Over-Optimization of Academic Publishing Metrics: Observing Goodhart's Law in Action

Michael Fire[1,2]* and Carlos Guestrin[1]

Abstract

The academic publishing world is changing significantly, with ever-growing numbers of publications each year and shifting publishing patterns. However, the metrics used to measure academic success, such as the number of publications, citation number, and impact factor, have not changed for decades. Moreover, recent studies indicate that these metrics have become targets and follow Goodhart's Law, according to which "when a measure becomes a target, it ceases to be a good measure." In this study, we analyzed over 120 million papers to examine how the academic publishing world has evolved over the last century. Our study shows that the validity of citation-based measures is being compromised and their usefulness is lessening. In particular, the number of publications has ceased to be a good metric as a result of longer author lists, shorter papers, and surging publication numbers. Citation-based metrics, such citation number and h-index, are likewise affected by the flood of papers, self-citations, and lengthy reference lists. Measures such as a journal's impact factor have also ceased to be good metrics due to the soaring numbers of papers that are published in top journals, particularly from the same pool of authors. Moreover, by analyzing properties of over 2600 research fields, we observed that citation-based metrics are not beneficial for comparing researchers in different fields, or even in the same department. Academic publishing has changed considerably; now we need to reconsider how we measure success.

Multimedia Links

► Interactive Data Visualization   ► Code Tutorials   ► Fields-of-Study Features Table

Keywords

Science of Science; Scientometrics; Goodhart's Law; Data Science; Big Data; Academic Publishing Metrics

[1] *Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA*
[2] *The eScience Institute, University of Washington, Seattle, WA*
*Corresponding author: fire@cs.washington.edu

Contents

Introduction

In the last century, the academic publishing world has changed drastically in volume and velocity. The volume of papers has increased sharply from about 174,000 papers published in 1950 to over 7 million papers published 2014 (see Figure 2). Furthermore, the speed in which researchers can share and publish their studies has increased significantly. Today's researchers can publish not only in an ever-growing number of traditional venues, such as conferences and journals, but also in electronic preprint repositories and in mega-journals that provide rapid publication times [1].

Figure 1: Measuring success in academic publishing

A comprehensive analysis of more than 120 million academic papers by 35 million authors published since the early 19th century reveals that the world of academic publishing is changing significantly across 19 major research domains and 2,600 research fields. Here we examine trends from 1990 to 2010, with a closer look at increasingly specialized fields within biology and genetics.

The number of papers has surged in recent decades, particularly in non-English languages. Titles have increased in length, as has the use of emphatic punctuation such as question marks and exclamation points.

The number of authors has risen, as has collaboration: Researchers who started their careers in the 1950s on average had a few coauthors over 25 years, but those who started in the 1990s average more than 60 over a smiliar period.

The number of ranked journals has exploded, surpassing 21,000 active journals in 2016 and adding several hundred new ones each year..
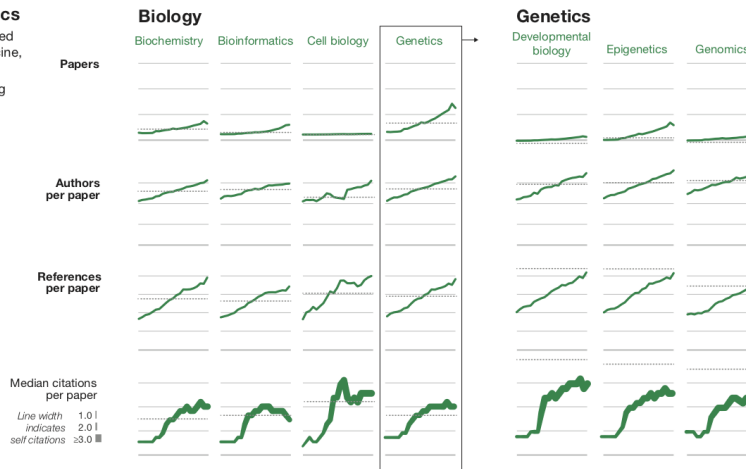
The number of references is growing across disciplines, even as papers' lengths have decreased.

Self-citations have increased along with citations in general, but papers' influences are not long lasting: 72% of papers did not have a single citation 5 years after publication.

ENGINEERING — Engineering

HUMANITIES — Art, History, Philosophy

LIFE SCIENCES — Medicine, Enviro. sci., Biology

MATHEMATICS — Computer sci., Mathematics

PHYSICAL SCIENCES — Chemistry, Geology, Materials sci., Physics

SOCIAL SCIENCES — Business, Economics, Geography, Political sci., Psychology, Sociology

Papers
Average
Authors per paper
References per paper
Median citations per paper

Line width 1.0 | indicates 2.0 | self citations ≥3.0 ▓

Trends within biology and genetics

The number of biology-related papers published per year has surged, mirroring trends in medicine, physics, computer science and mathematics.

But this rate of growth is not consistent among disciplines within biology, where research in genetics is growing quickly but cell biology continues at a relatively steady state.

Within genetics disciplines, epigenetics publishing is on the rise.

Biology
Biochemistry, Bioinformatics, Cell biology, Genetics

Genetics
Developmental biology, Epigenetics, Genomics

Papers
Authors per paper
References per paper
Median citations per paper

Line width 1.0 | indicates 2.0 | self citations ≥3.0 ▓

CHARTS: SEAN McNAUGHTON

Along with the exponential increase in the quantity of published papers, the number of ranked scientific journals has increased to over 20,000 journals (see Figures 15 and S17), and the number of published researchers has soared (see Figure S14). As part of this escalation, metrics such as the number of papers, number of citations, impact factor, h-index, and altmetrics are being used to compare the impact of papers, researchers, journals, and universities [2, 3, 4, 5]. Using quantitative metrics to rank researchers contributes to a hypercompetitive research environment, which is changing academic culture - and not in a positive direction [6].

Studies suggests that publication patterns have changed as a result of Goodhart's Law, according to which, "When a measure becomes a target, it ceases to be a good measure" [6, 7]. Goodhart's Law, and its closely related Campbell's Law [8], influence many systems in our everyday life, including economical [9], educational [8], and other decision-making systems [10]. As an example, Goodhart's Law can be found in the NYPD's manipulation of crime reports (the "measure") in order to improve crime statistics (the "target") [11]. Another example is found in the educational system, revealing that when "test scores become the goal of the teaching process, they both lose their value as indicators of educational status and distort the educational process in undesirable ways" [8].

Recent studies indicate that when measures become targets in academic publishing, the effectiveness of the measures can be compromised, and unwelcome and unethical behaviors may develop, such as salami publications [12], ghost authorships [13], p-hacking [14], metrics manipulation [15], and even faking of peer reviews [16].

If the influence of Goodhart's Law on academia is indeed significant, then it should be possible to observe that academic entities, such as researchers and journals, will over-optimize their own measures to achieve a desired target. Similar to the consequences of making test scores a target, chasing after certain measures in the academic publishing world to desperately win the battle of "*impact or perish*" [7] can have undesirable effects.

In this study, our main goal was to utilize new advances in data science tools to perform an in-depth and precise bottom-up analysis of academic publishing over the decades. Our comprehensive analysis ranged from micro to macro levels as we studied individual researchers' behaviors as well as behavioral changes within large research domains. Additionally, we wanted to uncover how and if Goodhart's Law has changed academic publishing.

To achieve our research goals, we developed an open-source code framework to analyze data from several large-scale datasets containing over 120 million publications, with 528 million references and 35 million authors,[1] since the beginning of the 19th century. This provided a precise and full picture of how the academic publishing world has evolved. In our analysis, we uncovered a wide variety of underlying

[1]The number of authors was estimated according to the unique full names in the Microsoft Academic Graph dataset (see Section 4.2)

changes in academia at different levels (see Figure 1):

- *Papers* – We observed that on average papers became shorter, yet other features, such as titles, abstracts, and author lists, became longer (see Figures 4, 9, S2, S5, S6, and S7). Furthermore, both the number of references and the number of self-citations considerably increased (see Figures S9 and 8). Moreover, the total number of papers without any citations at all increased sharply over time (see Figure 10).

- *Authors* – We noticed a sharp increase in the number of new authors. Moreover, we observed a negative relation between researchers' career ages and the number of publications, i.e., early career researchers tended to publish more quickly than those later in their careers (see Figure 12). Additionally, the average number of coauthors per author considerably increased over time (see Figure 13).

- *Journals* – We observed a drastic increase in the number of ranked journals, with several hundreds of new ranked journals published each year (see Figures 15, S17, and S18). In addition, we observed that journal ranking changed significantly as average citations per document and SCImago Journal Rank (SJR) measures increased, while the h-index measure decreased over time. Moreover, the percentage of papers with returning authors increased sharply in recent years (see Figure 19). For example, in Cell journal about 80% of all papers published in 2016 included at least one author who had published in the journal before, while in 1980 this rate stood at less than 40%.

- *Fields of Research* – We analyzed the properties of 19 major research domains, such as art, biology, and computer science, as well as the properties of 2600 subdomains, such photography, genetic diversity, and computer animation. Our analysis revealed that different domains had widely ranging properties (see Figures 21 and S24), even within subdomains. For example, different subdomains of biology, and even subdomains of genetics, had surprisingly different average numbers of citations (see Figures S33, and S34).

These observations support the hypothesis that commonly used measures, such as the author's number of papers, h-index, and citation number, have become targets. Researchers can increase their average number of papers by writing shorter papers with more coauthors. The h-index can be boosted by increasing the number of papers and self-citations. A higher citation number is attained with longer titles and abstracts, as well as more references (see Section 4.1 and Figure S13).

It is time to consider how we judge academic papers. Citation-based measures have been the standard for decades, but these measures are far from perfect. In fact, our study shows that the validity of citation-based measures is being

compromised and their usefulness is lessening. Goodhart's Law is in action in the academic publishing world.

The remainder of the paper is organized as follows: In Section 2, we provide an overview of related studies. In Section 3, we describe the datasets, methods, algorithms, and experiments used throughout this study. Next, in Section 4, we present the results of our study. Afterwards, in Section 5, we discuss the obtained results. Lastly, in Section 6, we present our conclusions.

2. Related Work

This research is a large-scale scientometrics study (also referred to as the "science of science" [17]). Scientometrics is the study of quantitative features and characteristics of scientific research. In this section, we give a short overview of the relevant scientometric papers to this study. We present studies that analyze changes in academic publications in recent years (see Section 2.1), and we provide an overview of common metrics that measure the impact of published papers (see Section 2.2).

2.1 Changes in Publication Trends

One prevalent and increasing trend is to publish papers in preprint repositories, such as arXiv, bioRxiv, and SSRN. For example, the use of arXiv surged from 4,275 papers in September 2006 to 10,570 papers in September 2017 [18]. Another common trend is to publish papers in mega-journals, such as PLOS ONE and Nature's Scientific Reports. Mega-journals are a new type of scientific journal that publishes peer-reviewed, open-access articles, where the articles have been reviewed for scientific trustworthiness, but not for scientific merit. Mega-journals accelerate review and publication times to 3-5 months and usually have high acceptance rates of over 50% [1]. In the first quarter of 2017, over 11,000 papers were published in PLOS ONE and Scientific Reports [19].

Another observable trend is that more and more papers are written by hundreds of authors. The recent Laser Interferometer Gravitational-Wave Observatory (LIGO) paper [20] has over 1000 authors [21]. Robert Aboukhalil measured this trend [22] and discovered that the average number of authors of academic papers has increased sharply since the beginning of the 20th century. While papers' average number of authors has gone up over time, not all the authors have significantly contributed to the paper. Honorary and ghost authors are prevalent. Wislar et al. found such evidence in biomedical journals [23], and similar findings were observed by Kennedy et al. [24] and by Vera-Badillo et al. [25]. The Economist recently published an article titled "Why research papers have so many authors" [26].

Lewison and Hartley [27] analyzed how papers' titles have changed over time. They discovered that titles' lengths have been increasing, along with the percentage of titles containing colons. Additionally, Gwilym Lockwood observed that "articles with positively-framed titles, interesting phrasing, and no wordplay get more attention online" [28].

Additionally, many studies have focused on how publication trends have changed over time, often focusing on specific geographical areas, various demographic characteristics, specific research domains, or specific journals. For example, Gálvez et al. [29] utilized the Science Citation Index to understand publication patterns in the developing world. Jagsi et al. [30] studied the gender gap in authorship of academic medical literature over 35 years. They discovered that the percentage of first and last authors who were women increased from 5.9% and 3.7% in 1970 to 29.3% and 19.3%, respectively, in 2004. Johnson et al. [31] studied publication trends in top-tier journals of higher education. Peter Aldhous analyzed publications in the National Academy of Sciences (PNAS) journal, to consider the influence of an "old boys' club" mentality [32].

Our study is greatly influenced by a recent study by Edwards and Roy [6], who observed that academia has become a hypercompetitive environment that can lead to unethical behaviors. The driving force behind such behaviors is to manipulate the metrics that measure the research's impact solely to increase the quantitative measures (and hence the status) of the research.

2.2 Success Metrics and Citation Trends

Over the years, various metrics have been proposed to measure papers, journal importance, and authors' impact. One of the most straightforward and commonly utilized measure is to simply count the researcher's number of publications. Another common metric is the citation number, either of a particular paper or the total citations received by all the author's papers. However, not all citations are equal [33]. Moreover, different research fields have different citation metrics, and therefore comparing them creates a problem: "The purpose of comparing citation records is to discriminate between scientists" [34].

One of the best-known and most-used measures to evaluate journals' importance is the impact factor, devised over 60 years ago by Eugene Garfield [4]. The impact factor measures the frequency in which an average article in a journal has been cited in a specific year. Over time, the measure has been used to "evaluate institutions, scientific research, entire journals, and individual articles" [35]. Another common metric to measure a researcher's output or a journal's impact is the h-index, which measures an author's or a journal's number of papers that have at least h citations each [3]. It has been shown that the h-index can predict academic achievements [36].

The above measures have been the standard for measuring academic publishing success. According to recent studies, and following Goodhart's Law, these metrics have now become targets, ripe for manipulation [6, 7, 37]. All types of manipulative methods are used, such as increasing the number of self-citations [15], increasing the number of publications by slicing studies into the smallest measurable quantum acceptable for publication [38], indexing false papers [39], and merging papers on Google Scholar [40]. Indeed, a recent

study by Fong and Wilhite [37], which utilized data from over 12,000 responses to a series of surveys sent to more than 110,000 scholars from eighteen different disciplines, discovered "widespread misattribution in publications and in research proposals." Fong and Wilhite's findings revealed that the majority of researchers disapprove of this type of metric manipulation, yet many feel pressured to participate; other researchers blandly state "that it is just the way the game is played" [37].

Due to many common metric shortcomings, various alternative measures have been proposed. For example, the q-index [15] and w-index [41] are alternatives to the h-index. Likewise, the SJR indicator [42] and simple citation distributions [43] are offered as alternatives to the impact factor. Senior employees at several leading science publishers called upon journals to restrain from using the impact factor and suggested replacing it with simple citation distributions [43, 44]. Similarly, the altmetric[2] was proposed as an alternative metric to the impact factor and h-index. The altmetric [45] is a generalization of article-level metrics and considers other aspects of the impact of the work, such as the number of downloads, article views, mentions in social media, and more. The altmetric measure has gained in popularity in recent years, and several large publishers have started providing this metric to their readers. Additionally, Semantic Scholar[3] offers various measures to judge papers and researchers' influence. A thorough report regarding potential uses and limitations of metrics was written by Wilsdon et al. [5]. Additionally, an overview of the changing scholarly landscape can be found in Roemer and Borchardt's study [2].

Even with their many known shortcomings [5, 34, 46, 47, 48], measures such as the impact factor, citation number, and h-index are still widely used. For example, the Journal Citation Reports publishes annual rankings based on journals' impact factors, and it continues to be widely followed.

