Assessing impact and quality from local dynamics of citation networks

Camille Roth, Jiang Wu, Sergi Lozano

► To cite this version:

Camille Roth, Jiang Wu, Sergi Lozano. Assessing impact and quality from local dynamics of citation networks. *Journal of Informetrics*, 2012, 6 (1), pp.111-120. 10.1016/j.joi.2011.08.005 . halshs-00778517

HAL Id: halshs-00778517

<https://shs.hal.science/halshs-00778517>

Submitted on 20 Jan 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Assessing impact and quality from local dynamics of citation networks

Camille Roth^{a,b,1}, Jiang Wu^{c,d}, Sergi Lozano^d

^a*CNRS, CAMS Centre d'Analyse et de Mathématique Sociales, CNRS/EHESS,
75013 Paris, France*

^b*Institut des Systèmes Complexes de Paris Ile-de-France (ISC-PIF)
75005 Paris, France*

^c*School of Information Management, Wuhan University, 430072, Wuhan, China*

^d*Chair of Sociology, in particular Modeling and Simulation, ETH Zurich
8092 Zurich, Switzerland*

Abstract

We show that essentially local dynamics of citation networks bring special information about the relevance/quality of a paper. Up to some rescaling, they exhibit universal behavior in citation dynamics: temporal patterns are remarkably consistent across disciplines, and uncover a prediction method for citations based on the structure of references only, at publication time. Above-average cited papers universally focus extensively on their own recent subfield – as such, citation counts essentially select what may plausibly be considered as the most disciplinary and normal science; whereas papers which have a peculiar dynamics, such as re-birthing scientific works – ‘rediscovered classics’ or ‘early birds’ – are comparatively poorly cited, despite their plausible relevance for the underlying communities. The “rebirth index” that we propose to quantify this phenomenon may be used as a complementary quality-defining criterion, in addition to final citation counts.

Keywords: citation networks, bibliometrics, qualitative assessment, citation dynamics

---

---

¹To whom correspondence should be addressed.

E-mail: roth@ehess.fr, Tel.: +33 1 4954 2032, Fax.: +33 1 4954 2109.

Permanent address: CAMS, CNRS/EHESS, 190 Avenue de France, 75013 Paris, France.

1. Introduction

Upon publication, a paper inserts itself into the network of existing references by building on and from the state-of-the-art [7]. It contributes to creating connections between bits of existing literature of various ages, and is in turn deemed relevant by being progressively embedded in the developing network of ongoing citations [15]. This continuous and cumulative process is rarely rendered by aggregate citation observations at a given date.

Yet, citations have often been appraised as a global and mainly static feature, plausibly dismissing the intrinsically dynamic nature of citation networks. In this respect, existing works on citation networks endeavour for instance at describing and reproducing through modeling their highly heterogeneous nature [8, 19, 20, 23] and universal features [18], or measuring the centrality of publication positions in static networks [6]. But even when addressing the growth of citation networks over time, a global approach has been adopted [1, 17].

Rather than considering (static) patterns at the global network level, we exhibit universal behaviors in local citation dynamics in both temporal directions, backward and forward. On the whole, we suggest that even very basic features of the immediate neighborhood reveal highly significant information on the relevance of a paper to its community, as long as it is being observed dynamically. There is another important motivation in looking at the local neighborhood only: irrespective of their age, papers should be observed on a similar period of time – the fully developed citation network indeed mixes papers from significantly different eras, where older papers are more likely to be more cited over a different period of time.

We incidentally aim at relating local citation dynamics to the interest a paper will eventually have triggered in its community. Our results therefore bear implications on the relationship between quality and citation metrics, currently widely used as a benchmark for scientific evaluation. While a highly cited paper is certainly of relevance to the scientific community, it is however quite common to debate how such an “impact” translates into attested quality [10, 16] – typically because of, among other factors, distinct [24] or biased citation practices [10, 14], excessive focus on “hubs” [6] or potentially flawed statistics derived from aggregate citation metrics [9, 21]. In this respect, we pursue two goals: exhibit robust structural predictors of the relevance of a paper (even *at* publication time) and, consequently, question the type of science which is captured by global/static citation metrics.

