

The Time Dimension of Science: Connecting the Past to the Future

Yian Yin^{1,2} and Dashun Wang^{1,2,3}

¹Northwestern Institute on Complex Systems and Data Science (NICO), Northwestern University, Evanston, IL 60208, USA

²McCormick School of Engineering and Applied Sciences, Northwestern University, Evanston, IL 60208, USA

³Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA

April 15, 2017

Abstract

A central question in science of science concerns how time affects citations. Despite the long-standing interests and its broad impact, we lack systematic answers to this simple yet fundamental question. By reviewing and classifying prior studies for the past 50 years, we find a significant lack of consensus in the literature, primarily due to the coexistence of retrospective and prospective approaches to measuring citation age distributions. These two approaches have been pursued in parallel, lacking any known connections between the two. Here we developed a new theoretical framework that not only allows us to connect the two approaches through precise mathematical relationships, it also helps us reconcile the interplay between temporal decay of citations and the growth of science, helping us uncover new functional forms characterizing citation age distributions. We find retrospective distribution follows a lognormal distribution with exponential cutoff, while prospective distribution is governed by the interplay between a lognormal distribution and the growth in the number of references. Most interestingly, the two approaches can be connected once rescaled by the growth of publications and citations. We further validate our framework using both large-scale citation datasets and analytical models capturing citation dynamics. Together this paper presents a comprehensive analysis of the time dimension of science, representing a new empirical and theoretical basis for all future studies in this area.

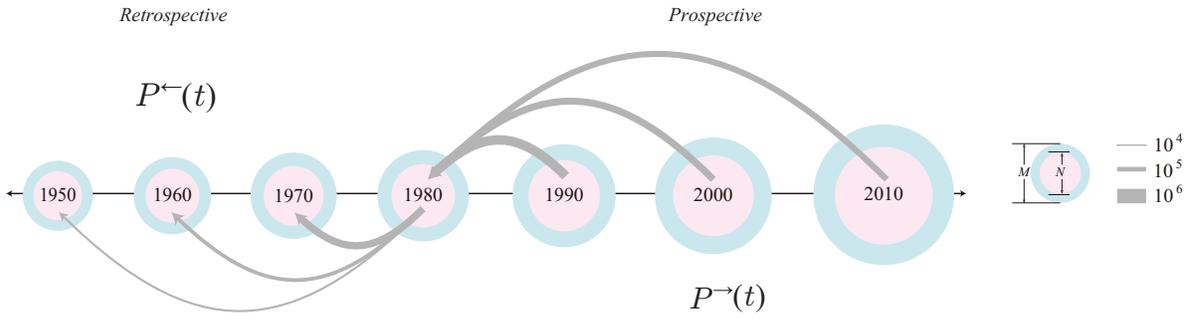
1 Introduction

The increasing availability of large-scale datasets that capture major activities in science has created an unprecedented opportunity to explore broad and important patterns underlying the scientific enterprise, catalyzing in a drastic fashion a recent multidisciplinary shift in our quantitative understanding of science (Redner, 2005; Guimera et al., 2005; Radicchi et al., 2008; Jones et al., 2008; Hirsch, 2005; Stringer et al., 2008; Newman, 2009; Petersen et al., 2014; Evans, 2008; Evans and Foster, 2011; Szántó-Várnagy et al., 2014; Wang et al., 2013; Uzzi et al., 2013; Sinatra et al., 2015; Ke et al., 2015; Sinatra et al., 2016; Guevara et al., 2016). Nowhere are these new advances more apparent than in the studies of citations (Redner, 2005; Bornmann and Daniel, 2008; Radicchi et al., 2008; Eom and Fortunato, 2011; Barabási et al., 2012; Moreira et al., 2015; Wang et al., 2013; Uzzi et al., 2013), owing largely to their widespread applications, from science policy to promotions (Lane, 2010; Radicchi et al., 2009; Shen and Barabási, 2014), hiring (Clauaset et al., 2015; Duch et al., 2012), assignment of grants (Lane and Bertuzzi, 2011; Li and Agha, 2015; Bromham et al., 2016) and prizes (Mazloumian et al., 2011). The aim of this paper is to carry out a comprehensive analysis on one of the most fundamental dimensions of citations: Time.

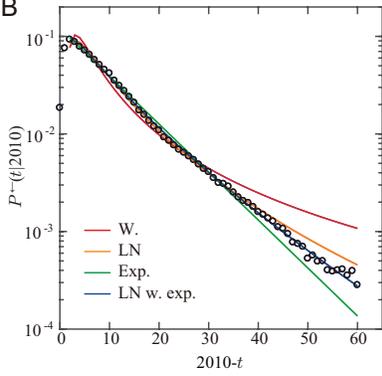
Time plays a central role in science. Indeed, while some papers stay relevant and continue to dominate the scientific discourse over a long period, most papers unfortunately ‘die out’, after collecting their fair share of citations (Wang et al., 2013; Stringer et al., 2010). Understanding the time dimension is critical for a wide

range of reasons. It not only helps us understand the rise and fall of scientific paradigms (Kuhn, 1962), tracing knowledge horizons and hotspots in science and technology (Sinatra et al., 2015; Orosz et al., 2016; Mukherjee et al., 2017), it is also critical for allocating investment (Szell and Sinatra, 2015; Ma et al., 2015; Bromham et al., 2016), identifying crucial yet initially unrecognized sleeping beauties (Van Raan, 2004; Ke et al., 2015), assessing and even predicting future impact of inventions and discoveries (Wang et al., 2013; Redner, 1998; Uzzi et al., 2013; Acuna et al., 2012; Newman, 2014; Sinatra et al., 2016). This has become ever more so with the increasing scale, cost and complexity of science (Van Noorden et al., 2014; Börner et al., 2004; Šubelj and Fiala, 2016; Michels and Schmoch, 2012). Indeed, the exponential growth of science (Price de Solla, 1963; Sinatra et al., 2015; Van Noorden et al., 2014) suggests that there is now much more work for scientists to learn from, build upon, and cite, which is further exacerbated by intensifying specialization in science and engineering disciplines (Jones, 2011) and the inevitable dominance of interdisciplinary research spanning institutional and international boundaries (Jones et al., 2008; Evans and Reimer, 2009; Guimera et al., 2005; Deville et al., 2014; Adams, 2013). Furthermore, because citation systems are commonly treated as models of complex interconnected systems, understanding the time dimension of citations will also help deepen our quantitative understanding of complex systems by tightening models and observations that are highly generalizable to broad areas.

