



RESEARCH ARTICLE

Individual and gender inequality in computer science: A career study of cohorts from 1970 to 2000

Haiko Lietz¹ , Mohsen Jadidi^{1,2}, Daniel Kostic¹ , Milena Tsvetkova³ ,
and Claudia Wagner^{1,4} 

¹Department of Computational Social Science, GESIS—Leibniz Institute for the Social Sciences, Cologne, Germany

²Institute for Computer Science, University of Koblenz, Koblenz, Germany

³Department of Methodology, The London School of Economics and Political Science, London, UK

⁴Department of Society, Technology, and Human Factors, RWTH Aachen University, Aachen, Germany

an open access  journal



Citation: Lietz, H., Jadidi, M., Kostic, D., Tsvetkova, M., & Wagner, C. (2024). Individual and gender inequality in computer science: A career study of cohorts from 1970 to 2000. *Quantitative Science Studies*, 5(1), 128–152. https://doi.org/10.1162/qss_a_00283

DOI:
https://doi.org/10.1162/qss_a_00283

Peer Review:
https://www.webofscience.com/api/gateway/wos/peer-review/10.1162/qss_a_00283

Supporting Information:
https://doi.org/10.1162/qss_a_00283

Received: 3 January 2022
Accepted: 28 August 2023

Corresponding Author:
Haiko Lietz
haiko.lietz@gesis.org

Handling Editor:
Vincent Larivière

Copyright: © 2024 Haiko Lietz, Mohsen Jadidi, Daniel Kostic, Milena Tsvetkova, and Claudia Wagner. Published under a Creative Commons Attribution 4.0 International (CC BY 4.0) license.



Keywords: careers, cohorts, computer science, gender, inequality, Matthew Effect

ABSTRACT

Inequality prevails in science. Individual inequality means that most perish quickly and only a few are successful, and gender inequality implies that there are differences in achievements for women and men. Using large-scale bibliographic data and following a computational approach, we study the evolution of individual and gender inequality for cohorts from 1970 to 2000 in the whole field of computer science as it grows and becomes a team-based science. We find that individual inequality in productivity (publications) increases over a scholar's career but is historically invariant, whereas individual inequality in impact (citations), albeit larger, is stable across cohorts and careers. Gender inequality prevails regarding productivity, but there is no evidence for differences in impact. The Matthew Effect is shown to accumulate advantages to early achievements and to become stronger over the decades, indicating the rise of a “publish or perish” imperative. Only some authors manage to reap the benefits that publishing in teams promises. The Matthew Effect then amplifies initial differences and propagates the gender gap. Women continue to fall behind because they continue to be at a higher risk of dropping out for reasons that have nothing to do with early-career achievements or social support.

1. INTRODUCTION

Half a century ago, Price diagnosed that the science system exhibits an “essential, built-in undemocracy,” meaning that academic achievements are strongly concentrated among a very limited number of persons or organizations. He observed inequality in the form of broad distributions of individual productivity and scientific impact and found this pattern to be stable as science grows, perpetuating a system where a “few giants” coexist with a “mass of pygmies” (Price, 1963, p. 53). The literature has found the broadness of these distributions to be a universal property of the science system (Albarrán, Crespo et al., 2011; Bradford, 1934; Lotka, 1926; Ruiz-Castillo & Costas, 2014) and has identified an endogenous process of reproduction as the main driving mechanism: the Matthew Effect (ME; Bol, de Vaan, & van de Rijdt, 2018; DiPrete & Eirich, 2006; Perc, 2014). In his explanations of advancement in academic careers, Merton (1968, 1988) referred to the ME as a cumulative-advantage process according to which “initial comparative advantages of trained capacity, structural location, and available

resources make for successive increments of advantage such that the gaps between the haves and the have-nots in science ... widen until dampened by countervailing processes" (Merton, 1988, p. 606). The larger the ME, the more "the rich get richer rendering the poor relatively poorer" (Page, 2015, p. 34).

Extreme individual inequality is problematic but could be considered fair if it is merit based (Starmans, Sheskin, & Bloom, 2017). However, for differences in merit and success to be considered fair, they should not be associated with ascribed characteristics such as gender, age, or ethnicity (Cole, 1979; Merton, 1973). Inequality among persons belonging to different groups, also known as horizontal inequality as opposed to individual, or vertical, inequality, is undesirable (Stewart, 2005). For example, gender is a prominent principle of distinction, and gender inequality in scientific productivity has been observed. This is known as the "productivity puzzle." For instance, research from the early days of science studies found that women produce about half as much as men (Cole & Singer, 1991; Cole & Zuckerman, 1984), particularly over the first decade of their careers (Long, 1992; Reskin & Hargens, 1979). More recent large-scale analyses show that each year, women are 20% more likely to drop out of science than men (Huang, Gates et al., 2020). In computer science, women on average publish less than men per year for the first several years of employment (Way, Larremore, & Clauset, 2016). Women are less likely to take prestigious author positions in publications (Holman, Stuart-Fox, & Hauser, 2018; West, Jacquet et al., 2013), yet they are more likely to perform better in the job market (Way et al., 2016). Gender inequality in impact has also been reported (Cole & Zuckerman, 1984; Larivière, Ni et al., 2013; Lincoln, Pincus et al., 2012).