3. Methods and Experiments

3.1 Datasets

3.1.1 The Microsoft Academic Graph (MAG) Dataset
In this study we primarily utilized the Microsoft Academic Graph (MAG) [49], which was released as part of the 2016 KDD Cup [50]. The large-scale MAG dataset contains scientific publication records of over 120 million papers, along with citation relationships among those publications as well as relationships among authors, institutions, journals, conferences, and fields of study. In addition, the MAG dataset contains every author's sequence number for each paper's authors list. Furthermore, the dataset contains field-of-study hierarchy ranking with four levels, L0 to L3, where L0 is the highest level, such as a research field of computer science, and L3 is the lowest level, such as a research field of decision tree [50].

Even though the MAG dataset contains papers that were published through 2016, we wanted to use years in which the data was the most comprehensive, so we focused our analysis on 120.7 million papers which were published through the end of 2014. Furthermore, we noted that the dataset contains many papers that are news items, response letters, comments, etc. Even though these items are important, they can affect a correct understanding of the underlying trends in scientific publications. Therefore, we focused our research on a dataset subset, which consists of over 22 million papers. This subset contains only papers which have a Digital Object Identifier (DOI) and at least 5 references. Additionally, while calculating various authors' properties, we primarily considered only the 22.4 million authors with unique author ID values in the selected papers' subset.

3.1.2 The AMiner Dataset
The AMiner open academic graph dataset [51] contains data from over 154 million papers. The dataset contains various papers' attributes, such as titles, keywords, abstracts, venues, languages, and ISSNs. In our study, we primarily utilized the AMiner dataset to analyze papers' abstracts, to estimate papers' lengths, and to compare results with those obtained using the MAG dataset in order to validate the existence of observed patterns in both datasets.

3.1.3 The SCImago Journal Rank Dataset
To better understand trends in journal publications, we used the SCImago Journal Ranking (SJR) open dataset [52].[4] This dataset contains details of over 23,000 journals with unique names between 1999 and 2016. For each journal, the SJR dataset contains the journal's SJR value, the number of published papers, the h-index, and the number of citations in each year. Additionally, the SJR dataset contains the best quartile (ranked from Q1 to Q4) of each journal. The quartile rank is typically used to compare and rank journals within a given subject category.

3.1.4 The Join Dataset
To match the MAG journal IDs with their correlated various ranking measures, such as h-index and SJR, we joined all three datasets in the following manner: First, we joined the MAG and AMiner datasets by matching unique DOI values. Then, we matched ISSN values between the MAG-AMiner joined dataset with the SJR dataset.

3.2 Code Framework
To analyze the above MAG and AMiner large-scale datasets, we developed an open source framework written in Python, which provided an easy way to query the datasets. The framework utilizes TuriCreate's SFrame dataframe objects [53] to perform big-data analysis on tens of millions of records to calculate how various properties have changed over time. For example, we used SFrame objects to analyze how the average number of authors and title lengths evolved. However, while

---

[2]/https://www.altmetric.com/
[3]/https://www.semanticscholar.org

---

[4]/https://www.scimagojr.com/journalrank.php

SFrame is exceptionally useful for calculating various statistics using all-papers features, it is less convenient and less computationally cost effective for performing more complicated queries, such as calculating the average age of the last authors in a certain journal in a specific year.

To perform more complex calculations, we loaded the datasets into the MongoDB database.[5] Next, we developed a code framework that easily let us obtain information on papers, authors, paper collections, venues, and research fields. The framework supports calculating complex features of the above object in a straightforward manner. For example, with only a few and relative simple lines of Python code, we were able to calculate the average number of coauthors per author in a specific year for authors who started their career in a specific decade. An overview of our code framework is presented in Figure S1.

To make our framework accessible to other researchers and to make this study completely reproducible, we have written Jupyter Notebook tutorials which demonstrate how the SFrame and MongoDB collections were constructed from the MAG, AMiner, and SJR datasets.

3.3 Analysis of Publication Trends

We used our developed code framework to explore how papers, authors, journals, and research fields have evolved over time. In the following subsections, we describe the specific calculations that were performed. Moreover, our Supplementary Materials section includes the precise code implementations which were used to obtain most of our results and to create the figures presented throughout this study.

3.3.1 Paper Trends

To explore how the quantity and structure of academic papers have changed over time, we performed the following: First, we calculated how many papers were published in the MAG dataset every year. Then, we detected the language of each paper's title and calculated the number of papers in each language. Next, we calculated the following paper features over time:

- Average number of words in titles and average number of characters per word (for papers with English titles)

- Percentage of titles that used question or exclamation marks (for papers with English titles)

- Average number of authors

- Percentage of papers in which authors appear in alphabetical order

- Average number of words in abstracts

- Average number of keywords

- Average number of references

- Length of papers

In addition, we utilized the papers with existing field-of-research values, matching the papers to their corresponding fields in order to identify each paper's top level (L0) research field. Using the top-level data, we were able to estimate the number of multidisciplinary papers that had more than one L0 research field. Afterwards, we calculated the percentage and total number of papers with no citations after 5 years, as well as the overall percentage of papers with self-citations over time.[6] Lastly, to better understand how citation patterns have changed across generations, we calculated the citation distributions after 10 years for each decade between 1950 and 2000.

Additionally, we selected all the papers in the Join dataset that had valid features[7] and were published between 1990 and 2009. Using the selected papers, we calculated the Spearman correlations among the title lengths, author numbers, reference numbers, overall lengths, and number of citations after 5 years. The results of the above described calculations are presented in Section 4.1. Moreover, the code implementation is provided in the "Part III - A: Analyzing Changing Trends in Academia – Paper Trends" Jupyter Notebook (see Section 7).

3.3.2 Author Trends

To study how authors' behaviors and characteristics have changed, we performed the following: First, we calculated how the number of new authors has changed over time. Second, for all authors who published their first paper after 1950, we divided the authors into groups according to each author's academic birth decade, i.e., the decade in which an author published his or her first paper. Next, for each group of authors with the same academic birth decade, we analyzed the following features:

- Average number of papers the authors in each group published n years after they began their careers, for $\forall n \in [0, 30]$. We performed these group calculations taking into account all papers, as well as only papers with at least 5 references.

- Average number of conference and journal papers each group published n years after they began their careers, for $\forall n \in [0, 30]$

- Average number of coauthors each group had n years after they began their careers, for $\forall n \in [0, 30]$

- Authors' median sequence number each group had n years after they began their careers, for $\forall n \in [0, 60]$. Additionally, we calculated the average percentage of times the authors in each group were first authors.

---

[5]/http://www.mongodb.com

[6]We define paper A as self-citing paper B if at least one of the authors of A is also an author of B.

[7]We selected only papers having English titles and abstracts, existing author lists, references, and valid lengths. Additionally, we checked if the paper's title contained question or exclamation marks.

The results of the above described calculations are presented in Section 4.2. Moreover, the code implementation is provided in the "Part III - B: Analyzing Changing Trends in Academia - Author Trends" Jupyter Notebook (see Section 7).

3.3.3 Journal Trends

To investigate how journal publication trends have changed over time, we used the SJR dataset to calculate the following features between 1999 and 2016:

- Number of journals with unique journal IDs that were active in each year

- Number of new journals that were published each year

- Average and maximal number of papers in each journal

Additionally, we utilized the SJR dataset to calculate how the journals' best quartile, average h-index, average SJR, and average citation number ($\frac{CitationNumber}{DocumentsNumber}$ (2 years)) metrics changed between 1999 and 2016.

Furthermore, we selected the 40 journals with the highest SJR values in 2016 and matched them to their corresponding journal IDs in the MAG dataset by matching each journal's ISSN and exact name in the MAG-AMiner joined dataset.[8] Then, for the matching journal IDs, we calculated the following features over time, for all papers that were published in the selected top journals:

- First and last authors' average career age

- Percentage of papers in which the first author had previously published in the one of the top journals

- Percentage of papers in which the last author had previously published in the one of the top journals

The results of the above described calculations are presented in Section 4.3. Moreover, the code implementation is provided in the "Part III - C: Analyzing Changing Trends in Academia - Journal Trends" Jupyter Notebook (see Section 7).

Additionally, for over 8,400 journals with at least 100 published papers with 5 references, we calculated the following features over time:

- Number of papers

- Number of authors

- Top keywords in a specific year

- First/last/all authors average or median academic age

- Average length of papers

- Percentage of returning first/last/all authors, i.e., those who had published at least one prior paper in the journal

We developed a website with an interactive interface, which visualizes how the above features changed for each journal (see Section 7).

3.3.4 Field-of-Research Trends

We utilized the MAG dataset field-of-study values and the hierarchical relationship between various fields to match papers to their research fields in various levels (L0-L3). Then, for each field of study in its highest hierarchical level (L0), we calculated the following features over time: number of papers, number of authors, number of references, and average number of citations after five years. Next, we focused on the field of biology, which is in the L0 level. For all the L1 subfields of biology, we repeated the same feature calculations as in the previous step. Afterwards, we focused on genetics. For all the L2 subfields of genetics, we repeated the same feature calculations as in the previous step.

Additionally, to better understand the differences in citation patterns of various fields of research, we performed the following: For each field of study with at least 100 papers published in 2009, we calculated the following features using only papers that were published in 2009 and had at least 5 references:

- Number of papers

- Number of authors

- Median and average number of citations after 5 years

- Maximal number of citations after 5 years

The full features of over 2600 L3 fields of study are presented in Table 1.

The results of the above described calculations are presented in Section 4.4. Moreover, the code implementation is provided in the "Part III - D: Analyzing Changing Trends in Academia - Research Fields" Jupyter Notebook (see Section 7).

4. Results

In the following subsections, we present all the results for the experiments which were described in Section 3.3. Additional results are presented in the Supplementary Materials.

4.1 Results of Paper Trends

In recent years there has been a surge in the number of published academic papers, with over 7 million new papers each year and over 1.8 million papers with at least 5 references (see Figure 2).[9] Additionally, by analyzing the language of the papers' titles, we observed a growth in papers with non-English titles (see Figure 3).

As described in Section 3.3.1, we analyzed how various properties of academic papers have changed over time to better understand how papers' structures have evolved. In this analysis, we discovered that papers' titles became longer, from an average of 8.71 words in 1900 to an average of 11.83 words in 2014 (see Figure 4). Moreover, the average number

---

[8]The top journal name was compared to the journal's name in the MAG dataset.

[9]There is a decline in the number of papers after 2014, probably due to missing papers in the MAG dataset, which was released in 2016.

Figure 2. The Number of Papers over Time. The total number of papers has surged exponentially over the years.


Figure 3. Papers with Top-9 Non-English Titles. Increasingly, more papers have non-English titles.

of characters per word increased from 5.95 characters per average title word in 1900 to 6.6 characters per average title word in 2014 (see Figure 4). Additionally, we observed that in recent years the percentage of papers with question or exclamation marks increased sharply, from less than 1% of all papers in 1950 to over 3% of all papers in 2013 (see Figure S2). Furthermore, the usage of interrobangs (represented by ?! or !?) also increased sharply, from 0.0005% in 1950 to 0.0037% in 2013 (see Figure S2).

We explored how the number and order of the authors list has changed over time. The number of authors for papers with at least 5 references more than tripled over the years, from an average of 1.41 authors to an average of 4.51 authors per paper between 1900 and 2014, respectively (see Figure S3). Also, the maximal number of authors for a single paper in each year increased sharply over time, especially in recent years (see Figure S4). In fact, some recent papers actually listed over 3000 authors. Moreover, we observed that the percentage of author lists ordered alphabetically decreased in recent years, from 43.5% of all papers published in 1950 to 21% of all papers published in 2014 (see Figure S5). Furthermore, we discovered that with a higher number of authors, it is less likely that the authors list will be ordered alphabetically (see Figure 5). For example, in 2014 only about 1% of papers with six authors were ordered alphabetically.

When calculating how the abstracts of papers have changed over time, we discovered that the abstract length increased from an average of 116.3 words in 1970 to an average of 179.8 words in 2014 (see Figure S6). Moreover, with each decade since 1950, the distributions shifted to the right, showing that papers with longer abstracts of 400 and even 500 words have become more common over time (see Figure 6). Additionally, we analyzed how the number of keywords in papers has changed. We discovered that both the number of papers containing keywords increased, as well as the average number of keywords per paper (see Figure S7).

By estimating the percentage and number of multidisciplinary papers over time, we discovered an increase in the number of multidisciplinary papers until 2010, followed by a sharp decrease (see Figures 7 and S8). After performing further analysis, we believe the decline in the number of multidisciplinary papers is a result of papers with missing keywords in the MAG dataset, such as papers that were published in the PLOS ONE journal. These papers have dynamically changing keywords in the online version, but not in the offline version.

By examining how the number of references has changed over time, we observed a sharp increase in the average number of references per paper (see Figure S9). In addition, by analyzing the reference number distributions grouped by publishing decade, we can observe that higher numbers of references have become increasingly common. For example, in 1960 few papers had over 20 references, but by 2010 many papers had over 20 references, and some over 40 references (see Figure S10).

We also examined how self-citation trends have changed,

Figure 4. Average Title Length over Time. A paper's average title length increased from 8.71 words to over 11.83 words. Moreover, the average word length increased from 5.95 characters to 6.6 characters per title word.

Figure 5. Percentage of Papers with Author Lists in Alphabetical Order, Grouped by the Number of Authors. The higher the number of authors, the less likely the authors will be organized alphabetically.
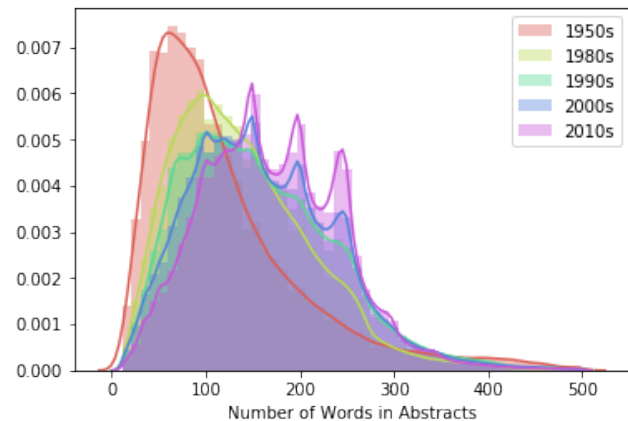

Figure 6. Distribution over Time of the Number of Words in Abstracts. Over time, papers' abstracts have tended to become longer.

and we observed that both the total number of self-citations and the percentage of papers with self-citations increased significantly (see Figure S12). Also, the average number of self-citations per paper, as well as the maximal number of self-citations in each year, increased sharply (see Figure 8). For example, about 3.67% of all papers in 1950 contained at least one self-citation, while 8.29% contained self-citations in 2014 (see Figure S12). Moreover, the maximal number of self-citations in a single paper increased sharply from 10 self-citations in a paper published in 1950 to over 250 self-citations in a paper published in 2013 (see Figure 8).

By using the AMiner dataset to analyze how papers' lengths have changed, we discovered that the average and median length of papers decreased over time (see Figure 9). The average length of a paper was 14.4, 10.1, and 8.4 pages in 1950, 1990, and 2014, respectively.

By analyzing citation patterns over time, we discovered

that the percentage of papers with no citations after 5 years decreased (see Figure S11). Nevertheless, still about 72.1% of all papers published in 2009, and 25.6% of those with at least 5 references, were without any citations after 5 years (see Figure S11). Moreover, the total number of papers without any citations increased sharply (see Figure 10).

Additionally, by analyzing the citation distributions of papers published in different decades, we discovered citation distributions changed notably over time (see Figure 11).