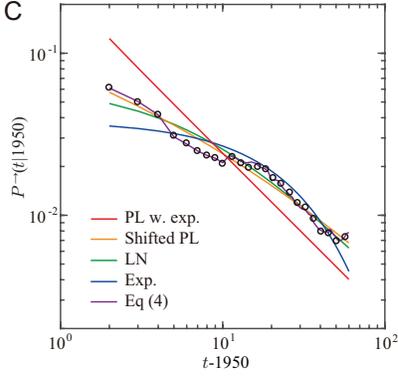
A



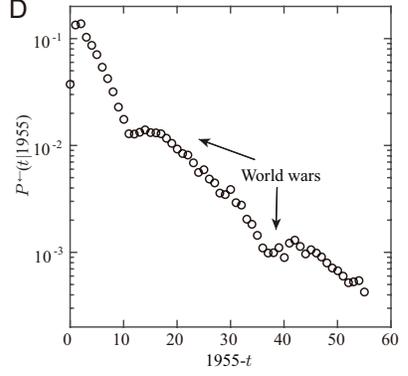
B



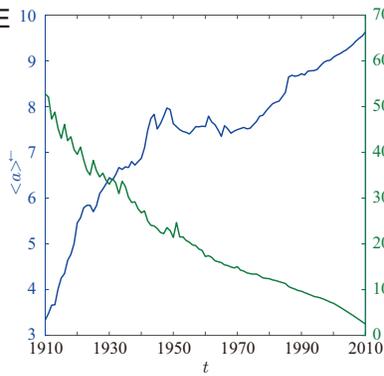
C



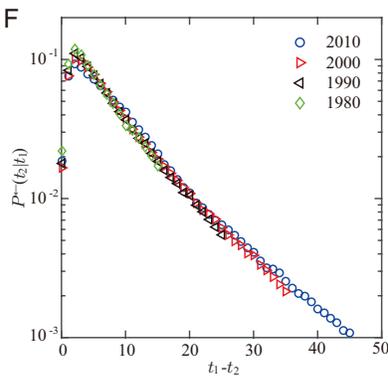
D



E



F



G

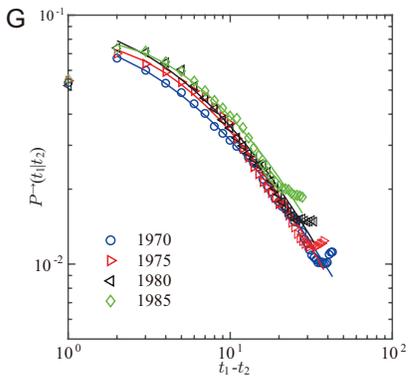


Figure 1 (*preceding page*): **A** The first approach, denoted by $P^{\leftarrow}(t)$, is also called ‘synchronous distribution’ (Nakamoto, 1988), ‘recent impact method’ (Vinkler, 2010) or ‘citations from’ approach as it measures the number of citations of a given age *from* a paper (Redner, 2005), or ‘retrospective citation’ approach as it takes a paper and looks retrospectively back in time (Burrell, 2002; Glänzel, 2004). Denoted with $P^{\rightarrow}(t)$, the second approach studies the age distribution of citations gained over time by a paper (or papers) published in a given year. Correspondingly, this method is also called ‘diachronous distribution’ (Nakamoto, 1988), ‘subsequent impact method’ (Vinkler, 2010) or ‘citations to’ approach (Redner, 2005), or ‘prospective citation’ approach (Burrell, 2002; Glänzel, 2004). Here the width of arrows is proportional to the logarithm of citations, the size of the outer (inner) circle corresponds to the logarithm of total number of references (publications). **B** Retrospective distribution for papers published in 2010, lognormal with exponential cutoff (prediction by Eq (6), LN w. exp.) outperforms Weibull (W., normalized likelihood ratio > 0 (599.3866), p -value $< 10^{-78016}$), lognormal (LN, normalized likelihood ratio > 0 (169.2839), p -value $< 10^{-6225}$) and exponential (Exp., normalized likelihood ratio > 0 (173.0337), p -value $< 10^{-6503}$). **C** Prospective distribution for papers published in 1950, prediction by Eq (4) outperforms lognormal (LN, normalized likelihood ratio > 0 (50.0941), p -value $< 10^{-546}$) shifted power law (Shifted PL, normalized likelihood ratio > 0 (59.0745), p -value $< 10^{-759}$), power law with exponential distribution (PL w. exp., normalized likelihood ratio > 0 (202.7706), p -value $< 10^{-8930}$) and exponential (Exp., normalized likelihood ratio > 0 (96.9191), p -value $< 10^{-2041}$). **D** Retrospective distribution for papers published in 1955, there are two dips due to world wars. **E** Average age of citations using retrospective and prospective approaches. Significant fluctuations are observed for average age of citations until 1970. **F** Retrospective distribution, $P^{\leftarrow}(t_2|t_1)$, in different years. **G** Prospective distribution, $P^{\rightarrow}(t_1|t_2)$, in different years. Solid lines show a lognormal fit.

It is, therefore, not surprising that this question has been extensively investigated over the past several decades (Redner, 2005; Evans, 2008; Wang et al., 2013; Börner et al., 2004; Burton and Kebler, 1960; Garfield and Sher, 1963; de Solla Price, 1965; Margolis, 1967; MacRae Jr, 1969; Krauze and Hillinger, 1971; Avramescu, 1979; Stinson and Lancaster, 1987; Nakamoto, 1988; Motylev, 1989; Matricciani, 1991; Egghe and Rao., 1992; Gupta, 1997; Tsay, 1999; Pollman, 2000; Burrell, 2002; Glänzel, 2004; Sanyal, 2006; Simkin and Roychowdhury, 2007; Larivière et al., 2008; Bouabid, 2011; Verstak et al., 2014; Golosovsky and Solomon, 2014; Parolo et al., 2015; Pan et al., 2016), being one of most prolific lines of inquiry in science of science studies. Yet, despite its broad impact and rich historical context, we lack systematic answers to a simple yet fundamental question: How does time affect citations? In this paper, we systematically investigate this question by leveraging large-scale citation datasets and recent models capturing citations dynamics. We start by conducting a comprehensive review of existing literature, which reveals a significant lack of consensus on this matter. The main reason of this lack of consensus is that empirical measurements of temporal effect in citations have taken two related yet distinct measurement approaches (Fig. 1A). The first approach considers papers that are cited by a publication and analyzes *retrospectively* the age distribution of these references (Redner, 2005; Burton and Kebler, 1960; Garfield and Sher, 1963; de Solla Price, 1965; Margolis, 1967; MacRae Jr, 1969; Krauze and Hillinger, 1971; Avramescu, 1979; Stinson and Lancaster, 1987; Nakamoto, 1988; Motylev, 1989; Matricciani, 1991; Egghe and Rao., 1992; Gupta, 1997; Tsay, 1999; Pollman, 2000; Burrell, 2002; Glänzel, 2004; Larivière et al., 2008; Verstak et al., 2014; Golosovsky and Solomon, 2014; Pan et al., 2016). The second approach, in contrast, studies *prospectively* the age distribution of citations that are gained over time by a paper (Redner, 2005; Wang et al., 2013; MacRae Jr, 1969; Krauze and Hillinger, 1971; Avramescu, 1979; Stinson and Lancaster, 1987; Nakamoto, 1988; Motylev, 1989; Glänzel, 2004; Sanyal, 2006; Simkin and Roychowdhury, 2007; Bouabid, 2011; Golosovsky and Solomon, 2014; Parolo et al., 2015). The subtle difference between the two approaches creates a dramatic yet largely understudied effect in temporal citation patterns (Table 1), further confounded by the exponential growth in both the number of papers and references cited by them (Sinatra et al., 2015; Van Noorden et al., 2014; Pan et al., 2016).

The temporal behavior of citations has been measured and reported by independent research groups, each using specific datasets and measurement details (Table 1). The lack of consensus, and the coexistence of the two approaches, raise an important question: What is the most appropriate functional form describing the time dimension of science? Here, by introducing a general theoretical framework, we provide systematic answers to this question, which is then validated both analytically and empirically through citation models and large-scale datasets. As such, the paper makes several contributions to this prolific direction in the science of science. First and foremost, it serves as a detailed and much needed survey by reviewing and repeating results from most major studies in this domain. We then used the Web of Science (WoS) dataset to systematically measure and test results obtained by prior studies. We examine these results in the context of growth in science, systematically testing the most appropriate model describing aging measured by the prevailing two approaches. More importantly, despite the concomitant development of both approaches in the literature, we lack any known connections between the two. Here we introduced a general framework that helps us uncover that the two different approaches are connected through precise mathematical relationships, allowing us to derive and even predict one approach from the other. We further derived our empirical results using citation models, providing theoretical support for the observed temporal behavior. We conclude by a brief discussion on scaling relationships between aging structure and citation impact of papers. Together, this paper presents a comprehensive study on the time dimension of science, providing a new empirical and theoretical basis for all further studies on this topic. These results are not only relevant for the emerging field of the science of science, they also help improve our understanding of complex systems, by pushing forward measurements and models in complex evolving networks.

2 Results

A frequently applied measure for the time dimension of science is the age distribution of citations. Since 1960s, there have been numerous studies reporting on this distribution by taking one or both of the two approaches. While both approaches focus on the age of citations, they measure quantities pertaining to rather different processes (Fig. 1A). Indeed, the retrospective approach, denoted by $P^-(t)$, measures how far back the paper looks in time to cite old references, hence quantifying the citing memory of a paper.

Paper	$P^{\leftarrow}(t)$	$P^{\rightarrow}(t)$	$P_N^{\leftarrow}(t)$	$P_r^{\rightarrow}(t)$
Burton and Kebler (1960)	Exp. v3	\times	\times	\times
Garfield and Sher (1963)	—	\times	\times	\times
de Solla Price (1965)	Exp.	\times	Const.	\times
Margolis (1967)	Exp.	\times	\times	\times
MacRae Jr (1969)	Exp.	Exp.	Exp.	\times
Krauze and Hillinger (1971)	Exp.	Exp.	Exp.	\times
Avramescu (1979)	Exp. v2	—	\times .	\times
Stinson and Lancaster (1987)	—	—	—	—
Nakamoto (1988)	Exp.	Exp.	\times	\times
Motylev (1989)	Exp. v3	—	—	\times
Matricciani (1991)	LN	\times	\times	\times
Egghe and Rao. (1992)	LN>W., Exp. v2, NB	\times	\times	\times
Gupta (1997)	LN	\times	\times	\times
Tsay (1999)	Exp.	\times	\times	\times
Pollman (2000)	Exp. v1> Exp.	\times	\times	\times
Burrell (2002)	LN>LL, W.	\times	\times	\times
Glänzel (2004)	—	PL	\times	\times
Börner et al. (2004)	W.	\times	\times	\times
Redner (2005)	Exp.	PL	PL	—
Sanyal (2006)	\times	LN	\times	\times
Simkin and Roychowdhury (2007)	\times	PL w. Exp.	\times	\times
Larivière et al. (2008)	—	\times	\times	\times
Evans (2008)	—	\times	\times	\times
Bouabid (2011)	\times	LL	\times	\times
Wang et al. (2013)	\times	LN	\times	\times
Verstak et al. (2014)	—	\times	\times	\times
Golosovsky and Solomon (2014)	—	—	\times	\times
Parolo et al. (2015)	\times	Exp. v1>PL w. const	\times	\times
Pan et al. (2016)	Exp.	—	—	\times

Table 1: Age distribution of citations: retrospective vs. prospective approaches. Not measured: \times . Measured but no functional form was offered: —. See Table 2 for explanations of abbreviations. $X > Y$: X outperforms Y during fitting for $P(t)$

Abbreviation	Full Name	Functional Form of $P(t)$
Const.	Constant	1
Exp.	Exponential	$e^{-\lambda t}$
Exp. v1	Exponential variant 1	$\beta e^{-\lambda t} + \gamma$
Exp. v2	Exponential variant 2	$e^{-\lambda t} - e^{-m\lambda t}$
Exp. v3	Exponential variant 3	$ae^{-x} + 2(1-a)e^{-2x}$
PL	Power law	$t^{-\alpha}$
PL. w. const	Power law with constant	$\beta t^{-\alpha} + \gamma$
Pow. w. Exp.	Power law with exponential cutoff	$t^{-\alpha} e^{-\lambda t}$
W.	Weibull	$(\frac{t}{\alpha})^{\beta-1} e^{-(\frac{t}{\alpha})^{\beta}}$
NB	Negative Binomial	$\frac{\Gamma(k+t)}{\Gamma(k)\Gamma(t+1)} p^k q^t$
LN	Lognormal	$\frac{1}{t} e^{-\frac{(\ln(t)-\mu)^2}{2\sigma^2}}$
LL	Log-logistic	$\frac{(x/\alpha)^{\beta-1}}{(1+(x/\alpha)^{\beta})^2}$

