

Random Drift versus Selection in Academic Vocabulary: An Evolutionary Analysis of Published Keywords

R. Alexander Bentley*

Department of Anthropology, Durham University, Durham, United Kingdom

Abstract

The evolution of vocabulary in academic publishing is characterized via keyword frequencies recorded in the ISI Web of Science citations database. In four distinct case-studies, evolutionary analysis of keyword frequency change through time is compared to a model of random copying used as the null hypothesis, such that selection may be identified against it. The case studies from the physical sciences indicate greater selection in keyword choice than in the social sciences. Similar evolutionary analyses can be applied to a wide range of phenomena; wherever the popularity of multiple items through time has been recorded, as with web searches, or sales of popular music and books, for example.

Citation: Bentley RA (2008) Random Drift versus Selection in Academic Vocabulary: An Evolutionary Analysis of Published Keywords. PLoS ONE 3(8): e3057. doi:10.1371/journal.pone.0003057

Editor: Colin Allen, Indiana University, United States of America

Received: June 13, 2008; **Accepted:** August 5, 2008; **Published:** August 27, 2008

Copyright: © 2008 Bentley. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: The author has no support or funding to report.

Competing Interests: The author has declared that no competing interests exist.

* E-mail: r.a.bentley@durham.ac.uk

Introduction

Ideally, science is the systematic process of testing multiple hypotheses, but as practiced by real people, it is also distinctly social. Within complex collaboration networks, academics compete for citations, particularly in our modern era of online citation databases that can ‘summarize’ an academic’s career at a single command [1–4]. They are therefore prone to copy ideas, and particularly buzzwords, from one another [2,3].

Diverse opinions exist as to what constitutes trendy ideas versus more meaningful research paradigms; the challenge is to evaluate this by some objective means. In other realms of fashion, ranked lists are increasingly a part of our world; from universities to Internet searches, downloads, book and music sales. Correspondingly, the design of algorithms needed to track ‘what’s hot and what’s not’ has itself become a hot topic in computer science [5]. Indeed, as journals are now ranked by their impact factor – increasingly a subject of study [6,7] – there is no reason why we cannot look at academic keywords the same way: rank them in order of popularity from year to year, and track the comings and goings of ‘what’s hot’ on such lists.

As the science of how attributes are passed on and modified through time [8], evolutionary theory is an ideal means to investigate these aspects of scientific process [9]. Previous work using evolutionary models has shown, counter-intuitively, that many patterns of change in cultural choices over time can be explained as random drift; i.e. the effect of chance on what happens to be copied, together with the occasional appearance of innovations [10–12]. Meaningful selection, as opposed to random copying, occurs when such choices are made on the basis of something inherent to the choice itself [13] – as with a ‘better mousetrap’ for example, or something inherently preferable to human tastes.

In knowledge production, ideas are not always adopted out of inherent superiority, but often merely because others are using

those ideas. In either case, the transmission process is evolutionary; predominantly one of adopting what others have done, with creative modifications contributing new ideas that eventually replace old ones through being adopted. ‘Ideas’ of course is a nebulous description, so this study focuses specifically on the evolution of keyword use in academic publishing.

By analyzing keyword frequencies as recorded in a citations database, one can characterize their replication in terms of a continuum between (a) random copying of fashionable buzzwords at one extreme (akin to random genetic drift), and (b) independent selection of keywords, based on inherent qualities, at the other (falsifying the neutral model). The question is one of degree, with variation expected along this basic continuum. Using random copying as the null hypothesis, one can simply seek to identify selection against the null without characterizing it specifically; although clearly the first hypothesis is that words are selected for usefully describing something real and relevant to the topic.

It may seem cynical to assume first that keywords are copied without much thought, but several studies suggest this [2,3,9,12] and even George Orwell thought as much in his famous 1946 essay, “Politics and the English language.” As the null hypothesis, random copying does not mean that the words themselves are chosen randomly, but that they are copied randomly from others who have already used them. The assumption is that randomly-copied keywords are *value-neutral*, in that no keyword is inherently more valuable than any other – the likelihood of any being chosen is simply proportional to its current popularity. This is in essence the neutral model of population genetics [8,14].

In previous simulations, the random copying, or neutral, model has been represented as follows: Start with a set of N individuals, which are replaced by N new individuals in each generation. Over successive generations, each of the N new individuals copies its variant from a randomly-selected individual in the previous generation, with exception of a small fraction, μ (<5%), of the N new individuals who invent a new variant in the current generation.

