

# The distribution of technological progress

Gianluca Carnabuci

Received: 10 January 2010 / Accepted: 24 January 2012 / Published online: 13 April 2012  
© Springer-Verlag 2012

**Abstract** The size distribution of the domains of US-patented technological knowledge obeys an exponential law, revealing a disproportionate concentration of progress among larger domains. Our analyses suggest that this phenomenon is explained by a combination of two factors. First, domains' trajectories of growth have inherently different potentials. Second, differences in domains' potentials are magnified by a mechanism—domains' self-hybridization—endogenous to the process of knowledge growth. Our results show that in addition to being stable, the observed distribution of technological progress is likely to arise under very general conditions.

**Keywords** Technological progress · Knowledge growth · Size distribution · Patents · Concentration

**JEL Classification** O030

## 1 Introduction

Understanding why certain areas of technological knowledge grow faster than others is of key importance for both economic theory and policy (Rosenberg 1976; Harberger 1988). Supported by a long tradition of empirical research (Gilfillan 1930; Rossman 1931), two main explanations became predominant among economists. On the one hand, so called demand-pull studies hold that the direction of technological progress reflects underlying market forces. Griliches (1957), for example, provides evidence for the importance of demand-pull dynamics by showing that the geographic diffusion of agricultural innovations is affected by market size. In a similar vein, Schmookler

---

G. Carnabuci (✉)  
University of Lugano, Lugano, Switzerland  
e-mail: Gianluca.carnabuci@usi.ch

(1966) analyzes the production of technological knowledge related to railway equipment, showing that the higher is an application's actual or potential demand the more inventive activity is directed toward it. On the other hand, the so called technology-push view posits that the "strength of science" determines the technological opportunities available across technology domains, thereby affecting the direction of technological progress. For example, [Rosenberg \(1976\)](#) and Nelson and Winter (1977) argue that changes in the scientific knowledge base of an industry have a major impact on which trajectories of technological innovation are pursued.

Rather than focusing on the factors affecting the growth of specific domains of technology, the present study adopts an inductive approach and examines large-scale data on the distribution of progress *across* technology domains. The paper is organized as follows. We begin by describing our empirical data, which comprises all technological knowledge patented in the US over a quarter of a century. We then use this data to statistically analyze the distribution of technological progress, showing that the growth of technological knowledge tends to be disproportionately concentrated among a few large technology domains. We move on to demonstrate that this empirical finding reflects an enduring characteristic of technological progress. Inspired by the notion that knowledge grows by hybridization of existing ideas ([Usher 1954](#); [Romer 1994](#); [Weitzman 1998](#)), we then identify a causal mechanism that may be responsible for the observed distribution of technological progress, and we test its hypothesized effect. We conclude by discussing the scope and implications of our results.

## 2 Data and measures

Digitalized patent archives offer a rich data source to study the progress of codified technological knowledge. In this paper, we use the NBER Patent and Patent Citations Dataset, comprising data on all patents (over two million) granted by the USPTO between 1975 and 1999, and all citations between them (see [Hall et al. 2001](#) for a detailed description of the data). Based on consolidated methods (e.g., [Griliches 1984](#); [Jaffe et al. 2000](#); [Carnabuci and Bruggeman 2009](#); [Carnabuci 2010](#)), we operationalize the data in three ways. First, patents indicate novel, non-trivial, and usable inventions. We thus use patent counts to measure the advancement of technological knowledge (e.g., [Scherer 1965](#))<sup>1</sup>. Second, USPTO patents must cite all previous patents that the focal invention draws from, the so-called *prior art*. Accordingly, patent citations may be regarded as an indication of the previous knowledge hybridized into the focal invention (e.g., [Jaffe et al. 2000](#)). Third, USPTO officers must partition all patents into one of over four hundred internationally standardized technological categories, also called primary patent classes<sup>2</sup>. We follow prior research and use the primary patent classes

---

<sup>1</sup> Although citation-based measures are known to yield more accurate estimates of technological output, here we prefer to use simple patent counts. As will become clear, this choice provides a more straightforward way to disentangle causal dynamics in the context of the analyses presented. However, all the statistical analyses presented in this paper have also been performed using citation-weighted counts, leading to equivalent results.