By showing that the most cited papers are those citing extensively the field, having recent references, to the contrary of papers having little references from *within* the field, citing farther in the past, or being cited late in the future, particularly those being resurrected later, we quantitatively show how citation counts essentially select what may be considered as the most disciplinary and normal science.

We pay special attention to rebirthing articles, i.e. those ones showing two separate ‘lives’ or citation bursts. By using an innovative *Rebirth index*, we perform an impact analysis not dependent on threshold-based definitions (in contrast with similar studies on related citation profiles like *Sleeping Beauties* [12]). We find that rebirth is a notion distinct from impact, in the sense that low impacts corresponds both to high and no rebirth. We therefore suggest that this index could be used as a complementary quality-defining criterion, in addition to final citation counts.

In Section 2 we introduce the methodology developed, which includes the processing of bibliometric data and the proposal of several indices quantifying both simple and rebirthing citation dynamics. Section 3 examines the relationship between the local citation dynamics of papers, which are characterized by these indices, and their relative citations counts (i.e. normalized with respect to other works with the same disciplinary and temporal characteristics). Finally, Section 4 summarizes and discusses our work.

2. Material and Methods

We first describe the dataset, then various variables which will be used throughout the paper.

Dataset

The database we use is based on Thomson Reuters Scientific “ISI Web of Science”, spanning over 1960-2009 for four general *a priori* categories: computer science (denoted by “cs”), economics (“eco”), engineering (“eng”) and physics (“phys”). The choice of these fields was primarily guided by the wish to have datasets of mixed size (as evidenced from Tab. 1), with mixed recency as constituted disciplines (“cs” being younger than “phys”), as well as at least one field from social sciences (“eco”) and one from applied sciences (“eng”). To ease the presentation and comparison of results, a number of four fields seemed to be a reasonable compromise. Apart from that, specific field choices – i.e. “phys” or “eng” rather than “medicine” or “maths” – were

somewhat arbitrary; leaving to further research the task of exploring other fields and possibly finding outliers.

Introducing temporal metrics requires to observe papers some time before and after publication. We therefore need to restrict the dataset to papers having sufficient history and, in some cases, to papers having sufficient dynamic information both in terms of citations or references. We suggest that 10 years of history ahead and before provides a sufficient temporal resolution; by doing so, we still consider about 30 years of data by focusing on papers published between 1970 and 1999.

Going further, we consider several suitable constraints: “ALL” represents papers having at least 10 years of history before and after publication date, among which “CMIN” represents the subset of papers having at least 10 citations after 10 years, and “RMIN” the subset of papers having at least 10 references at most 10 years old. The rationale for defining CMIN and RMIN lies in the fact that average citation or reference ages cannot be defined when papers have little or no citations; most importantly, dynamics of link arrival are likely to be sketchy and strongly discretized on few points. Again, we arbitrarily assume that 10 points provide a sufficient resolution on the dynamics, while it does not jeopardize statistical significance: these constraints result in diverse size reduction effects on each dataset as shown on Tab. 1 (see Fig. A.1 too). More precisely, ALL only aims at ensuring sufficient temporal coverage before and after publication and includes 100% of papers published over 1970-1999. From this baseline, CMIN will be used for citation dynamics-related variables and does not deal with the majority of papers which have few or no citations (keeping from 6% (“eng”) to 26% (“phys”) of ALL papers). Similarly, RMIN will be used for reference dynamics and conserves between 2% (“cs”) and 23% (“phys”) of ALL papers, i.e. the minority of papers having a significant span of references from the field.

Nonetheless, we checked that our results do hold qualitatively using thresholds of 5 and 20 for CMIN and RMIN.² In the remainder, we then use a threshold of 10 as a decent trade-off between temporal resolution (the higher the threshold, the better) and statistical significance and representativity.

---

²For a threshold of 5, CMIN keeps from two to three times more papers than for a threshold of 10, while RMIN conserves from four to five times more papers.