The literature on individual and gender inequality in science is abundant, but we identify two major research gaps in it. The first relates to cohort design and data availability. Older analyses tend to have sound cohort designs but are often restricted in the amount of data (number and size of cohorts) that were studied. For example, Zuckerman and Merton (1972) only analyze one cohort, and Allison, Long, and Krauze (1982) analyze three cohorts. More recent computational analyses tend to study large amounts of data but are often restricted regarding cohort design. For example, Penner, Pan et al. (2013) aggregated scientists who started their careers in the same decade. Petersen, Fortunato et al. (2014a) group authors into one cohort that published their first paper in a competitive journal within the same 15 years. These cohorts are heterogeneous with respect to career age and do not include unsuccessful scientists and early career researchers. Previous research, however, has shown that life-course approaches are important because dropouts can partially explain gender inequality. Specifically, productivity inequality almost vanishes when women and men are compared for the same career ages (Azoulay & Lynn, 2020; Jadidi, Karimi et al., 2018; Huang et al., 2020), although, even when the survival bias is removed, women still have fewer publications than men when they become a professor (Aksnes, Rorstad et al., 2011; Lutter & Schröder, 2016).

The second research gap relates to recognizing the scientific field's growth and transformation. Price (1963) argued that recruiting more people into science implies that less talented people will enter. Zuckerman and Merton (1972) hypothesized that this leads to larger differences between the most and least talented, suggesting that inequality should be higher in more recent cohorts than in older ones. Early work on chemistry cohorts found inequality in productivity (publications) and scientific impact (citations) to increase as a cohort ages (Allison et al., 1982). Yet, using full-scale bibliographic databases, scholars found impact inequality decreases over time (Larivière, Gingras, & Archambault, 2009; Pan, Petersen et al., 2018; Petersen & Penner, 2014) as the academic system transitions from a scholar-centered to a globalized, interdisciplinary, team- and project-based mode of knowledge production (Gibbons,

1994). Both findings are plausible and can be explained by changes in the academic system: The increased tendency to publish papers with multiple authors (Petersen, Pavlidis, & Semendeferi, 2014b; Wuchty, Jones, & Uzzi, 2007) may function as a social multiplier that potentially increases inequality, whereas a higher number of references per paper decreases the number of uncited papers, which may decrease inequality (Pan et al., 2018; Wallace, Larivière, & Gingras, 2009).

In this paper, we take a cohort-design approach to the problem of individual and gender inequalities in academia and their origins. Using bibliographic data on the whole field of computer science, we define cohorts from 1970 to 2000 and study the careers of authors over 15 years. Computer science presents an ideal case study because we can observe it since its early days as it grows and evolves from an individual-based to a team-based science. The field is relatively young, growing, in ongoing transformation, and a driver of the digital revolution. Last but not least, it concerns potentially large gender disparities because only one out of five computer scientists is female (Lee, Karimi et al., 2019). We find that individual inequality in productivity is slightly increasing over academic careers. In contrast, individual inequality in impact is stable. These trends are invariant as the field grows and matures. Gender inequality exists, but impact inequality finds an explanation in productivity inequality, which is a result of higher dropout rates for women. The ME is shown to increase historically. Over the decades, we expose the emergence of an imperative to “publish or perish” and the citation-based consequences of early-career achievements as well as early-career social capital. By shedding light on the mechanisms behind individual and gender inequality we motivate science policy interventions to mentor women in social networking.

In the next section, we distill from the literature an evolutionary theory of careers in competitive fields, with the ME at its center. This theory guides our analysis. Then, we present our research design, discuss our results in detail, and conclude our work. For readability, materials and methods are placed at the end (Section 6).

2. THE MATTHEW EFFECT IN THE CENTER OF A THEORY OF CAREERS

The ME is a feedback mechanism that generates inequality. On the one hand, the ME implies *cumulative advantage*. For example, getting a more prestigious job entails an increase in productivity (Allison & Long, 1990). Departmental prestige helps careers because prestige operates and reproduces in networks. As a scholar climbs up the career ladder, she or he advances into the core of a field and becomes part of a reproductive vortex that makes it increasingly hard to *not* benefit from collective dynamics (Burris, 2004; Clauset, Arbesman, & Larremore, 2015; Way, Morgan et al., 2019). Cores harbor the few positions that strongly influence how a field reproduces (Fuchs, 2001). Padgett and Powell (2012) introduce the concept of autocatalytic feedback to model these dynamics.

Inversely, the ME also takes the form of a *cumulative disadvantage*. This has sustained the hypothesis that success either comes early or not at all (Zuckerman & Merton, 1972). As a consequence, the ME makes it increasingly difficult for an individual to stay in academia (Cole & Cole, 1973). Young scientists must overcome a “barrier” to excel (Petersen, Jung et al., 2011). If positive feedback does not set in early in a career, the respective scholar requires motivation to be productive for the love of the work or some amount of tenacity (Huber, 2002). Surprisingly, though most computer scientists are most productive in their fifth year after hiring, there is a huge variance in productivity career patterns (Way, Morgan et al., 2017). And success can come at any time in a career, but it depends on persistence, ability to excel, and, last but not least, luck (Sinatra, Wang et al., 2016).

Cumulative advantage and disadvantage both imply that past achievement to some extent predicts current achievement. Thus, empirical research on the ME typically quantifies the size of the effect and even attempts to establish a scaling law (Jeong, Nédá, & Barabási, 2003; Perc, 2014; Ronda-Pupo & Pham, 2018). Career reinforcement via the ME manifests as increasing returns to the average number of citations per paper as an author becomes more productive (Costas, Bordons et al., 2009). For highly cited authors, staying in academia twice as long means being up to 2.8 times more productive and being up to eight times more impactful. Below a certain citation threshold, the ME operates via the author's reputation as measured by their cumulative citation record, but above that threshold, mainly via publication visibility (Petersen et al., 2014a). Overall, studies predicting the success of scholars or publications have found that current productivity and impact (Acuna, Allesina, & Kording, 2012; Dong, Johnson, & Chawla, 2015; Mazlounian, 2012; Penner et al., 2013), combined with an intrinsic "fitness," or ability and quality (Wang, Song, & Barabási, 2013), and mediated through networks (Sarigöl, Pfitzner et al., 2014) are positively correlated with future success. The observation that the early career of a scientist is predictive of her or his later success and gains in predictive power diminish as more career ages are used for prediction provides further evidence for the ME (Mazlounian, 2012; Penner et al., 2013; Wang et al., 2013).