<sup>2</sup> The 418 domains have about 120.000 subclasses that are less tractable due to their number and are less reliable due to higher arbitrariness of categories.

to indicate technology domains (Henderson et al. 2005; Carnabuci and Bruggeman 2009).

### 3 Size distribution

We regard the amount of inventions produced within a given technology domain over a given time interval as its size. Figure 1 shows a scatter plot of domains' size, as measured by the number of patents granted within the 25 years we observed, against domains' rank, in linear (A) and semi-logarithmic (B) scale. The data follow a straight line in the semi-logarithmic plot, indicating that domains' sizes are distributed exponentially.

This impression is confirmed by Fig. 2, which reports a probability plot of the observed size distribution against the theoretical exponential distribution. Ordinary





Fig. 1 Distribution of domains' size by rank in linear (a) and semi-logarithmic (b) plot

Fig. 2 Plot of the observed probability distribution (gray line) against the expected exponential probability distribution (black line)



Least Square regression of the logarithmic transform of domains' rank on domains' size makes it possible to establish statistically how closely the data fit an exponential function. The regression yielded an  $R^2$  of 0.963, indicating very close fit.

To examine how the observed size distribution changed over time, we time-partitioned the data into five subsequent cross-sections (based, respectively, on the patents granted in the intervals 1975 through 1979, 1980 through 1984, 1985 through 1989, 1990 through 1994, and 1995 through 1999), and we again regressed the log of domains' rank on domains' size for each of the time intervals. The data turned out to be well described by an exponential distribution during each of the time intervals, with  $R^2$  deviating only slightly from the one estimated for the entire observation period. This suggests that the observed exponential distribution reflects a stable trait in the growth of US-patented knowledge. In order to provide evidence that this stability is likely to persist in the longer run, in the following sections we will investigate the generative mechanism behind the observed size distribution by analyzing how domains' growth rates have been varying in relation to domains' size.

#### 4 Toward a heavy-tail distribution?

There is increasing evidence that two distributions, power-law, and lognormal, are empirically ubiquitous. Thus, it seems relevant to investigate if the observed exponential distribution of technological progress in fact reflects a transition toward one of these two distributions. Both the power-law and the log-normal distributions are highly skewed, with tails heavier than in the exponential, and apply to systems wherein concentration of resources is extreme. A classic example is the distribution of income in capitalist societies, which Pareto showed to obey a power-law in the upper tail (Pareto 1897). Other well-known cases are the size distribution of cities (Gabaix 1999), firms (Gibrat 1931), human sexual contacts (Liljeros et al. 2001), the World Wide Web (Albert et al. 1999), usage of words (Zipf 1949), and citations received by scientists (Lotka 1926). The ubiquity of power-law and log-normal distributions led scholars to investigate their generative processes in detail. Not unexpectedly, it turned out that the two distributions are inherently related, and represent the equilibrium solutions of a wide-ranging growth mechanism (Mitzenmacher 2004). The distinctive feature of this mechanism, which in the literature has been interchangeably termed *law of proportionate effect* (Gibrat 1931), *cumulative advantage* (Price 1980), *Gibrat's law* (Simon 1955) and *Matthew Effect* (Merton 1968), is that the size of an object affects positively its subsequent growth.

##### 4.1 Gibrat's model

In his seminal book, Gibrat was the first to provide a general mathematical model of growth by cumulative advantage (Gibrat 1931). His model shows that given a fixed population, the size of its units will converge to a log-normal distribution when growth dynamics are driven by a stochastic multiplicative component. Formally, let  $s$  denote the size of a population unit,  $t$  denote time, and  $\varepsilon$  be a stochastic term drawn from a random distribution; then Gibrat's law of proportionate effects can be expressed as:

$$s_t - s_{t-\Delta t} = \varepsilon_t s_{t-\Delta t} \tag{1}$$

Gibrat assumed that  $\varepsilon_t$  is independent of  $s_t$ ; that  $\varepsilon_t$  features no significant inter-temporal correlation; and, that there is no interaction between units. Given that:

$$s_t = s_0(1 + \varepsilon_1)(1 + \varepsilon_2) \cdots (1 + \varepsilon_t), \tag{2}$$

the logarithm of  $s_{it}$  follows a random walk. Because in a short period of time  $\varepsilon_{it}$  can be regarded as small, the approximation  $\ln(1 + \varepsilon_t) = \varepsilon_t$  is justified. Then, taking logs:

$$\ln s_t \cong \ln s_0 + \sum_{t=1}^T \varepsilon_t \tag{3}$$

Due to central limit theorem, after a sufficiently long time the logarithm of units' size will converge to a normal distribution; that is, units' size will converge to a log-normal distribution.