Name	Description	Paper count			
		cs	eco	eng	phys
RAW	Papers from the raw dataset	634 592	268 785	2 056 282	2 543 769
ALL	At least 10 years of forward and backward history	258 303	142 935	1 049 626	1 365 774
CMIN	At least 10 citations after 10 years, among ALL	21 167	15 595	66 465	349 792
RMIN	At least 10 references at most 10 years old, among ALL	5 381	8 610	17 720	317 106

Table 1: Constraints applied to the data.

Figure 1: Dynamic, first neighborhood of a sample paper published in 1990. At left, references (solid links); at right, citations (dashed links).

History of citation and reference

We use the data presented above to examine the relationship between the *temporal ordering* of links arriving to and from a given paper, and its *relative citation count* with respect to its context. By *relative citation count* we mean that we consider the number of citations received by a paper published in a given year with respect to the mean number of citations that all papers published in the same field have received after an identical period of time [2, 5, 18]. By *temporal ordering of links*, we mean that we work, in the broad sense, on arrival times of links connected to a given paper, distinguishing between links originating from a paper (*references*) and those pointing to a paper (*citations*), see Fig. 1.

Most basic features relate to the number of references and citations over the considered time window. Various elaborate dynamics-related metrics are

derivable from first neighborhood linking patterns — from the mean temporal gap from publication time (i.e. average age of references or citations) to metrics indicating in a more sophisticated manner how much a citation profile experiences a late surge of interest or not (called below the “rebirth index”).

For each paper i , from the dynamic citation network we define a citation vector c_i (resp. reference vector r_i) describing temporally the number of incoming links (resp. outgoing links), such that $(c_i)_j$ is the number of incoming links or citations at j years after publication (resp. $(r_i)_j$ denotes the number of outgoing links or references j years old before publication). Because of the above thresholds, j ranges from 0 to 10 for a given c_i , and from -10 to 0 for r_i . In the toy example presented in Fig. 1, for, say, paper “0” from “eng”, we thus have: $c_0 = (0, 0, 1, 0, 0, 1, 0, 0, 2, 0, 1)$ and $r_0 = (0, 2, 0, 1, 1, 0, 0, 0, 0, 0, 0)$.

The raw *citation count* C_i and *reference count* R_i of paper i are simply the sum of its raw citation history and, respectively, reference history:

$$C_i = \sum_{j=0}^{10} (c_i)_j \quad R_i = \sum_{j=-10}^0 (r_i)_j \quad (1)$$

The relative citation count of i is thus the ratio between C_i and the mean of $C_{i'}$ for all papers i' published on the same year as i . Here, $C_0 = 5$ and $R_0 = 4$.

Dynamics-related variables

To compute citation dynamics-related variables, we normalize each history vector c_i and r_i twice: values are first adjusted such that $(\hat{c}_i)_j$ denotes the ratio between $(c_i)_j$ and the total number of papers published on year j , and eventually $\hat{c}_i$ is normalized such that its coefficients sum to 1 — we denote twice-normalized vectors with a tilde: $\tilde{c}_i$. Adjusted proportions presented on Fig. 4 are averages of $\tilde{c}_i$ for all papers i in a given field. This kind of normalisation makes it possible to compare patterns across different ‘historical periods’. In our example, $\hat{c}_0 = (0, 0, \frac{1}{44\,560}, 0, 0, \frac{1}{60\,673}, 0, 0, \frac{2}{72\,680}, 0, \frac{1}{74\,693})$ as there are respectively 44 560, 60 673, 72 680 and 74 693 papers published in “eng” in 1992, 1995, 1998 and 2000. Then, $\tilde{c}_0 \simeq (0, 0, 0.28, 0, 0, 0.21, 0, 0, 0.34, 0, 0.17)$. A similar computation goes for $\tilde{r}_0$.