Lastly, using the properties of over 3.29 million papers published between 1950 and 2009, we discovered positive correlations among the papers' citation numbers after 5 years and the following features: (a) title lengths ($\tau_s = 0.1$); (b) author numbers ($\tau_s = 0.22$); (c) abstract lengths ($\tau_s = 0.26$); (d) keyword numbers ($\tau_s = 0.15$); (e) reference numbers ($\tau_s = 0.48$); (e) paper lengths ($\tau_s = 0.13$); and (f) use of question or
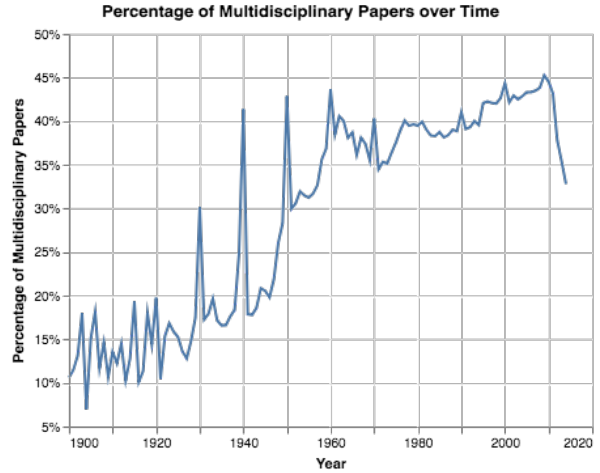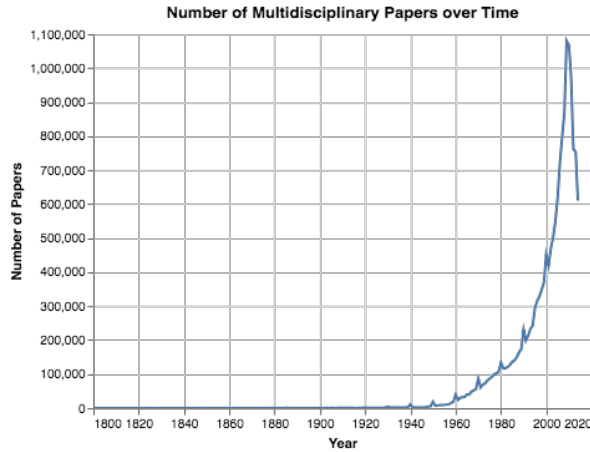
Figure 7. The Number and Percentage of Multidisciplinary Papers over Time. Between 1900 and 2010, both the number and percentage of multidisciplinary papers increased over time.
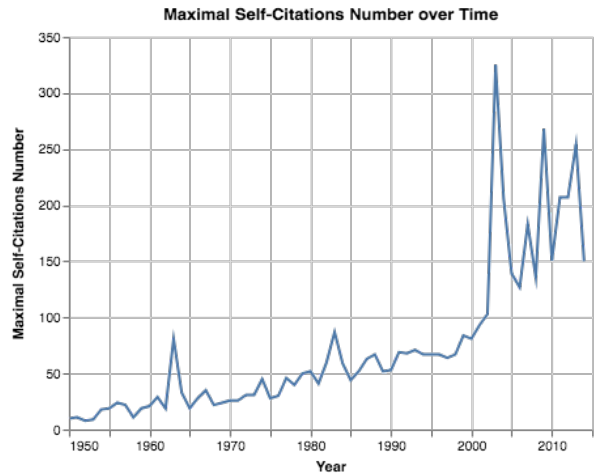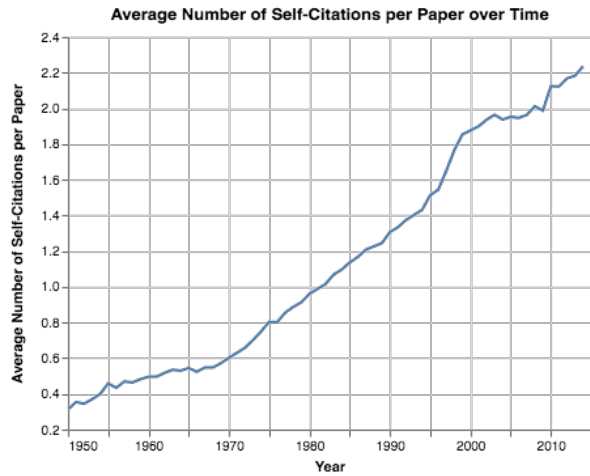
Figure 8. The Average and Maximal Number of Self-Citations. Both the average and maximal number of self-citations increased over time.
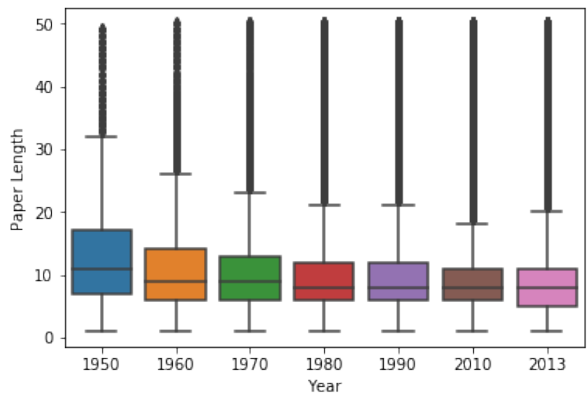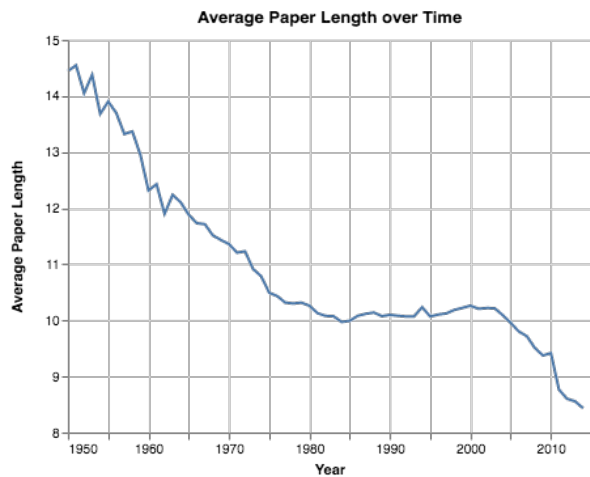
Figure 9. Papers' Lengths. Both the papers' average and median lengths decreased over time.

Figure 10. **Total Number of Papers with No Citations after 5 Years.** The number of papers with increased sharply over time.

Figure 11. Citation Distributions over Time. The citation distributions of different decades show notable changes.

Figure 12. Average Number of Papers by Authors' Academic Birth Decades. With each decade, the rate of paper publication has increased.

exclamation marks ($\tau_s = 0.022$) (see Figure S13).[10]

4.2 Results of Author Trends

By analyzing the number of new authors each year, we discovered a sharp increase over time, with several million new authors publishing each year in recent years (see Figure S14).[11] Additionally, when analyzing the trends grouped by the authors' academic birth decades, we discovered a significant increase in the average number of published papers for the younger birth decades (see Figure 12). For example, researchers who started their careers in 1950 published on average 1.55 papers in a time period of 10 years, while researchers who started their careers in 2000 published on average 4.05 papers in the same time frame. Furthermore, we observed that authors who started their careers after 1990 tended to publish more in conferences in the first years of their career than their more senior peers who started their careers in the 1950s or 1970s (see Figure S15). For example, researchers who started their careers in the 1970s published on average about 2 conference papers and 1.65 journal papers after 10 years; researchers who started their careers in the 2000s published about 4 conference papers and 2.59 journal papers in the same time frame.

We can also observe that the average number of coauthors has considerably increased over the decades (see Figure 13). Moreover, we can notice that researchers who started their careers in the 1950s and 1970s had on average only few coauthors over a period of 25 years, while researchers who started their careers in the 1990s had over 60 coauthors in the same career length of 25 years (see Figure 13).

Lastly, by exploring how author sequence numbers evolved, we discovered that with seniority, the researchers' median sequence number increased (see Figure S16). Additionally, with seniority, the percentage of published papers with the researcher listed as the first author decreased (see Figure 14).

---

[10]Similar correlation values were obtained by calculating the correlations for papers published in a specific year.

[11]It is possible that the same author has several MAG author IDs.

Figure 13. Average Number of Coauthors by Academic Birth Decade. The average number of coauthors has considerably increased over the decades.

Figure 14. Percentage of Times Researcher was First Author. We can observe that over time on average the percentage of senior researchers as first authors declined. Moreover, in the same time intervals, the percentage of times recent generations of researchers were first authors declined compared to older generations.

Moreover, by looking at the decade researchers started their careers, we can see a sharp decline in the percentages of first authors (see Figure 14). Overall, early career researchers are publishing more in their careers but appear as first authors much less than in previous generations.

4.3 Results of Journal Trends

By analyzing journal trends using the SJR and MAG datasets, we discovered that the number of journals increased significantly over the years, with 20,975 active ranked journals in 2016 (see Figure 15). Furthermore, we observed that hundreds of new ranked journals were published each year (see Figures S17 and S18). In addition, we discovered that the number of published papers per journal increased sharply, from an average of 74.2 papers in 1999 to an average of 99.6 papers in 2016 (see Figure 15). We also observed that in recent years, journals that publish thousands of papers have become more common. For example, in 2016, according to the SJR dataset, 197 journals published over 1000 papers each.

By exploring how various metrics have changed over time, we discovered the following: First, over the last 18 years, the number of papers published in Q1 and Q2 journals more than doubled, from 550,109 Q1 papers and 229,373 Q2 papers in 1999, to 1,187,514 Q1 papers and 554,782 Q2 papers in 2016 (see Figure 16). According to the SJR dataset, in 2016, 51.3% of journal papers were published in Q1 journals and only 8.66% were published in Q4 journals. Second, the average h-index decreased over recent years from an average value of 37.36 and median value of 23 in 1999 to an average value of 31.3 and median value of 16 in 2016 (see Figure S19). Third, we noted that the SJR and the average number of citations measures both increased considerably during the last 18 years (see Figures 17 and S20).

We selected the top 40 journals with the highest SJR values in 2016. We matched the journals' titles and ISSNs with the data in our Join dataset to identify the journal IDs in the MAG dataset. Using this method, we identified 30 unique journal IDs in the MAG dataset that published 110,825 papers with over 5 references. Next, we performed an in-depth analysis on these top-selected journal papers to better understand how various properties changed over time. We discovered that between 2000 and 2014, the number of papers in these journals more than doubled, and the number of authors increased significantly (see Figure S21).[12]

Additionally, by calculating the average academic career ages of first and last authors, we discovered that in recent years the average academic age has increased notably (see Figure 18). Moreover, when looking at first and last authors who previously published in one of the selected top-30 journals, we discovered that over time the percentage of returning authors increased substantially (see Figure 19). By 2014, 46.2% of all published papers in top-30 selected journals were published by last authors who had published at least one paper in a top-30 selected journal before (see Figure 19).

By calculating the number of papers, number of authors, authors' average age, and percentage of returning authors in each selected top-30 journal, we observed the following: (a) the number of published papers per year increased considerably in the vast majority of the journals (see Figure S22);[13] (b) the average career ages of last authors in the vast majority of the selected journals considerably increased (see Figure S23), like in Cell journal where the last authors' career ages increased from about 4.5 years in 1980 to about 20 years in 2014 (see Figure S23); and (c) the percentage of returning authors in the vast majority of the selected journals increased drastically, like in Nature Genetics where in 86.6% of 2014 papers, at least one of the authors had published in the journal before (see Figure 20).

---

[12]The total number of authors each year was determined by summing the number of authors in each published paper.

[13]Due to missing references in the MAG dataset, there are decline in the number of papers in Nature (1990s), and in Science (before 2008).

Figure 15. Number of Active Journals over Time. Over a period of 18 years, from 1999 to 2016, both the number of active journals and the papers per journal increased greatly.


Figure 16. Journals' Quartile Number of Papers over Time. The number of papers published in Q1 journals has vastly increased.

4.4 Results of Fields-of-Research Trends

By matching each paper to its L0 field of study and analyzing each field's properties, we discovered substantial differences in these properties. Namely, we observed the following:

- A large variance in the number of published papers in each field. For example, 231,756 papers were published in the field of biology in 2010, but only 5,684 were published that year in the field of history (see Figures 21 and S24).

- A considerable variance in the average number of paper authors among the various research fields. For example, the number of authors in 2010 ranged from an average of 2.28 authors in the field of political science to an average of 5.39 authors in medicine (see Figure S25).

- A variance in the papers' average number of references in different fields. For example, in 2010, the average reference number in the fields of material science and

engineering was less than 24, while in the fields of biology and history it was over 33 (see Figure S26).

- A big variance in each L0 field's average and median number of citations after 5 years. For example, for 2009 papers in the fields of computer science and political science, the median citation number after 5 years was 4 citations. In biology and environmental science, the median citation number after 5 years was 9 and 13 citations, respectively (see Figure 22).

By repeating the above analysis for the L1 subfields of biology and for the L2 subfields of genetics, we uncovered similar differences among fields of study. Namely, we observed the following for subfields in the same hierarchal level: (a) significant variance in the average number of papers (see Figures S27 and S28); (b) notable variance in the average number of authors (see Figures S29 and S30); (c) noteworthy variance in the average number of references (see Figures S31 and S32); and (d) vast variance in median citation numbers (see Figures S33 and S34).

Lastly, by analyzing various features of 2,673 L3 fields of study, we observed a huge variance in the different properties (see Table 1 and Figure S35). For example, several fields of study, such as gallium (chemistry), ontology (computer science), and presentation of a group (mathematics), had median citation numbers of 2, while other fields of study, such as microRNA and genetic recombination (biology), had median citation numbers of over 47 and 50.5, respectively (see Table 1 and Figure S35).

5. Discussion

By analyzing the results presented in Section 4, the following can be noted:

First, we can observe that the structure of academic papers has changed in distinct ways in recent decades. While the average overall length of papers has become shorter (see Figure 9), the title, abstract, and references have become longer (see

Table 1. L3 Fields-of-Study Features in 2009

Parent Field of Study	Field of Study Name	Median Citations After 5 Years	MAX Citations After 5 Years	Number of Papers	Average Author Number
Engineering	**Structural material**	61.0	1250	174	6.14
Biology	**Genetic recombination**	50.5	451	196	6.07
Biology	**Nature**	48.0	5660	4162	6.28
Biology	**microRNA**	47.0	3076	1691	6.24
Biology	**Induced pluripotent stem cell ...**	39.0	987	213	6.53
Economics	**Signalling**	39.0	695	1030	5.87
Biology	**Genome evolution**	35.5	392	140	5.04
Biology	**Non-coding RNA**	35.0	1414	375	5.39
Biology	**Post-transcriptional modification ...**	34.0	1414	315	5.49
Biology	**Autophagy**	34.0	789	381	5.71
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
Mathematics	**Finite impulse response**	2.0	167	337	3.0
Computer Science	**Pixel**	2.0	380	2484	3.27
Computer Science	**Ontology**	2.0	616	733	3.35
Computer Science	**Mesh networking**	2.0	62	274	3.43
Computer Science	**Camera resectioning**	2.0	43	114	3.13
Computer Science	**Session Initiation Protocol ...**	2.0	116	100	3.6
Chemistry	**Gallium**	2.0	73	484	3.43
Mathematics	**Presentation of a group**	2.0	91	706	3.22
Mathematics	**Spiral**	2.0	80	122	3.65
Mathematics	**Block code**	2.0	54	281	2.83

Figure 17. The Average Number of Citations ($\frac{Cites}{Docs}$ (2 years)) over Time. The average number of citations values have almost doubled in the last 18 years; additionally, their distributions have changed considerably.

Figure 18. Top-Selected Journals' Average First and Last Authors Ages. Both the first and last authors' average ages have increased sharply.

Figure 19. Percentage of Papers with Returning First or Last Authors. The percentage of returning first or last top-journal authors increased considerably.

Section 4.1 and Figures 4, 6, S3, S6, S9, and S10). Also, the number of papers that include keywords has increased considerably, as has the average number of keywords in each paper (see Figure S7). Furthermore, the average and median number of authors per paper has increased sharply (see Figures S3 and S4).

These results support Goodhart's Law as it relates to academic publishing: the measures (e.g., number of papers, number of citations, h-index, and impact factor) have become targets, and now they are no longer good measures. By making papers shorter and collaborating with more authors, researchers are able to produce more papers in the same amount of time. Moreover, we can observe that the majority of changes in papers' properties are correlated with papers that receive higher numbers of citations (see Figure S13). Authors can use longer titles and abstracts, or use question or exclamation marks in titles, to make their papers more appealing. Thus more readers are attracted to the paper, and ideally they will cite it, i.e., academic clickbait [28]. These results support our hypothesis that the citation number has become a target. Consequently, the properties of academic papers have evolved in order to win-to score a bullseye on the academic target.