In subsequent studies, Gibrat's model has been extended to show that the mechanism of cumulative advantage may also lead to another commonly observed heavy-tail distribution: the power-law distribution. For example, [Simon and Bonini \(1958\)](#) showed that sizes converge to a power-law distribution if entry dynamics and the principle of *preferential attachment* are introduced in Gibrat's model; [Kesten \(1973\)](#) showed that a power-law distribution results from integrating Gibrat's purely multiplicative mechanism with an additive term; in [Levy and Solomon \(1996\)](#), a power-law distribution is obtained by imposing a lower-bound size on Gibrat's model.

#### 4.2 Testing for cumulative advantage

Based on Gibrat's model, one can derive a straightforward method to empirically test whether processes of growth by cumulative advantage are at work that may lead to heavy-tail distributions. As shown in (3), cumulative advantage occurs when units' growth rates are unrelated to their current size and only depend on the sum of idiosyncratic shocks. Then, the following logarithmic specification can be used to test for growth processes driven by cumulative advantage:

$$\ln s_{it} - \ln s_{it-1} = \alpha + \beta \ln s_{it-1} + u_{it} \tag{4}$$

where  $i$  and  $j$  are units, and  $u_{it}$  is a random variable that satisfies the following conditions:

$$E(u_{it} | s_{it-s}, s > 0) = 0$$

$$E(u_{it}u_{jt} | s_{it-s}, s > 0) = \begin{cases} \sigma^2 & i = j, t = \tau \\ 0 & \text{otherwise} \end{cases}$$

**Table 1** Fixed-effects regression of domains' size on domains' growth rates

	Dependent variable: domains' growth rate			
	Coefficient	Std. error	<i>t</i> Value	Sig.
Intercept	1.337	0.132	10.06	0.000
Domains' size	−0.212	0.022	−9.31	0.000
Number of units	417			
Periods	4			
Number of observations	1653			
<i>F</i>	86.65			
Prob. > <i>F</i>	0.0000			

In Eq. (4), if  $\beta < 0$ , Gibrat's law can be rejected, and the process is mean reverting<sup>3</sup>. If  $\beta \cong 0$ , Gibrat's law is confirmed, and cumulative advantage is at work. The case  $\beta > 0$  can be regarded as a transient phase of a Gibrat-like process, whereby the process of cumulative advantage leads to explosive growth patterns.

The model specification described by (4) has been widely used to test Gibrat's law. However, the model assumes that unit effects are homogeneous—i.e., that  $\alpha_i = \alpha$  or, equivalently, that  $\sigma^2(\alpha_i) = 0$ . Hence, under the hypothesis that Gibrat law does not hold, all units are assumed to converge to the same steady-state size. Because this assumption seems unwarranted in the case of technology domains, we use the Breusch and Pagan Lagrange multiplier (1980) to test the homogeneity assumption that  $\sigma^2(\alpha_i) = 0$ . The test clearly rejects the null hypothesis of homogeneity ( $\chi^2 = 60.20$ ; Probability >  $\chi^2 = 0.000$ ); thus, the use of model (4) for the analysis of our data would overestimate the parameter of interest,  $\beta$ , whenever  $\beta < 0$  (Breusch and Pagan 1980). We therefore reformulate (4) as a variable-intercept model:

$$\ln s_{it} - \ln s_{it-1} = \alpha_i + \beta \ln s_{it-1} + u_{it} \quad (5)$$

Domains' progress is likely to be interdependent, as the knowledge generated in one domain may spawn progress in others, resulting in autocorrelated residuals and biased parameter estimates. To eliminate this problem, we apply a spatial autocorrelation model (Anselin 1998), where proximities among domains are measured by the row-normalized aggregate citation patterns among them (Carnabuci 2010). Table 1 reports the results, which show that over the observation period domains' growth rates decreased significantly with domains' size ( $\beta = -0.21$ ,  $p < 0.001$ ). This is in stark contradiction with the hypothesis of cumulative advantage, and indicates that domains' observed growth patterns are not compatible with the emergence of heavy-tail size distributions.

