. Typically, the Pareto exponent is around 1.5 for wealth and between 1.5 and 3 for income (recall that a lower Pareto exponent means a higher degree of inequality in a distribution). Indeed, power laws and random growth processes are rapidly becoming a central tool to analyze inequalities of income and wealth (Piketty and Zucman 2014; Atkinson, Piketty, and Saez 2011; Benhabib, Bisin, and Zhu 2011; Lucas and Moll 2014; Gabaix, Lasry, Lions, and Moll 2015; Toda and Walsh 2015).

What Causes Power Laws?

Here I sketch the two main mechanisms that generate powers laws: 1) random growth models, which generate a power law and 2) transfer of that power law via matching and optimization (that is, the variable generated by random growth is used as an input begetting another power law in another output variable).

Random Growth

The basic mechanism for generating power laws is proportional random growth (Champernowne 1953; Simon 1955). Suppose that we start with an initial distribution of firms, and they grow and shrink randomly with independent shocks, and they satisfy Gibrat's law (Gibrat 1931), which starts with the assumption that all firms have the same expected growth rate and the same standard deviation of growth rate. However, these basic assumptions do not assure a steady state distribution, because they imply a distribution that over time becomes lognormal with larger and larger variance. However, things change altogether if we add to the model some friction that guarantees the existence of a steady state distribution. Suppose for instance that there is also a lower bound on size, so that a firm size cannot go below a given threshold. Then, as it turns out, the model yields a steady-state distribution, and it is a power law, with some exponent ζ that depends on the details of the growth process.

However, while this mechanism generates a power law, the exponent need not be equal to 1. Why might an exponent of 1 arise?

I give one explanation in Gabaix (1999; for a more thorough review, see Gabaix 2009). Suppose that the size of the assumed friction (the lower bound) becomes very small, and that we have a given exogenous population size to allocate in the system (between the different cities or firms). Then, the exponent ζ becomes 1, rather than any other value. One intuition behind this result is as follows: the exponent cannot be below 1 because then the distribution would have infinite mean

(see footnote 3). Indeed, an exponent just above 1 is the smallest consistent with a finite total population. As the friction becomes very small, the exponent becomes the fattest that is consistent with a finite population.

This insight can explain why we observe lots of Zipf's laws with the exponent ζ equal to 1 in the real economy: because a number of economic variables are well-represented by an underlying pattern of proportional random growth with a small friction and some adding-up constraint for the total size of the system.

Of course, this explanation so far is mechanical; that is, it just points out that a certain process leads to a certain distribution. A social scientist will want to know why these variables exhibit proportional random growth in the first place. Fortunately, once the basic mechanics are clarified, one can suggest good economic reasons as to why they might plausibly exist.

The simplest microfoundation suggested by power laws and Gibrat's law is the following: cities and firms largely exhibit constant returns to scale, perhaps with small deviations from that benchmark, and lots of randomness. In this spirit, many fully economic models for the random growth of cities and firms have been proposed since the 2000s which add more economics to this random growth mechanism (for example, Rossi-Hansberg and Wright 2007; Luttmer 2007). The shocks to this system come from shocks to productivities or amenities. A certain death rate for firms may arise from the processes of creative destruction, and the minimum size can come from a fixed cost that the firm needs to pay to be alive.

As Gibrat's rule of proportionate growth is most naturally consistent with constant return to scale economies, it is an interesting question as to how it fits with the many standard economic theories of firm size that emphasize increasing returns to scale as well as with economies and diseconomies of agglomeration. One possibility is that the effects of increasing returns to scale and the economies and diseconomies of scale are not that big (see the discussion in Rozenfeld, Rybski, Gabaix, and Makse 2009). Another possibility is that these effects are to some extent offset by other large compensating factors, like urban amenities or geography. Yet another possibility is that some of these theories assume that shocks have permanent effects, but some shocks do mean-revert: for example, Japanese cities in large part reverted to their previous sizes after World War II bombing (Davis and Weinstein 2002). Writing richer theories of city size promises to be a fruitful task for future researchers.

Likewise, for the income distribution, the details of the underlying mechanism—say, the extent of luck, the distribution of thrift, and the varying responsiveness to incentives—are very important for a variety of questions, and microfounded models are important. Still, to write sensible theories on this basis one needs to keep in mind the core mechanics of these models, which is proportional random growth leading to power laws.

Matching and Economics of Superstars

Another manifestation of power laws is in the extremely high earnings of top earners in areas of arts, sports, and business. Rosen (1981) suggests a qualitative explanation for this pattern with the "economics of superstars." In Gabaix and

Landier (2008), we present a tractable, calibratable model of this phenomenon along the following lines: Suppose that lots of firms, of different sizes, compete to hire the talents of chief executive officers. In this model, the talent of a chief executive officer (CEO) is given by how much (in percentages) that person is expected to increase the profits of the firm. Competition implements the efficient outcome, which is that the largest firm will be matched with the best CEO in the economy, the second largest firm with the second best CEO, and so on (as in Terviö 2008).[5]

One might think it hopeless to derive a quantitative theory from this starting point, because the distribution of executive talent is very hard to observe. However, we can draw on extreme value theory (which is a branch of probability) to obtain some properties of the tail of the distribution of talent, without knowing the distribution itself. One implication is that, given adjacent chief executive officers in the ordering of talent, the approximate difference in talent between these two CEOs varies like a power law of their rank. The exponent depends on the distribution, but the power law functional form holds for essentially any reasonable distribution (in a way that can be made precise). Given this, in Gabaix and Landier (2008), we work out the pay of CEO number n, who manages a firm of size $S(n)$. We denote by $S(n^*)$ the size of a reference firm, which is the size of the median firm in the Standard and Poor's 500. $D(n^*)$ is a constant that depends on model parameters, like the scarcity of talent. The pay of CEO number n is:

$$w(n) = D(n^*)\, S(n^*)^{1-b} S(n)^b.$$

In this approach, one calibrates $b = 1/3$ (empirically, the exponent tends to be in the [0.3, 0.4] range). For instance, if a firm is eight times bigger than the median firm (so $S(n) = 8\ S(n^*)$), then the CEO of that larger firm earns twice ($8^{1/3}$) as much as the median CEO. But if the size of all firms is multiplied by 8 (so $S(n)$ and $S(n^*)$) are multiplied by 8, the pay of all CEOs is increased by 8.