Citation and reference age. We define a variable related to the citation delay of paper i as the expectation of the distribution across time of its normalized

citation history. We denote it by $\langle c_i \rangle$. We similarly define the reference age $\langle r_i \rangle$ on the normalized reference history distribution. Formally, we have:

$$\langle c_i \rangle = \sum_{j=0}^{10} j \cdot (\tilde{c}_i)_j \quad \langle r_i \rangle = \sum_{j=-10}^0 j \cdot (\tilde{r}_i)_j \quad (2)$$

While these variables are computed on citation and reference counts normalized by the number of papers published on each corresponding year, note that the results are qualitatively stable if we consider non-normalized mean citation and reference age, i.e. actual mean delay of citation, irrespective of the volume of papers published on each respective year. In the case of paper “0”, $\langle c_0 \rangle = 6.03$.

Rebirth index. Metrics introduced above are the most basic ones, but we can also study the concept of quality of a publication by focusing on more sophisticated dynamic patterns. A significant example of this are the so-called ‘Sleeping Beauties’ or ‘Early Birds’ papers. In the literature [3, 4, 22], such publications are characterized by making a significant contribution that triggers no or very little interest upon publication, but is recognized by the community only long time afterwards. These works are based on a binary operationalization of the idea of delayed recognized articles, i.e. papers belong or not to this category.

Building upon this literature, in order to further facilitate the systematical characterization of papers having a second life after a period of sleep, here we propose a second-order dynamic pattern named “*Rebirth index*” and denoted as ρ . This index aims at describing how likely it is for a publication to have two periods of citations, that is, two distinct modes in their citation history distribution. In comparison to previous approaches, it allows for a continuous analysis, i.e. different degrees of delayed recognition. It also focuses on papers which, upon publication, exhibit at least a first surge of interest rather than no impact at all. Notice that this feature allows to avoid relying on an a priori “rebirthing” threshold. Our Rebirth index is defined as follows for a given paper i of normalized citation history $\tilde{c}_i$:

$$\rho_i = \frac{\sum_{t, t' \in \{0, \dots, 10\}^2} \tilde{c}_i^2(t) \tilde{c}_i^2(t') \frac{|t - t'|}{10}}{\sum_{t, t' \in \{0, \dots, 10\}^2} \tilde{c}_i^2(t) \tilde{c}_i^2(t')} \quad (3)$$

Figure 2: Behavior of citation profiles with respect to three typical rebirth index ρ values (from left to right: low to high). Samples were taken from the ‘*phys*’ dataset.

in such a way that ρ_i gives more weight to plateaus (i.e. irregularities in $\tilde{c}_i(t)$). In other words, ρ is higher when a paper exhibits a first burst of citations, then a kind of stagnation period and, finally, a second burst.

Some examples representing typical possible scenarios can be found on Fig. 2. Considering paper “0” again, $\rho_0 = 0.31$. See also Appendix A for more details regarding how c_i , $\hat{c}_i$, $\tilde{c}_i$ and ρ_i vectors and indices may be computed.

3. Results

Universal scaling in the preference towards the recent past

We first focus on the age of forward and backward links, i.e. citations and references. Their average ages are, remarkably, similarly distributed across all datasets, as shown on Fig. 3. Besides, both distributions are extremely modal, and their mean does not seem to evolve in time: in other words, papers published in either 1970 or 1999 are, on average, getting cited after almost exactly the same time.

For a given paper, the probability of being cited after y years or of featuring references y' years old is represented on Fig. 4. Notice that probabilities are normalized in order to effectively allow comparison across historical periods: as said in the above section, we normalize quantities of references or citations on a given year by the number of papers published on that year in the given field. On a paper-by-paper basis, reference lists are more likely to include papers that are around three years old. The same goes for citations: papers are generally getting the highest rate of citation within around 3-5 years of their original publication. These findings are, actually, neither surprising nor new. De Solla Price already discussed the concept of *Immediacy Factor* in Ref. [7], using a limited sample of data. Moreover these results are

Figure 3: Left: Distributions of papers having given mean ages of citations or references, respectively for CMIN and RMIN papers (bins of width 0.2). Right: temporal evolution of the *average* mean ages of citations or references (put simply, it is the temporal evolution of the mean of distributions featured on the left). Legend: *cs*, black solid; *eco*, blue dashed; *eng*, gray dotted; *phys*, red dot-dashed.