In sum, the ME has become central to an empirically oriented evolutionary *theory of careers* in competitive fields that is taking shape at the intersection of the social and computational sciences. It is a *field theory* (Bourdieu, 1988) because the academic fields, as spacetimes that delimit agents' social positions and interactions, are the loci that harbor the ME (White, Owen-Smith et al., 2004). Emerging from collective action, field structure acts as a memory in which advantages accumulate and lead to institutionalization (Flack, 2017; Pan et al., 2018; Petersen & Penner, 2014). This field-endogenous feedback process operates behind (i.e., it reinforces or impedes) life-course factors such as creativity, self-perceptions, dispositions, access to resources, and environmental conditions (Cole & Cole, 1973; Cole & Singer, 1991; Padgett & Powell, 2012). *Competition* for ideas, positions, and funds ensues. Careers are tournament-like endeavors (Sørensen, 1986) to improve one's rank in the academic "pecking order" (Chase, 1980). Ranks translate to positions in networks, and upward or downward mobility resembles approaching or withdrawing from network cores (Burris, 2004; Clauset et al., 2015). Only a few make it up those "chains of opportunity"; for most, the way is down (White, 1970). As an *evolutionary theory*, it looks for path dependence and the long-term consequences of initial conditions (Cole & Singer, 1991; Wray, 2011). Small differences in ability, persistence, or luck accumulate and lock a career into an upward or downward path (Petersen, Riccaboni et al., 2012; Way et al., 2019). As this is a collective phenomenon, good ideas can fail if they are put forth at the "wrong time" (Bornholdt, Jensen, & Sneppen, 2011; Newman, 2009), but if the time is "right," success breeds success in an avalanche-like way (Mazlounian, Eom et al., 2011).

This theory also prepares the ground for understanding gender inequality as cogenerated by the ME (DiPrete & Eirich, 2006; Long & Fox, 1995). Some or many of the career factors exemplified above are likely to be gender correlated and thus generate outcome differences that increase over the career as they interact with the ME (Cole & Singer, 1991; Xie & Shauman, 1998). For example, absence from the job market (e.g., because of motherhood) leads to disadvantages that accumulate (Cole, 1979; DiPrete & Eirich, 2006). And women's disadvantages grow early in a career (Long, 1992; Reskin & Hargens, 1979).

3. RESEARCH DESIGN

We adopt an integrated modeling approach to study individual and gender inequality in an academic field. By "integrated" we mean that we are interested in both explaining and

predicting inequality (Hofman, Watts et al., 2021). We study 15-year careers in the entire computer science discipline for cohorts from 1970 to 2000. Starting with descriptive modeling, we explore the evolution of individual and gender inequality in productivity and impact over the career within cohorts and between cohorts over time. In an explanatory modeling step, we then present the ME as a plausible mechanism that generates the patterns of individual inequality we observe. Finally, in a predictive modeling step, we inquire how accurately the early career predicts total-career achievements. We identify the meritocratic and nonmeritocratic early-career factors that predict whether an author drops out of the field and how successful they become eventually. Explanations of individual and gender inequality then derive from the assumption that the ME accumulates the early advantages from these career factors.

As the main data source, we use DBLP, a comprehensive collection of computer science papers that were published in major and minor computer science outlets (Ley, 2009). We study cohorts from 1970 to 2000, where an author belongs to a cohort if they have published their first paper in the given year. For each cohort, we study careers over 15 years, including the start year. We measure productivity in terms of the number of publications because those are the vehicles of academic communication (Merton, 1968) and scientific impact in terms of the number of citations, a widely used measure (Aksnes, Langfeldt, & Wouters, 2019; Merton, 1988). For the details of our methods, we refer to Section 6 at the end of the paper. Selected results obtained from the DBLP data set (Jadidi et al., 2018; Way et al., 2016) were reported above in Section 1. Our cut of the DBLP data set consists of 1.9 million publications from 1970 to 2014 that are authored by 1.1 million authors. Of those, about 300,000 authors started their careers between 1970 and 2000 and are counted as cohort members. There are 6.6 million citations among those publications, which we use for the impact analyses.

Figures 1A and 1B show that cohorts grow exponentially with time and that the field is becoming a team science in the process. Individual inequality at the most aggregate level (all publications and citations accumulated over an author's 15-year career, aggregated for all cohorts) is depicted via broad probability distributions. The citation distribution is broader than the productivity distribution; that is, inequality in impact is larger than inequality in productivity (Figure 1C). Correspondingly, the Gini coefficient, our measure of individual inequality, is larger for impact (0.83) than for productivity (0.68). This is not surprising because authors are physically constrained about the number of projects they can work on during any year but




Figure 1. Description of the field of computer science. (A) The size of cohorts increases exponentially with time for both males and females. (B) The average team size, measured by the number of authors per paper, increases over time. (C) Distributions of productivity (cumulative number of papers P per author at career age 15) and impact (cumulative number of citations C per author at career age 15) are broad. The lines are best fits to the data: a truncated power law ($P(15)$) and a stretched exponential ($C(15)$). (D) The number of authors decreases with the number of years during which they publish persistently after the beginning of their careers (early career persistence). Female scientists show equal persistence in their early careers but after 4 years they are less likely to persist. (E) The fraction of authors in a cohort that drop out of academia (for 10 years in a row) decreases but is more or less constant since the mid-1980s. Women drop out more than men.

there are no such restrictions when it comes to the number of citations their work receives. The last two plots show early career *persistence*, that is, the number of career years during which an author publishes consecutively from the beginning of their career, and the dropout rate per cohort. Most authors persist for only 1 year before they become inactive (for at least a year) or drop out of computer science. Long persistence is decreasingly likely, especially for female scientists (Figure 1D). Dropout rates decrease for subsequent cohorts but women continue to be more likely to drop out than men (Figure 1E).