<sup>3</sup> Note that the dependent variable in Eq. (4) is expressed as a difference of logarithms; hence, it approximates percentage change.

## 5 Converging to a common size?

The negative value of  $\beta$  in model (5) implies that a process of mean-reversion is afoot in the distribution of US-patented knowledge across technology domains. This reversion, however, means neither that domains have been converging to a common size, nor that they will. As indicated by the Breusch Pagan Lagrange multiplier test of homogeneity, there are significant domain-specific effects, revealing permanent differences across domains' growth paths, which in model (5) are accounted for by a variable-intercept specification. Hence, the estimated negative relationship between domains' current size and subsequent growth rate expressed by  $\beta$  should be interpreted as evidence that domains tend toward different steady-state sizes.

Model (5) lends itself to a substantive interpretation of why domains cannot converge to a common size. The model says that growth rates decline similarly across all domains as these grow larger, while domains' initial growth rates are significantly different. A metaphor for this is a group of cars consuming a similar amount of fuel per kilometer, but beginning their trip with different amounts of fuel in their tank. Likewise, all domains progressively exhaust their growth potential, but they differ in their initial endowment. Differences in domains' inherent growth potential can be explained in several ways, the most obvious one being that possibilities of technological developments are, at least in part, determined by idiosyncratic technical and physical constraints (Nelson and Winter 1982; Rosenberg 1994; Mokyr 2002).

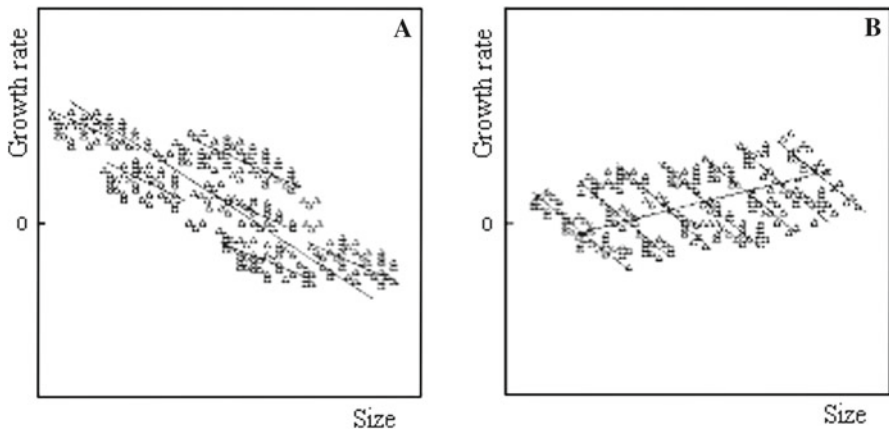
## 6 Domains' steady-state sizes

In the last two sections we described the overall picture: on the one hand, domains' growth dynamics are incompatible with the emergence of heavy-tail size distributions while, on the other, there are permanent differences across domains that prevent them from converging to a common size. Now, we would like to turn to a more detailed view, and ask whether we should expect the observed concentration of technological progress to increase or decrease as domains approach their steady-state size. To get a more nuanced inspection of the available data, we propose to complement the results obtained under model (5) with the estimation of the following model:

$$\frac{1}{T} \sum_t (\ln s_{it} - \ln s_{it-1}) = \lambda + \tau \frac{1}{T} \sum_t \ln s_{it-1} + \frac{1}{T} \sum_t \varepsilon_{it} + v_i \quad (6)$$

where  $T$  is the number of observed time points,  $\lambda$  and  $\tau$  are parameters to be estimated, and  $\varepsilon_{it}$  is an error term. In model (6), longitudinal observations are collapsed into domains' means across time, thus the only variance left is between domains. For this reason, this model is also referred to as the *between-effects* estimator. As it results in substantial loss of information, the between-effects estimator is not much used in practice. However, in combination with model (5), it can shed light on the nature of domains' observed size differences. Figure 3 helps illustrate how.