The neutral model is simple to simulate, yet has been shown to provide richly complex results that produce at least three useful predictions relevant to cultural drift [10,12,15]:

(1) If individual variants are tracked through the generations, their frequencies (relative popularities) will change in a stochastic manner, as opposed to a directed manner or completely random manner. More specifically, the neutral model predicts that the only source of change in variant frequencies over time is random sampling, such that [14]:

$$V = \frac{v(1-v)}{N} \quad (1)$$

where V is the variance in frequencies over time (see methods), and $v \leq 1$ is the relative frequency of the variant as fraction of N , the maximum possible number of variant copies per generation. For small v , $v(1-v) \approx v$, which after rearranging eq. (1) indicates that $NV/v \approx 1$.

(2) Like many processes of proportional advantage (under random copying the chance of being copied is proportional to current frequency), the variant frequencies exhibit a long-tailed distribution, which for small values of μ follows a power law form [10,12]. This is one of the less diagnostic predictions, as a variety of mechanisms can generate power law and similar distributions [16]. Nonetheless, the distribution is useful as a null expectation. Among the possible departures from this null, selective bias for novelty (e.g., some maximum threshold of popularity) should truncate the tail (high end) of the variant frequency distribution [17,18]. Alternatively, there might be a conformist bias resulting in a ‘winner take all’ distribution, whereby one word has a higher frequency than predicted by the power law for the rest of the words [17,18]. Within journal publication, an example of this is the bias towards authors seeking to publish in a few high-impact journals, apparently at a level that exceeds merely the effect of proportionate advantage [19].

(3) There is continual turnover in the variant pool [15]. If the variants are ranked in order of decreasing frequency, the turnover z in that list over successive generations (time) depends much more strongly on μ than on N , such that:

$$z \approx \sqrt{\mu} \quad (2)$$

where z is measured as the fraction of turnover in the list (e.g., two items replaced in a Top 10 list would be 20% turnover). In contrast to random copying, under selection the population size N should correlate positively with the turnover rate in the ranked list of most popular variants [15].

Using these three predictions as the null model, it is easier to identify selection, which is effectively demonstrated by departures from these patterns, dependent on the kind of selection operating.

In applying this to keyword use, let N represent the number of keywords in a given time period (rather than the number of articles, which vary in their number of keywords). This ensures that each ‘individual’ corresponds with exactly one variant. The invention fraction μ is then the fraction of those words in each time interval that are appearing for the first time.

As listed on the WoS database, the four case studies presented here provide a test of differences of keyword use among published articles within older paradigms versus younger ones, and within

the physical sciences versus the social sciences. In order to define these case studies, we need a working definition of a sub-field of academic publishing. If belabored, this could be quite a difficult task – many definitions would be too subjective, variable or broad.

A way forward is to define a scientific ‘paradigm’ [20] as comprising the scientific papers that were in some way inspired by a certain highly-influential paper. We thus can define each academic paradigm as the set of all papers that cited a certain highly-cited paper. The citing papers may occur in a range of different journals, but they will all share the defining characteristic of citing the highly-influential work.

Results

Consider four highly-cited, seminal works, two from the natural sciences and two from the social sciences. These works were chosen on the basis of being among the most highly-cited in their respective fields, and also to be paired in similar areas of social science and physical science, respectively. To see the effect of time, from the pair in each category we include one work about 30 years old and the other about ten years old. This provides two comparisons: older versus younger fields of study, and social sciences versus physical sciences.

From the physical sciences we have a paper by Barabási and Albert (PS99 for ‘physical sciences, 1999’) in 1999 [21], which introduced a quantitative model of ‘scale-free networks’ and has been cited over 2,000 times, and one by Witten and Sander (PS81) from 1981 [22], which introduced the physics model of ‘diffusion limited aggregation’, and has been cited over 1,300 times. From the social sciences, there is a paper by Nahapiet and Ghoshal (SS98) in 1998 [23], cited over 460 times, which reviewed the influential concept of ‘social capital’, and a 1977 book by Bordieu (SS77), cited over 2,700 times, which introduced such influential concepts as ‘agency’ and ‘structuration’ into the social sciences [24].