4. RESULTS

4.1. Individual Inequality Over Careers and Cohorts

In the first, *descriptive*, modeling step, we explore the evolution of individual and gender inequality regarding productivity and impact. If the ME is in place, how would inequality change over the career? Intuitively, one might expect that inequality should increase if the rich get richer and that an increase in productivity inequality should directly translate to an increase in impact inequality. This is what Allison et al. (1982) find in their aforementioned study of the chemistry cohorts from the 1950s and 1960s. But they also find that the method of counting publications and citations—window vs. cumulative counting—is decisive. They find stable impact inequality for cumulative counting; increases are found only for window counting. Here, we report results using cumulative counting but also include plots with 3-year window counting in the Supplementary material. Our measure of individual inequality is the Gini coefficient.

For cumulative counting, we find that productivity inequality is increasing over career years (Figure 2A) but impact inequality is larger but mostly stable after an initial decrease (Figure 2B). We study several modifications to validate this finding. The change around career year 4 that can be seen in almost all figures is because the career of an author starts with the first publication (i.e., in career year 1 every author has at least one publication, but even those authors that eventually become highly cited may still have zero citations). As we saw in Figure 1E, many authors drop out of academia early on, but their publication and citation counts influence the Gini coefficients. We introduce the convention that we observe a *dropout* if an author is absent for at least 10 consecutive years. When we remove dropouts¹ (Figure 2E), author careers are more comparable and productivity inequality drops, but a trend of increasing inequality remains. When this filter is applied to measuring impact, the inequality level also drops but the trend toward stable inequality does not change (Figure 2F).

In computer science, the order of authors is typically important. The first author usually did the most valued part of the work. Hence, in our analysis, attributing publications only to first authors serves the purpose of studying scholars of heightened importance². When we add the first-author filter to the removal of dropouts, the trend for increasing productivity inequality completely vanishes (Figure 2G), and the trend for stable impact inequality remains (Figure 2H). When we employ window counting, the Gini coefficients are systematically higher and the trends less pronounced due to the ceiling effect, but qualitatively almost the same. A difference is, though, that productivity inequality still increases slightly over a career when dropout and first-author filters are applied (Figure S1 in the Supplementary material). In sum, when all authors are considered, productivity inequality is increasing over career ages while impact inequality,

¹ Results are qualitatively similar for absences of five and 10 consecutive years.

² In our cut of the DBLP data set, 69% of all publications have author lists that are not alphabetically sorted. As an author ranking by importance can be alphabetic by chance, the fraction where the author ranking is indicative of importance will be even higher.




Figure 2. Individual inequality in productivity and impact as a function of career ages, depicted for seven cohorts between 1970 and 2000. We count publications and citations cumulatively ($P(t)$ and $C(t)$, defined in Section 6.4). First two columns: assigning publications to all authors. Last two columns: assigning publications only to first authors. Second row: authors are filtered that have not published for 10 consecutive years (most likely left academia).

albeit larger, is mostly stable. Rising inequality in productivity is an effect of considering the full workforce and disappears for comparable authors, at least for cumulative counting.

Now turning to the historical analysis of cohorts, we address Zuckerman and Merton's (1972) hypothesis that recruiting more people into science will lead to larger differences between the most and the least talented. Our results do not entirely support this hypothesis: We do not see a remarkable increase in inequality over cohorts for impact, but we observe an upward trend for productivity (Figures 3A and B). Removing dropouts reduces inequality levels but no longer softens the increasing trend for productivity (Figure 3E). Also considering first authorships preserves all trends, this time also for productivity. The increase in productivity inequality does not vanish when only comparable authors are considered (Figure 3G). The results are not as evident for window counting of publications and citations because the Gini coefficients are much closer to 1 (Figure S2 in the Supplementary material). In sum, though the field has grown exponentially, similar levels of impact inequality can be observed for authors that started their careers in 1970 and 2000. When it comes to productivity, however, inequality appears to be increasing over time in parallel to the field's transition from an individual- to team-based science.

4.2. Gender Inequality Over Careers and Cohorts

Increasing individual inequality in science is not necessarily problematic if the evaluation is based solely on merit rather than on functionally irrelevant factors such as gender, race, nationality, age, or class. Due to its societal importance, we focus on gender inequality.




Figure 3. Individual inequality in productivity and impact as a function of cohorts, depicted for career ages 3, 5, 10, and 15. We count publications and citations cumulatively ($P(t)$ and $C(t)$, defined in Section 6.4). First two columns: assigning publications to all authors. Last two columns: assigning publications only to first authors. Second row: authors are filtered that have not published for 10 consecutive years (most likely left academia).

Figure 4 shows a systematic comparison of the cumulative productivity and impact distributions of male and female computer scientists with the same filters applied as in the previous figure. Positive values (red) indicate that the distribution of men is dominant, that is, men are more productive or their work has a higher scientific impact. Negative values (blue) reveal that the distribution of women is dominant.