in line with those of Glänzel and collaborators in [12, 13]. However, as far as we know, no previous work has generalized such an observation by comparing temporal citation profiles across scientific fields, as we are doing here. The cross-field similarity shown in Fig. 4 can indeed be stressed by re-scaling the time variable of plots. As shown in the inset of the figure, once re-scaled the different citation patterns almost overlap. This result suggests the generality of the underlying citation dynamics, just differentiated from one field to the other by temporal factors like the duration of publishing processes (economics and, to a lesser extent, computer science being comparatively slower than engineering which itself has a slower pace than physics).

References and future impact, recent and normal science

In order to get a deeper understanding of the relationship between temporal linking aspects and paper impact, we examined how linking patterns are correlated with relative citation counts. Results are summarized on the panel on Figure 5 (three first columns) and suggest essentially that:

1. Papers with a larger number of references *from the field* are very significantly more likely to get cited *within the field* — up to 4 times more than the average for papers having about 10 references, and 4 times less for papers having none.
2. Papers with a low (but not too low) average reference age are more likely to be more cited; in particular, papers with a high mean reference age (papers mostly citing “old publications”) are significantly less cited.

Figure 4: Adjusted proportions of reference links pointing back to $y < 0$ years before publication, or respectively, citation links being $y > 0$ years old. Adjusted proportions are averages of $\tilde{c}_i$ for all papers i in a given field. (Observations made on papers from **RMIN** and respectively **CMIN**). Inset: The same adjusted proportions after time re-scaling. More concretely, $t_{\text{eco}} = t_{\text{cs}} = 1.37 \cdot t_{\text{eng}} = 1.95 \cdot t_{\text{phys}}$ for references and $t_{\text{eco}} = 1.2 \cdot t_{\text{cs}} = 1.37 \cdot t_{\text{eng}} = 2.6 \cdot t_{\text{phys}}$ for citations. Profile overlapping is highly significant. Colors are set as defined in Fig. 3.

Figure 5: Relationship between relative citation counts (expressed in percentages above or below the average of the given field) and various citation dynamics variables. Dots indicate bin means, error bars correspond to confidence intervals ($p = .95$).

This regularity is particularly striking in physics, and much less obvious in computer science; it is also the only feature which is not similar across all fields.

3. Papers with a medium citation age are those which are generally more cited. This last feature possibly indicates that papers with highest relative impact are having a longer career, yet not too long: lately cited papers might get cited too late to raise sufficient mainstream interest.

Interestingly, features 1 and 2 may globally both act as predictors, at publication time, of the likely future impact of a paper in its field. In any event, on the whole, these results emphasize that the most discipline-focused papers and those citing essentially the most recent science are getting the highest interest from the community. Notice that these are traditional markers of “normal science”, but could also correspond to review articles.³ By contrast, interdisciplinary papers and/or papers referring back to older times are raising a markedly weaker interest.

Rebirth as a complementary quality criterion

As observed above, papers being cited late (higher average citation age) are receiving less citations from the community. Such kind of papers are likely to belong to an especially interesting type (e.g. ‘sleeping beauties’), yet their empirical occurrence has also been shown to be rare [11, 12]. More precisely, from these previous studies, which also address the validity of citation-based indicators and the role of time windows when analyzing rebirthing papers [in particular 12], “non-impact” on the early years after publication appears to be a good predictor of future “non-impact” — in other words, from papers which exhibit “non-impact” initially, few eventually surge. A similar effect is observable in our study: in “phys”, for instance, papers with an average citation age over 7.5 years (corresponding to less-than-average citations, from Fig. 5) represent only 0.37% of CMIN, or about a thousand papers (to be compared with the 1.37 million papers present in ALL).

However, these works were based on an *a priori* threshold where a “sleeping beauty” is defined to be such that after s steps we have c citations. By

---

³Unfortunately, a comprehensive examination of the dataset in order to identify separately these two classes of publication was out of the range of this work.