Table 2: Distributions appeared in Table 1

The prospective approach $P^{\rightarrow}(t)$, on the other hand, measures how a paper is cited by others in the future, effectively measuring how the research community remembers a paper over time. Here, we review all major studies on this topic and summarize them in Table 1, asking specifically: what approach did they take? And what results did they obtain? Most early studies have been focused on the retrospective approach, a tradition dating back at least to Burton and Kebler (Burton and Kebler, 1960). Out of the seventeen papers that reported explicitly the functional form of $P^{\leftarrow}(t)$ in Table 1, twelve of them supported the exponential model or its variants to be the best fit to data. Since MacRae compared the retrospective and prospective citation age distributions in 1969 (MacRae Jr, 1969), more and more studies have started to focus on the prospective approach to measuring age distribution of citations. Comparing the two columns under $P^{\leftarrow}(t)$ and $P^{\rightarrow}(t)$ in Table 1 reveals that the prospective distributions are typically characterized by broader distributions than retrospective distributions, ranging from power law to power law with an exponential cutoff to lognormal distributions. Note that the different results reported in Table 1 do not necessarily imply that these papers contradict with each other, because they strongly depend on measurement details as well as the corpus of papers on which these studies were carried out. The remarkably diverse conclusions summarized in Table 1, however, raise a tantalizing question: When it comes to citation age distributions, for both $P^{\leftarrow}(t)$ and $P^{\rightarrow}(t)$, what are the most appropriate functional forms that capture their temporal behavior? To answer this question, we take the Web of Science dataset, arguably the most authoritative citation index, to systematically measure both retrospective and prospective distributions.

2.1 Empirical Measurements

To understand $P^{\leftarrow}(t)$, we first test the three most frequently applied distributions: Weibull, exponential and lognormal distributions (Table 2), asking which one best fits our data. We use maximum likelihood estimation in this paper to determine the fit (Sec. S1). Since the tail part of the distribution is of the most interest and the citations in the first year after publication can be largely affected by publication date, we follow recent approaches by Parolo *et al.* (Parolo et al., 2015) and focus on data after the peak for the fit (typically occurring around 2 years). More specifically, We take all papers published in the same year (e.g. year 2010), trace retrospectively back in time, and measure the number of citations from the focal year to each earlier time unit. We find among the three functions tested, lognormal distribution provides the best fit for $P^{\leftarrow}(t)$ in terms of log-likelihood estimation, outperforming both Weibull and exponential distributions (Fig. 1B). However, none of them provides a satisfactory fit, showing an evident deviation for the tail part, hinting for alternative models that can fit $P^{\leftarrow}(t)$ more accurately. Indeed, we use a lognormal distribution with exponential cutoff, for reasons that will be explained later, and find it provides a better fit to the data (outperforming other distributions from likelihood ratio test with $p - value < 10^{-6225}$).

For the prospective distribution $P^{\rightarrow}(t)$, we tested four types of functions reported in Table 1, from exponential distribution to lognormal to power law with and without exponential cutoff (Table 2). We also used shifted power law, i.e. $P(t) \sim (t_0 + t)^\alpha$, to fit our data. We calculated $P^{\rightarrow}(t)$ using the same set of papers as Figure 1B, finding that lognormal function outperforms other distributions (Fig. 1C). Interestingly, power law with and without exponential cutoff yield similar results, because the obtained exponential cutoff approaches its lower bound, playing a minor role in distinguishing the two distributions. Indeed, as we show later, the prospective distribution follows Eq (4), which not only gives a better fitting (outperforming other distributions from likelihood ratio test with $p - value < 10^{-546}$), but also captures all the fluctuations in empirical data.

Both distributions can be affected by significant external shocks to the system. Indeed, repeating our analysis for year 1955 reveals two ‘dips’ in the retrospective distribution around 1942 and 1919, respectively (Fig. 1D), documenting interruptions created by wars. To quantify the extent to which the system is affected by external forces, we calculate the average age of all citations both from and to papers published at time t , denoted by $\langle a \rangle^{\leftarrow}(t)$ and $\langle a \rangle^{\rightarrow}(t)$, respectively. Figure 1E shows that $\langle a \rangle^{\leftarrow}(t)$ shows larger fluctuations prior to 1970. Hence in this paper, we measure P^{\leftarrow} for data from 1970 to 2013. For P^{\rightarrow} , we use data after the Second World War, i.e. from 1948 to 2013.

Figure 1E also demonstrates that the aging structure of citations changes over time (Sinatra et al., 2015; Verstak et al., 2014; Pan et al., 2016). To further understand temporal evolution of the aging structures, we calculate $P^{\leftarrow}(t)$ and $P^{\rightarrow}(t)$ for different years t_0 . If we group paper samples based on their publication year t_0 , mathematically this can be described as $P^{\leftarrow}(t|t_0)$ and $P^{\rightarrow}(t|t_0)$. To this end we assume t_1 is a more recent

year than t_2 ($t_1 > t_2$). Then $P^{\leftarrow}(t_2|t_1)$ measures the probability that a citation from t_1 points to a reference published at time t_2 . Similarly $P^{\rightarrow}(t_1|t_2)$ measures the probability that a paper published in t_1 is later cited at t_2 . We plot both distributions for different years (Fig. 1F,G). For $P^{\leftarrow}(t_2|t_1)$, we find papers published in different years after 1980s follow a similar citation pattern, with recent papers having a stronger tendency to cite ‘old papers’ (Sec. S2), which supports earlier findings by using Google Scholar dataset (Verstak et al., 2014). Comparing $P^{\leftarrow}(t_2|t_1)$ (Fig. 1F) with average reference ages for papers published in different years (Fig. 1E), we find this increasing preference for old papers can be traced back to the 1970s, after leaving an unstable regime in the 1960s.

Most important, we find in many cases $P^{\rightarrow}(t)$ is not a monotonically decreasing function. Indeed, we observed ubiquitous ‘upward tails’ characterizing temporal behavior of citations. For example, papers published in 1970 received 7.87% more citations in 2013 than they did in 2006. This indicates, as papers get older, they collect more instead of fewer citations. Repeating this analysis for different years, we find the ‘upward tails’ characterizing $P^{\rightarrow}(t_1|t_2)$ is nearly universal (Fig. 1G). The upward tails offer the strongest demonstration that none of the functions we have discussed so far (Table 1) are adequate in describing $P^{\rightarrow}(t)$. This phenomenon has been reported in prior works (Sinatra et al., 2015), but has received scarce attention in the literature. Next we show, understanding these tails not only helps us understand the role of growth in science in driving age distributions of citations, but also enables us to uncover the correct functional forms describing this temporal behavior.

2.2 Accounting for Growth

$P^{\leftarrow}(t)$ and $P^{\rightarrow}(t)$ are determined by both aging of citations and growth of the citation system. Indeed, if we look further back in time, there are inherently fewer papers to cite, hence affecting $P^{\leftarrow}(t)$. Similarly, the continuous expansion of science impacts how an existing paper is remembered over time as new citations enter the system at an increasing rate. Let us denote with $N(t)$, the number of papers published at time t . Extensive literatures documented that $N(t)$ is best approximated by an exponential function: $N(t) \sim e^{\beta t}$ (Wang et al., 2013; Sinatra et al., 2015; Van Noorden et al., 2014) (Fig. 2A). Besides the sheer increase in the number of papers, the average number of references contained in each paper, $m(t)$, also increases with time (Fig. 2A inset). Indeed, if we define $M(t)$ as the number of citations from papers published at time t , i.e., $M(t) \equiv N(t)m(t)$, to the leading order the growth of M can also be approximated by an exponential function (Fig. 2A). To account for the growth of science, we need to normalize the obtained retrospective distributions by the number of papers available for citations

$$P_N^{\leftarrow}(t_2|t_1) = \frac{P^{\leftarrow}(t_2|t_1)/N(t_2)}{\int_0^{t_1} P^{\leftarrow}(t_2|t_1)/N(t_2)dt_2}, \quad (1)$$

where $P_N^{\leftarrow}(t_2|t_1)$ captures the fact that $P^{\leftarrow}(t_2|t_1)$ is normalized by the number of papers available for reference at time t_2 , $N(t_2)$.