Second, we observed that over time fewer papers list authors alphabetically, especially papers with a relatively high number of authors (see Section 4.1 and Figures 5 and S5). These results may indicate the increased importance of an author's sequence number in a paper, which may reflect the author's contribution to the study. This result is another signal of the rising importance of measures that rate an individual's research contribution.

Third, from matching papers to their L0 fields of study, we observed that the number of multidisciplinary papers has increased sharply over time (see Figure 7). It is important to keep in mind that these results were obtained by matching keywords to their corresponding fields of study. Therefore, these results have several limitations: First, not all papers contain

keywords. Second, the dataset may not extract keywords from papers in the correct manner. For example, we found some papers contained keywords in their online version but not in their offline version (see Section 4.1). It is also possible that in some fields it is less common to use keywords. Therefore, the papers' keywords may be missing in the datasets, and the presented results may be an underestimate of the actual number of multidisciplinary studies. Nevertheless, we observed a strong trend in increasing numbers of multidisciplinary papers.

Fourth, from seeing sharp increases in both the maximal and average number of self-citations (see Section 4.1 and Figures 8, S11, 11, and S12), it is clear that citation numbers have become a target for some researchers who cite their own papers dozens, or even hundreds, of times. Furthermore, we can observe a general increasing trend for researchers to cite their previous work in their new studies. Moreover, from analyzing the percentage of papers without citations after 5 years, we can observe that a huge quantity of papers – over 72% of all papers and 25% of all papers with at least 5 references – have no citations at all (see Figure S11). Obviously, many resources are spent on papers with limited impact. The lack of citations may indicate that researchers are publishing more papers of poorer quality to boost their total number of papers. Additionally, by exploring papers' citation distributions (see Figure 11), we can observe that different decades have very different citation distributions. This result indicates that comparing citation records of researchers who published papers in different time periods can be challenging.

Fifth, by exploring trends in authors (see Section 4.2 and Figures 12, 13, 14, S14, S15, and S16), we observed an exponential growth in the number of new researchers who publish papers. We also observed that young career researchers tend to publish considerably more than researchers in previous generations, using the same time frames for comparison (see Figure 12). Moreover, young career researchers tend to publish their work much more in conferences in the beginning of their careers than older researchers did in previous decades (see Figure S15). We also observed that young career researchers tend to collaborate considerably more in the beginning of their careers than those who are older (see Figure 13). Furthermore, we see that the average percentage of researchers as first authors early in their career is considerably less than those in previous generations (see Figure 14). In addition, authors' median sequence numbers typically increase over time, and the rate is typically faster for young career researchers (see Figure S16). These results emphasize the changes in academia in recent years. In a culture of "publish or perish," researchers publish more by increasing collaboration (and being added to more author lists) and by publishing more conference papers than in the past. However, as can be observed by the overall decline of researchers as first authors, young career researchers may be publishing more in their careers but contributing less to each paper. The numbers can be misleading: a researcher who has 5 "first author" claims but has published 20 papers may be less of a true contributor than one with 4 "first author" claims and 10 published papers.

Sixth, by analyzing journal trends (see Section 4.3), we see a rapid increase in the number of ranked active journals in recent years (see Figure 15). Moreover, on average, journals publish more papers than in the past, and dozens of journals publish over 1,000 papers each year (see Figure 15 and S17). With the increase in the number of active journals, we observed rapid changes in impact measures: (a) the number of papers published in the first and second quartiles (Q1 and Q2) has increased sharply, and today the vast majority of papers are published in these quartiles (see Figure 16); (b) the journals' average and median h-index has decreased sharply (see Figure S18); and (c) both the SJR and the average number of citations has increased considerably (see Figures 17 and S20). With these significant changes, it is clear that some measures, such as the use of quartiles and the h-index, are rapidly losing meaning and value. Moreover, with the abundance of journals, researchers can "shop around" for a high impact journal and submit a rejected paper from one Q1 journal to another Q1 journal, time after time, and then start the review process again. These repeated reviews for the same paper wastes time, and in the long run the burden of reviewing papers several times may affect the quality of the reviews.

There are compelling reasons to change the current system. We need to think about making all reviews open and online. We should consider the function of published journals; for that matter, is it even necessary to have journals in a world with over 20,000 journals that publish hundreds or even thousands of papers each year? We need to seriously evaluate the measures we use to judge research work. If all these measures have been devalued to being merely targets, they are no longer effective measures. Instead, they should be adapted to meet our current needs and priorities.

Seventh, by focusing on trends in selected top journals, we can observe that these journals have changed considerably in recent years (see Figures 18, 19, 20, S21, and S22). The number of papers in the selected journals has increased sharply, along with the career age of the authors and the percentage of returning authors. The number of submissions to top journals, like Nature, have increased greatly in recent years [54]; however, many of these journals mainly publish papers in which at least one of the authors has previously published in the journal (see Figure 19 and 20). We believe that this situation is also a result of Goodhart's Law. The target is the impact factor, and so researchers are vigorously seeking journals with high impact factors. Therefore, the yearly volume of papers sent to these top journals has considerably increased, and overwhelmed by the volume of submissions, editors at these journals may choose safety over risk and select papers written by only well-known, experienced researchers.

Eighth, by analyzing how features evolve in the various L0 fields of study using the MAG dataset, we can observe that different fields have completely different sets of features (see Figures 21, 22, 21, S25, S26, and Table 1). While some fields have hundreds of thousands of papers published yearly, others

have only thousands published yearly (see Figures 21and S22). Moreover, similar large differences are reflected in other examined fields' features, such as the average number of references and the average and median citation numbers (see Figures 22 and S35).

Lastly, by examining over 2600 research fields of various scales (see Table 1 and Figure S35), we observed vast diversity in the properties of papers in different domains – some research domains grew phenomenally while others did not. Even research domains in the same subfields presented a wide range of properties, including papers' number of references and median number of citations per research field (see Table 1 and Figures S31, S32, S33, and S34). These results indicate that using measures such as citation number, h-index, and impact factor are useless when comparing researchers in different fields, and even for comparing researchers in the same subfield, such as genetics. These results emphasize that using citation-based measures for comparing various academic entities is like comparing apples to oranges, and is to "discriminate between scientists." [34]. Moreover, using these measures as gauges to compare academic entities can drastically affect the allocation of resources and consequently damage research. For example, to improve their world ranking, universities might choose to invest in faculty for computer science and biology, rather than faculty for less-cited research fields, such as economics and psychology. Moreover, even within a department, the selection of new faculty members can be biased due to using targeted measures, such as citation number and impact factor. A biology department might hire genetic researchers in the field of epigenetics, instead of researchers in the field of medical genetics, due to the higher average number of citations in the epigenetics field. Over time, this can unfairly favor high-citation research fields at the expense of other equally significant fields.

6. Conclusions

In this study, we performed a large-scale analysis of academic publishing trends, utilizing data on over 120 million papers and over 20,000 journals. By analyzing this huge dataset, we can observe that over the last century, especially the last few decades, published research has changed considerably, including the numbers of papers, authors, and journals; the lengths of papers; and the average number of references in specific fields of study.

While the research environment has changed, the measures to determine the impact of papers, authors, and journals have not changed. Measures based on citations, such as impact factor and citation number, were used 60 years ago, in a time before preprint repositories and mega-journals existed and before academia became such a hypercompetitive environment. Most important, however, is that these measures have degenerated into becoming purely targets. Goodhart's Law is clearly being illustrated: when a citation-based measure becomes the target, the measure itself ceases to be meaningful, useful, or accurate.

Our study's extensive analysis of academic publications reveals why using citation-based metrics as measures of impact are wrong from the core: First, not all citations are equal; there is a big difference between a study that cites a paper that greatly influenced it and a study that cites multiple papers with only minor connections. Many of the impact measures used today do not take into consideration distinctions among the various types of citations. Second, it is not logical to measure a paper's impact based on the citation numbers of other papers that are published in the same journal. In the academic world, there are over 20,000 journals that publish hundreds or even thousands of papers each year, with papers written by hundreds or even thousands of authors. It is even less logical to measure a researcher's impact based on a paper coauthored with many other researchers according to the journal in which it is published. Third, as we demonstrated in Section 4.4, it is wrong to compare studies from different fields, and even to compare papers and researchers within the same parent field of study, due to the many differences in the median and average number of citations in each field (see Table 1).

As we have revealed in this study, to measure impact with citation-based measures—that have now become targets—clearly has many undesirable effects. The number of papers with limited impact has increased sharply (see Figure 10), papers may contain hundreds of self-citations (see Figure 8), and some top journals have become "old boys' clubs" that mainly publish papers from the same researchers (see Figures 18 and 19). Moreover, using citation-based measures to compare researchers in different fields may have the dangerous effect of allocating more resources to high-citation domains, shortchanging other domains that are equally important.

We believe the solution to the above issues is to utilize data-science tools and release new and open datasets in order to develop new measures that will more accurately determine a paper's impact in a specific research field. Certain metrics have been proposed, but the key is to wisely and carefully evaluate new measures to ensure that they will not follow Goodhart's Law and end up merely as targets. Researchers do valuable work. Communicating the work to others is vital, and correctly assessing the impact of that work is essential.

7. Data and Code Availability

One of the main goals of this study was to create an open source framework, which provided an easy way to query the datasets. Our code framework, including tutorials, is available at the project's website, which also gives researchers the ability to interactively explore and better understand how various journals' properties have changed over time (see Figure S36).

Acknowledgments

References

[1] Bo-Christer Björk. Have the "mega-journals" reached the limits to growth? *PeerJ*, 3:e981, 2015.

[2] Robin Chin Roemer and Rachel Borchardt. From bibliometrics to altmetrics: A changing scholarly landscape. *College & Research Libraries News*, 73(10):596–600, 2012.

[3] Jorge E Hirsch. An index to quantify an individual's scientific research output. *Proceedings of the National academy of Sciences*, 102(46):16569–16572, 2005.

[4] Eugene Garfield. The agony and the ecstasy—the history and meaning of the journal impact factor, 2005.

[5] James Wilsdon. *The metric tide: Independent review of the role of metrics in research assessment and management*. Sage, 2016.

[6] Marc A Edwards and Siddhartha Roy. Academic research in the 21st century: Maintaining scientific integrity in a climate of perverse incentives and hypercompetition. *Environmental Engineering Science*, 34(1):51–61, 2017.

[7] Mario Biagioli. Watch out for cheats in citation game. *Nature News*, 535(7611):201, 2016.

[8] Donald T Campbell. Assessing the impact of planned social change. *Evaluation and program planning*, 2(1):67–90, 1979.

[9] Paul Mizen. *Central Banking, monetary theory and practice: essays in honour of Charles Goodhart*, volume 1. Edward Elgar Publishing, 2003.

[10] K Alec Chrystal, Paul D Mizen, and PD Mizen. Goodhart's law: its origins, meaning and implications for monetary policy. *Central banking, monetary theory and practice: Essays in honour of Charles Goodhart*, 1:221–243, 2003.

[11] Chris Francescani. Nypd report confirms manipulation of crime stats, 2012.

[12] Vesna Šupak Smolčić. Salami publication: Definitions and examples. *Biochemia medica: Biochemia medica*, 23(3):237–241, 2013.

[13] Jerome Schofferman, F Todd Wetzel, and Christopher Bono. Ghost and guest authors: You can't always trust who you read. *Pain Medicine*, 16(3):416–420, 2015.

[14] Megan L Head, Luke Holman, Rob Lanfear, Andrew T Kahn, and Michael D Jennions. The extent and consequences of p-hacking in science. *PLoS biology*, 13(3):e1002106, 2015.

[15] Christoph Bartneck and Servaas Kokkelmans. Detecting h-index manipulation through self-citation analysis. *Scientometrics*, 87(1):85–98, 2010.

[16] Charlotte J Haug. Peer-review fraud—hacking the scientific publication process. *New England Journal of Medicine*, 373(25):2393–2395, 2015.

[17] Santo Fortunato, Carl T Bergstrom, Katy Börner, James A Evans, Dirk Helbing, Staša Milojević, Alexander M Petersen, Filippo Radicchi, Roberta Sinatra, Brian Uzzi, et al. Science of science. *Science*, 359(6379):eaao0185, 2018.

[18] arXiv. arxiv monthly submission rates, (Accessed: 9-July-2018). /https://arxiv.org/stats/monthly_submissions.

[19] Phil Davis. Scientific reports overtakes plos one as largest megajournal - the scholarly kitchen, (Accessed: 9-July-2018). /https://scholarlykitchen.sspnet.org/2017/04/06/scientific-reports-overtakes-plos-one-as-largest-megajournal/.

[20] Benjamin P Abbott, Richard Abbott, TD Abbott, MR Abernathy, Fausto Acernese, Kendall Ackley, Carl Adams, Thomas Adams, Paolo Addesso, RX Adhikari, et al. Observation of gravitational waves from a binary black hole merger. *Physical review letters*, 116(6):061102, 2016.

[21] Davide Castelvecchi. Ligo's unsung heroes. *Nature News*, 2017.

[22] Robert Aboukhalil. The rising trend in authorship. *The winnower*, 2:e141832, 2014.

[23] Joseph S Wislar, Annette Flanagin, Phil B Fontanarosa, and Catherine D DeAngelis. Honorary and ghost authorship in high impact biomedical journals: a cross sectional survey. *Bmj*, 343:d6128, 2011.

[24] Maureen Shawn Kennedy, Jane Barnsteiner, and John Daly. Honorary and ghost authorship in nursing publications. *Journal of Nursing Scholarship*, 46(6):416–422, 2014.

[25] Francisco Emilio Vera-Badillo, Marc Napoleone, Monika K Krzyzanowska, Shabbir MH Alibhai, An-Wen Chan, Alberto Ocana, Arnoud J Templeton, Bostjan Seruga, Eitan Amir, and Ian F Tannock. Honorary and ghost authorship in reports of randomised clinical trials in oncology. *European Journal of Cancer*, 66:1–8, 2016.

[26] The Economist. Why research papers have so many authors. *The Economist*, 2016. /http://www.economist.com/news/science-and-technology/21710792-scientific-publications-are-getting-more-and-more-names-attached-them-why.

[27] Grant Lewison and James Hartley. What's in a title? numbers of words and the presence of colons. *Scientometrics*, 63(2):341–356, 2005.

[28] Gwilym Lockwood. Academic clickbait: articles with positively-framed titles, interesting phrasing, and no wordplay get more attention online. *The Winnower*, 3, 2016.

[29] Antonio Gálvez, Mercedes Maqueda, Manuel Martinez-Bueno, and Eva Valdivia. Scientific publication trends and the developing world: what can the volume and authorship of scientific articles tell us about scientific progress in various regions? *American scientist*, 88(6):526–533, 2000.

[30] Reshma Jagsi, Elizabeth A Guancial, Cynthia Cooper Worobey, Lori E Henault, Yuchiao Chang, Rebecca Starr, Nancy J Tarbell, and Elaine M Hylek. The "gender gap" in authorship of academic medical literature—a 35-year perspective. *New England Journal of Medicine*, 355(3):281–287, 2006.

[31] Matthew R Johnson, Nick J Wagner, and Jonathan Reusch. Publication trends in top-tier journals in higher education. *Journal of Applied Research in Higher Education*, 8(4):439–454, 2016.

[32] Peter Aldhous. Scientific publishing: The inside track. *Nature*, 2014.

[33] Erjia Yan and Ying Ding. Weighted citation: An indicator of an article's prestige. *Journal of the American Society for Information Science and Technology*, 61(8):1635–1643, 2010.

[34] Sune Lehmann, Andrew D Jackson, and Benny E Lautrup. Measures for measures. *Nature*, 444(7122):1003, 2006.

[35] Eugene Garfield. The meaning of the impact factor. *International Journal of Clinical and Health Psychology*, 3(2), 2003.

[36] Jorge E Hirsch. Does the h index have predictive power? *Proceedings of the National Academy of Sciences*, 104(49):19193–19198, 2007.

[37] Eric A Fong and Allen W Wilhite. Authorship and citation manipulation in academic research. *PloS one*, 12(12):e0187394, 2017.