In this way, the equation creates a "dual scaling" or double power law—because there is a scaling in both average firm size and own firm size. This approach has three implications:

1) Cross-sectional prediction. In a given year, the compensation of a CEO is proportional to the size of the firm to the power of 1/3, $S(n)^{1/3}$, an empirical relationship sometimes called Roberts' (1956) law.

2) Time-series prediction. When the size of all large firms is multiplied by λ (perhaps over a decade), the compensation at all large firms is multiplied by λ. In particular, the pay at the reference firm is proportional to the average market cap of a large firm.

3) Cross-country prediction. Suppose that CEO labor markets are national rather than internationally integrated. For a given firm size S, CEO compensation varies

[5] The microfoundations of matching models and internal organizations are also interesting, and full of power laws (Garicano and Rossi-Hansberg 2006; Geerolf 2015). The availability of microdata makes these detailed microeconomic models even more testable.

across countries with the market capitalization of the reference firm, $S(n^*)^{2/3}$, using the same rank n^* of the reference firm across countries.

It turns out that all three predictions seem to hold empirically since the 1970s. This theory thus points to the increase in firm size as the cause for the increase in CEO pay.

In this way, power laws and extreme value theory are the natural language for drawing quantitative lessons from the "economics of superstars." This formulation explains why very small differences in talent give rise to very large differences in pay: in our calibration in Gabaix and Landier (2008), differences in talent are small and bounded, but differences in pay are unboundedly large. This is what happens when very large firms compete to hire the services of CEOs: small differences of talent, affecting unboundedly large firms, give rise to unboundedly large differences in pay.

The same logic should apply to other markets with superstar characteristics: apartments with a large view of Central Park in New York City, and also top athletes in sports and the price of famous works of art. As far as I know, a systematic quantitative exploration of those issues remains to be done.[6]

This line of thinking leads to a fresh way of thinking about pay–performance sensitivity for chief executive officers. In the classic paper by Jensen and Murphy (1990), they define pay–performance sensitivity as how many dollars the compensation (or wealth) of a CEO changes for a given dollar change in firm value. They find that pay–performance sensitivity is very small: CEOs earn "only" \$3 extra when their firm increases by \$1,000 in value, and so they conclude that corporate governance may not work well. In contrast, in Edmans, Gabaix, and Landier (2009), we propose a different way to think of the benchmark incentives, resting on scaling arguments. Suppose that to motivate the CEO, it is percent/percent incentives, not dollar/dollar incentives that matter: namely, for a 1 percent increase in firm value, the CEO's wealth should increase by k percent, where k is independent of firm size (this relationship is derived from preferences that are multiplicative in effort and consumption).[7] Then, if our earlier expression for the pay of a CEO based on the size of the firm holds, the pay–performance sensitivity in the sense of Jensen and Murphy should decrease as (Firm size)$^{2/3}$. This pattern holds true empirically, as illustrated in Figure 5. Hence, thinking in terms of scaling leads to new thinking about pay–performance sensitivity.

It is thus interesting to study models where assignment and incentives are optimally and jointly determined, a task I carried out with coauthors in Edmans,

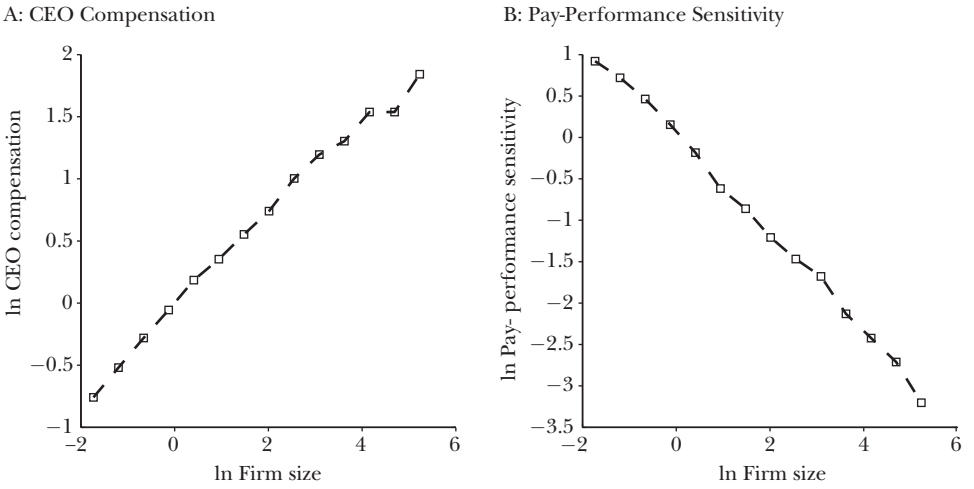
[6] Behrens, Duranton, and Robert-Nicoud (2013) propose a theory of Zipf's law based on matching.

[7] The percent/percent incentive is constant ($d \ln w / dr = k$, where w is CEO pay, and r the firm return), so the Jensen–Murphy pay–performance sensitivity measure is,

$$PPS \coloneqq \frac{dw}{dS} = \frac{dw}{Sdr} = \frac{w}{S}\frac{d\ln w}{dr} = \frac{k' S^b}{S} k = k'' S^{b-1}$$

using $w = k' S^b$ from Roberts' law, and for firm-independent constants k', k''.