Figure 6: *Left*: distribution of papers according to ρ for all datasets (markers located at the middle of each bin of width 0.05). *Right*: absolute number of papers having a given rebirth index (histogram bins of width 0.1).

contrast, our re-birth index ρ aims at discriminating papers with two periods of citations, among those which have a minimal citation count (CMIN). Its continuous nature additionally allows for a more complete study of these cases. First we have computed the distribution of ρ values across the four discipline datasets, as shown in Fig. 6. In all four cases ρ appears to be modally distributed around ≈ 0.25 , with a small majority of cases falling within the $[0.2, 0.3]$ range, while a still significant proportion of papers exhibits lower values ($[0, 0.2[$, around 30% of papers on average) and higher values ($\rho > 0.3$, around 15% on average).

Then, we have calculated the relative impact of articles as a function of their rebirth index value, which can be found in Fig. 5. These results can be interpreted in terms of low, mid and high rebirthing characteristics without pre-defining such categories and specific range values:

- *Low values of ρ* : These are short-lived papers whose impact is generally limited to its publication time. It is not surprising that they got fewer citations than average.
- *Mid values of ρ* : These are papers that have been receiving citations, more or less continuously, during 10 years. This means that either they made an interesting contribution to the community, or the authors have been continuously supporting them through self-citation. At least in the former case, it is normal (and deserved) that they have a high relative citation count.
- *High values of ρ* : These papers all had a first surge, then a plateau,

then a boost again, 7-10 years after publication. These papers are not getting much success in terms of eventual total citation impact. This could sound counterintuitive: papers which are rediscovered later and, therefore, experiencing a “second life”, are plausibly helpful for a future field.

Put shortly, we rather show that papers with two periods of life after a first surge of interest, which are plausibly interesting in some respect as they exhibit delayed (re-)recognition, have lower-than-average citation counts compared with other CMIN papers. In other words, from papers which exhibit some impact, the ones that have two citation periods (possibly pioneers during their first period, at least with respect to the second period) are not well retributed by citation counts.

Revisiting these results in light of the debate about the appropriateness of the citation count as a impact (and frequently, by extension, quality) measure for publications, we find them contradictory. On one side the first two cases are in agreement with the role of citation count as a quality discriminator. On the other, the important role of ‘rediscovered’ or ‘early birds’ papers is not recompensed by a high citation rate. This latter point is crucial, because in these cases citation count would fail as a quality measure; the rebirth index, subsequently, would be an appropriate complementary and novel quality measure. Put differently, we show here that rebirth and impact are two distinct concepts, in that low impacts corresponds both to high and no rebirth. In this respect, the rebirth index could be used as a quality-defining criterion complementing final citation counts and able to distinguish for instance, among papers with lower-than-average final citation counts, those which are eventually rebirthing from the others.

4. Concluding remarks

It is remarkable that the above dynamic temporal signatures for quite diverse fields – from physics to economics – are significantly similar for all datasets. In most cases, they exhibit a similar behavior; often, identical trends with comparable orders of magnitudes in the y-axis and suggest a universal behavior that even seems to be weakly dependent on a given time period. We suggest that this similarity, which resembles and complements the static universality of citation distributions across disciplines identified by Radicchi and co-authors in [18], may be helpful in generalizing the present results and conclusions.

These correlations provide possibly useful predictive proxies regarding the citation future of a given paper, based only on its immediate temporal neighborhood. In other words, it is possible to statistically infer the relative interest a paper will trigger in a given community, based on its linkage profile towards its direct past neighborhood.

We also propose an index, ρ , quantifying rebirthing of a paper, i.e. the extent to which a paper has a second period of citation after several years of stagnation. Despite such kind of papers are usually quite relevant for their corresponding communities, we find that they are normally weakly cited. The same applies to papers with a higher mean citation age, possibly ‘sleeping beauties’ being cited late, which have a lower final citation count. In a broader view, second-order dynamic indices like ρ can make it possible to detect papers with interesting metrics but little number of citations — so as to eventually diverge from traditional considerations equating quality with citation counts. As such, we aim to contribute to challenging raw citation counts by emphasizing the *type of impact*, often equated with quality, that it specifically appraises.