[38] Nature Editorial. The cost of salami slicing. *Nature Materials*, 4:1, 2005.

[39] Emilio Delgado López-Cózar, Nicolás Robinson-García, and Daniel Torres-Salinas. The google scholar experiment: How to index false papers and manipulate biblio-metric indicators. *Journal of the Association for Information Science and Technology*, 65(3):446–454, 2014.

[40] René Van Bevern, Christian Komusiewicz, Rolf Niedermeier, Manuel Sorge, and Toby Walsh. H-index manipulation by merging articles: Models, theory, and experiments. *Artificial Intelligence*, 240:19–35, 2016.

[41] Qiang Wu. The w-index: A measure to assess scientific impact by focusing on widely cited papers. *Journal of the American Society for Information Science and Technology*, 61(3):609–614, 2010.

[42] Matthew E Falagas, Vasilios D Kouranos, Ricardo Arencibia-Jorge, and Drosos E Karageorgopoulos. Comparison of scimago journal rank indicator with journal impact factor. *The FASEB journal*, 22(8):2623–2628, 2008.

[43] Vincent Lariviere, Veronique Kiermer, Catriona J MacCallum, Marcia McNutt, Mark Patterson, Bernd Pulverer, Sowmya Swaminathan, Stuart Taylor, and Stephen Curry. A simple proposal for the publication of journal citation distributions. *BioRxiv*, page 062109, 2016.

[44] Ewen Callaway. Beat it, impact factor! publishing elite turns against controversial metric. *Nature*, 2016.

[45] Steffan Arthur Griffin, Christopher W Oliver, and Andrew Murray. 'altmetrics'! can you afford to ignore it?, 2017.

[46] Per O Seglen. Why the impact factor of journals should not be used for evaluating research. *BMJ: British Medical Journal*, 314(7079):498, 1997.

[47] Aidan Byrne. Comment: Measure for measure. *Nature*, 2017.

[48] Frederick Hecht, Barbara Kaiser Hecht, and Avery A. Sandberg. The journal "impact factor": a misnamed, misleading, misused measure. *Cancer genetics and cytogenetics*, 104 2:77–81, 1998.

[49] Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-June Paul Hsu, and Kuansan Wang. An overview of microsoft academic service (mas) and applications. In *WWW*, 2015.

[50] KDD Cup 2016. Sigkdd: Kdd cup 2016: Whose papers are accepted the most: towards measuring the impact of research institutions, 2016. /http://www.kdd.org/kdd-cup/view/kdd-cup-2016/Data.

[51] Jie Tang, Jing Zhang, Limin Yao, Juan-Zi Li, Li Zhang, and Zhong Su. Arnetminer: extraction and mining of academic social networks. In *KDD*, 2008.

[52] Declan Butler. Free journal-ranking tool enters citation market. *Nature*, 451:6–6, 2008.

[53] Yucheng Low, Joseph Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, and Joseph M. Hellerstein. Graphlab: A new framework for parallel machine learning. In *UAI*, 2010.

[54] Nature editorial criteria and processes., (Accessed: 15-July-2018). /https://www.nature.com/nature/for-authors/editorial-criteria-and-processes.

1. Supplementary Materials - Additional Results

A.1 Additional Results

Figure 20. Average Percentage of Return Authors in Top-Selected Journals over Time. In most journals the number of papers with at least one author who previously published in the journal increased sharply. In many of the selected journals the percentage of papers with returning authors was above 60%, and in some cases above 80%.

Figure 21. L0 Fields-of-Study Number of Papers over Time. The numbers of papers in each field of study have increased drastically.

Figure 22. L0 Field-of-Study Median Citation Number after 5 Years. There is notable variance among the L0 fields-of-study median citation numbers.

Figure S1. Overview of the Code Framework. The datasets are loaded into SFrame objects and MongoDB collections. The SFrame objects are used mainly to obtain general insights by analyzing tens of millions of papers and author records. The MongoDB collections are used to construct Paper and Author objects that can be used to analyze more complicated statistics for specific venues and research fields with usually hundreds of thousands of records.

Figure S2. Percentage of Titles with Question or Exclamation Marks. The percentage of papers with question or exclamation marks in their titles increased over time, as well as the percentage of titles with interrobangs (represented by ?! or !?).

Figure S3. Average Number of Authors over Time. There has been a rise in the average number of authors, especially in recent decades.

Figure S4. Maximal Number of Authors over Time. In recent years the maximal number of authors per paper increased sharply from 520 authors in 2000 to over 3100 authors in 2010.

Figure S6. Average Length of Abstracts. Since 1970 there has been an increase in abstracts' average number of words.

Figure S5. Percentage of Author Lists in Alphabetical Order. There has been a decline in the number of author lists organized in alphabetical order.
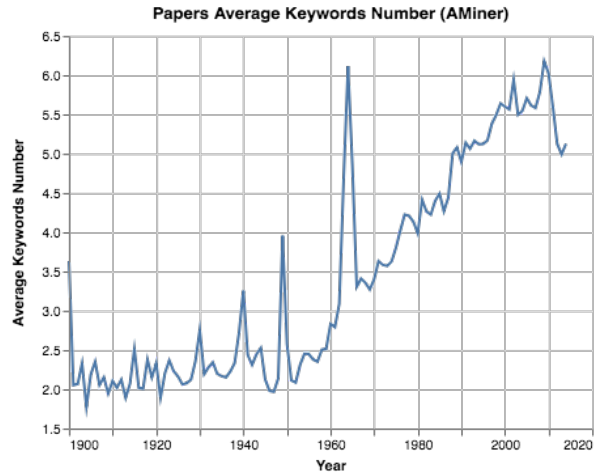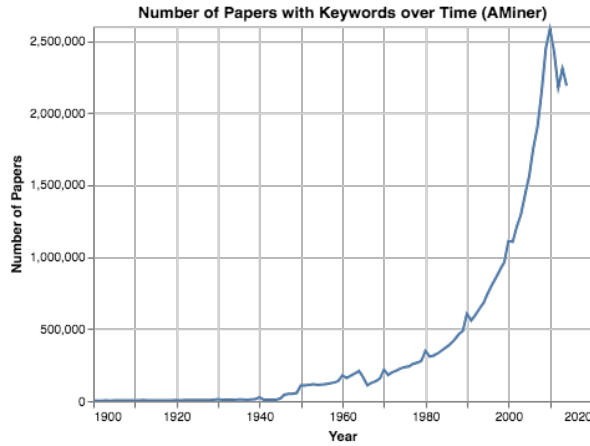
Figure S7. Keyword Trends. Both the number of papers with keywords has increased, as well as the average number of keywords per paper.

Figure S8. Average Number of Fields of Study over Time. Over time both the average number of L0 and L1 fields of studies per paper considerably increased. We believe the drop in the average number of L0 and L1 fields is a direct results of the drop in the number of papers with keywords in the same years (see Section 4.1).

Figure S9. Average Number of References over Time.
Over time, the average number of references sharply increased.
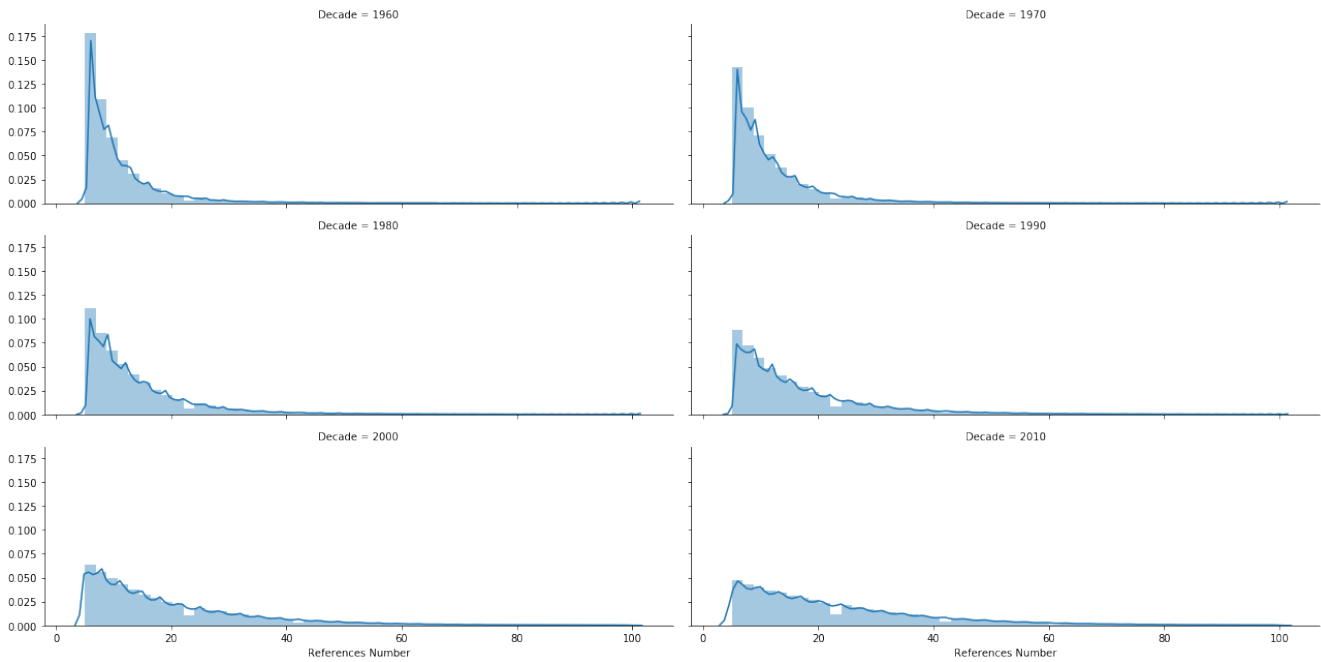
Figure S10. Distributions over Time of References in Papers. Over time, papers with a relatively high number of references have become more common.

Figure S11. Papers with No Citations after 5 Years. Papers with no citations after 5-years decreased; nevertheless, in 2009 over 72.1% of all published papers had no citations after 5 years.
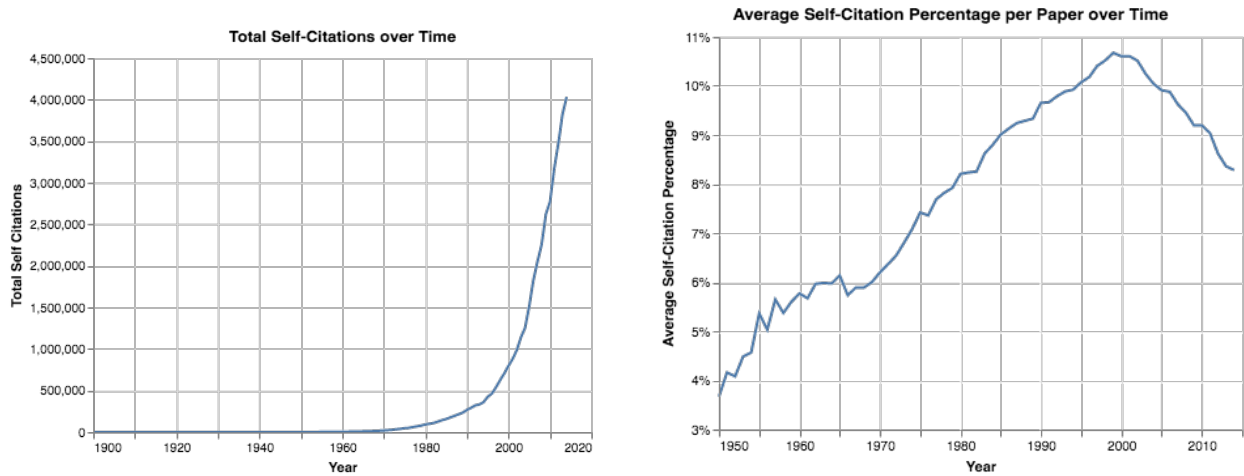
Figure S12. Total Number of Self-Citations and Percentage of Papers with Self-Citations. We can observe that over time both the total number of self-citations as well as the percentage of papers with self-citations increased significantly.

Figure S13. Spearman Correlation Heat Map for Papers' Properties. We can observe positive correlations among papers' various structural properties and the papers' total number of citations after 5 years.

Figure S14. New Authors over Time. The number of authors, with unique MAG author IDs, who published their first paper each year.

Figure S15. Authors Average Number of Conference and Journal Papers over Time. The average publication rate of both journal and conference papers increased with every decade.

Figure S16. Authors' Median Sequence Number over Time. We can see that over time the median sequence numbers increased; i.e., senior researchers tend to have higher sequence numbers.

Figure S18. Number of New Journals by Year. Hundreds of new ranked journals are being published each year.

Figure S17. Number of Journals over Time according to the MAG Dataset. There has been a drastic increase in the number of journals since the 1960s.
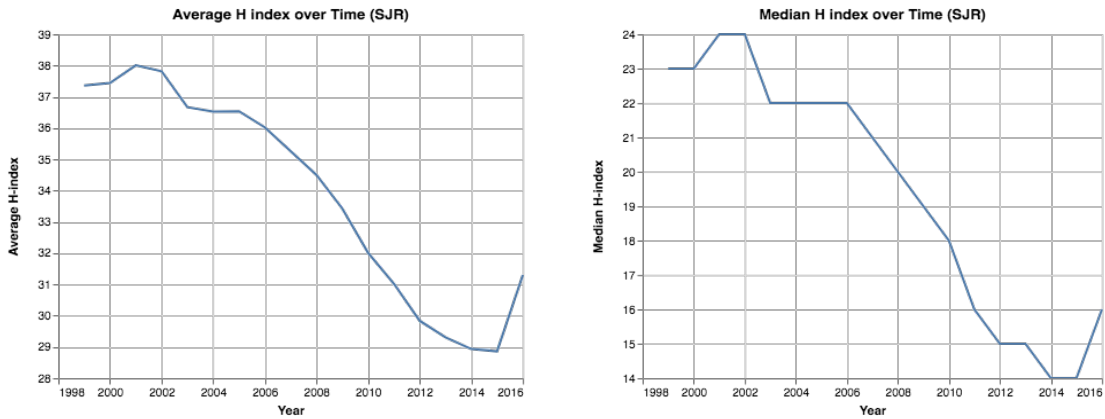
Figure S19. Journals' H-Index Average and Median Values. We can notice that over time both the average and median values of the journals' h-index measures decreased.

Figure S20. SJR Values over Time. We can observe that over time both the average and median SJR values increased.