$M(t_1)$ new citations point towards $\int_0^{t_1} N(t)dt$ papers in the system. Thus, for the prospective distribution, we normalize $P^{\rightarrow}(t)$ by the likelihood to be cited by new references, $r(t_1) \equiv \frac{M(t_1)}{\int_0^{t_1} N(t)dt}$, yielding

$$P_r^{\rightarrow}(t_1|t_2) = \frac{P^{\rightarrow}(t_1|t_2)/r(t_1)}{\int_{t_2}^T P^{\rightarrow}(t_1|t_2)/r(t_1)dt_1} \quad (2)$$

Note that $r(t_1)$ reduces to $\beta m(t_1)$ in the condition of $N(t) \sim e^{\beta t}$ (Fig. 2A inset). Here we use $N(t)$ to normalize retrospective distribution, which is equivalent to a null model that assumes every paper has the same chance to be cited. The prospective distribution, on the other hand, is normalized by $r(t)$, indicating the probability for a paper to be cited at a given time is driven by both growth of papers and references. Only a modest amount of literature measured temporal effect on citations by taking into account growth factors (Table 1). To find the best functional form for $P_r^{\rightarrow}(t_1|t_2)$, we take all papers published in a year, measure the number of citations they receive in subsequent years, and normalize the distributions by $r(t)$. For $P_N^{\leftarrow}(t)$, we select papers from year 2013, and normalize their $P^{\leftarrow}(t)$ by $N(t)$. Although $P^{\rightarrow}(t)$ is characterized by a broader distribution than $P^{\leftarrow}(t)$, once normalized however, we find both distributions are best fitted by the

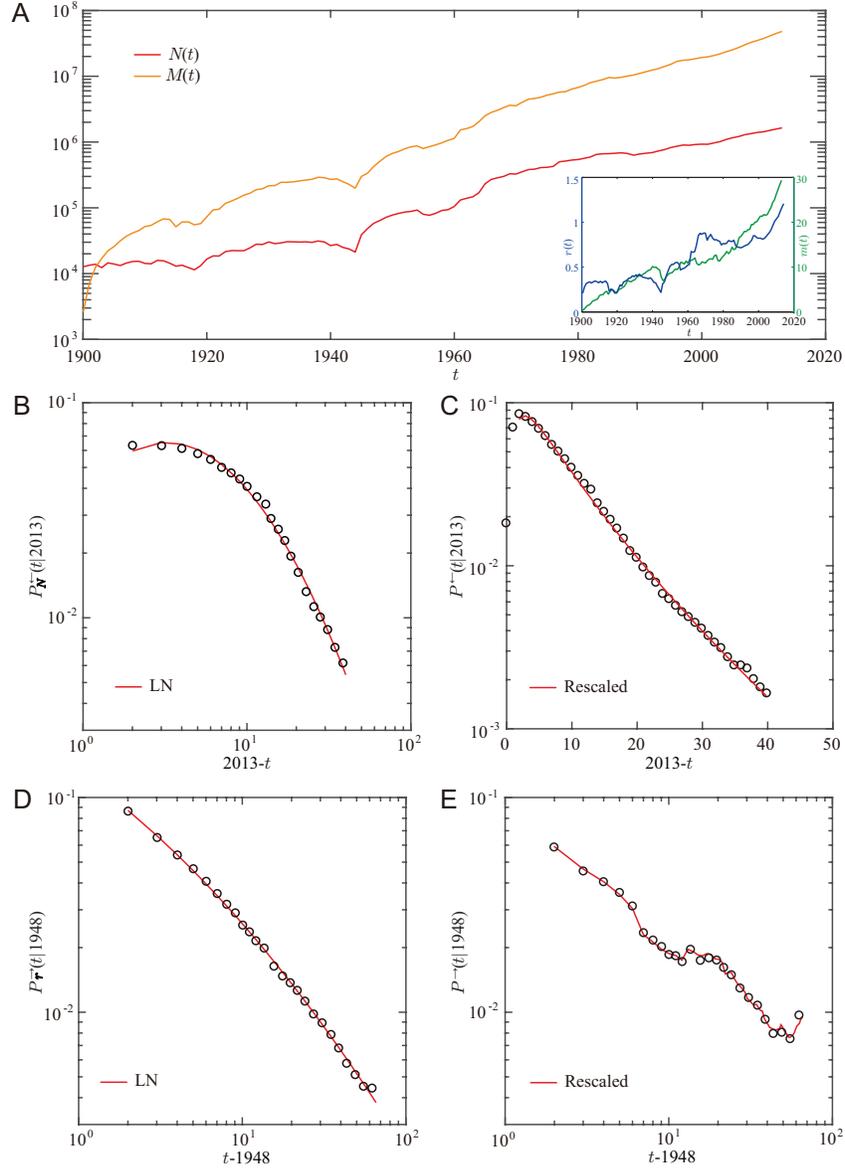


Figure 2: **A** Number of papers published and citations generated in every year grows exponentially. Inset: The growth in the average number of references from papers published every year. **B** Retrospective distribution normalized by N in 2013, with lognormal showing the best fit. **C** Retrospective distribution in 2013. The fitting is recovered from normalized distribution, i.e. lognormal with exponential cutoff is the best fit. **D** Prospective distribution normalized by r for papers published in 1948, with lognormal showing the best fit. **E** Prospective distribution in 1948, the fitting is recovered from normalized distribution. The best fit reproduces the upward tail.

same function: lognormal distribution (Fig. 2BD). That is,

$$P_N^{\leftarrow}(t_2|t_1) \sim \frac{1}{\sqrt{2\pi}\sigma(t_1 - t_2)} e^{-(\ln(t_1-t_2)-\mu)^2/2\sigma^2} \quad (3)$$

$$P_r^{\rightarrow}(t_2|t_1) \sim \frac{1}{\sqrt{2\pi}\sigma(t_1 - t_2)} e^{-(\ln(t_1-t_2)-\mu)^2/2\sigma^2} \quad (4)$$

Indeed, $P^{\leftarrow}(t)$ is simultaneously driven by two factors: the obsolescence of citations, $P_N^{\leftarrow}(t)$, and the growth of science, $N(t)$. Since the latter leads to an exponential decay when we look backwards, the true aging function can be hidden from the rapid growth of science when measuring P^{\leftarrow} directly. To this end, we use (1) and (3) to recover the functional form of P^{\leftarrow} , obtaining

$$P^{\leftarrow}(t_2|t_1) \sim \frac{1}{\sqrt{2\pi}\sigma(t_1 - t_2)} e^{-(\ln(t_1-t_2)-\mu)^2/2\sigma^2 - \beta(t_1-t_2)} \quad (5)$$

where μ and σ are estimated from P_N^{\leftarrow} . Note that β characterizes the exponential growth of publications, which is a fixed system-wide parameter. Therefore, (5) has the same number of independent parameters as other fitting functions analyzed in this paper. Yet, when fitted to retrospective distributions, (5) offers a superior fit than all functions discussed in Table 1 (Fig. 2C). Furthermore, the exponential part dominates over the lognormal factor in the long term, explaining why it appears to be approximated by an exponential function in most earlier cases.

Similarly, we recover the mathematical form of P^{\rightarrow} from P_r^{\rightarrow} , obtaining

$$P^{\rightarrow}(t_1|t_2) \sim \frac{r(t_1)}{\sqrt{2\pi}\sigma(t_1 - t_2)} e^{-(\ln(t_1-t_2)-\mu)^2/2\sigma^2} \quad (6)$$

(6) not only provides an excellent fit to empirical data (Fig. 2E), it also predicts accurately the emergence of upward tails: since $r(t_1)$ increases with time, it competes with aging of citations, compensating the decay in (in some case even increasing) $P^{\rightarrow}(t)$. Using empirical data of $r(t)$, we successfully reproduce these tails, documenting the validity of (6) (Fig. 2E). Note that when $t_1 - t_2$ is small, $\frac{\partial \ln(P^{\rightarrow}(t_1|t_2))}{\partial t_1}$ can be approximated as $\frac{\ln(t_1-t_2)-\mu+\sigma^2}{\sigma^2(t_1-t_2)}$, explaining why lognormal appears to fit P^{\rightarrow} well when $t_1 - t_2$ is small. Moreover, (6) indicates that, for growth factor to dominate over the aging factor, $t_1 - t_2$ needs to be large enough, explaining why upward tails are only evident for old papers.

Taken together, our results show that the time dimension of science is different from any result reported in Table 1: First of all, both distributions follow lognormal distributions after appropriate normalization. $P^{\leftarrow}(t)$ can be modeled by a lognormal distribution with exponential cutoff, where the cutoff is caused by the exponential growth of papers. $P^{\rightarrow}(t)$ is the product of a lognormal function and $r(t)$, which lacks clear mathematical forms. A recent literature (Pan et al., 2016) suggests $r(t)$ to be an exponential function, which means $P^{\rightarrow}(t_1|t_2)$ will be dominated by the growth of $r(t)$ and increases infinitely as t_1 increases. However, if we assume $r(t)$ to grow with a slower rate, e.g. following a linear or power law function (Ucar et al., 2014; Krapivsky and Redner, 2005; Vazquez, 2001; Allen et al., 1994), the upward tail is only temporary and $P^{\rightarrow}(t_1|t_2)$ will be eventually dominated by the lognormal decay from $P_r^{\rightarrow}(t)$.

To test the robustness of (5) and (6), we further repeat our empirical measurement for thirty years (See Fig. S2,3 in Supplementary Information), finding a universal pattern of citation decay across decades. Next we show, (5) and (6) not only offer more accurate mathematical forms to fit the distributions, they also allow us to build a conceptually new framework which is closely related to several citation models.

2.3 Validations

Given the new results obtained in (5) and (6), we must ask: how do we validate these results? Next we show two types of validations, by deriving both empirical and analytical evidence that supports the validity of (5) and (6).

To validate (5) empirically, we use normalized retrospective distribution from different years and rescaled them with respect to μ and σ . (5) and (6) predicts that both distributions collapse into standard lognormal distribution, i.e. $Z = \frac{1}{\sqrt{2\pi}X} e^{-(\ln X)^2/2}$, after a proper rescaling (Sec. S3). We test this hypothesis empirically

for P_N^{\leftarrow} , finding excellent data collapse across different years (Fig. 3A). We repeat this test for P_r^{\rightarrow} using (6), finding again an excellent data collapse on the same curve (Fig. 3B). The data collapses documented in Figure 3 offer strong empirical evidence for (5) and (6).