Figure 5
CEO Pay and CEO Pay–Performance Sensitivity versus Firm Size



Source: The data and methodology is from Edmans, Gabaix, and Landier (2009), for the years 1994–2008. Years are lumped together by reporting $\ln \frac{w_{it}}{\overline{w}_t}$ (panel A) and $\ln \frac{PPS_{it}}{\overline{PPS}_t}$ (panel B) versus $\ln \frac{S_{it}}{\overline{S}_t}$ (horizontal axis in both panels), where $\overline{x}_t$ indicate the median value of x_{it} in year t.
Notes: Left panel: The CEO compensation is the ex ante one, including Black–Scholes value of options granted. The slope is about 1/3, a reflection of Roberts's law: Pay $\sim$ Size^b with $b \simeq 1/3$. Right panel: The pay–performance sensitivity (PPS) is the Jensen–Murphy measure: by how many dollars does the CEO wealth change, for a given dollar change in firm value. The slope is about $-2/3$, so that PPS $\sim$ Size^{b-1} with $b \simeq 1/3$. The congruence between the scalings is predicted by the Edmans, Gabaix, Landier (2009) model.

Gabaix, and Landier (2009) and Edmans and Gabaix (2011). One result is that pay in the aggregate is all about rewarding talent, not about paying for risk and incentives (which affect pay in the cross-section, not in the aggregate). In the aggregate, the reward to talent fully governs the level of expected pay; incentive issues are quite secondary and simply pin down the form of pay, like what fraction is in fixed or variable pay, not its level. For instance, if some firms are riskier than others, they need to reward their CEOs more (a cross-sectional effect). But if all firms become riskier, the level of pay does not budge (there is no aggregate effect). Hence from that perspective, the rise in pay is all about talent, not incentives.

Optimization and Transfer of Power Laws

Optimization provides a useful way to obtain power laws. For instance, the Allais–Baumol–Tobin rule for the demand for money (which scales as $i^{-1/2}$, where i is the interest rate), is a power law (Allais 1947; Baumol and Tobin 1989). The first scaling relation in economics—and, not coincidentally, the first nontrivial empirical success in economics—may be Hume's thought experiment that doubling the

money supply should lead, after a time, to a doubling of the price level—a basic theory that has stood the test of time.[8]

Power laws have very favorable aggregation properties: taking the sum of two (independent) power law distributions gives another power law distribution. Likewise, multiplying two power laws, taking their max or their min, or a power, gives again a power law distribution. This partly explains the prevalence of power laws: they survive many transformations along with the addition of noise.

Granularity: Aggregate Fluctuations from Microeconomic Shocks

I now turn to an application of power laws: developing a better sense of the origins of aggregate fluctuations in GDP, exports, and the stock market.

Basic Ideas

Where do aggregate fluctuations come from? In Gabaix (2011), I propose that idiosyncratic shocks to firms (or narrowly defined industries) can generate aggregate fluctuations. A priori, many economists would say that this is not quantitatively plausible: there are millions of firms and their idiosyncratic variations should tend to cancel each other out, so the resulting total fluctuations should be very small. However, when the firm size distribution is fat-tailed, this intuition no longer applies, and random shocks to the largest firms can affect total output in a noticeable way.[9]

Empirically, the existence of a power law distribution for firm size suggests that economic activity is indeed very concentrated amongst firms. For instance, di Giovanni and Levchenko (2012) find: "In Korea, the 10 biggest business groups account for 54% of GDP and 51% of total exports. . . . The largest one, Samsung, is responsible for 23% of exports and 14% of GDP." In a setting like this, it seems more plausible that idiosyncratic shocks to firms would affect macroeconomic activity. Likewise, in Japan, the top 10 firms account for 35 percent of exports, and in the United States, the sales of the top 50 firms represent about 25 percent of output (Gabaix 2011). In this view, economic activity is not made of a smooth continuum of firms, but it is made of incompressible "grains" of activities—we call them "firms"—whose fluctuations do not wash out in the aggregate. One plain reason is that some of the firms are very big, and a further reason is that initial shocks can be intensified by a variety of generic amplification mechanisms, such as endogenous changes to hours worked.

Is this granular hypothesis relevant empirically? In Gabaix (2011), I find that idiosyncratic shocks to large firms explain about one-third of GDP fluctuations in

[8] Hume's scaling says $\textit{Price level} = a(\textit{Money supply})^1$, with a proportionality factor a that depends on GDP, a very simple "power law" with exponent of 1.

[9] If there are N firms and a distribution of firms where the central limit theorem applies, the effect of a random shock on total fluctuations should decay as $1/\sqrt{N}$. In a fat-tailed distribution, the standard central limit theorem no longer applies. Instead, a power-law variant holds called the Lévy central limit theorem. In this setting, the effect of random shocks on total GDP fluctuation decays as $1/\ln N$.

the US economy. Di Giovanni, Levchenko, and Mejean (2014) find that they explain over half the fluctuations in France. Further support is given in Foerster, Sarte, and Watson (2011) for industrial production, and in di Giovanni and Levchenko (2012) for exports. The exploration of this theme continues.

This analysis offers two payoffs. First, it may help us to better understand the origins of aggregate fluctuations. Second, these large idiosyncratic shocks may suggest some useful instruments for macroeconomic policy. For instance, Amiti and Weinstein (2013) start from the fact that banking is very concentrated, such that idiosyncratic bank shocks may have strong ripple effects in the aggregate economy; they then seek to quantify these banking channels. They find that idiosyncratic bank shocks can explain 40 percent of aggregate loan and investment fluctuations.

Another implication of granularity is to emphasize the potential importance of networks (Acemoglu, Carvalho, Ozdalar, and Tahbaz-Salehi 2012; Carvalho 2014). Those large firm-level shocks propagate through networks, which create an interesting amplification mechanism and a way to observe the propagation effects. Networks are a particular case of granularity rather than an alternative to it: if all firms had small sales, the central limit theorem would hold and idiosyncratic shocks would all wash out. Networks offer a way to visualize and express the propagation of idiosyncratic firm shocks. Indeed, ongoing research is finding more precise evidence for the explanatory power of this perspective (Kelly, Lustig, and van Nieuwerburgh 2013; Acemoglu, Akcigit, and Kerr 2015).