Both Fig. 3a and b represent sets of hypothetical domains of technological knowledge whose growth paths display both domain-specific effects (the data points are



**Fig. 3** **a** and **b** Display two hypothetical size-growth *scatter plots* featuring negative within-domain effects. The between-domain effects are positive in **a**, but negative in **b**

clustered within domains), and a mean-reverting process towards domain-specific steady-state sizes (the clusters follow a negative slope). In other words, both figures represent situations that are perfectly compatible with the estimates obtained for US-patented technology domains, i.e.,  $\alpha_i \neq \alpha$ , and  $\beta < 0$  in model (5). The difference between A and B is in the patterns observed *between* domains, represented in each figure by the long line interpolating the clusters, and whose slope is captured by parameter  $\tau$  in model (6). As can be seen, in A this line has a negative slope, thus  $\tau < 0$ , whereas in B its slope is positive, thus  $\tau > 0$ . This means that in A, domains that are on average larger over the observation period display a lower average growth rate; in contrast in B, larger domains also display a higher average growth rate. As we will try to explain, this implies that in A, larger domains have on the average less potential for future hybridization and growth, whereas the opposite is true for B.

To illustrate why this is the case, let me again refer to Fig. 3. Notice that that the zero-point on the Y-axis indicates the point where domains do not grow, i.e., where they reach their steady-state size<sup>4</sup>. As can be seen, in A a part of the observed size differences between larger and smaller domains is due to larger domains having exploited, on average, more of their growth potential. While larger domains are bound to grow yet relatively modestly or even to shrink to reach their steady-state size, smaller domains will reach their steady-state size only after growing a great deal more. Thus, one can expect size differences to diminish as domains approach their steady-state sizes, leading to a decrease in the concentration of technological progress. In B the situation is reversed. Here the domains that have been observed to be smaller during the observation period are, on average, closer to or larger than their steady-state

<sup>4</sup> Accordingly, when domains are above this point their sizes grow, and when they are below their sizes shrink. Clearly, this representation abstracts away from time.



**Table 2** Between-effects regression of domains' size on domains' growth rates

	Dependent variable: domains' growth rate			
	Coefficient	Std. error	<i>t</i> Value	Sig.
Intercept	-0.284	0.055	-5.15	0.000
Domains' size	0.062	0.009	7.02	0.000
Number of units	417			
Periods	4			
Number of observations	1653			
<i>F</i>	49.28			
Prob. > <i>F</i>	0.0000			

size. Accordingly, in this case the concentration of technological progress among large domains will increase as domains approach their steady-state size<sup>5</sup>.

The estimates<sup>6</sup> of model (6) for the US-patent data are reported in Table 2, showing that domains' average size is positively associated with domains' average growth rate (i.e.,  $\tau > 0$ ) and the association is highly significant. It follows that the growth paths of the domains of US-patented knowledge conform to the situation sketched in Fig. 3b, thus technological progress should become even more concentrated among large domains as domains approach their steady-state size.

## 7 Explaining the concentration of technological progress

The analyses presented so far indicate that the observed concentration of knowledge progress among larger domains reflects a continuing feature in the advancement of US-patented knowledge. Without delving deeper into the causes of this phenomenon, however, it is impossible to establish its generality beyond the empirical case observed. According to model (5), domains' steady-state sizes are determined by two parameters: a domain-specific intercept, and the relationship between domains' size and growth. The first of these two parameters, as we have argued, can be interpreted as the growth potential inherent in domains' path of development: clearly, domains' inherent growth potentials have a bearing on their steady-state size. The second of these two parameters—the size-growth relationship—determines the degree to which domains' inherent growth potentials are realized. To see why, consider the following two limiting cases. If domains' growth rates decay infinitely fast as a function of domains' size, then the size of all domains drops to zero regardless of domains'

<sup>5</sup> To get back to the car metaphor introduced earlier, in A larger domains are like cars that have proceeded a longer way but have less fuel left in their tank. In B, in contrast, larger domains are like cars that have gone a longer way, and still have more fuel in their tank than other cars. By the time all cars are left without fuel, in A the cars that were lagging behind will have caught up to some degree, whereas in B they will be left even further behind.