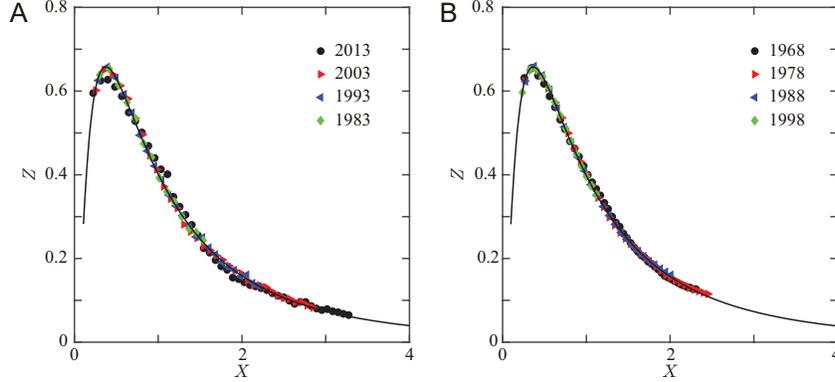


Figure 3: **Empirical support for $P^{\leftarrow}(t)$ and $P^{\rightarrow}(t)$.** **A** $P_N^{\leftarrow}(t)$ rescaled with respect to μ and σ . Distributions from different years collapse on the probability density function (PDF) curve of standard lognormal distribution. **B** $P_r^{\rightarrow}(t)$ rescaled with respect to μ and σ . Similar to A, distributions from different years collapse on the PDF curve of standard lognormal distribution.

What do existing models predict for $P^{\leftarrow}(t)$ and $P^{\rightarrow}(t)$? To this end, we test three commonly used citation models with ascending levels of complexity: Price’s model (Price, 1976), Bianconi-Barabasi model (Bianconi and Barabási, 2001), and Wang-Song-Barabasi model (Wang et al., 2013).

Price’s model assumes that each new paper cites m references and the probability for an old paper to gain citations is proportional to its current citations plus a constant c_0 (Sec. S4.1). We calculate $P^{\leftarrow}(t)$ and $P^{\rightarrow}(t)$ analytically for this model, finding

$$P^{\leftarrow}(t_2|t_1) \sim \frac{1}{t_1} \left(1 - \frac{\Delta t}{t_1}\right)^{-\frac{m}{c_0+m}}, \quad (7)$$

where $\Delta t = t_1 - t_2$. (7) predicts that $P^{\leftarrow}(t_2|t_1)$ increases with Δt . That is, if we measure the citation age distribution following the retrospective approach at a given time t_1 , the distribution increases as Δt increases. This is due to the cumulative advantage mechanism. Indeed, if we look further back in time, older papers are more likely to be characterized by a higher degree hence more likely to be cited. For prospective approach, we obtain

$$P^{\rightarrow}(t_1|t_2) \sim A(t_2 + \Delta t)^{-c_0/(c_0+m)} \quad (8)$$

where $A = \frac{1}{T^{m/(c_0+m)} - t_2^{m/(c_0+m)}}$. Therefore the age distribution decays as a power law $\Delta t^{-c_0/(c_0+m)}$ as we look further in time. Interestingly, Price’s model also predicts that $P^{\leftarrow}(t_2|t_1)$ collapse into a universal function for different t_1 after rescaling $x \equiv t_2/t_1$ (Fig. 4A):

$$P^{\leftarrow}(t_2|t_1) = \frac{c_0}{c_0 + m} \frac{1}{t_1} x^{-m/(c_0+m)} \quad (9)$$

$P^{\rightarrow}(t_1|t_2)$ can also be rescaled into a universal form as (Fig. 4B)

$$P^{\rightarrow}(t_1|t_2) = \frac{m}{c_0 + m} A t_2^{-c_0/(c_0+m)} x^{c_0/(c_0+m)} \quad (10)$$

The Bianconi-Barabasi (BB) model builds on Price’s model by adding a fitness parameter to each paper, η_i (Sec. S4.2). For retrospective approach, we find

$$P^{\leftarrow}(t_2|t_1) \sim t_1 \langle \eta x^{-\eta/C} \rangle, \quad (11)$$

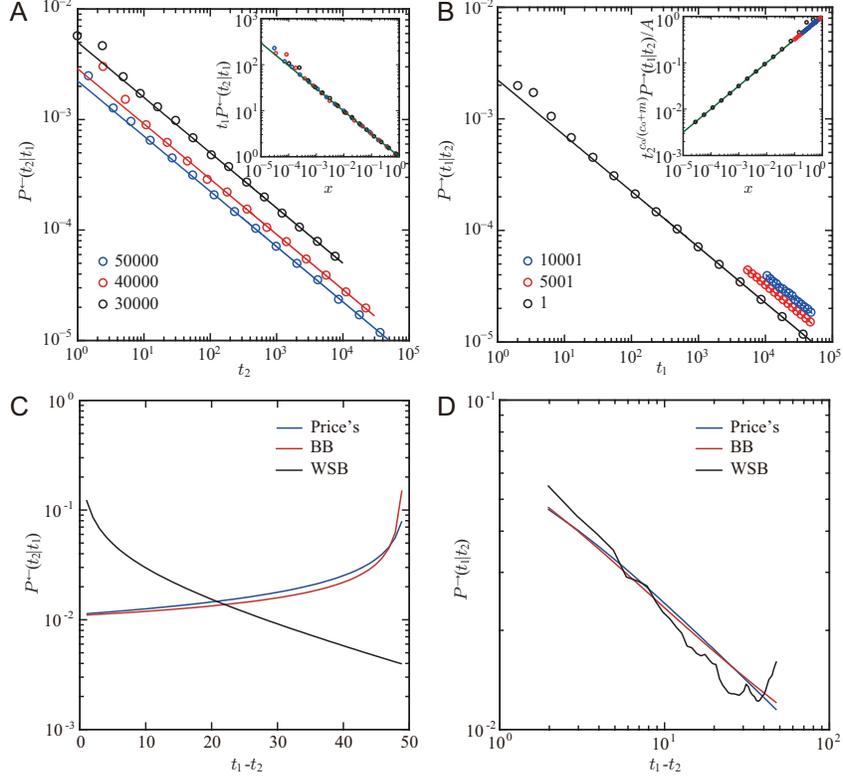


Figure 4: **Theoretical support for $P^{\leftarrow}(t)$ and $P^{\rightarrow}(t)$.** **A** Price's model's prediction of $P^{\leftarrow}(t_2|t_1)$ for different time units. Different circles are our simulations for different t_1 . Solid line shows our analytical calculation (7). Inset shows $P^{\leftarrow}(t_2|t_1)$ for different years collapse into a single curve, solid line showing prediction of (9). **B** Price's model's prediction of $P^{\rightarrow}(t_1|t_2)$ for different years. Different circles are our simulations for different t_2 . Solid line shows our analytical calculation (8). Inset shows $P^{\rightarrow}(t_1|t_2)$ for different years collapse into a single curve, solid line showing prediction of (10). **C** $P^{\leftarrow}(t_2|t_1)$ predicted by Price's model, Bianconi-Barabasi model and Wang-Song-Barabasi model. The WSB model predicts the correct function over time consistent with our empirical results. **D** $P^{\rightarrow}(t_1|t_2)$ predicted by Price's model, Bianconi-Barabasi model and Wang-Song-Barabasi model. Price's model and BB model predict approximately power law decay, while the WSB model is able to reproduce the upward tails given $r(t)$.

where $x \equiv t_2/t_1$, same as the rescaled time we previously defined for Price's model and C is a normalization constant. Hence the BB model too predicts a retrospective distribution that increases with time. For prospective distribution, we find $P^{\rightarrow}(t_1|t_2)$ decays as

$$P^{\rightarrow}(t_1|t_2) \sim \langle \eta x^{1-\eta/C} / t_2 \rangle / (\langle (T/t_2)^{\eta/C} \rangle - 1). \quad (12)$$

Eqs (11) and (12) also indicate that the shape of both distributions depends on the distribution of fitness. We calculate both distributions by assuming an exponential fitness distribution, i.e. $P(\eta) \sim e^{-\eta/\eta_0}$, a commonly used functional form (Kong et al., 2008). We find that $P^{\rightarrow}(t_1|t_2)$ decays as a power law of $(\Delta t)^{-1}$, while $P^{\leftarrow}(t_2|t_1)$ again increases with time, albeit with a slower rate (Fig. 4CD). Both distributions again collapse into universal forms if we use $x = t_2/t_1$:

$$P^{\leftarrow}(t_2|t_1) = \frac{1}{C t_1 \eta_0} \left(\frac{1}{\eta_0} + \frac{1}{C} \ln x \right)^{-2} \quad (13)$$

$$P^{\rightarrow}(t_2|t_1) = \frac{\alpha_2 x}{C t_2 (1 - \alpha_2 \eta_0)} \left(\frac{1}{\eta_0} + \frac{1}{C} \ln x \right)^{-2} \quad (14)$$

where $\alpha_2 = 1/\eta_0 - \ln(T/t_2)/C$.

Finally we repeat our calculations for the Wang-Song-Barabasi (WSB) model (Wang et al., 2013), which builds on the BB model but further incorporates an aging mechanism with the exponential growth of the system (Sec. S4.3). We find prospective distribution follows a lognormal distribution:

$$P^{\rightarrow}(t_2|t_1) \sim \frac{1}{\sqrt{2\pi\sigma\Delta t}} e^{-(\ln(\Delta t)-\mu)^2/2\sigma^2}, \quad (15)$$

and the retrospective distribution follows lognormal distribution with exponential cutoff, i.e.

$$P^{\leftarrow}(t_1|t_2) \sim \frac{1}{\sqrt{2\pi\sigma\Delta t}} e^{-(\ln(\Delta t)-\mu)^2/2\sigma^2-\beta\Delta t}. \quad (16)$$

Somewhat surprisingly, these theoretical predictions are in remarkable agreement with our empirical results (Fig. 4CD). The accuracy of the WSB model is partly due to its ability to incorporate both aging and growth factors into the model. Together these results provide strong theoretical support for (5) and (6).