Finally, these findings touch an essential feature of most reflexive social systems, especially those where quality and evaluation issues are crucially influencing agent behavior. Revealing these patterns, and thus assuming that the average age in the reference list of a given paper says something about its quality, it would be interesting to see how scientific communities might subsequently behave: will authors, in the future, tend to adapt their reference list in such a way that they all exhibit this optimal value; making, for one, the present finding immediately obsolete?

Acknowledgments.

We thank Dirk Helbing for his support, Marius Beceanu for fruitful discussions, and Marc Barthélémy for his useful comments. We also warmly thank two anonymous referees for their constructive remarks and suggestions. This work was partially supported by the Future and Emerging Technologies programme FP7-COSI-ICT of the European Commission through project QLectives (grant no.: 231200). We acknowledge Thomson Reuters Scientific ISI as required by the terms of use of ISI WoS data.

Appendix A. Basic quantitative features of the dataset

The dynamics of publication, absolute citation and reference counts *within* each dataset is shown on Figs. A.1 and A.2. Note that when a paper is marked as having R references or C citations, the figure corresponds to within-field references and citations *as described in the field-focused dataset*. In other words, a *cs* paper with a reference count of 0 is still likely to have references, but none marked as belonging to the field of computer science, as tracked in the ISI dataset.

Figure A.1: *Left:* yearly number of papers published. *Middle and right:* proportion of papers respectively having at least C citations and R references. Colors are set as defined in Fig. 3 of the main text.

Figure A.2: Left: average count of citations after 10 years $\langle C^{year} \rangle$ for ALL (below) and CMIN (above). Right: average count of references not older than 10 years for $\langle R^{year} \rangle$ for ALL (above) and RMIN (below). Colors are set as defined in Fig. 3 of the main text.

References

- [1] Boerner, K., Maru, J. T., Goldstone, R. L., 2004. The simultaneous evolution of author and paper networks. *PNAS* 101 (Suppl 1), 5266–5273.
- [2] Bornmann, L., Daniel, H.-D., 2009. Universality of citation distributions-a validation of radicchi et al.’s relative indicator. *Journal of the American Society for Information Science and Technology* 60.
- [3] Braun, T., Glänzel, W., Schubert, A., 2010. On sleeping beauties, princes and other tales of citation distributions... *Research Evaluation* 19 (3), 195–202.
- [4] Burrell, Q. L., 2005. Are “sleeping beauties” to be expected? *Scientometrics* 65, 381–389.
- [5] Castellano, C., Radicchi, F., 2009. On the fairness of using relative indicators for comparing citation performance in different disciplines. *Archivum Immunologiae et Therapiae Experimentalis* 57 (2), 85–90.
- [6] Chen, P., Xie, H., Maslov, S., Redner, S., 2007. Finding scientific gems with Google’s PageRank algorithm. *Journal of Informetrics* 1, 8–15.
- [7] de Solla Price, D. J., 1965. Networks of scientific papers. *Science* 149 (3683), 510–515.
- [8] de Solla Price, D. J., 1976. A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science* 27 (5–6), 292–306.
- [9] Egghe, L., 2010. The hirsch-index and related impact measures. *Annual Review of Information Science and Technology* 44, 65–114.
- [10] Garfield, E., 1979. Is citation analysis a legitimate evaluation tool? *Scientometrics* 1 (4), 359–375.
- [11] Glänzel, W., Garfield, E., 2004. The myth of delayed recognition. *The Scientist* 18 (11), 8.