Figure 1 shows the temporal change in N , the total number of keywords per year, and in $N\mu$, the number of new keywords per year, for paradigms about 10 years old (Figure 1a) and 30 years old (Figure 1b). A new keyword was one which had not appeared in the record beforehand, with records starting in 1994 for the older works and date of publication (1998, 1999) for the younger paradigms.

Table 1 shows additional statistics for each paradigm averaged from 2002 to 2006, the sample period applicable to all four case studies (the newer case studies do not have enough data before 2002). In each case, the quantities N and $N\mu$ parallel each other (Figure 1), indicating a consistent and substantial invention fraction μ between 15 and 30% in all cases (Table 1). Within the older pair and the younger pair of paradigms, the invention fraction μ was higher for the social science than for the physical science case (Table 1). This is true even though the comparison differs in the number of words: N is larger for PS99 than SS98, but lower for PS81 than SS77.

In addition to a higher innovation fraction for the social science paradigms, there is also a marked difference in the turnover in keywords. Consider the top 5 keywords, in terms of popularity, over the years in each case study (below the top 5, keywords start to become insufficient in their numbers of appearances). As the best way to view overall trends in turnover, Figure 2 shows the cumulative turnover in the top 5 keywords, expressed as a fraction (e.g., 4 words having passed through the top 5 = 80% turnover). In the physical science paradigms, the turnover in the top 5 keywords leveled off to virtually no turnover in the last several years. At the other end of the spectrum, the keywords in the social science cases show a high and steady turnover throughout the sampling period




Figure 1. Keywords, total and new, among paradigms about (a) 10 and (b) 30 years old. Social science cases are shown in red and physical sciences in black. Solid curves show the total number of keywords N per year, and the dashed curve shows number of new keywords $N\mu$ introduced per year.

doi:10.1371/journal.pone.0003057.g001

(Figure 2). In the case of SS77, this turnover persisted long after its publication, and many years beyond which PS81 had leveled off.

Whereas the continual turnover in SS77 and SS98 is consistent with random copying with innovation, the cessation of turnover in PS81 and especially PS99 suggests selection. As Figure 3a shows, the selective sorting of the keyword frequencies for PS99 was strong enough that even the keyword *networks* (highlighted in red) occupies a distinct frequency ranking from the singular *network* (blue), while other entries are similarly locked into their positions among the top 5. Although this pattern of selection is not as strong in the older physical science paradigm (PS81), the blue versus black lines in Figure 3c show apparent groupings of words by selected frequencies. In contrast, both the older and younger social science cases (SS77 and SS98) appear more stochastic in their histories of individual word frequencies (Figure 3 b, d), and with each at a relatively low frequency compared to the network science case (Figure 3 a, c). In the SS77 case, the ratio $N\mu/v$ increases moving down the rankings (Table 2), which suggests a possible conformist bias, in that the more frequent words have been preferentially selected (e.g. red curve in Figure 3d).

As described above, the ratio $N\mu/v$ can be used to characterize keyword variability, allowing comparison across cases studies for

the period 2002–2006 (Table 2). Averaged over the five keywords, $N\mu/v$ differs more by age of the paradigm than by subject matter, being higher for the younger (~ 2.3) than the older (1.3–1.4) paradigms. Within each age pair, however, the physical sciences paradigm has the larger standard deviation in the mean value of $N\mu/v$ (Table 2). This reflects certain keywords in the physical science paradigms whose popularity changed directionally, apparently due to selection. In the PS99 case, the word ‘complex’ (word 3, score = 5.8) appears to have been selected for, as it doubled in frequency from 2002 to 2006 beyond what would be expected from random drift. Also in the PS99 case, *networks* (word




Figure 2. Cumulative turnover in the top 5 keywords. Social science cases are shown in red and physical sciences in black. Turnover refers to words making a first appearance in the top 5. For the older paradigms (SS77; PS81), symbols are squares and the count begins at 1994, for the newer articles (PS99; SS98) symbols are circles and the count begins the year after publication.

doi:10.1371/journal.pone.0003057.g002

Table 1. Average values, from 2002–2006.