2.4 Connecting the Two Approaches

The retrospective and prospective approaches are studied concomitantly in the literature, often pursued as separate lines of inquiry, which raise an interesting question: Are these two distributions connected? Indeed, after all, they both measure the age of citations. Should we have measured the age distribution of *all* citations, both approaches would result in the same distribution, as if it did not matter whether we look retrospectively or prospectively in time. Yet despite the large amount of work on this topic (Table 1), little is known if there exists any relationship between the two approaches. To shed light on this question, we calculate $P(t_1, t_2)$, measuring the probability that a randomly selected citation connects a paper published in t_1 with a paper published in t_2 . Because the difference between the two approaches lies in whether we fix the citing or cited papers, we can derive their citation age distribution from $P(t_1, t_2)$. Next we show that expressing P^{\leftarrow} or P^{\rightarrow} in terms of $P(t_1, t_2)$ allows us to not only connect the two approaches, but also derive one approach from the other. Indeed, $P(t_1, t_2)$ can be expressed as

$$P(t_1, t_2) \sim P^{\leftarrow}(t_2|t_1)M(t_1) = P^{\rightarrow}(t_1|t_2)L(t_2) \quad (17)$$

where $L(t_2) = \int_{t_2}^T P^{\leftarrow}(t_2|t_1)M(t_1)dt_1$ measures the total number of citations papers published at t_2 receive. Therefore, (17) documents a precise mathematical connection between $P^{\rightarrow}(t_1|t_2)$ and $P^{\leftarrow}(t_2|t_1)$, mediated by $P(t_1, t_2)$, indicating one can obtain $P^{\rightarrow}(t)$ from $P^{\leftarrow}(t)$ or vice versa. Indeed, using (17) we obtain a rescaling formula for $P^{\rightarrow}(t)$:

$$P^{\rightarrow}(t_1|t_2) = \frac{P^{\leftarrow}(t_2|t_1)M(t_1)}{\int_{t_2}^T P^{\leftarrow}(t_2|t_1)M(t_1)dt_1}, \quad (18)$$

which allows us to calculate the prospective distribution based solely on retrospective distributions (Sec. S5). Taking papers published in 1980 as an example, we collected $P^{\leftarrow}(1980|t)$ for $t > 1980$ and rescaled them using (18). We then compare the rescaled results with real $P^{\rightarrow}(t)$, finding an excellent match between the two, reproducing not only the decay rate but also the upward tail (Fig. 5A).

Similarly, we can solve for $P^{\leftarrow}(t)$ through proper rescaling of $P^{\rightarrow}(t)$ using (17) (Sec. S5). To this end we need to calculate $L(t)$ in terms of $P^{\rightarrow}(t)$ from

$$\int_0^T P^{\rightarrow}(t_1|t_2)L(t_2)dt_2 = M(t_1). \quad (19)$$

(19) is the Fredholm integral equation of the first kind (Wazwaz, 2011), which lacks analytical solution. We can, however, consider a discretized version, obtaining the rescaling formula for $P^{\leftarrow}(t)$:

$$P^{\leftarrow}(t_2|t_1) = \frac{P^{\rightarrow}(t_1|t_2)[(\mathbf{P}^{\rightarrow})^{-1}\mathbf{M}](t_2)}{\mathbf{M}(t_1)} \quad (20)$$

where

$$\begin{aligned} \mathbf{P}^{\rightarrow} &= (P^{\rightarrow}(i|j)\Theta(i-j))_{1 \leq i, j \leq n} \\ \mathbf{M} &= [N(1)m(1) \ \cdots \ N(T)m(T)]^T \end{aligned}$$

To test its validity, we use (20) to numerically calculate the rescaled P^{\rightarrow} , and compare it with real P^{\leftarrow} . More specifically, we take papers published in 1980, and measure all prospective distributions for every year between 1950 to 1980, i.e. $P^{\rightarrow}(t_1|t_2)$ for t_2 ranging from 1950 to 1980. We then plug them into (20) to calculate the corresponding retrospective distributions. The obtained result yields excellent agreement with our empirical measurement (Fig. 5B).

The remarkable agreement documented in Figure 5AB raises an interesting question: does our rescaling method have predictive power as well? To test this, we predict the prospective distribution in year 1990 by using retrospective distributions and the number of new citations prior to 1990. That is, we estimate values of $M(t)$ and $P^{\leftarrow}(t|t_0)$ for $t = 1990$ and t_0 in [1990, 2010] in advance, seeking to obtain $P^{\rightarrow}(t|t_0)$ following (18).

Figure 2A indicate that $M(t)$ is the product of $N(t)$, and $m(t)$. $P^{\leftarrow}(t|t_0)$, on the other hand, follows a lognormal distribution with exponential cutoff, whose parameters in (5) can be time-dependent. To test the feasibility of our approach, we impose a minimum set of assumptions to measure $M(t)$ and $P^{\leftarrow}(t|t_0)$, to check the degree to which the predicted $P^{\rightarrow}(t|t_0)$ matches the real distributions. More specifically, for $M(t)$, we use an exponential function $\hat{M}(t)$ to fit the growth of references. We use a lognormal distribution with exponential cutoff to fit $P^{\leftarrow}(t|t_0)$.

To test the validity of our rescaling formula, we tried three models to predict P^{\rightarrow} (Sec. S6.1). Model 1: we relax M and predict P^{\rightarrow} based on simulated \hat{M} and real P^{\leftarrow} . Next, we build model 2 by relaxing P^{\leftarrow} and hence using real M and simulated \hat{P}^{\leftarrow} . Model 3 uses both simulated \hat{M} and \hat{P}^{\leftarrow} as approximations. Comparing with real data, we witness strong predictive power from all three models (Fig. 5C).

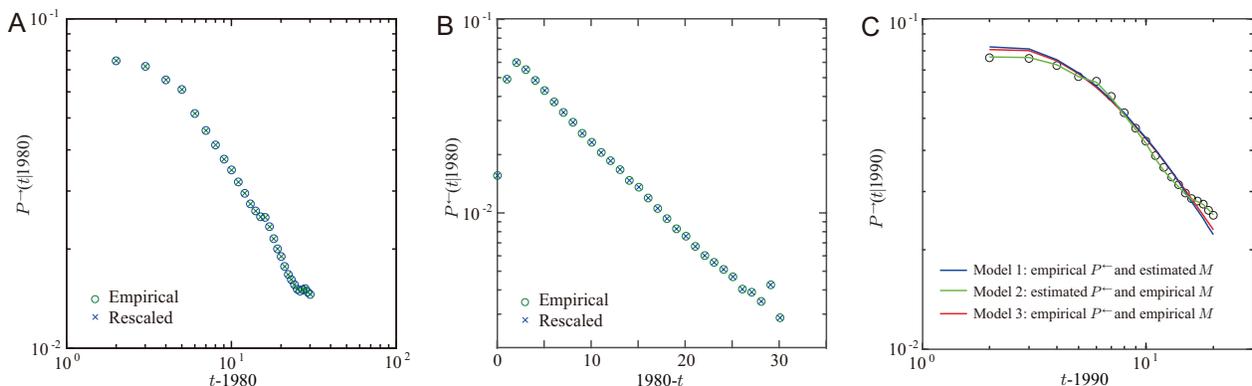


Figure 5: **Validations of (18) and (20).** **A** Using $P^{\leftarrow}(1980|t)$ data to calculate $P^{\rightarrow}(t|1980)$ through rescaling formula (18). **B** Using $P^{\rightarrow}(1980|t)$ data to calculate $P^{\leftarrow}(t|1980)$ through rescaling formula (20). **C** We use history data to estimate $M(t)$ and $P^{\leftarrow}(1990|t)$ and predict $P^{\rightarrow}(t|1990)$ through rescaling formula (18). Model 1 is based on estimated M and empirical P^{\leftarrow} , Model 2 is based on empirical M and estimated P^{\leftarrow} , while Model 3 is based on estimated M and estimated P^{\leftarrow} . All three models show great predictive power.

Our rescaling formula is also consistent with the derived functional form reported in preceding sections of this paper. Indeed, provided that P^{\leftarrow} follows (5), (18) predicts that P^{\rightarrow} follows the same functional form as (6), which is consistent with our empirical measurement (Sec. S6.2).

2.5 Redner's Ansatz

How does the citation aging structure of a paper connect to its citations? To answer this question, we consider an interesting ansatz raised by Redner (Redner, 2005): there appears to exist a power law relationship between average citation age and total citations a paper receives. Redner took all papers with 500 or fewer citations published in the American Physics Society (APS) dataset, and found that the average age of all citations to a paper, i.e. $\langle a \rangle^{\rightarrow}$, grows with the number of citations its receives, c , following $\langle a \rangle^{\rightarrow} \sim c^{\alpha}$. We repeated this analysis on the WoS dataset, finding $\langle a \rangle^{\rightarrow}$ and c indeed follow a power law scaling ($\alpha \approx 0.13$) in the small to medium c regime ($c < 500$) (Fig. 6A), thus independently confirming Redner's ansatz. We also find

significant deviations in the large c regime ($c > 500$), where $\langle a \rangle^\rightarrow$ appears to increase with c at a faster rate than power law.