- [12] Glänzel, W., Schlemmer, B., Thijs, B., 2003. Better late than never? On the chance to become highly cited only beyond the standard bibliometric time horizon. *Scientometrics* 58, 571–586.
- [13] Glänzel, W., Schubert, A., 1995. Predictive aspects of a stochastic model for citation processes. *Information Processing & Management* 31, 69–80.
- [14] Greenberg, S., 2009. How citation distortions create unfounded authority: analysis of a citation network. *BMJ* 339 (2680).
- [15] Hu, X., Rousseau, R., Chen, J., 2011. On the definition of forward and backward citation generations. *Journal of Informetrics* 5 (1), 27–36.
- [16] MacRoberts, M. H., MacRoberts, B. R., 1996. Problems of citation analysis. *Scientometrics* 36 (3), 435–444.
- [17] Menczer, F., 2004. Evolution of document networks. *PNAS* 101 (S1), 5261–5265.
- [18] Radicchi, F., Fortunato, S., Castellano, C., 2008. Universality of citation distributions: Toward an objective measure of scientific impact. *PNAS* 105, 17268–17272.
- [19] Redner, S., 1998. How popular is your paper? An empirical study of the citation distribution. *European Physical Journal B* 4, 131.
- [20] Seglen, P. O., 1992. The skewness of science. *Journal of the American Society for Information Science* 43 (9), 628–638.
- [21] The PLoS Medicine Editors, 06 2006. The impact factor game. *PLoS Med* 3 (6), e291.
- [22] van Raan, A. F. J., 2004. Sleeping beauties in science. *Scientometrics*, 467–472.
- [23] Wallace, M. L., Larivière, V., Gingras, Y., 2009. Modeling a century of citation distributions. *Journal of Informetrics* 3 (4), 296–303.
- [24] Zitt, M., Ramanana-Rahary, S., Bassecoulard, E., 2005. Relativity of citation performance and excellence measures: From cross-field to cross-scale effects of field-normalisation. *Scientometrics* 63, 373–401.

Appendix A. Pseudo-code for the computation of c_i , $\hat{c}_i$, $\tilde{c}_i$ and ρ_i

---

Algorithm 1:

---

input : A list of publication years $(Y_i)_{i \in \{1, \dots, n\}}$ and a list of citation sets $(\text{Cit}_i)_{i \in \{1, \dots, n\}}$, such that $\forall i, Y_i \in \mathbb{N}$ is the publication year of i and $\text{Cit}_i \subset \{1, \dots, n\}$ is the list of papers citing i .

output: A list of c_i , $\hat{c}_i$, $\tilde{c}_i$ and ρ_i indices.

```

 $t_{\max} \leftarrow 10$ ;
for  $y \in \{\text{set of possible years}\}$  do
   $S_y \leftarrow 0$ ;           // define  $S_y$  as the number of publications on year  $y$ 
for  $i \leftarrow 1$  to  $n$  do
   $S_{Y_i} \leftarrow S_{Y_i} + 1$ ;
for  $i \leftarrow 1$  to  $n$  do
   $\kappa \leftarrow 0$ ;
  for  $j \leftarrow 0$  to  $t_{\max}$  do
     $c_i(j) = |\{k \in \text{Cit}_i | Y_k = Y_i + j\}|$ ;
     $\hat{c}_i(j) = c_i(j) / S_{Y_i + j}$ ;    //  $c_i(j)$  is the number of citations  $j$  years
    after publication,  $\hat{c}_i(j)$  is normalized by the number of
    publication on that year
     $\kappa \leftarrow \kappa + \hat{c}_i(j)$ ;
  for  $j \leftarrow 0$  to  $t_{\max}$  do
     $\tilde{c}_i(j) = \hat{c}_i(j) / \kappa$ ;           // normalize  $\hat{c}_i$  coefficients to sum to 1
   $\rho^\top \leftarrow 0$ ;  $\rho^\perp \leftarrow 0$ ;
  for  $j \leftarrow 0$  to  $t_{\max}$  do
    for  $k \leftarrow 0$  to  $t_{\max}$  do
       $r \leftarrow \tilde{c}_i(j)^2 * \tilde{c}_i(k)^2$ ;
       $\rho^\perp \leftarrow \rho^\perp + r$ ;
       $\rho^\top \leftarrow \rho^\top + r * |k - j| / t_{\max}$ ;
   $\rho_i \leftarrow \rho^\top / \rho^\perp$ ;

```