Paradigm	N	Vocabulary	$N\mu$	μ
SS77	2089 $\pm$ 255	1209 $\pm$ 100	433 $\pm$ 14	21% $\pm$ 2%
PS81	1201 $\pm$ 230	635 $\pm$ 88	197 $\pm$ 34	16% $\pm$ 1%
SS98	885 $\pm$ 416	431 $\pm$ 159	224 $\pm$ 68	28% $\pm$ 7%
PS99	2660 $\pm$ 1525	979 $\pm$ 408	511 $\pm$ 135	23% $\pm$ 8%

Parameters include the number of keywords N , different keywords or ‘vocabulary’, newly appearing keywords $N\mu$, and the invention fraction μ .

doi:10.1371/journal.pone.0003057.t001




Figure 3. Frequencies of the top 5 keywords of 2005. Shown are the four paradigm case studies, including: (a) newer physical sciences (PS99); (b) newer social sciences (SS98); (c) older physical sciences (PS81); and (d) older social sciences (SS77). Logarithmic y-axes.
doi:10.1371/journal.pone.0003057.g003

1) declined steadily as *network* (word 2) increased, such that their variability scores are near 2. By contrast, the words in the SS98 case do not show such directionality in their change (Figure 3b), and the high variability scores for four of the five words (Table 2) is due to their fluctuating frequencies over the time interval (Figure 4b). Curiously, in the PS81 case, the word *aggregation* (word 2, score = 0.3) was considerably less variable than *diffusion* (word 5, score = 2.7) even though the seminal paper [22] was about diffusion-limited aggregation.

Finally, consider the cumulative frequency distributions of keywords for two selected time-slices, years 2001 and 2005 (Figure 4). All show essentially a power law form, which could be

consistent with either the neutral model but also a variety of models of proportionate advantage [16]. What is more revealing is the degree of change in the power law exponent (slope on the log-log plot) over this 4-year time span. In three cases, the slope is nearly the same for 2005 as for 2001, but for PS99, the slope is considerably less for 2005. The decreasing slope for PS99 correlates with a decreasing invention fraction μ over this time span (Table 1), and reflects the diminishing probability for any new keyword to ever reach the top 5.

The frequency distributions in Figure 4 enable the identification of copying biases. Although all four paradigms yield essentially power law distributions, two cases – PS81 and SS98 – show marked departures from a power law in the truncations of the tail (Figure 4b and 4c). In these two cases, the top 3 or 4 words of 2005 have nearly the same frequencies rather than following the power law. In the PS81 case, those three words (*aggregation*, *fractal*, *growth*) as mentioned above were grouped together for the entire sampling period (Figure 3c), whereas for the SS98 case, the top four words (*social*, *capital*, *knowledge*, *management*) show no obvious grouping (Figure 3b). Indeed, even though the seminal article for SS 98 [23] reviewed the “social capital” concept, those two words only parallel each other for the first two years (Figure 3b). A further difference is that the 3 words in the PS 81 case are consistently above 10% frequency (Figure 3c), whereas those for SS98 never even reach that level (Figure 3b). One possibility is that the departures from the power law, though they look similar, reflect

Table 2. Values of NV/v for the top 5 words, 2002–2006, tracked in Figure 3.

Paradigm	Word 1	Word 2	Word 3	Word 4	Word 5	Ave.
SS77	0.82	1.13	1.10	1.46	1.75	1.25 (16)
PS81	1.52	0.34	0.71	1.55	2.69	1.36 (41)
SS98	2.73	2.75	0.97	2.66	2.19	2.26 (34)
PS99	2.10	1.82	5.85	1.28	0.59	2.33 (92)

Numbers in parentheses give standard error on the trailing digits.

doi:10.1371/journal.pone.0003057.t002




Figure 4. Cumulative frequency distributions of all keywords. Double logarithmic axes. The paradigms are (a) newer physical sciences (PS99), (b) newer social sciences (SS98), (c) older physical sciences (PS81) and (d) older social sciences (SS77). Using the maximum likelihood method [16], the estimated power-law exponents for 2001 and 2005, respectively, are as follows: PS99: 2.11, 2.00; SS98: 2.11, 2.02; PS81: 2.09, 2.05; SS77: 2.18, 2.09. Errors (by jackknife estimate) on these exponents are <0.01. doi:10.1371/journal.pone.0003057.g004

different effects for PS81 versus SS98 – selection in the former case, and anti-conformity in the latter.

Discussion

By treating academic keywords as discrete elements of evolution, this study finds that different academic niches – as defined by sets of publications which share a single seminal article in their cited reference lists – can show markedly different evolutionary patterns. From the case studies considered, it appears that some academic fields are characterized by a high degree of drift, resulting in continual and unpredictable change in vocabulary, whereas in others words appear under selection, such that the predominant vocabulary becomes increasingly crystallized and unchanging over time.