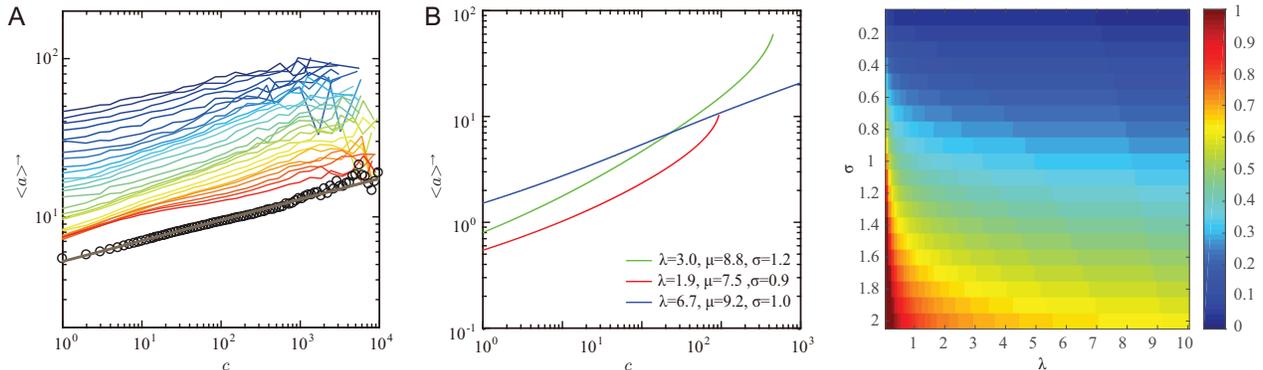


Figure 6: **Redner’s ansatz.** **A** Average age growth with citations of all papers on the WoS dataset (black circle) can be well fitted by power law (grey line, exponent is 0.13), we also see similar scaling when papers are grouped according to their ages (colored lines). Blue lines correspond to earlier years (1900-1902), while red lines correspond to later years (1978-1980). We also observe a decrease of average citation age for recent years since these papers can only be cited by more recent papers. **B** Simulation of average citation growth for different (λ, μ, σ) using Wang-Song-Barabasi model, the power law scaling is observed in a long range. **C** Varying λ and σ parameters in the WSB model can result in a wide range of scaling exponent. Color indicates the scaling exponent measured as slopes in log-log plots similar to A and B.

To test the robustness of this ansatz, we grouped papers based on their publication times and measure the average citation age for papers with different citations within each group, finding again a power law scaling between c and $\langle a \rangle^\rightarrow$ (Fig. 6A). Next, we show that this power law scaling can be explained by the inherent citation dynamics individual papers follow. To approximate the citation dynamics of a paper, we use the WSB model to generate citations of a paper over time and measure their relationship with the average citation age (Sec. S7). As we vary its (λ, μ, σ) parameters, we find that across a wide range of c values, $\langle a \rangle^\rightarrow$ can be approximated by power law relationships. The model also successfully captures the acceleration phase of the relationship in large c region (Fig. 6B). The scaling exponent of the power law relationship α depends on the specific (λ, μ, σ) parameters, hence varying (λ, μ, σ) through the parameter space allows us to reproduce the scaling behavior with a wide range of exponents (Fig. 6C).

3 Discussion

Understanding the temporal behavior of citations is of fundamental importance. In this paper, we carried out a large-scale quantitative analysis on how time affects citations. We began by reviewing and comparing the most important papers on this topic, finding significant inconsistency among prior studies in terms of both their measurement approaches and conclusions. To systematically measure citation age distributions, we used the Web of Science dataset to test different functional forms through maximum likelihood estimation. We found neither of the two prevailing approaches taken by current studies can be well captured by existing models. Moreover, we found ubiquitous upward tails characterizing P^\rightarrow , indicating that citation age distributions are affected by growth of papers and citations. To account for such impact, we further defined and calculated normalized distributions for both approaches. We found both normalized distributions can be well captured by lognormal functions, which not only allows us to fit age distributions more accurately, but also explains upward tails observed in empirical data.

Given that all the distributions reported in Table 1 are inherently quite similar (Fig. 1), here instead of focusing on visually how distributions fit the data, we aimed to look at this question through a theoretical lens. Indeed, we situated the observations within a theoretical framework connecting both distributions through mathematical relationships and network models, allowing us to rescale both distributions and reason about the most appropriate functional forms. Indeed, the main virtue of this paper lies in the theoretical

framework that helps us connect the retrospective and prospective approaches. We validated the newly derived distributions both empirically and theoretically. Through theoretical and empirical validations, we show that either distribution can be mathematically obtained from another distribution because of deeper underlying connections between the two approaches.

We then placed our results in the context of citation models, analytically calculating the citation age distributions for both approaches in three citation model systems. Compared with our empirical results, we find the Wang-Song-Barabasi model provides an accurate prediction for the functional forms of citation age distributions we obtained, offering theoretical support for our empirical findings. On the other hand, our results also provide additional evidence supporting the accuracy of the WSB model in characterizing the temporal behavior of citations. We concluded the paper with a discussion on the Redner’s ansatz that postulates a power law scaling between citations and average citation ages, finding it again can be well captured by the WSB model.

Taken together, this paper provides systematic answer to the mathematical form of citation age distributions. We believe the results reported in this paper provide a new empirical and theoretical basis for future studies on temporal analyses of citations, and will play an increasingly important role as our quantitative understanding of science deepens. Future directions include extending our analyses to examining their dependence on contextual factors, such as disciplines, countries or publishing venues, remains unknown, which could offer a more holistic understanding of the time dimension of science.

4 ACKNOWLEDGMENTS

The authors wish to thank Zhen Lei, Chaoming Song, and A.-L. Barabási for helpful discussions. This work is supported by the Air Force Office of Scientific Research under award number FA9550-15-1-0162 and FA9550-17-1-0089.

References

- Acuna, D. E., Allesina, S., and Kording, K. P. (2012). Future impact: Predicting scientific success. *Nature*, 489(7415):201–202.
- Adams, J. (2013). Collaborations: The fourth age of research. *Nature*, 497(7451):557–560.
- Allen, B., Qin, J., and Lancaster, F. (1994). Persuasive communities: A longitudinal analysis of references in the philosophical transactions of the royal society, 1665-1990. *Social Studies of Science*, 24(2):279–310.
- Avramescu, A. (1979). Actuality and obsolescence of scientific literature. *Journal of the American Society for Information Science*, 30(5):296–303.
- Barabási, A.-L., Song, C., and Wang, D. (2012). Publishing: Handful of papers dominates citation. *Nature*, 491(7422):40–40.
- Bianconi, G. and Barabási, A.-L. (2001). Competition and multiscaling in evolving networks. *EPL (Europhysics Letters)*, 54(4):436.
- Börner, K., Maru, J. T., and Goldstone, R. L. (2004). The simultaneous evolution of author and paper networks. *Proceedings of the National Academy of Sciences*, 101(suppl 1):5266–5273.
- Bornmann, L. and Daniel, H.-D. (2008). What do citation counts measure? a review of studies on citing behavior. *Journal of Documentation*, 64(1):45–80.
- Bouabid, H. (2011). Revisiting citation aging: a model for citation distribution and life-cycle prediction. *Scientometrics*, 88(1):199–211.
- Bromham, L., Dinnage, R., and Hua, X. (2016). Interdisciplinary research has consistently lower funding success. *Nature*, 534(7609):684–687.

- Burrell, Q. L. (2002). Modelling citation age data: Simple graphical methods from reliability theory. *Scientometrics*, 55(2):273–285.
- Burton, R. E. and Kebler, R. (1960). The “half-life” of some scientific and technical literatures. *American documentation*, 11(1):18–22.
- Clauset, A., Arbesman, S., and Larremore, D. B. (2015). Systematic inequality and hierarchy in faculty hiring networks. *Science Advances*, 1(1):e1400005.
- de Solla Price, D. J. (1965). Networks of scientific papers. *Science*, 149(3683):510–515.
- Deville, P., Wang, D., Sinatra, R., Song, C., Blondel, V. D., and Barabási, A.-L. (2014). Career on the move: Geography, stratification, and scientific impact. *Scientific reports*, 4.
- Duch, J., Zeng, X. H. T., Sales-Pardo, M., Radicchi, F., Otis, S., Woodruff, T. K., and Amaral, L. A. N. (2012). The possible role of resource requirements and academic career-choice risk on gender differences in publication rate and impact. *PloS one*, 7(12):e51332.
- Egghe, L. and Rao, I. K. R. (1992). Citation age data and the obsolescence function: Fits and explanations. *Information Processing & Management*, 28(2):201–217.
- Eom, Y.-H. and Fortunato, S. (2011). Characterizing and modeling citation dynamics. *PloS one*, 6(9):e24926.
- Evans, J. A. (2008). Electronic publication and the narrowing of science and scholarship. *science*, 321(5887):395–399.
- Evans, J. A. and Foster, J. G. (2011). Metaknowledge. *Science*, 331(6018):721–725.
- Evans, J. A. and Reimer, J. (2009). Open access and global participation in science. *Science*, 323(5917):1025–1025.
- Garfield, E. and Sher, I. H. (1963). New factors in the evaluation of scientific literature through citation indexing. *American Documentation*, 14(3):195–201.
- Glänzel, W. (2004). Towards a model for diachronous and synchronous citation analyses. *Scientometrics*, 60(3):511–522.
- Golosovsky, M. and Solomon, S. (2014). Uncovering the dynamics of citations of scientific papers. *arXiv preprint arXiv:1410.0343*.
- Guevara, M. R., Hartmann, D., Aristarán, M., Mendoza, M., and Hidalgo, C. A. (2016). The research space: using career paths to predict the evolution of the research output of individuals, institutions, and nations. *Scientometrics*, 109(3):1695–1709.
- Guimera, R., Uzzi, B., Spiro, J., and Amaral, L. A. N. (2005). Team assembly mechanisms determine collaboration network structure and team performance. *Science*, 308(5722):697–702.
- Gupta, B. (1997). Analysis of distribution of the age of citations in theoretical population genetics. *Scientometrics*, 40(1):139–162.
- Hirsch, J. E. (2005). An index to quantify an individual’s scientific research output. *Proceedings of the National academy of Sciences of the United States of America*, 102(46):16569–16572.
- Jones, B. F. (2011). As science evolves, how can science policy? In *Innovation Policy and the Economy, Volume 11*, pages 103–131. University of Chicago Press.
- Jones, B. F., Wuchty, S., and Uzzi, B. (2008). Multi-university research teams: shifting impact, geography, and stratification in science. *science*, 322(5905):1259–1262.
- Ke, Q., Ferrara, E., Radicchi, F., and Flammini, A. (2015). Defining and identifying sleeping beauties in science. *Proceedings of the National Academy of Sciences*, 112(24):7426–7431.