The Great Moderation: A Granular Post Mortem

This granular perspective also offers a way to understand the time variations in economic volatility. Say that granular or "fundamental" volatility is the volatility that would come only from idiosyncratic sectoral- or firm-level shocks. By construction, when the economy is more diversified, or when the large sectors are in less-volatile industries, fundamental volatility is lower. In Carvalho and Gabaix (2013), we find that fundamental volatility is quite correlated with actual volatility, again consistent with the idea that firm- or industry-specific shocks are an important driver of aggregate fluctuations in the United States and other high-income economies. In addition, policy may dampen or amplify those primitive granular shocks but is not (typically) the primary driver. For instance, in the case of the Great Recession, the primitive shock is a shock to a narrow sector—real estate finance—which was then propagated to the rest of the economy via interesting economic and policy linkages.

This perspective offers an additional narrative for some events of the US economy in recent decades. The US economy experienced what was often called a Great Moderation of lower volatility from the mid-1980s through the mid-2000s, which is often credited at least in part to greater stability of monetary policy. However, from a granularity perspective, the long and large decline of volatility can be traced back to the long and large decline in fundamental volatility at the same period—which came about in part because of the shrinkage of a handful of heavy-manufacturing sectors, whose demise made the economy more diversified. The burst of economic volatility in the 1970s can be attributed in part to the increased importance of a

single sector—the energy sector—which itself can be traced to the rise of oil prices. From this view, the growth in the size of the financial sector is an important determinant of the increase in fundamental volatility—and of actual volatility—in the 2000s. Rather than relying on abstract shocks, a granularity perspective helps us to understand concretely the (proximate) origin of macroeconomic developments.

Linked Volatility of Firms and National Economies

The volatility of the growth rate of firms varies by size, with larger firms tending to have a smaller proportional standard deviation than smaller firms. The volatility of national economies also varies by size, with larger economies tending to have less volatility. Intriguingly, this relationship for firms looks much the same as it does for national economies.

Stanley et al. (1996) study how the volatility of the growth rate of firms changes with size, looking at data for all publicly traded US manufacturing firms between 1975 and 1991. To do this, they calculate the standard deviation $\sigma(S)$ of the growth rate of firms' sales, S, and regress its log against log size. They find an approximately linear relationship, displayed in Figure 6: $\ln \sigma^{\text{firms}}(S) = -\alpha \ln S + \beta$. This means that a firm of size S has volatility proportional to $S^{-\alpha}$ with $\alpha = 0.15$.[10] Lee et al. (1998) conduct the same analysis for fluctuations in the gross domestic product (GDP) of 152 countries for the period 1950–1992 and also find a volatility proportional to $S^{-\alpha'}$ with $\alpha' = 0.15$ (Koren and Tenreyro 2013). These size/volatility relationships, for firms and for countries, are both plotted in Figure 6. The slopes are indeed very similar. This may be a type of "universality." This may be explained by a combination of granularity and power laws: if aggregate fluctuations come from microeconomic shocks, and firm sizes follow Zipf's law, then the identical scaling of Figure 6 should hold true (Gabaix 2011).

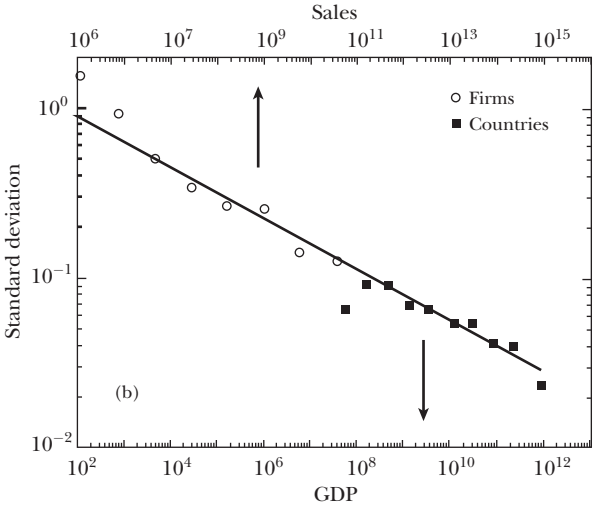
Origins of Stock Market Crashes?

We saw earlier illustrations of power laws in stock market fluctuations. What causes these fluctuations? The power law distribution of firms might actually explain the power law distribution of stock market crashes. In Gabaix, Gopikrishnan, Plerou, and Stanley (2003, 2006), we develop the hypothesis that stock market crashes are due to large financial institutions selling under pressure in illiquid markets (see also Levy and Solomon 1996; Solomon and Richmond 2001). This may account not only for large crashes, but also for the whole distribution of mini-crashes described by the power law.

The power law perspective begins by noting that large institutions are roughly Zipf-distributed, so when they trade, they could have a very large price impact. However, the institutions trade intelligently—and working out the optimal trading strategy of the large institutions gives an explanation for the empirically found

[10] This suggests that large firms are a little more diversified than small firms. However, this diversification effect is weaker than if a firm of size S were composed of S independent units of size 1, which would predict $\alpha = 1/2$. See Riccaboni et al. (2008).