There is a general pattern for gender inequality in productivity (Figure 4A): It seems to accumulate and is more prevalent in the later career stages. If there are differences in productivity it is always men publishing more. This gender productivity gap exists in almost all cohorts. For gender inequality in impact, the picture is less clear (Figure 4B). Female and male dominance both exist sporadically in cohorts. In four cohorts, women are statistically more likely to have more citations than men; for the 1982 cohort even for 10 consecutive career years. There is no cohort in which gender inequality shifts signs, which means, it is always one cohort's gender that is dominant. In total, gender inequality is more pronounced for productivity than for impact. For cumulative numbers of publications, 55% of 465 cohort-age pair differences are statistically significant; for cumulative numbers of citations, 19% are significant. That means, the productivity gap does not automatically translate into an impact gap. However, whenever there is an impact gap, it can be explained by a productivity gap: Significant differences in citation are strongly correlated with differences in publications ($r = 0.91$; $p \leq 0.001$). That means, as Azoulay and Lynn (2020) found, the productivity gap is the puzzle to solve.

When we limit authors to first authors, we get a step closer to solving this puzzle (Figure 4C): The magnitude of the gender gap becomes smaller (only 28% of cohort-career year pairs are




Figure 4. Gender inequality for productivity and impact as a function of cohorts and career ages. We compare the cumulative publications distribution $P_{\text{gender}}(t)$ and cumulative citations distribution $C_{\text{gender}}(t)$ of male and female scientists in the same cohort at the same career age t and test differences between these distributions. Color marks the effect size (Cliff's d). Positive values (red) indicate that men dominate women, and negative values (blue) reveal that women dominate men. Effects are only shown if they are significant ($p \leq 0.05$) according to a Mann–Whitney U test. Details in Section 6.5. Publications are assigned to all authors (A, B) or first authors only (C, D). In general, effects decrease with cohort and increase with career age.

significantly different). This is particularly the case for the more recent cohorts. The observation that larger differences in productivity between male and female scientists diminish when only first-author contributions are counted suggests that, as team sizes increased in computer science, male scientists boosted their productivity more via collaborations than female scientists. Applying the first author filter makes the impact gap a purely male phenomenon but also a phenomenon of the 1970s (Figure 4D).

Accounting for dropouts removes any pattern (Figures 4E to H). Neither does gender inequality increase with career ages nor does it persist on the historical scale or single out any gender. Any significant inequality is likely just noise. In sum, gender inequality exists. Although it appears to be diminishing on the timescale of cohorts, it is more persistent on the career timescale. Importantly, however, gender inequality practically disappears in recent cohorts when authors with comparable careers are studied.

4.3. The Role of the Matthew Effect

In the *explanatory* modeling step that now follows, we inquire if reproductive feedback operates in the field as an underlying mechanism and to what extent it can generate the patterns of individual inequality we observe. The ME states that present achievement (productivity or impact) depends on past achievement and that resulting advantages can accumulate over time. Our guiding theory describes this feedback process as a vortex, an autocatalytic mechanism that fuels itself (Padgett & Powell, 2012). When the ME is fully operational—formally, when it is linear—it generates power law distributions that signal the absence of a

characteristic scale (Albert & Barabási, 2002). In our case of computer science, productivity and impact distributions are broad but not pure power laws. Distributions for individual cohorts are much like the (truncated power law and stretched exponential) distributions that we measure when all cohorts are lumped together (Figure 1C). These deviations can result from a damped (sublinear) ME and other mechanisms and factors that interact with the ME but also from sampling and finite-size effects intrinsic to the DBLP database.

We quantify the strength of the reproductive feedback of the field that a cohort experiences in a career age by regressing the number of publications or citations in a career age on the corresponding cumulative number in the previous career age (details in Section 6.6). We interpret two parameters. The *exponent* of the scaling relationship quantifies the strength of reproductive feedback. An exponent that is larger than zero over time is indicative of a cumulative advantage. The *lower cutoff* states at and above which number of publications or citations the advantage accruing from past selection unfolds. It resembles the boundary of the basin of attraction of the feedback dynamics: Once an author crosses it, she or he gets attracted by the reproductive vortex and advantages can accumulate. Examples of the fitting procedure are depicted in Figures 5A and F. They show that scaling relationships are plausible fits to the data.

Our results show that the ME is a plausible explanation for productivity and impact inequality because all exponents are larger than zero. For an average cohort, the strength of the ME is stable over an author's career, allowing for a constant cumulative advantage. This holds for both productivity and impact, as there are no discernible trends in Figures 5B and G. To enter the productivity basin of attraction (i.e., to reap benefits), an author must produce a certain number of publications that is constant over career ages (nondiscernible trend in Figure 5C). However, getting one's publications cited becomes increasingly difficult as careers




Figure 5. Matthew Effect. A, F: measurement of the strength of a cohort's reproductive feedback as the exponent that relates an author's number of papers produced, or citations received, in a career age (y-axis) to the respective cumulative numbers in the previous career age (x-axis), shown for the 2000 cohort and the last career age. Exponents show as slopes of the continuous lines on the log-log plot. Dotted lines indicate that feedback fully unfolds only above a lower cutoff. B, G: for an average cohort, potential individual advantages from feedback are constant along the career path, for both productivity and impact. C, H: for an average cohort, the number of citations required to take advantage of feedback increases along the career path. D, I: for an average career age, potential individual advantages from feedback increase historically, but more so for productivity. E, J: for an average career age, the numbers of citations and publications required to take advantage of feedback increase historically. All columns but the first: shaded areas are bounded by minimal and maximal, lines show means.

progress as the lower cutoff increases with career age (Figure 5H). In other words, regarding productivity, it is equally possible for an early- and late-career author to benefit from autocatalytic feedback, but regarding impact, moving early is advantageous.