Among the cases presented, the social science paradigms showed the stronger patterns of random copying with invention, including constant turnover in the keywords of highest frequency, and the stochastic ups-and-downs of individual word frequencies over time. In contrast, the physical science paradigms showed a rejection of the neutral model, particularly in the cessation of turnover in the top keywords over time.

Here is where one naturally seeks an explanation of *why* the physical science paradigms demonstrated the greater degree of keyword selection. Some physical scientists might like to argue that social scientists invent meaningless jargon, whereas social scientists might claim that they are more creative and free with their language, and their rivals have no imagination. Clearly there are different proclivities, which may have fed back on each other over time. For example, many physics journals (e.g., *Physical Review*) force authors to choose keywords from lists, whereas most social science journals still allow authors to invent keywords without limitations. Whether this *caused* the differences, or merely reflects habits that were already there, now becomes a new layer of inquiry. The point is, evolutionary analysis has provided this basis for new investigations – as so often in science, the finding itself raises more questions.

New questions are also revealed by the scale of analysis; a certain group of keywords might be selected, yet drifting within the group. One can then ask why certain keywords group together, and also, given that they are selected as a group, what their selective ‘fitness’ is relative to other groups [28]. Similarly, in a different study, while choices of baby names for the entire United States are indistinguishable from random copying [12], different ethnic groups certainly select from different pools of names [25], and it remains to be studied whether random drift would predominate again within these groups.

In addition to these particular points, this study is meant to demonstrate how a similar evolutionary analysis could be performed on any cultural dataset comprising discrete elements. This evolutionary analysis contrasts with the increasing representation of knowledge growth as networks e.g., [1,4, 26], with the individuals (e.g. authors) as ‘nodes’, and their interactions (e.g. cited references) as ‘links’. A particular challenge for network analysis, however, is change, because a network implies a structure to interactions – the connections of today determine what will happen tomorrow, such that change must be modeled as a modification of the existing network. However, in fashionable realms, yesterday might be less important than tomorrow, and interactions of influence may differ completely from one day to the next. Change can be the essence of the process, rather than just a modification.

For this reason, evolutionary theory can often naturally account for change that may be seen as exceptionally noteworthy in a network model [26,27]. A recent network analysis [29], for example, tracked co-authorships and mobile phone calls to show that, in order to have longevity, small groups require stability in their composition, whereas large groups last a bit longer with some degree of turnover in their membership. This, in fact, is a basic prediction of the genetic drift model: small populations are destroyed by drift, large populations can tolerate it and even find it adaptive. The crucial difference is that in the network analysis [29] mutation was measured as absolute number of changes, whereas the random copying model defines mutation μ as a fraction of N . Hence the random copying model would have predicted the network result, in that coherence disintegrates more quickly with one mutation per time step in a population of 4 versus a population of 100, for example, because the former is a much higher mutation rate.

Change, in fact, is central to evolutionary theory. The use of some basic evolutionary analyses, with parallels in population genetics, can be used to characterize different forms of innovation and transmission of discrete cultural elements. Identifying what proceeds in predictable directions, as opposed to drifting upon the tides of fashion, would be of great utility in understanding the evolution of knowledge. It is wasted effort to try to predict the future of randomly drifting fashionable buzzwords [2,30], but one might hope to predict selected elements, such as valid new scientific terms. The kind of evolutionary analysis used here is generally applicable to any case study where popularity can be presented in the form of frequencies and ranked lists over time.

Methods

The data used in this analysis were taken from Thompson Scientific’s ‘Web of Science’ (WoS) database, which covers articles thousands of journals in science and engineering, social sciences, arts and humanities. Among the wealth of information provided, each journal article description in the WoS database contains the title, keywords and abstract, references cited, and a list of all papers in other journals that have cited the paper to date.

For each of the sets of articles within each defined paradigm, the keywords data from the WoS database were taken only from titles and keywords chosen by the authors (not the WoS ‘Keywords plus’ which is an automated condensation of the cited references), and then sorted by publication year. The following common words were removed from the data: *a, an, and, as, by, for, from, in, its, of, on, the, to, using* and *with*. Aside from these, no other common words were present in high enough frequencies to significantly affect the patterns discussed below.