<sup>6</sup> As with Model (5), we correct for the potential network autocorrelation in the residuals by means of a network disturbance model.

inherent growth potential. In contrast, if domains' growth rates do not decay at all with domains' size (as is the case with processes of cumulative advantage), then domains' inherent growth potentials will be perpetually amplified.

We posit that there exists a mechanism, integral to the process of knowledge growth, which intervenes in the size-growth relationship so as to systematically boost concentration in knowledge progress. Our argument is the following. Most scholars agree that the engine of technological progress is knowledge hybridization, whereby new ideas are generated by combining existing ones (Usher 1954; Romer 1994; Weitzman 1998). As the bulk of knowledge hybridizations occur within rather than between domains (Nelson and Winter 1982; Dosi 1982), the number of self-hybridizations made within domains, and hence the progress domains make, should increase proportionally to domains' size<sup>7</sup>. If this is the case, then domains' self-hybridizations will counter the tendency of growth rates to decay with domains' size and, thereby, will magnify the size differences induced by domains' inherent growth potentials. For that reason, we expect the mechanism of self-hybridization to reinforce the concentration of knowledge progress.

In statistical jargon, our argument is that domains' self-hybridizations mediate the size-growth relationship. This hypothesis is corroborated if two conditions are met. The first is that domains' size and self-hybridizations are positively correlated. This is clearly the case with our data: on average, during the 25 years observed, the Pearson correlation coefficient between the log of domains' size—as measured by number of patents—and the log of domains' self-hybridizations—as measured by number of self-citations—has been 0.942. The second condition is that the effects of size and self-hybridizations are counteracting, resulting in a stronger negative effect of size on growth when the effect of self-hybridizations are controlled for. A straightforward way to test whether this condition is met is to augment model (5) with the log of domains' self-citations. Due to the very high correlation between domains' size and self-hybridization, however, the resulting model specification exhibited high multicollinearity (highest VIF = 18.2). Therefore, we preferred to augment model (5) with the raw count of self-citations instead. This reduced the correlation between the regressors to 0.599, thereby eliminating the problem of multicollinearity (highest VIF = 1.54)<sup>8</sup>. Table 3 reports the results of the analysis: net of the positive effect of self-hybridizations, the negative effect of domains' size on domains' growth is more than 20% stronger (compare Tables 1 and 3). This indicates that also the second condition necessary for the corroboration of our hypothesis is met.

## 8 Conclusions

In this paper, we showed that a feature in the growth of US-patented knowledge is that it concentrates disproportionately among larger domains. In other words, the

<sup>7</sup> In our data, for example, domains' self-citations—indicating self-hybridizations—are roughly half of all citations.

<sup>8</sup> Notice that we also estimated the model in which both domains' size and self-hybridization are logarithmically transformed, obtaining qualitatively identical results.

**Table 3** Variable-intercept (fixed-effects) regression of domains' size and domains' self-hybridization, on domains' growth rates

	Dependent variable: domains' growth rate			
	Coefficient	Std. error	<i>t</i> Value	Sig.
Intercept	1.616	0.143	11.26	0.000
Domains' size	-0.256	0.024	-10.59	0.000
Domains' self-hybridizations	2.04e-4	4.16e-06	4.90	0.000
Number of units	417			
Periods	4			
Number of observations	1653			
<i>F</i>	56.13			
Prob. > <i>F</i>	0.0000			

bulk of progress takes place in a handful of large technology domains, while most other domains contribute surprisingly little to technological progress. Interestingly, this pattern appears to reflect a stable and enduring dynamics, as our analyses indicate that we can be reasonably confident that it will hold in a foreseeable future. Furthermore, we showed that the concentration of technological progress observed in the data is explained by the combination of two factors. On the one hand, domains' trajectories of growth have inherently different growth potentials. On the other hand, the differences in domains' inherent growth potentials are magnified by the endogenous mechanism of domains' self-hybridization. As it is hard to conceive of a process of knowledge growth in which either of these two factors would not be at work, it is reasonable to expect the observed concentration of progress to manifest itself under very general conditions. Future attempts at explaining the direction of technological progress should account for this empirical tendency.