Figure S21. Top Journals' Number of Papers and Authors over Time. We can observe that both the number of papers and authors increased sharply in recent years.
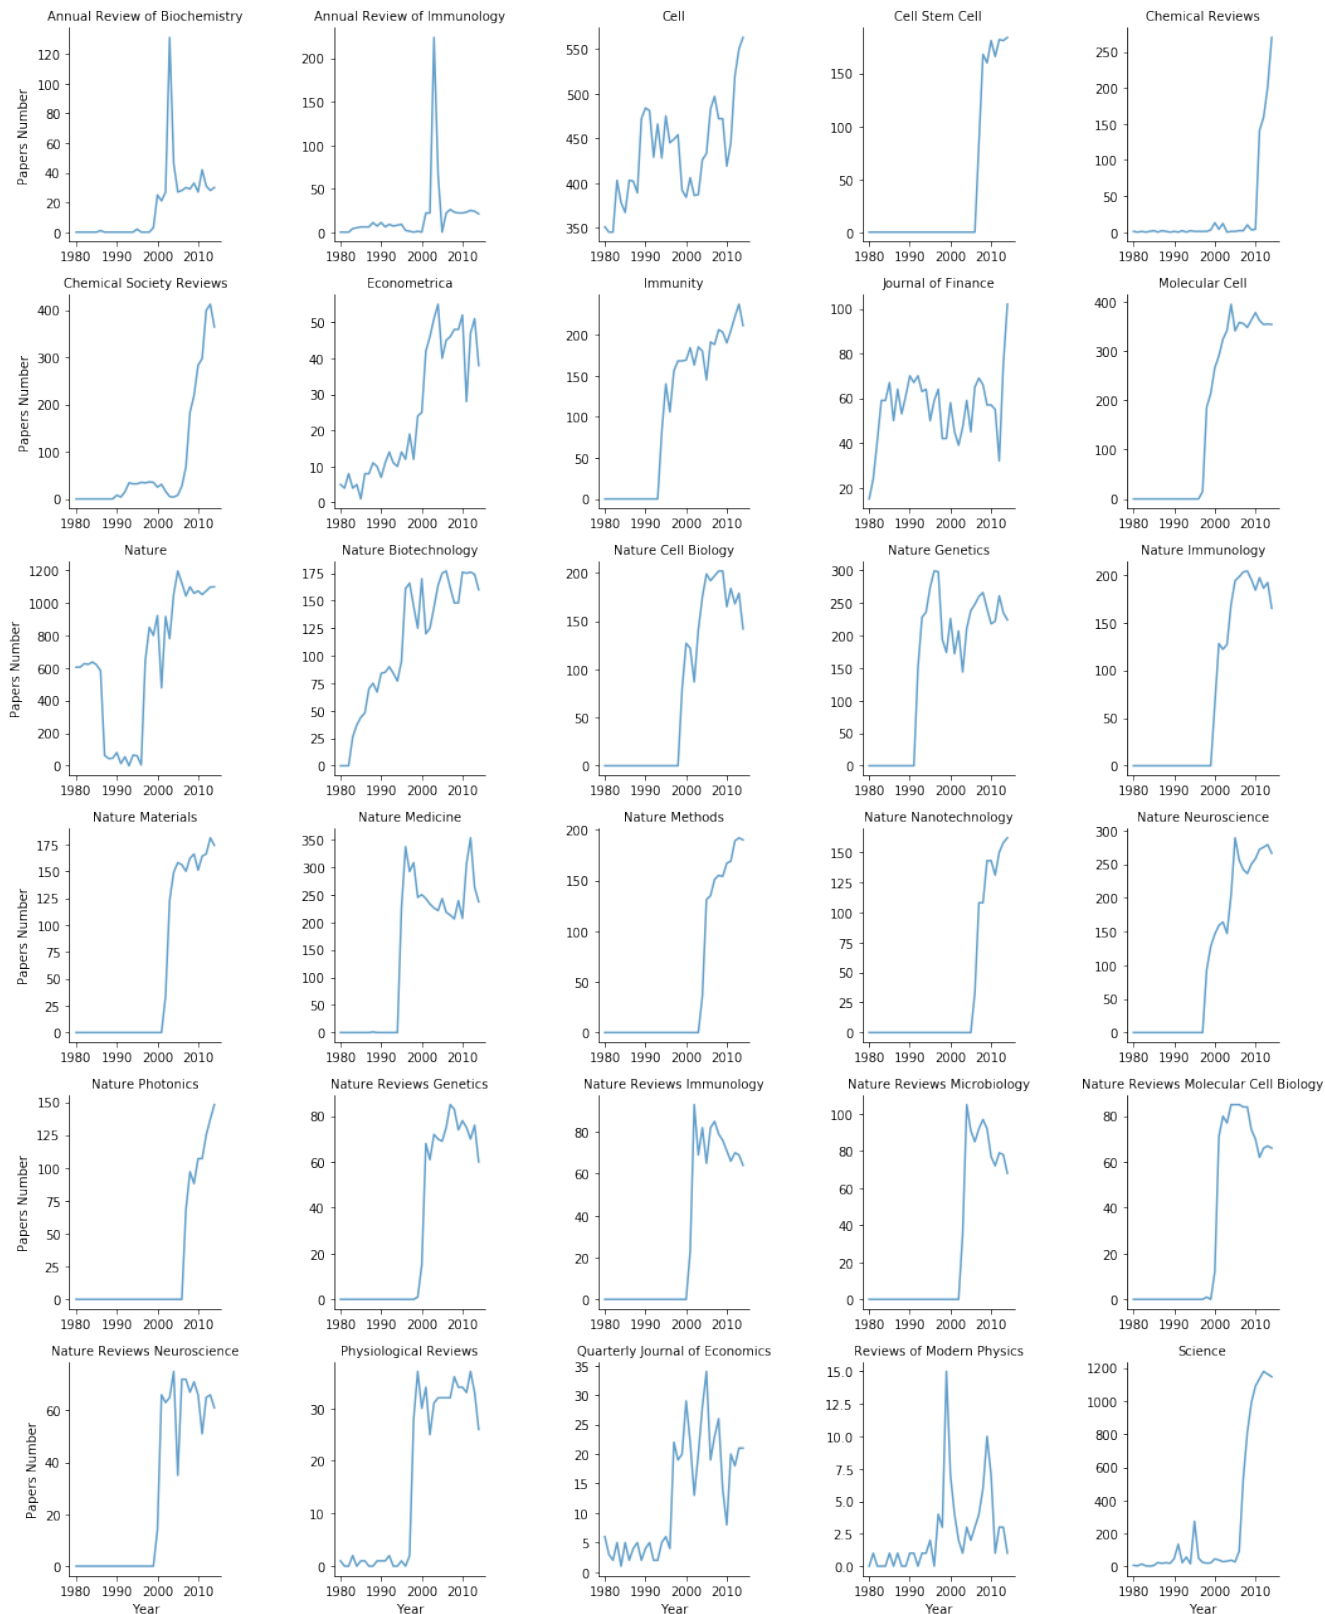
Figure S22. Top Selected Journals' Number of Papers over Time. It can be noted that in the vast majority of the selected journals the number of published papers with at least 5 references increased considerably over time.
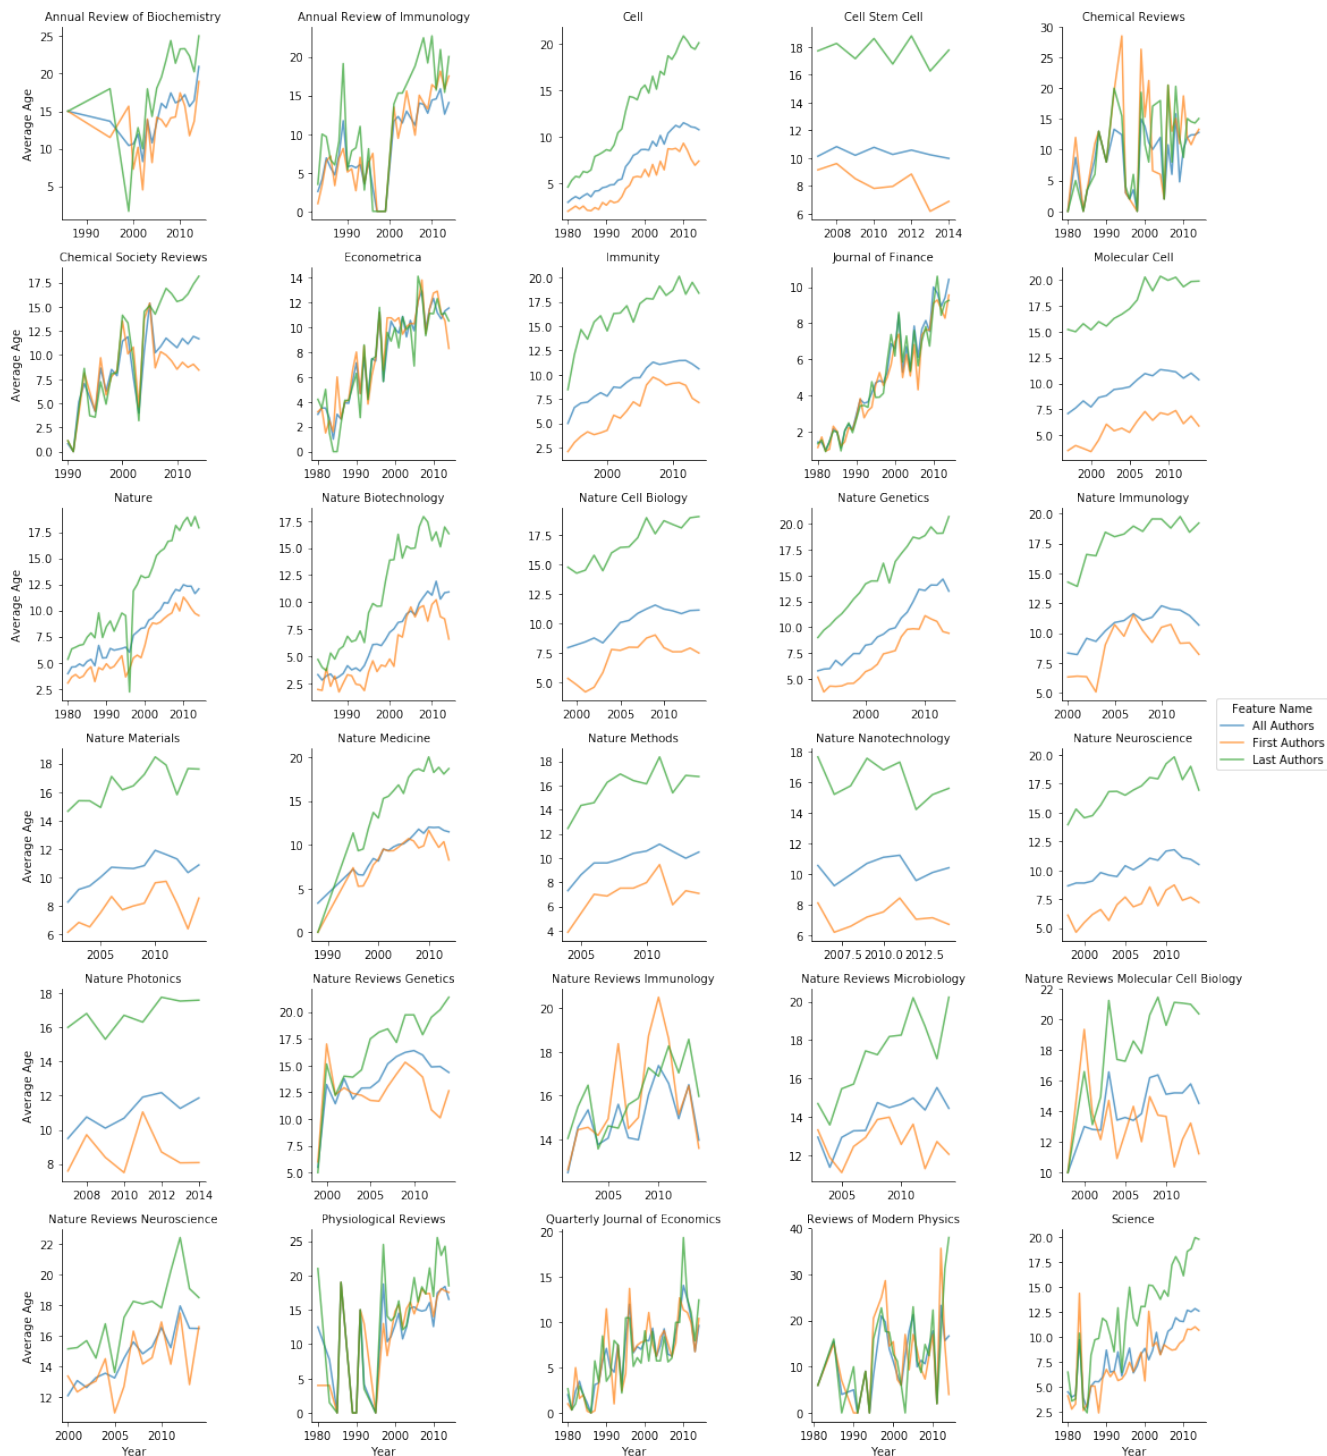
Figure S23. Top Selected Journals Average Author Career Age over Time. It can be noted that in the vast majority of the selected journals, the average age of authors, especially last authors, increased greatly over time.

Figure S24. L0 Fields-of-Study Number of Papers over Time. We can observe the large diversity in the number of papers published in each L0 research field.

Figure S25. L0 Fields-of-Study Average Authors Number. We can observe a variation in the average number of authors across the various research fields.