Figure 6
Standard Deviation of the Distribution of Annual Growth Rates



Source: Lee et al. (1998). The firm data are taken from the Compustat for the years 1974–93, the GDP data from Summers and Heston (1991) for the years 1950–92.
Notes: We see that $\sigma(S)$ decays with size S with the same exponent for both countries and firms: $\sigma(S) \sim S^{-\alpha}$, with $\alpha \simeq 1/6$. The size is measured in sales for the companies (top axis) and in GDP for the countries (bottom axis).

power law exponents of trading (an exponent of 3 for returns and 1.5 for volume, as presented earlier). More specifically, large institutions want to moderate their price impact. So, they take more time to execute their large trades. At the optimum, it turns out the price impact they achieve is proportional to the square root of the size of the trade, both in the theory and in the data. Also, in order to moderate total transaction costs, large institutions need to trade proportionally less than smaller ones. In the resulting equilibrium, the distribution of trades is less fat-tailed than the distribution of firms (with an exponent of 1.5), and the distribution of returns is even less fat-tailed (with an exponent of 3)—the factor of 2 coming from the 1/2 exponent of the square root price impact. This theory can be summarized more qualitatively: when large institutions sell under time pressure, they make the market fall and even crash.

One can speculate that this type of mechanism might have been at work in a variety of well-known events, whose origins had one or just a few primitive large traders (which suggests interesting ramifications that are under-researched). The Long Term Capital Management crash in summer 1998 was clearly due to one large fund, with repercussions for large markets (in particular bond markets). The rapid unwinding of very large stock positions by Société Générale after the Kerviel rogue trader scandal caused European stock markets to fall by 6 percent in January 21, 2008, which led the Fed to decrease its rates by 0.75 percentage points. Similarly, it seems the so-called "flash crash" of May 2010 was due to one trader. There is even

tentative evidence that a similar process unwound at the onset of the great stock market crashes of 1929 and 1987 (see the discussion in Gabaix et al. 2006; Kyle and Obizhaeva 2013). This research potentially brings us closer to understanding the origins of stock market movements, with a useful blend of narrative concreteness and general underlying mechanisms leading to clean predictions in the forms of power laws.[11]

Power Laws Outside Economics

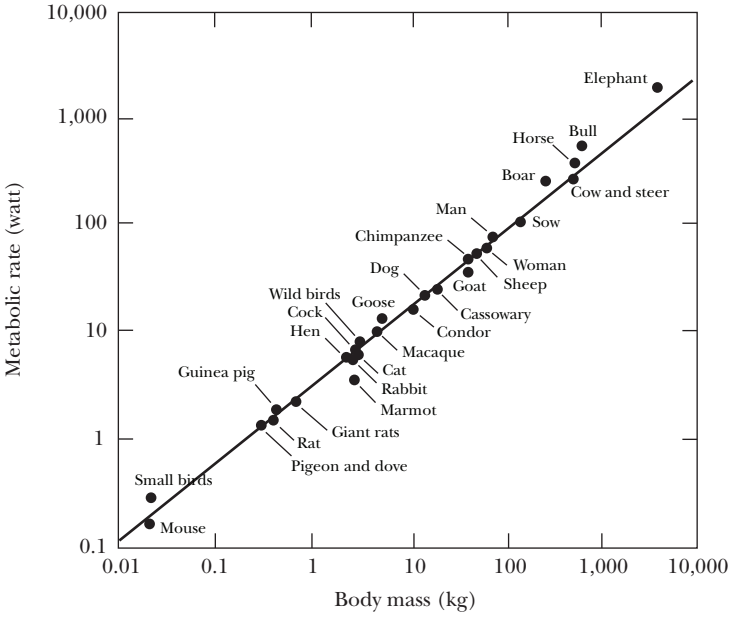
Power laws turn out to be a useful tool for analysis in many areas. For example, natural phenomena like forest fires and rivers have a power law structure (Turcotte 1997). There are concepts here underlying the power law structure that have not (yet) found widespread use in economics. For example, the concept of self-organized criticality considers how a dynamic system can converge to a "critical" state where power law fluctuations occur (for example, a sand pile with many avalanches). The concept of percolation began as the study of how a fluid would filter through a random medium, but has also found applications in other areas like how immunization affects the spread of epidemics. For some applications of these ideas in economics, see Bak, Chen, Scheinkman, and Woodford (1993) and Nirei (2006).

Networks are full of power laws: for instance (probably because of random growth), the popularity of websites, as measured by the number of sites linking to a website, can be represented by a power law (Barabasi and Albert 1999).

Biology is replete with intriguing and seemingly universal relations of the power law type. For instance, the energy that an animal of mass M requires to function is proportional to $M^{3/4}$—rather than M as a simple "constant return to scale" model would predict, as illustrated in Figure 7. West et al. (1997) have proposed the following explanation: If one wants to design an optimal network system to send nutrients to the animal, one designs a fractal system; the resulting efficiency generates the $M^{3/4}$ law. There is an interesting lesson: a priori, lots of things could matter for energy consumption—for example, climate, predator or prey status, thickness of the fur—and they probably do matter to a limited extent. However, in its essence, an animal is best viewed as a network in which nutrients circulate at maximum efficiency. Understanding the power laws forces the researcher to forget, in the first pass, about the details. Likewise, this research shows similar laws for a host of variables, including life expectancy (which scales as $M^{1/4}$). Here a possible interpretation is that the animal is constructed optimally, given engineering constraints that are biologically determined.

[11] Another related perspective on power laws in the stock market concerns potential disasters. Barro and Jin (2011) document a power law distribution of macroeconomic disasters, which may explain abrupt shocks to valuations and indeed many puzzles in finance (Gabaix 2012; Kelly and Jiang 2014). This offers a potentially fruitful direction of research, even if it is still just a hypothesis. See Fu et al. (2005) and Riccaboni et al. (2008).

Figure 7
Metabolic Rate as a Function of Mass across Animals



Source: West, Brown, and Enquist (2000).
Notes: "Metabolic rate" is energy requirement per day. The slope of this log-log graph is 3/4: the metabolic rate of an animal of mass M is proportional to $M^{3/4}$ (Kleiber's law).

Perhaps surprisingly, this type of mechanism, generating power laws from maximum fitness, doesn't seem to have been much studied in economics. For instance, the economy resembles a network with power-law-distributed firms: does this pattern arise from optimality as opposed to randomness? It would be nice to know. Likewise, Zipf's law holds for the usage frequency of words. The simplest explanation is via random growth (as the popularity of words follow a random growth process). However, this law might instead reflect an "optimal" organization of mental categories, perhaps in some tree-like structure? Again, one would like to know.