- Kong, J. S., Sarshar, N., and Roychowdhury, V. P. (2008). Experience versus talent shapes the structure of the web. *Proceedings of the National Academy of Sciences*, 105(37):13724–13729.
- Krapivsky, P. and Redner, S. (2005). Network growth by copying. *Physical Review E*, 71(3):036118.
- Krauze, T. K. and Hillinger, C. (1971). Citations, references and the growth of scientific literature: A model of dynamic interaction. *Journal of the American Society for Information Science*, 22(5):333–336.
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago press.
- Lane, J. (2010). Let’s make science metrics more scientific. *Nature*, 464(7288):488–489.
- Lane, J. and Bertuzzi, S. (2011). Measuring the results of science investments. *Science*, 331(6018):678–680.
- Larivière, V., Archambault, É., and Gingras, Y. (2008). Long-term variations in the aging of scientific literature: From exponential growth to steady-state science (1900–2004). *Journal of the American Society for Information Science and technology*, 59(2):288–296.
- Li, D. and Agha, L. (2015). Big names or big ideas: Do peer-review panels select the best science proposals? *Science*, 348(6233):434–438.
- Ma, A., Mondragón, R. J., and Latora, V. (2015). Anatomy of funded research in science. *Proceedings of the National Academy of Sciences*, 112(48):14760–14765.
- MacRae Jr, D. (1969). Growth and decay curves in scientific citations. *American Sociological Review*, pages 631–635.
- Margolis, J. (1967). Citation indexing and evaluation of scientific papers. *Science*, 155(3767):1213–1219.
- Matricciani, E. (1991). The probability distribution of the age of references in engineering papers. *Professional Communication, IEEE Transactions on*, 34(1):7–12.
- Mazloumian, A., Eom, Y.-H., Helbing, D., Lozano, S., and Fortunato, S. (2011). How citation boosts promote scientific paradigm shifts and nobel prizes. *PloS one*, 6(5):e18975.
- Michels, C. and Schmoch, U. (2012). The growth of science and database coverage. *Scientometrics*, 93(3):831–846.
- Moreira, J. A., Zeng, X. H. T., and Amaral, L. A. N. (2015). The distribution of the asymptotic number of citations to sets of publications by a researcher or from an academic department are consistent with a discrete lognormal model. *PloS one*, 10(11):e0143108.
- Motylev, V. (1989). The main problems of studying literature aging. *Scientometrics*, 15(1-2):97–109.
- Mukherjee, S., Romero, D., Jones, B., and Uzzi, B. (2017). The nearly universal link between the age of past knowledge and tomorrow’s breakthroughs in science and technology. *Science Advances*.
- Nakamoto, H. (1988). Synchronous and diachronous citation distribution.
- Newman, M. (2009). The first-mover advantage in scientific publication. *EPL (Europhysics Letters)*, 86(6):68001.
- Newman, M. (2014). Prediction of highly cited papers. *EPL (Europhysics Letters)*, 105(2):28002.
- Orosz, K., Farkas, I. J., and Pollner, P. (2016). Quantifying the changing role of past publications. *Scientometrics*, pages 1–25.
- Pan, R. K., Petersen, A. M., Pammolli, F., and Fortunato, S. (2016). The memory of science: Inflation, myopia, and the knowledge network. *arXiv preprint arXiv:1607.05606*.
- Parolo, B., Della, P., Pan, R. K., Ghosh, R., Huberman, B. A., Kaski, K., and Fortunato, S. (2015). Attention decay in science. *Available at SSRN 2575225*.

- Petersen, A. M., Fortunato, S., Pan, R. K., Kaski, K., Penner, O., Rungi, A., Riccaboni, M., Stanley, H. E., and Pammolli, F. (2014). Reputation and impact in academic careers. *Proceedings of the National Academy of Sciences*, 111(43):15316–15321.
- Pollman, T. (2000). Forgetting and the ageing of scientific publications. *Scientometrics*, 47(1):43–54.
- Price, D. d. S. (1976). A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science*, 27(5):292–306.
- Price de Solla, D. (1963). Little science, big science. *New York, NY: Columbia UP*.
- Radicchi, F., Fortunato, S., and Castellano, C. (2008). Universality of citation distributions: Toward an objective measure of scientific impact. *Proceedings of the National Academy of Sciences*, 105(45):17268–17272.
- Radicchi, F., Fortunato, S., Markines, B., and Vespignani, A. (2009). Diffusion of scientific credits and the ranking of scientists. *Physical Review E*, 80(5):056103.
- Redner, S. (1998). How popular is your paper? an empirical study of the citation distribution. *The European Physical Journal B-Condensed Matter and Complex Systems*, 4(2):131–134.
- Redner, S. (2005). Citation statistics from 110 years of physical review. *Physics Today*, 58:49.
- Sanyal, S. (2006). Effect of citation patterns on network structure. *arXiv preprint physics/0611139*.
- Shen, H.-W. and Barabási, A.-L. (2014). Collective credit allocation in science. *Proceedings of the National Academy of Sciences*, 111(34):12325–12330.
- Simkin, M. V. and Roychowdhury, V. P. (2007). A mathematical theory of citing. *Journal of the American Society for Information Science and Technology*, 58(11):1661–1673.
- Sinatra, R., Deville, P., Szell, M., Wang, D., and Barabási, A.-L. (2015). A century of physics. *Nature Physics*, 11(10):791–796.
- Sinatra, R., Wang, D., Deville, P., Song, C., and Barabási, A.-L. (2016). Quantifying the evolution of individual scientific impact. *Science*, 354(6312):aaf5239.
- Stinson, E. R. and Lancaster, F. (1987). Synchronous versus diachronous methods in the measurement of obsolescence by citation studies. *Journal of information science*, 13(2):65–74.
- Stringer, M. J., Sales-Pardo, M., and Amaral, L. A. N. (2008). Effectiveness of journal ranking schemes as a tool for locating information. *PLoS One*, 3(2):e1683.
- Stringer, M. J., Sales-Pardo, M., and Amaral, L. A. N. (2010). Statistical validation of a global model for the distribution of the ultimate number of citations accrued by papers published in a scientific journal. *Journal of the American Society for Information Science and Technology*, 61(7):1377–1385.
- Šubelj, L. and Fiala, D. (2016). Publication boost in web of science journals and its effect on citation distributions. *arXiv preprint arXiv:1601.01203*.
- Szántó-Várnagy, Á., Pollner, P., Vicsek, T., and Farkas, I. J. (2014). Scientometrics: Untangling the topics. *National Science Review*, page nwu027.
- Szell, M. and Sinatra, R. (2015). Research funding goes to rich clubs. *Proceedings of the National Academy of Sciences*, 112(48):14749–14750.
- Tsay, M.-Y. (1999). Library journal use and citation age in medical science. *Journal of documentation*, 55(5):543–555.
- Ucar, I., López-Fernandino, F., Rodríguez-Ulibarri, P., Sesma-Sanchez, L., Urrea-Micó, V., and Sevilla, J. (2014). Growth in the number of references in engineering journal papers during the 1972–2013 period. *Scientometrics*, 98(3):1855–1864.

- Uzzi, B., Mukherjee, S., Stringer, M., and Jones, B. (2013). Atypical combinations and scientific impact. *Science*, 342(6157):468–472.
- Van Noorden, R., Maher, B., Nuzzo, R., et al. (2014). The top 100 papers. *Nature*, 514(7524):550–553.
- Van Raan, A. F. (2004). Sleeping beauties in science. *Scientometrics*, 59(3):467–472.
- Vazquez, A. (2001). Statistics of citation networks. *arXiv preprint cond-mat/0105031*.
- Verstak, A., Acharya, A., Suzuki, H., Henderson, S., Iakhiaev, M., Lin, C. C. Y., and Shetty, N. (2014). On the shoulders of giants: The growing impact of older articles. *arXiv preprint arXiv:1411.0275*.
- Vinkler, P. (2010). *The evaluation of research by scientometric indicators*. Elsevier.
- Wang, D., Song, C., and Barabási, A.-L. (2013). Quantifying long-term scientific impact. *Science*, 342(6154):127–132.
- Wazwaz, A.-M. (2011). *Linear and nonlinear integral equations: methods and applications*. Springer Science & Business Media.