**Acknowledgments** This research was supported by the Netherlands Organization for Scientific Research, NWO grant no. 401-01-098.

## References

- Albert R, Jeong H, Barabasi A-L (1999) Diameter of the world-wide web. *Nature* 401:130–131
- Anselin L (1998) *Spatial econometrics: methods and models*. Kluwer Academic Publishers, Dordrecht
- Breusch T, Pagan A (1980) The Lagrange multiplier test and its applications to model specifications in econometrics. *Rev Econ Stat* 47:239–253
- Carnabuci G (2010) The ecology of technological progress: how symbiosis and competition affect the growth of technology domains. *Soc Forces* 88:2163–2187
- Carnabuci G, Bruggeman JP (2009) Knowledge specialization, knowledge brokerage, and the uneven growth of technology domains. *Soc Forces* 88:607–641
- Dosi G (1982) Technological paradigms and technological trajectories. *Res Policy* 11:147–162
- Gabaix X (1999) Zipf's Law for cities: an explanation. *Q J Econ August*:739–767
- Gilfillan SC (1930) Inventiveness by nation: a note on statistical treatment. *Geogr Rev* 20(2):301–304
- Griliches Z (1957) Hybrid corn: an exploration in the economics of technological change. *Econometrica* 25:501–522
- Griliches Z (1984) *R&D, patents, and productivity*. University of Chicago Press, Chicago

- Hall B, Jaffe A, Trajtenberg M (2001) The NBER patent citations data file: lessons, insights and methodological tools. NBER, Cambridge
- Harberger AC (1988) A vision of the growth process. *Am Econ Rev* 88:1–32
- Henderson R, Jaffe A, Trajtenberg M (2005) Patent citations and the geography of knowledge spillovers: a reassessment: comment. *Am Econ Rev* 95:450–460
- Jaffe A, Trajtenberg M, Fogarty MS (2000) Knowledge spillovers and patent citations: evidence from a survey of inventors. *Am Econ Rev* 90(2):215–218
- Kesten H (1973) Random difference equations and renewal theory for products of random matrices. *Acta Math* 131:207–248
- Levy M, Solomon S (1996) Power-laws are logarithmic Boltzmann laws. *Int J Mod Phys C* 595:65–72
- Liljeros FL et al (2001) The web of human sexual contacts. *Nature* 411:907–908
- Lotka A (1926) The frequency distribution of scientific productivity. *J Wash Acad Sci* 16:317–323
- Merton RK (1968) The Matthew effect in science. *Science* 159:56–63
- Mitzenmacher M (2004) A brief history of generative models for power law and lognormal distributions. *Int Math* 2:226–251
- Mokyr J (2002) The gifts of Athena: historical origins of the knowledge economy. Princeton University Press, Princeton
- Nelson RR, Winter SG (1977) In search of a useful theory of innovation. *Res Policy* 6:36–76
- Nelson R, Winter S (1982) An evolutionary theory of economic change. Belknap Press of Harvard University Press, Cambridge
- Pareto V (1897) *Cours d'économie politique*. F. Rouge, Lausanne
- Price DDS (1980) A general theory of bibliometric and other cumulative advantage processes. *J Am Soc Inf Sci Technol* 27:292–306
- Romer PM (1994) Economic growth and investment in children. *Daedalus* 123:141–154
- Rosenberg N (1976) *Perspectives on technology*. Cambridge University Press, New York
- Rosenberg N (1994) *Exploring the black box: technology, economics, and history*. Cambridge University Press, New York
- Scherer FM (1965) Firm size, market structure, opportunity and the output of patented innovations. *Am Econ Rev* 55:1097–1125
- Schmookler j (1966) *Invention and economic growth*. Harvard University Press, Cambridge, MA
- Simon HA (1955) On a class of skew distribution functions. *Biometrika* 42:425–440
- Simon HA, Bonini CP (1958) The size distribution of business firms. *Am Econ Rev* 4:607–617
- Usher AP (1954) *A history of mechanical inventions*. McGraw-Hill, London
- Weitzman ML (1998) Recombinant growth. *Q J Econ* 113:331–360
- Zipf GK (1949) *Human behavior and the principle of least effort*. Addison-Wesley Press, Cambridge