Although the strength of the ME is stable over a computer scientist's career, it does increase at the historical timescale of cohorts. Nowadays, the ME is strong for both impact and productivity (the exponents in Figures 5D and I are ≈ 1 for the 2000 cohort start year). However, whereas the 1970 cohort already experienced a strong effect from past citations (exponent ≈ 0.8), the effect of the past number of publications started weak (≈ 0.3). In other words, although getting cited has long been endowed with a strong reinforcement effect, increasing returns for productivity became prominent only recently. At the same time, the lower cutoff for reinforcement to set in has been growing historically, particularly so for productivity (Figure 5J). As the field grew and transitioned towards team-based science, this is likely the result of how limited resources get distributed among an increasing number of scholars. For the authorship practice, this mechanism has a name: "publish or perish" (Garfield, 1996).

There is a correspondence between the ME and the inequalities described in Figures 2 and 3: The ME is persistently stronger for impact than for productivity and thus generates individual inequalities that are persistently higher, over both careers and cohorts³. Looking at cohorts for average career ages, a large cumulative advantage for productivity corresponds to a modest increase in inequality, and a modest cumulative advantage for impact corresponds to stable inequality. For careers, there is no meaningful correspondence. We suggest that this is because, for a certain year, the ME is always computed for authors that have been active in that year; that is, that they have either published or got cited. This means that, in Figures 5B, C, G, and H, the values for small career ages result from all authors and the values for large career ages tend to result from authors with dropouts removed. In sum, the ME increases, strongly so for productivity where it resembles the imperative to "publish or perish." Reproductive feedback is stronger and creates larger inequality for impact than for productivity but, alone, is not capable of explaining the individual inequality patterns we observe.

4.4. The Role of Gender

Following the leads from the descriptive and explanatory analyses, we proceed with the final analytical step: out-of-sample *predictions* of dropout and future success. We study the effect of the cohort, gender, and variables from two classes of constructs. First, our theory of careers states that inequality results from differences in the early career that are then amplified by the ME, which we have found to operate. We aim to gain more insights into how the ME operates by using a set of variables on the early-career achievements of authors. Second, we study inequality as the field becomes a more team-based science. We have found that it is historically increasingly important to produce publications to benefit from the ME. Hence, we use a set of variables on social support which, we hypothesize, makes it easier to write papers. The constructs are fully described and operationalized in Section 6.7 and summarized in Table 1. We use regression models with standardized variables so that the coefficient size indicates the relative importance of the predictor. The models also employ variable selection. When new variables are added to the model, this can result in decreasing effect sizes for previously considered variables. If that happens, there is collinearity among the variables and the variable that predicts better has a larger coefficient.

³ As we have measured the ME cumulatively we do not discuss a correspondence with the window-counted inequalities.

Table 1. Independent variables used in the prediction models. The variables characterize authors in their early career ages $[1, t_e]$, with the end of the early career chosen to be $t_e = 3$. The variables are used to predict author dropout and success. Details are given in Section 6

Variable	Description
Baseline	
Cohort	Year in which cohort members started publishing
Gender	
Male	Dummy
Female	Dummy
Undetected	Dummy
Early achievement	
Productivity	Cumulative number of publications authored in the early career
Productivity (1st author)	Cumulative number of publications authored in the early career as a first author
Impact	Cumulative number of citations received in the early career
Top source	Smallest $h5$ -index-based quartile rank of all journals and conference proceedings an author has published in in the early career
Social support	
Collaboration network	Number of distinct coauthors in the early career
Senior support	Largest h -index of all coauthors in the early career
Team size	Median number of authors of all publications produced in the early career

4.4.1. Predicting dropout

Table 2 shows the results of logistic regression models that use dropout as a binary dependent variable. The most important factor for predicting dropout is early career productivity. Scientists that publish much in the first 3 years of their careers, not necessarily as a first author, are less likely to drop out. This is not surprising, given that dropout is defined as the absence of publications for 10 consecutive career years. Publishing in a top source early on is the second strongest predictor of not dropping out. Having a publication in a top journal or conference proceedings represents symbolic capital, a reputation signal for academic worthiness, which likely influences the career path. These differences in early-career achievements are then amplified by the ME, contributing to the individual inequalities we have diagnosed. Social support has effects that seem contradictory at first glance. On the one hand, coauthoring publications in large teams is positively correlated with dropout. On the other hand, having larger collaboration networks decreases the likelihood to drop out. This tells us that being one author among many is not automatically an achievement. It is the well-connected authors that stay in the field. Having early senior support (a coauthor with a high h -index) has negligible influence on dropping out. Similarly, dropout is not associated with the initial impact of early-career publications. Having many early citations does not make it more likely for an author to stay in computer science. Women are more likely to drop out than an average computer scientist

Table 2. Dropout prediction. Each column corresponds to a separate logistic regression model that predicts whether (1) or not (0) an author dropped out of computer science (described in Section 6.9). Coefficients are reported as means (with standard deviations in brackets) from 10-fold cross-validation over 292,443 observations. Mean coefficient s away from zero indicate that the effects are sizable, and low standard deviations indicate that the effects are robust. Goodness-of-fit measures (F1 and average precision) are also means across all folds

	Model 1	Model 2	Model 3	Model 4
	Baseline	+ Gender	+ Early achievement	+ Social support
Cohort	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Female		0.03 (0.01)	0.06 (0.01)	0.06 (0.01)
Male		−0.06 (0.01)	−0.07 (0.00)	−0.08 (0.00)
Undetected		0.03 (0.00)	0.02 (0.00)	0.02 (0.00)
Productivity			−0.56 (0.00)	−0.54 (0.00)
Productivity (1st)			−0.26 (0.00)	−0.22 (0.01)
Impact			0.01 (0.00)	0.01 (0.00)
Top source			−0.22 (0.01)	−0.21 (0.01)
Collaboration network				−0.08 (0.00)
Senior support				−0.01 (0.00)
Median team size				0.22 (0.01)
Intercept	0.00	0.30	0.00	0.00
F1	0.44	0.44	0.67	0.68
Average precision	0.58	0.60	0.75	0.76

and adding early achievements and social support even increases this effect. Interestingly, the cohort has no effect. These dropout-related patterns are historically invariant. That means that large author teams did not guard against dropping out before and in the team era.