Conclusion

All economists should become familiar with power laws and the basic mechanisms that generate them because power laws are everywhere. One place to teach power laws is in the macro sequence when discussing models with heterogeneous agents and sectors. The future of power laws as a subject of research looks very healthy: when datasets contain enough variation in some "size"-like factor, such as income or number of employees, power laws seem to appear almost invariably. In addition, power laws can guide the researcher to the essence of a phenomenon.

For instance, consider city size. Lots of things might conceivably be important for city size: specialization, transportation cost, congestion, positive externalities in human capital, and others. The power law approach concludes that while those things exist, they are not the essence of what determines the distribution of city sizes: the essence is random growth with a small friction. To generate the random growth, a judicious mix of the traditional ingredients in the study of cities may be useful, but to orient understanding, one should first think about the essence, and after that about the economic underpinnings.

Many open questions remain about the prevalence and explanation of power laws, and in many of these areas, new data have recently become available. Along with the earlier examples of the distribution of income, wealth, firm size, and city size, here is a sampling of some other questions.

In the study of international trade, the "gravity equation" suggests that the trade flow between any two countries is proportional to the GDP of the two countries (a result which can be derived from a simple model of economies with constant returns to scale) and declines with distance, a finding which seems intuitive. However, the relationship between volume of trade and distance seems to decline with the inverse of distance to the power 1—and there doesn't seem to be any obvious intuitive reason for this particular scale factor. As one possible underlying reason, Chaney (2013) proposes an ingenious model linking this distance to the probability of forming a link in a random growth model of networks, which, under Zipf's law, generates the appropriate coefficient. Similar scaling holds for migration (Levy 2010; Levy and Goldenberg 2014), raising similar questions. For a discussion of power laws in trade, useful starting points include Helpman, Melitz, and Yeaple (2004) and Eaton, Kortum, and Kramarz (2011).

Why is aggregate production in a high-income economy (roughly) Cobb–Douglas with a capital share about 1/3? Jones (2005) generates the functional form but not the particular exponent. Perhaps finding a way to generate this exponent will suggest a deeper understand of the causes of technical progress.

In the study of networks and granularity, how big is the volatility generated from idiosyncratic shocks propagated and amplified in networks?

Is random growth the fundamental origin of power law relationships for the distribution of cities and firms? Or is there potentially some other very different underlying force, like the economics of superstars, or efficiency maximization? Although there is evidence that Gibrat's law seems to roughly hold (that is, the mean and variance of growth rates of a given city or firm is roughly independent of its size, see Ioannides and Overman 2003; Eeckhout 2004), the issue is not settled, as the literature hasn't fully differentiated between permanent shocks and transitory ones, and plain measurement error.

As more of the huge datasets often referred to as "big data" become available, it will be important to characterize and order them. Scaling questions are a natural way to do that and have met with great success in the natural sciences.

A reader seeking a gentle introduction to power law techniques might begin with Gabaix (1999) and then move on to more systematic exposition in Gabaix

(2009), which contains many other pointers. For extreme value theory, Embrechts, Klüppelberg, and Mikosch (1997) is very pedagogical. For networks, Jackson (2008) and Newman (2010) are now standard references. Mantegna and Stanley (2000) contains an accessible introduction to the field of econophysics, while Sornette (2006) contains many interesting physics mechanisms generating scaling.

■ *For excellent research assistance, I thank Chenxi Wang and Jerome Williams, and for comments, I thank David Autor, Chang-Tai Hseih, Chad Jones, Bryan Kelly, Moshe Levy, Timothy Taylor, and David Weinstein.*

References

Acemoglu, Daron, Ufuk Akcigit, and William Kerr. 2015. "Networks and the Macroeconomy: An Empirical Exploration." *NBER Macroeconomics Annual.*

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica* 80(5): 1977–2016.

Allais, Maurice. 1947. *Economie et Intérêt.* Paris: Imprimerie Nationale.

Amiti, Mary, and David Weinstein. 2013. "How Much Do Bank Shocks Affect Investment? Evidence from Matched Bank-Firm Loan Data." NBER Working Paper 18890.

Atkinson, Anthony B, Thomas Piketty, and Emmanuel Saez. 2011. "Top Incomes in the Long Run of History." *Journal of Economic Literature* 49(1): 3–71.

Axtell, Robert L. 2001. "Zipf Distribution of US Firm Sizes." *Science* 293(5536): 1818–20.

Bak, Per, Kan Chen, José Scheinkman, and Michael Woodford. 1993. "Aggregate Fluctuations from Independent Sectoral Shocks: Self-organized Criticality in a Model of Production and Inventory Dynamics." *Ricerche Economiche* 47(1): 3–30.

Barabási, Albert-László, and Réka Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286(5439): 509–512.

Barro, Robert J., and Tao Jin. 2011. "On the Size Distribution of Macroeconomic Disasters." *Econometrica* 79(5): 1567–89.

Baumol, William J., and James Tobin. 1989. "The Optimal Cash Balance Proposition: Maurice Allais' Priority." *Journal of Economic Literature* 27(3): 1160–62.

Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud. 2013. "Productive Cities: Sorting, Selection, and Agglomeration." *Journal of Political Economy* 122(3): 507–553.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu. 2011. "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents." *Econometrica* 79(1): 123–57.

Bouchaud, Jean-Philippe, J. Doyne Farmer, and Fabrizio Lillo. 2009. "How Markets Slowly Digest Changes in Supply and Demand." In *Handbook of Financial Markets: Dynamics and Evolution*, edited by Klaus Reiner Schenk-Hoppe and Thorsten Hens, 57–160. North Holland.

Carvalho, Vasco. 2014. "From Micro to Macro via Production Networks." *Journal of Economic Perspectives* 28(4): 23–48.