For each yearly sequence of keyword frequencies, the variance V in frequencies over n years was calculated as $V_n = \sum \frac{v - \bar{v}}{n-1}$, where $\bar{v}$ is the average frequency of the keyword over the n year time interval.

Acknowledgments

Thanks to Mark Thomas, Peter Schauer, Mikhail Simkin and Tony Chemero for helpful comments.

Author Contributions

Conceived and designed the experiments: RAB. Performed the experiments: RAB. Analyzed the data: RAB. Wrote the paper: RAB.

References

- Guimerà R, Uzzi B, Amaral LAN (2005) Team assembly mechanisms determine collaboration network structure and team performance. *Science* 308: 697–702.
- Bentley RA (2006) Academic copying, archaeology and the English language. *Antiquity* 80: 196–201.
- Simkin MV, Roychowdhury VP (2003) Read before you cite! *Complex Systems* 14: 269.
- Wuchty S, Jones BF, Uzzi B (2007) The increasing dominance of teams in production of knowledge. *Science* 316: 1036–1039.
- Hayes B (2008) The Britney Spears problem. *American Scientist* 96: 274–279.
- Hirsch J (2006) An index to quantify and individual's scientific research input. *Proc Natl Acad Sci USA* 102: 16569–16572.
- Stringer MJ, Sales-Pardo M, Amaral LAN (2008) Effectiveness of journal ranking schemes as a tool for locating information. *PLoS ONE* 3: e0001683.
- Nowak MA (2006) *Evolutionary dynamics* Harvard University Press.
- Hull DL (2001) *Science and selection* Cambridge University Press.
- Bentley RA, Hahn MW, Shennan SJ (2004) Random drift and culture change. *Proc Roy Soc B* 271: 1443–1450.
- Shennan SJ, Wilkinson JR (2001) Ceramic style change and neutral evolution: A case study from Neolithic Europe. *American Antiquity* 66: 577–594.
- Hahn MW, Bentley RA (2003) Drift as a mechanism for cultural change: an example from baby names. *Proc Roy Soc B* 270: S1–S4.
- Shennan SJ (2008) Canoes and cultural evolution. *Proc Natl Acad Sci USA* 105: 3175–3176.
- Gillespie JH (2004) *Population genetics: A concise guide* The Johns Hopkins University Press.
- Bentley RA, Lipo CP, Herzog HA, Hahn MW (2007) Regular rates of popular culture change reflect random copying. *Evolution and Human Behavior* 28: 151–158.
- Newman MEJ (2005) Power laws, Pareto distributions and Zipf's law. *Contemporary Physics* 46: 323–351.
- Bentley RA, Shennan SJ (2003) Cultural evolution and stochastic network growth. *American Antiquity* 68: 459–485.
- Mesoudi A, Lycett SJ (in press) Random copying, frequency-dependent copying and culture change. *Evolution and Human Behavior*.
- Ioannidis JPA (2006) Concentration of the most-cited papers in the scientific literature: analysis of journal ecosystems. *PLoS ONE* 1: e0000005.
- Kuhn TS (1962) *The structure of scientific revolutions* University of Chicago Press.
- Barabási A-L, Albert R (1999) Emergence of scaling in random networks. *Science* 286: 509–512.
- Witten T, Sander LM (1981) Diffusion-limited aggregation, a kinetic critical phenomenon. *Physical Review Letters* 47: 1400–1403.
- Nahapiet J, Ghoshal S (1998) Social capital, intellectual capital and the organizational advantage. *Academy of Management Review* 23: 242–266.
- Bordieu P (1977) *Outline of a theory of practice* Cambridge University Press.
- Fryer RG, Levitt SD (2004) The causes and consequences of distinctively black names. *Quarterly Journal of Economics* 119: 767–805.
- Humphries MD, Kevin Gurney K (2008) Network 'small-world-ness': A quantitative method for determining canonical network equivalence. *PLoS ONE* 3: e0002051.
- Minnhagen P, Bernhardsson S (2008) The blind watchmaker network: Scale-freeness and evolution. *PLoS ONE* 3: e0001690.
- Simkin MV, Roychowdhury VP (2007) A mathematical theory of citing. *J Am Soc Information Science and Technology* 58: 1661–1673.
- Palla G, Barabási A-L, Vicsek T (2007) Quantifying social group evolution. *Nature* 446: 664–667.
- Salganik MJ, Dodds PS, Watts DJ (2006) Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311: 854–856.