4.4.2. Predicting success

Next, we study which factors predict whether success in terms of scientific impact increases, on average, after the first 3 career years (Table 3). The dependent variable is the increase in citations until career age 15. First of all, the cohort has a small positive effect on success. This can be an effect of the exponential growth of the field: As more publications are produced and reference lists become longer, more citations are made and accumulated (Pan et al., 2018). Early-career productivity is a requirement for success, and its effect is on par with that of early-career impact. This confirms an observation made in the literature, namely that total success is well predictable from early success (Mazloumian, 2012; Penner et al., 2013; Wang et al., 2013). As the ME is path-dependent this is further evidence for cumulative advantage as an underlying mechanism. In contrast to dropout, early senior support is an important factor for success. But similar to dropout prediction, publishing in large teams exhibits a negative effect on citation success, and the size of the early collaboration network has a weakly positive effect. Publishing in a top source is associated with success but gives way to senior support when added to the model.

Table 3. Success prediction. Each column corresponds to a separate linear regression model that aims to predict $C_i^+(15)$, the increase in citations an author gains after the early career of the first 3 career years (described in Section 6.9). Coefficients are reported as means (with standard deviations in brackets) from 10-fold cross-validation over 292,443 observations. Mean coefficients away from zero indicate that the effects are sizable, and low standard deviations indicate that the effects are robust. Goodness-of-fit measures (mean squared error and adjusted R^2) are also means across all folds. In the first four models, there are 292,443 observations. In model 5, dropouts are removed, which causes the number of observations to drop to 119,113

	Model 1	Model 2	Model 3	Model 4	Model 5
	Baseline	+ Gender	+ Early achievement	+ Social support	Dropouts removed
Cohort	0.08 (0.00)	0.08 (0.00)	0.07 (0.00)	0.05 (0.00)	0.04 (0.00)
Female		-0.07 (0.11)	-0.09 (0.01)	-0.10 (0.01)	-0.04 (0.01)
Male		0.01 (0.02)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Undetected		0.00 (0.00)	0.06 (0.01)	0.04 (0.01)	0.03 (0.01)
Productivity			0.99 (0.01)	0.88 (0.01)	0.45 (0.01)
Productivity (1st)			0.44 (0.01)	0.45 (0.01)	0.33 (0.01)
Impact			1.00 (0.01)	0.91 (0.01)	0.37 (0.01)
Top source			0.19 (0.01)	0.00 (0.00)	0.05 (0.00)
Collaboration network				0.05 (0.01)	0.02 (0.01)
Senior support				0.70 (0.01)	0.40 (0.01)
Median team size				-0.09 (0.01)	-0.05 (0.01)
Intercept	-158.94	-157.86	-131.24	-94.95	-77.45
Mean squared error	40.17	40.17	31.95	31.54	7.28
Adjusted R^2	0.01	0.01	0.21	0.22	0.22

Finally, being female is a negative predictor for an increase in citations, and the effect increases when we account for early achievements and social support. This adds to works that have also found gender inequality in impact (Cole & Zuckerman, 1984; Larivière et al., 2013; Lincoln et al., 2012). To check if this effect can again be explained by the gender difference in career persistence, we construct a last model that predicts success with dropouts removed. The result is reported in the last column of Table 3. Removing dropouts makes the absolute values of all effects smaller (except for publishing in a top source). This means that persistent authors are more homogeneous in their characteristics. The markedly smaller error also shows that their success is easier to predict. However, the gender effect does not fully go away: After removing dropouts, women are still slightly less successful at career age 15 than an average computer scientist.

5. DISCUSSION

5.1. Summary and Conclusion

We studied individual and gender inequalities in computer science, their changes over author careers as well as over the field's transformation from a sole-scholar to team-based science, and their origins. We found that individual inequality in productivity increases during the

careers of an average cohort but, contrary to what has been previously suggested (Allison et al., 1982), it does not translate to an increase in impact inequality. The increase in productivity inequality can be a result of comparing scholars with different persistence (all authors vs. only those that kept producing papers) and status (all authors vs. only first authors of a paper), but this explanation is not robust to changing the counting method. The inequality patterns of cohorts from 1970 to 2000 are different from those of chemistry cohorts in the 1960s and 1970s. As computer science exhibits an exponential influx of personnel, we have also checked if it leads to an increase in individual inequality on the historical time scale of cohorts (Zuckerman & Merton, 1972). We found such an effect but only a small one for publications. The inequality patterns for impact are largely the same since the 1970s.

Regarding gender inequality, we found that men produce more publications than women, particularly towards the end of their careers, though this phenomenon was more pronounced in the past. This gender gap, known as the “productivity puzzle” (Cole & Singer, 1991; Cole & Zuckerman, 1984; Long, 1992; Reskin & Hargens, 1979), disappears once we remove drop-outs and focus on first authors only. Regarding reports of gender inequality in impact (Cole & Zuckerman, 1984; Larivière et al., 2013; Lincoln et al., 2012), we found that men having more publications does not automatically entail having more citations, but more citations find their explanation in more publications.

To understand individual inequality, we quantified the ME. We found that it is stable over an author’s career regarding productivity and impact. Although the number of publications above which nonlinear benefits accrue is also stable as the career progresses, this becomes increasingly difficult regarding the number of citations, indicating an early-citation advantage. The ME for publications and citations has increased historically to similar levels, but the climb has been much steeper for productivity, indicating the rise of the imperative to “publish or perish.”