Carvalho, Vasco, and Xavier Gabaix. 2013. "The Great Diversification and Its Undoing." *American Economic Review* 103(5): 1697–1727.

Champernowne, D. G. 1953. "A Model of Income Distribution." *Economic Journal* 63(250): 318–51.

Chaney, Thomas. 2013. "The Gravity Equation in International Trade: An Explanation." NBER Working Paper 19285.

Davis, Donald R., and David E. Weinstein. 2002. "Bones, Bombs and Break Points: The Geography of Economic Activity." *American Economic Review* 92(5): 1269–89.

di Giovanni, Julian, and Andrei A. Levchenko. 2012. "Country Size, International Trade, and Aggregate Fluctuations in Granular Economies." *Journal of Political Economy* 120(6): 1083–1132.

di Giovanni, Julian, Andrei Levchenko, and Isabelle Mejean. 2014. "Firms, Destinations, and

Aggregate Fluctuations." *Econometrica* 82(4): 1303–40.

Dittmar, Jeremiah. 2011. "Cities, Markets, and Growth: The Emergence of Zipf's Law." http://www.jeremiahdittmar.com/files/Zipf_Dittmar.pdf.

Eaton, Jonathan, Samuel Kortum, and Francis Kramarz. 2011. "An Anatomy of International Trade: Evidence from French Firms." *Econometrica* 79(5): 1453–98.

Edmans, Alex, and Xavier Gabaix. 2011. "The Effect of Risk on the CEO Market." *Review of Financial Studies* 24(8): 2822–63.

Edmans, Alex, Xavier Gabaix, and Augustin Landier. 2009. "A Multiplicative Model of Optimal CEO Incentives in Market Equilibrium." *Review of Financial Studies* 22(12): 4881–4917.

Eeckhout, Jan. 2004. "Gibrat's Law for (All) Cities." *American Economic Review* 94(5): 1429–51.

Embrechts, Paul, Claudia Klüppelberg, and Thomas Mikosch. 1997. *Modelling Extremal Events for Insurance and Finance.* New York: Springer.

Foerster, Andrew, Pierre-Daniel Sarte, and Mark Watson. 2011. "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production." *Journal of Political Economy* 119(1): 1–38.

Fu, Dongfeng, Fabio Pammolli, Sergey V. Buldyrev, Massimo Riccaboni, Kaushik Matia, Kazuko Yamasaki, and H. Eugene Stanley. 2005. "The Growth of Business Firms: Theoretical Framework and Empirical Evidence." *PNAS* 102(52): 18801–06.

Gabaix, Xavier. 1999. "Zipf's Law for Cities: An Explanation." *Quarterly Journal of Economics* 114(3): 739–67.

Gabaix, Xavier. 2009. "Power Laws in Economics and Finance." *Annual Review of Economics* 1(1): 255–93.

Gabaix, Xavier. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica* 79(3): 733–72.

Gabaix, Xavier. 2012. "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-finance." *Quarterly Journal of Economics* 127(2): 645–700.

Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 129(4): 1661–1710.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley. 2003. "A Theory of Power-Law Distributions in Financial Market Fluctuations." *Nature* 423(6937): 267–70.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley. 2006. "Institutional Investors and Stock Market Volatility." *Quarterly Journal of Economics* 121(2): 461–504.

Gabaix, Xavier, and Rustam Ibragimov. 2011. "Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents." *Journal of Business Economics and Statistics* 29(1): 24–39.

Gabaix, Xavier, and Augustin Landier. 2008. "Why Has CEO Pay Increased So Much?" *Quarterly Journal of Economics* 123(1): 49–100.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2015. "The Dynamics of Inequality." NBER Working Paper 21363.

Garicano, Luis, and Estaban Rossi-Hansberg. 2006. "Organization and Inequality in a Knowledge Economy." *Quarterly Journal of Economics* 121(4): 1383–1435.

Geerolf, François. 2015. "A Static and Microfounded Theory of Zipf's Law for Firms and of the Top Labor Income Distribution." Working Paper, UCLA.

Gibrat, Robert. 1931. *Les inégalités économiques.* Recueil Sirey.

Gopikrishnan, Parameswaran, Vasiliki Plerou, Luis A. Nunes Amaral, Martin Meyer, and H. Eugene Stanley. 1999. "Scaling of the Distribution of Fluctuations of Financial Market Indices." *Physical Review E* 60(5): 5305–16.

Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple. 2004. "Export versus FDI with Heterogeneous Firms." *American Economic Review* 94(1): 300–316.

Ioannides, Yannis M., and Henry G. Overman. 2003. "Zipf's Law for Cities: An Empirical Examination." *Regional Science and Urban Economics* 33(2): 127–37.

Jackson, Matthew O. 2008. *Social and Economic Networks.* Princeton University Press.

Jensen, Michael C., and Kevin J. Murphy. 1990. "Performance Pay and Top-Management Incentives." *Journal of Political Economy* 98(2): 225–65.

Jones, Charles I. 2005. "The Shape of Production Functions and the Direction of Technical Change." *Quarterly Journal of Economics* 120(2): 517–49.

Kelly, Bryan, and Hao Jiang. 2014. "Tail Risk and Asset Prices." *Review of Financial Studies* 27(10): 2841–71.

Kelly, Bryan, Hanno Lustig, and Stijn Van Nieuwerburgh. 2013. "Firm Volatility in Granular Networks." NBER Working Paper 19466.

Koren, Miklós, and Silvana Tenreyro. 2013. "Technological Diversification." *American Economic Review* 103(1): 378–414.

Krugman, Paul. 1996. "Confronting the Mystery of Urban Hierarchy." *Journal of the Japanese and International Economies* 10(4): 399–418.

Kyle, Albert, and Anna Obizhaeva. 2013. "Large Bets and Stock Market Crashes." AFA 2013 San Diego Meetings Paper. Available at SSRN: http://ssrn.com/abstract=2023776.