Figure S26. L0 Fields-of-Study Average References Numbers. We can observe variance among the reference numbers in different fields.
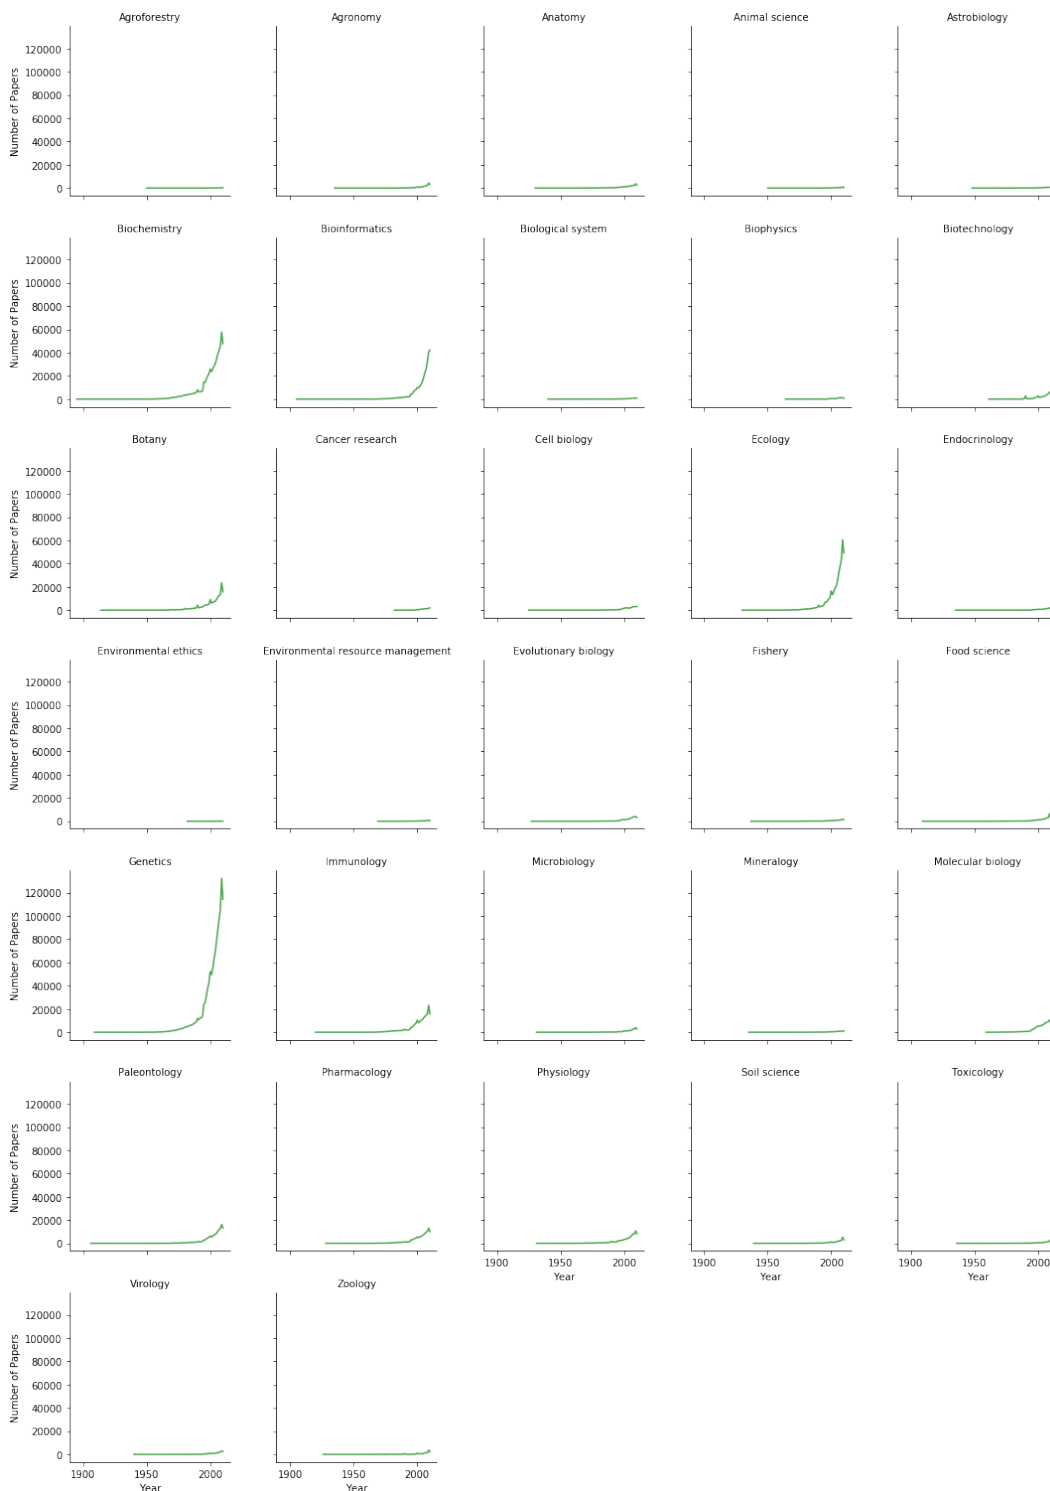
Figure S27. Biology L1-Subfields Number of Papers over Time. We can observe a big variance in the number of papers over time in the various biology subfields.
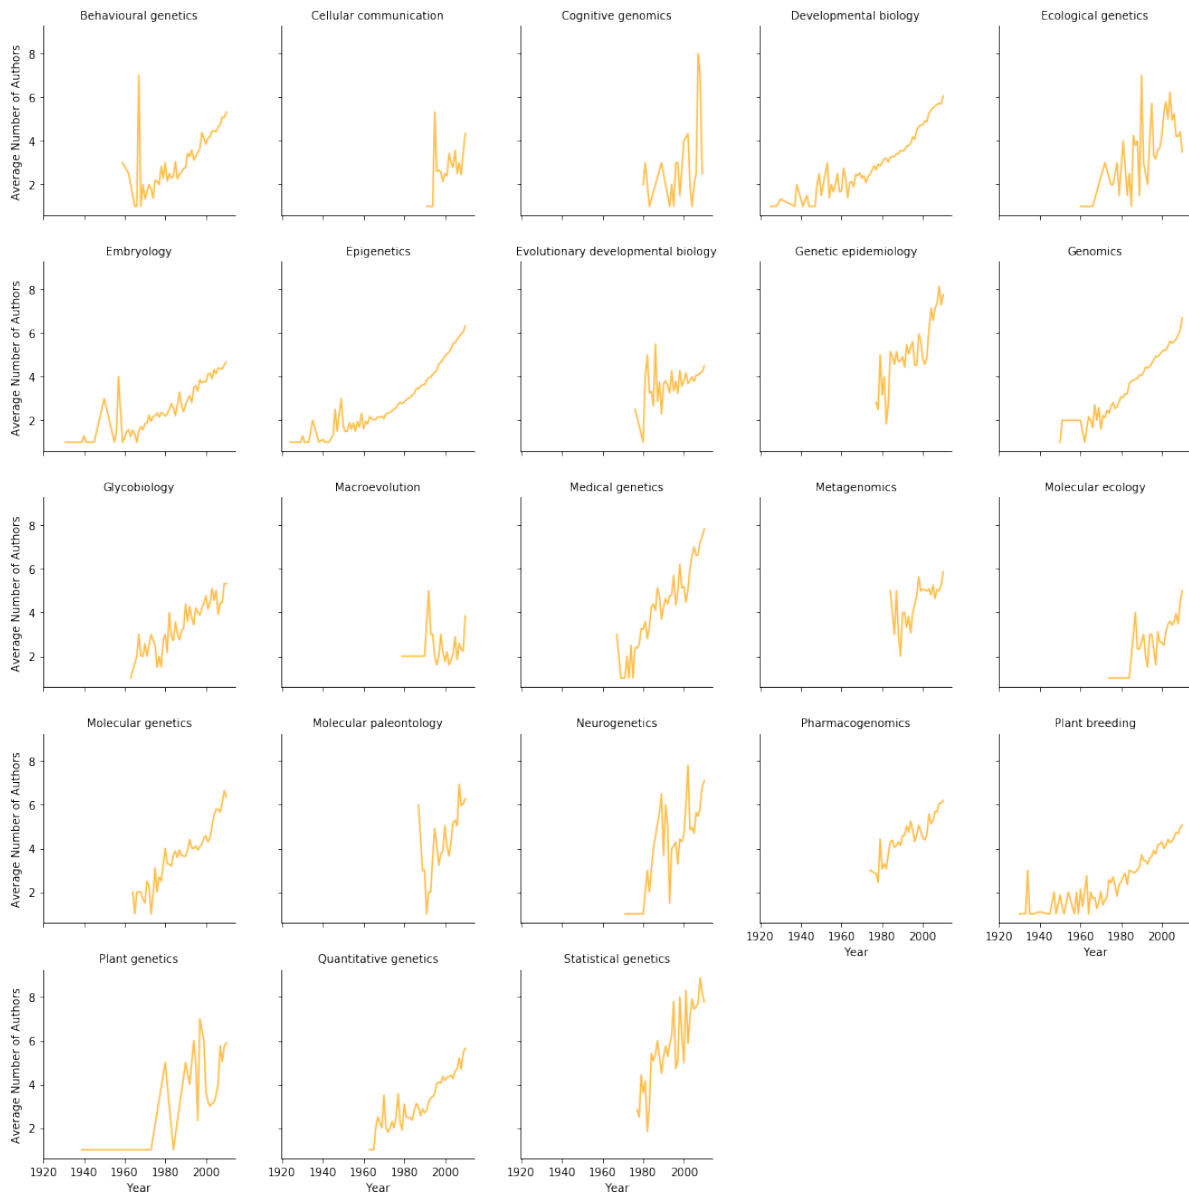
Figure S28. Genetics L2-Subfields Number of Papers over Time. We can observe a big variance in the number of papers over time in the various genetics subfields.
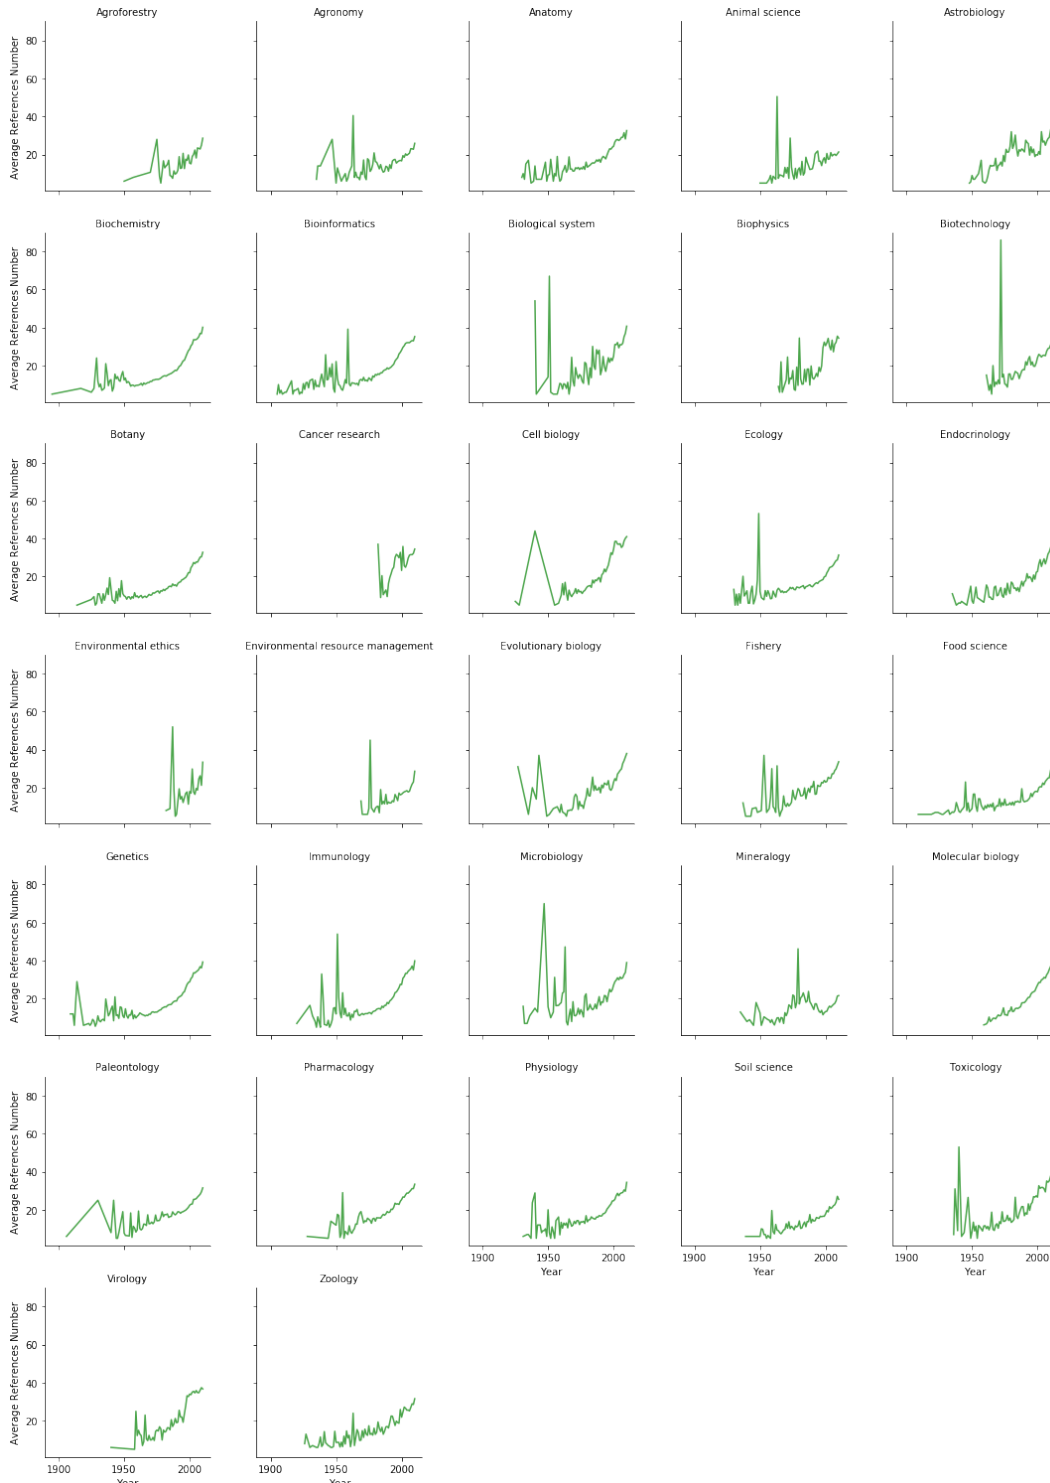
Figure S29. Biology L1-Subfields Average Number of Authors over Time. We can observe a variance in the average number of authors over time in the various biology subfields.
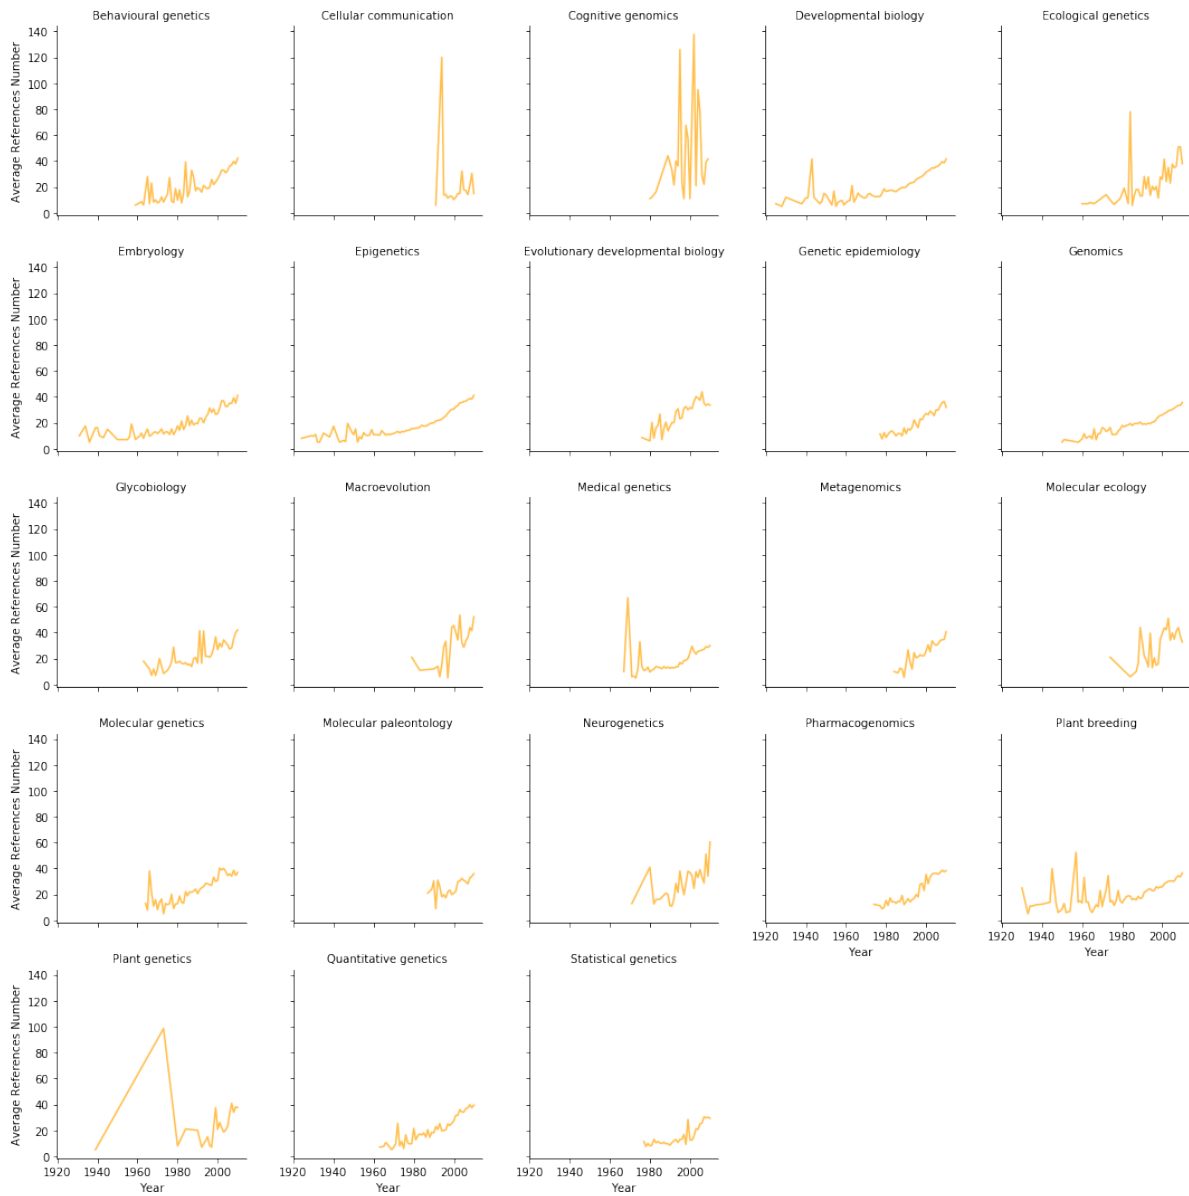
Figure S30. Genetics L3-Subfields Average Number of Authors over Time. We can observe a significant variance in the average number of authors over time in the various genetics subfields.

Figure S31. Biology L1-Subfields Average Number of References over Time. We can observe a variance in the average number of references over time in the various biology subfields.