Kyle, Albert, and Anna Obizhaeva. 2014. "Market Microstructure Invariance: Theory and

Empirical Tests." Available at SSRN: http://ssrn.com/abstract=1687965.

Lee, Youngki, Luís A. Nunes Amaral, David Canning, Martin Meyer, and H. Eugene Stanley. 1998. "Universal Features in the Growth Dynamics of Complex Organizations." *Physical Review Letters* 81(15): 3275–78.

Levy, Moshe. 2010. "Scale-free Human Migration and the Geography of Social Networks." *Physica A: Statistical Mechanics and Its Applications* 389(21): 4913–17.

Levy, Moshe, and Jacob Goldenberg. 2014. "The Gravitational Law of Social Interaction." *Physica A* 393: 418–26.

Levy, Moshe, and Sorin Solomon. 1996. "Power Laws are Logarithmic Boltzmann Laws." *International Journal of Modern Physics C* 7(4): 595–601.

Lucas, Robert E., and Benjamin Moll. 2014. "Knowledge Growth and the Allocation of Time." *Journal of Political Economy* 122(1): 1–51.

Luttmer, Erzo G. J. 2007. "Selection, Growth, and the Size Distribution of Firms." *Quarterly Journal of Economics* 122(3): 1103–44.

Mantegna, Rosario N., and H. Eugene Stanley. 2000. *An Introduction to Econophysics: Correlations and Complexity in Finance.* Cambridge University Press.

Newman, Mark E. J. 2010. *Networks: An Introduction*. Oxford University Press.

Nirei, Makoto. 2006. "Threshold Behavior and Aggregate Fluctuation." *Journal of Economic Theory* 127(1): 309–322.

Pareto, Vilfredo. 1896. *Cours d.économie politique.* Librairie Droz.

Persky, Joseph. 1992. "Retrospectives: Pareto's Law." *Journal of Economic Perspectives* 6(2): 181–92.

Piketty, Thomas, and Gabriel Zucman. 2014. "Capital is Back: Wealth-Income Ratios in Rich Countries, 1700–2010." *Quarterly Journal of Economics* 129(3): 1255–1310.

Plerou, Vasiliki, Parameswaran Gopikrishnan, Luis A. Nunes Amaral, Martin Meyer, and H. Eugene Stanley. 1999. "Scaling of the Distribution of Price Fluctuations of Individual Companies." *Physical Review E* 60(6): 6519–29.

Plerou, Vasiliki, Parameswaran Gopikrishnan, and H. Eugene Stanley. 2005. "Quantifying Fluctuations in Market Liquidity: Analysis of the Bid-Ask Spread." *Physical Review E* 71(4): 046131.

Riccaboni, Massimo, Fabio Pammolli, Sergey V. Buldyrev, Linda Ponta, and H. Eugene Stanley. 2008. "The Size Variance Relationship of Business Firm Growth Rates." *PNAS* 105(50): 19595–19600.

Roberts, David R. 1956. "A General Theory of Executive Compensation Based on Statistically Tested Propositions." *Quarterly Journal of Economics* 70(2): 270–94.

Rosen, Sherwin. 1981. "The Economics of Superstars." *American Economic Review* 71(5): 845–58.

Rossi-Hansberg, Esteban, and Mark L. J. Wright. 2007. "Urban Structure and Growth." *Review of Economic Studies* 74(2): 597–624.

Rozenfeld, Hernán D., Diego Rybski, Xavier Gabaix, and Hernán A. Makse. 2011. "The Area and Population of Cities: New Insights from a Different Perspective on Cities." *American Economic Review* 101(5): 2205–25.

Samuelson, Paul A. 1969. "The Way of an Economist." *International Economic Relations: Proceedings of the Third Congress of the International Economic Association*, pp. 1–11. London: MacMillan.

Simon, Herbert A. 1955. "On a Class of Skew Distribution Functions." *Biometrika* 44(3/4): 425–40.

Solomon, Sorin, and Peter Richmond. 2001. "Power Laws of Wealth, Market Order Volumes and Market Returns." *Physica A: Statistical Mechanics and its Applications* 299(1): 188–97.

Sornette, Didier. 2006. *Critical Phenomena in Natural Sciences.* Springer Science & Business.

Stanley, Michael H. R., Luís A. N. Amaral, Sergey V. Buldyrev, Shlomo Havlin, Heiko Leschhorn, Philipp Maass, Michael A. Salinger, and H. Eugene Stanley. 1996. "Scaling Behaviour in the Growth of Companies." *Nature* 379(6568): 804–06.

Summers, Robert, and Alan Heston. 1991. "The Penn World Table (Mark 5): An Expanded Set of International Comparisons, 1950–1988." *Quarterly Journal of Economics* 106(2): 327–68.

Terviö, Marko. 2008. "The Difference that CEOs Make: An Assignment Model Approach." *American Economic Review* 98(3): 642–68.

Toda, Alexis, and Kieran Walsh. 2015. "The Double Power Law in Consumption and Implications for Testing Euler Equations." *Journal of Political Economy* 123(5): 1177–1200.

Turcotte, Donald L. 1997. *Fractals and Chaos in Geology and Geophysics.* Cambridge University Press.

US Bureau of the Census. 2012. *Statistical Abstract of the United States: 2012*. http://www.census.gov/compendia/statab/2012edition.html.

West, Geoffrey B., James H. Brown, and Brian J. Enquist. 1997. "A General Model for the Origin of Allometric Scaling Laws in Biology." *Science* 276(5309): 122–26.

West, Geoffrey B., James H. Brown, and Brian J. Enquist. 2000. "The Origin of Universal Scaling Laws in Biology." In *Scaling in Biology*, edited by J. H. Brown, and G. B. West, 87–112. Oxford University Press.

Zipf, George Kingsley. 1949. *Human Behavior and the Principle of Least Effort.* Cambridge: Addison-Wesley.