Figure S32. Genetics L2-Subfields Average Number of References over Time. We can observe a significant variance in the average number of references over time in the various genetics subfields.
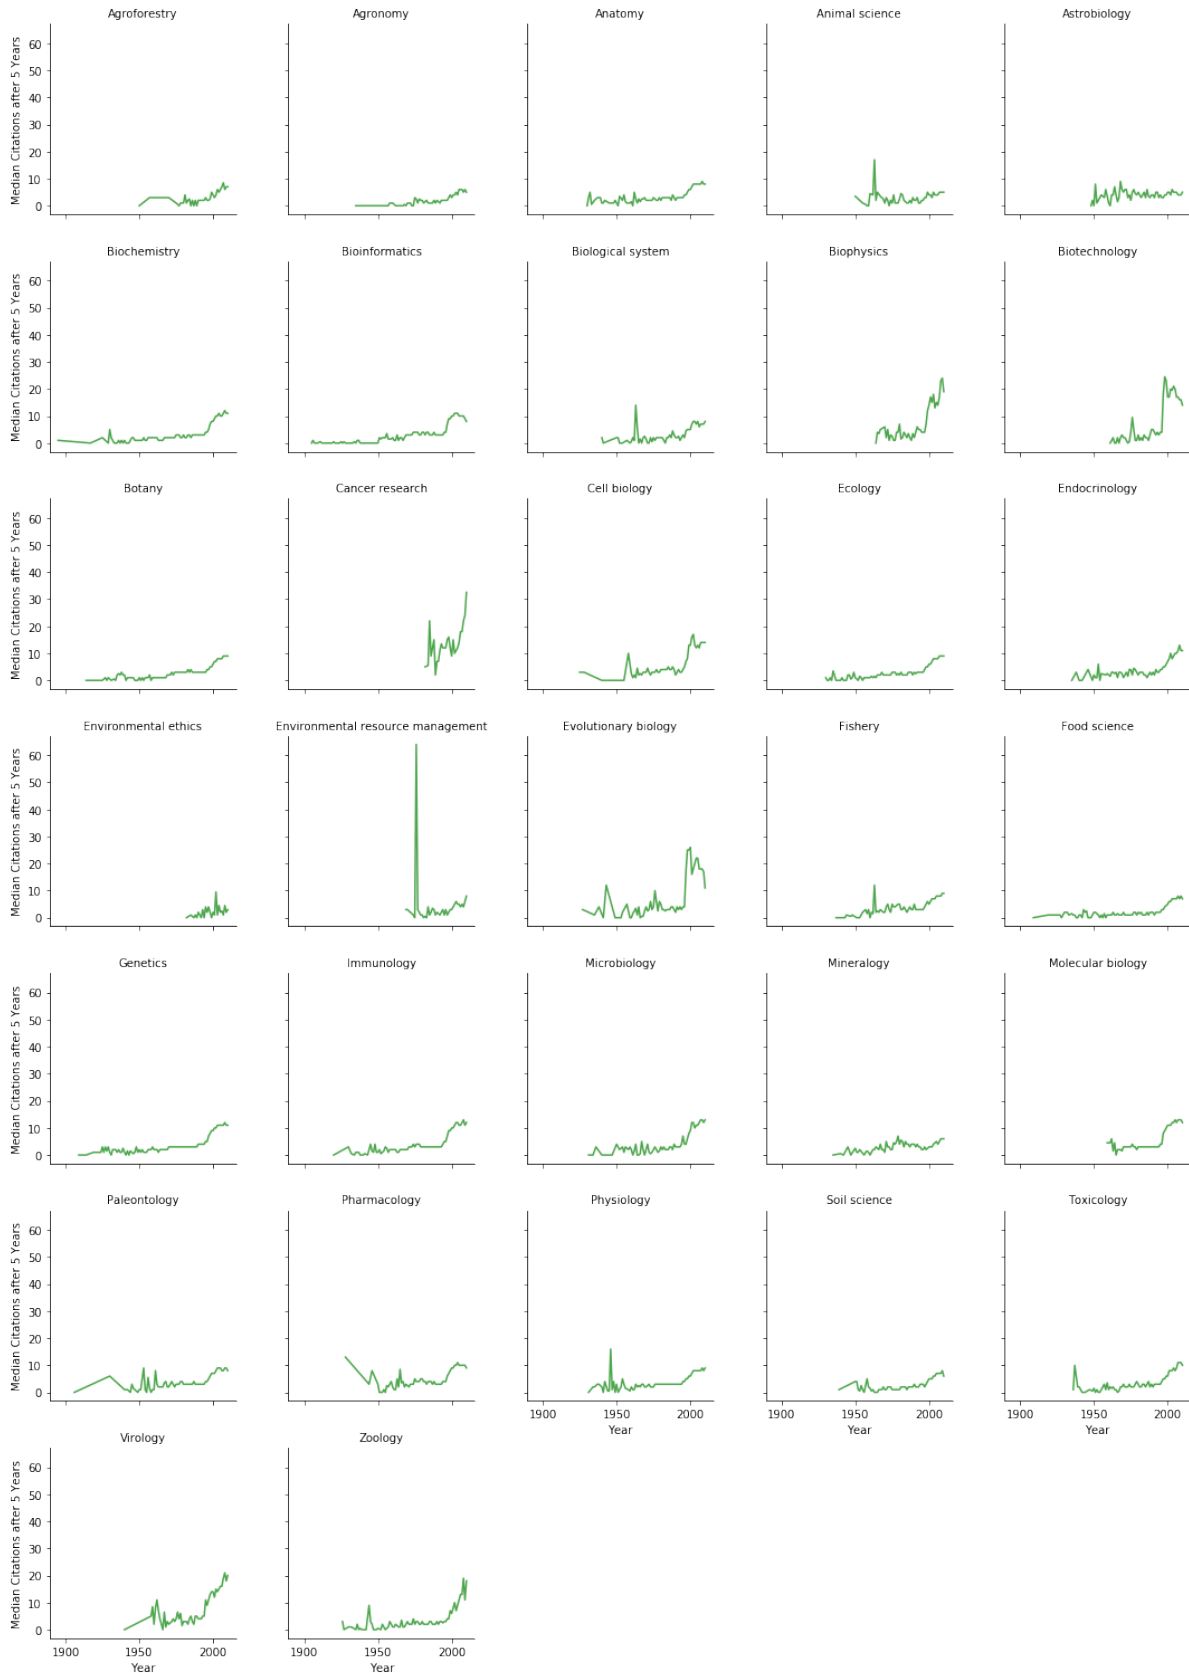
Figure S33. Biology L1-Subfields Median Number of 5-Year Citations over Time. We can observe a variance in the median number of citations over time in the various biology subfields.
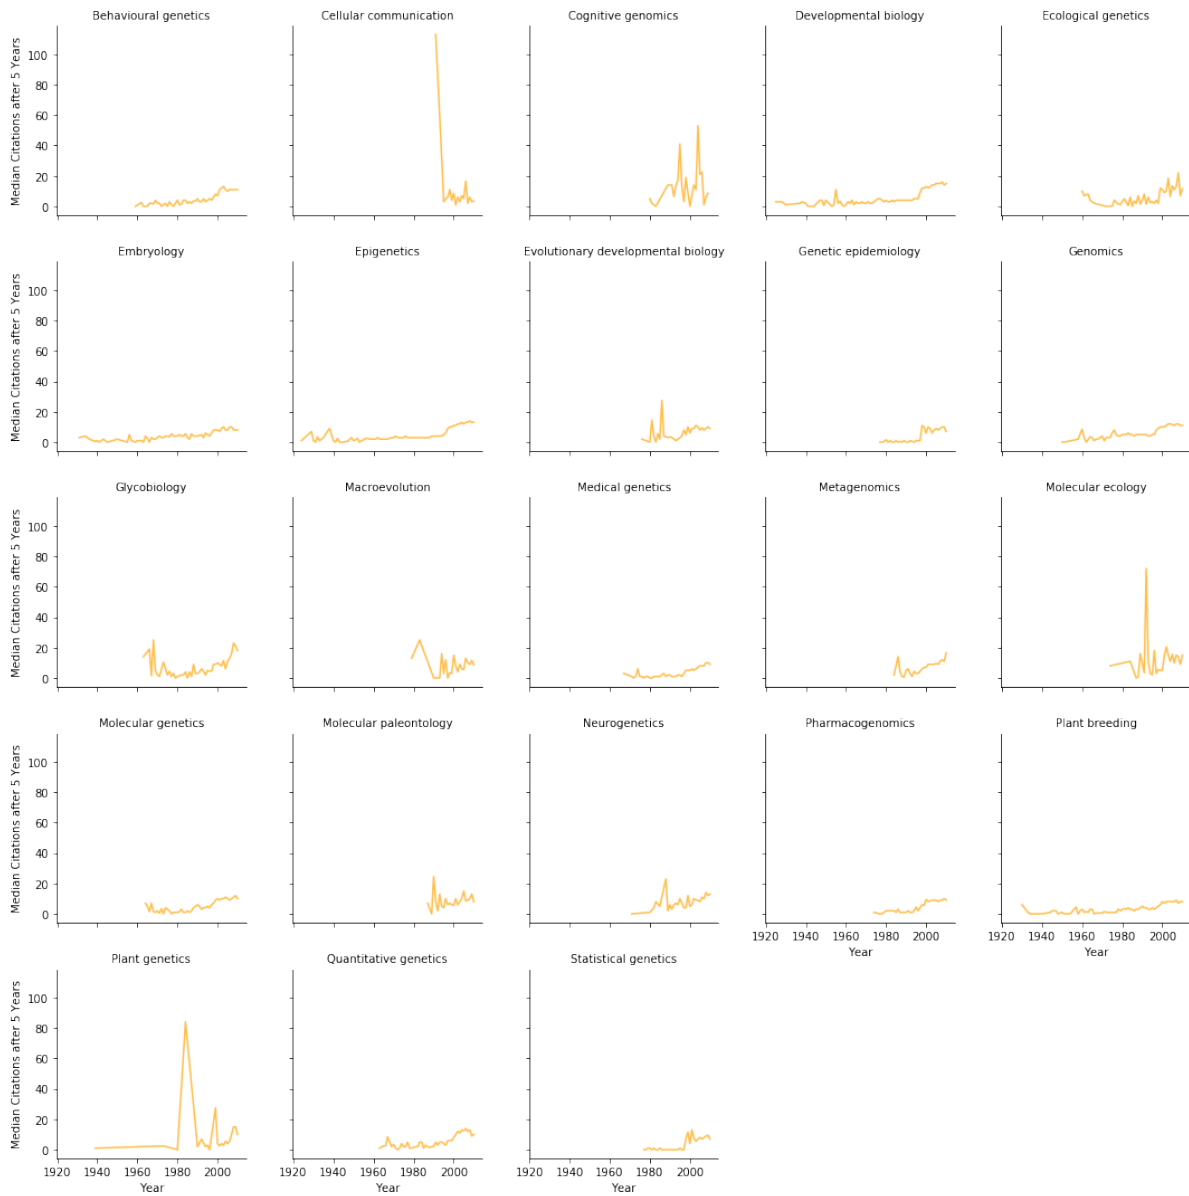
Figure S34. Genetics L2-Subfields Median Number of 5-Year Citations over Time. We can observe a significant variance in the median number of citations over time in the various genetics subfields.

Figure S35. L3 Fields-of-Study Median 5-Year Citation Distributions by Parent Fields. We can observe the high variance among the L3 fields-of-study median citation numbers.
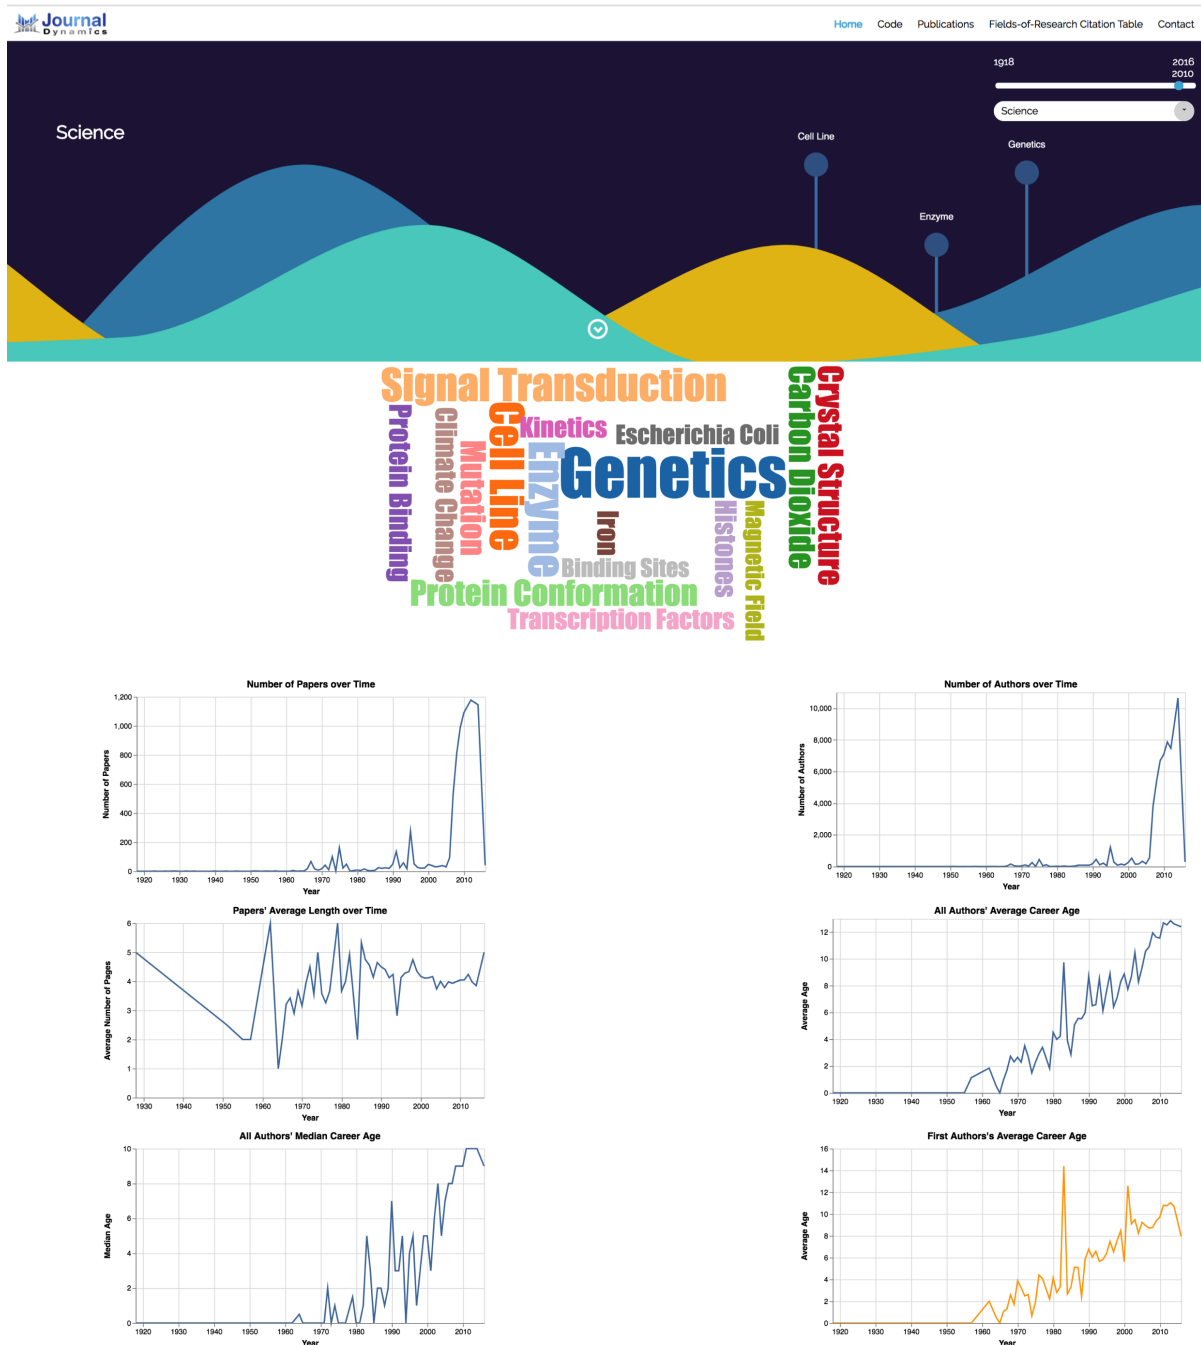
Figure S36. Interactive Website. We have developed an interactive website that makes it possible to view and interact directly with the study's data.