Using regression models to inquire about the importance of early-career achievements and social support in shaping total-career outcomes, we found that early productivity is the best predictor for not dropping out of computer science and for impact-based success. Although staying in the field is a condition for success, success and having successful coauthors are not conditions for staying in the field. As our result for the ME suggested, success turned out to be correlated with early-citation success. Publishing in top journals and conference proceedings in the early career is predictive of staying in the field as well as success. Authors with a large social support network are more likely to stay and be successful, stressing the importance of social capital, but authors that are part of large coauthor collectives are more likely to drop out or remain unsuccessful, potentially because more of the same in terms of team structure hinders creativity (Guimerà, Uzzi et al., 2005; Uzzi & Spiro, 2005). This supports the argument that the transition to team science is a reason why we observe a historical increase in individual inequality for productivity.

Finally, we found that women are not only more likely to drop out of the field but also somewhat less successful after 15 years than an average computer scientist. Both effects cannot be explained with less impressive early-career achievements or lower social support. Whereas gender inequalities could be explained when only comparable authors were analyzed, this also does not explain the impact gender effect: It is softened but does not go away when dropouts are removed. The gender effects can potentially be explained by differences in network structure. We have found in previous work that, on average, “female” collaboration networks are smaller and more cohesive than the networks of their male counterparts (Jadidi et al., 2018). If this type of embedding is a disadvantage, the latter would accumulate due to

the ME. For example, if men manage to have more publication counts than women due to having different social capital (Way et al., 2016), this can explain our finding that the productivity gap between women and men is smaller when only publications authored as first authors are counted.

In conclusion, we have contributed to the reconstruction of the chain of events that results in individual and gender inequality in computer science. Teams become more important. These help scholars increase the number of publications, which serves their career. But being part of a team is not enough, as only some authors manage to reap benefits from team science and the cumulative advantage of early-career achievements. Most scholars drop out of the field, especially women due to reasons we cannot measure. Differences in dropout entail differences in individual and gender inequality. Taken together, in the 45 years we have studied, computer science has increasingly become a competitive field. At the end of a career that works like a tournament, senior female computer scientists even fall behind in terms of citation impact to some extent.

These conclusions in the context of computer science likely carry over to other male-dominated fields that have experienced growth and transformation from individual-based to project- and team-based science, such as statistics, applied mathematics, and engineering. Nevertheless, our findings should be replicated for other academic fields to establish the extent to which the trends in individual and gender inequality we detected are sensitive to the rates of growth, dropout, or entry of women.

Our findings are relevant for science policy measures that aim at more gender equality. The way the field operates on autocatalytic feedback and natural differences between women and men, such as motherhood, but also small behavioral differences in the ways women and men embed into collaboration networks, can have large career consequences. Regarding behavioral differences, our results suggest that mentorship programs with the goal of “broadening and institutionalizing women’s support networks” (Abbate, 2017, p. 175) are promising. It is not unexpected that small policy interventions can have large consequences, too.

5.2. Methodological Considerations

We close our paper with a few methodological considerations. The first set of considerations relates to data. Inquiries into inequality and the ME date back to the 1970s when it was only possible to study small cohorts. Much of this research was done using one of two carefully constructed bibliographic chemistry cohort data sets (Allison et al., 1982). On the contrary, we use a large-scale data set on the complete trails of computer science. The use of bibliographic traces has allowed us to reconstruct and study scholarly careers in historical comparison. Although formal communication just represents the observable aspect of academic careers, it is undeniably an important part, because academic careers are subject to collective field dynamics that work on what is observable.

The ability to model processes with behavioral data comes at the cost of a reduced ability to model individual perception. Gender is an example. Although we maintain that our inferred gender variable is a true gender variable because authors are free to choose which name they put on a paper, the variable does not allow us to differentiate between the various types of socially constructed gender. Hence, we can only contribute insights into the social construction of binary gender. In particular, we show indirectly how a structural mechanism that accumulates achievements—the ME—contributes to generating gender disparities, even after we account for early achievements and social support. Augmenting behavior with data

on cognitive states (e.g., whether computer scientists dropped out on free terms, or because of structural constraints or even discrimination) would allow for deeper insights into the origins of inequality.

Limitations to gender disambiguation from names forced us to remove most Asian authors from the analysis. However, the proportion of Asian scholars in computer science has been increasing since the 1970s as a result of the passing of the 1965 Immigration Act in the United States and the increasing internationalization of science. We are thus missing a larger proportion of one type of authors. Yet, the DBLP data has the converse problem too, because its coverage of publications increases over time. We recognize that both of these biases might be influencing the historical trends we observe, and hence our conclusions should be interpreted with caution.

Operationalizing scientific impact via the number of citations is straightforward because it very well captures that impact is a collective phenomenon. On the other hand, citation scores are not unobtrusive measures anymore. Citations have become a currency in science: Scholars try to improve their scores, and the databases we use for research are also used to compute a scholar's market value. This adds an exploitative angle to the ME that cannot be disentangled from its systemic effects (Clauset, Larremore, & Sinatra, 2017; Xie, 2014). And yet, strategic human behavior is the result of structural constraints as well as incentives and, hence, just as much part of the problem of inequality.

The second set of considerations relates to methods. We followed the integrative modeling approach (Hofman et al., 2021) and found the combination of descriptive, explanatory, and predictive modeling insightful. With large-scale behavioral data, exploratory description is necessary because existing knowledge may not translate into meaningful research questions or hypotheses for testing: Past small-scale studies may not have captured new phenomena, and new large-scale studies may not generalize due to preprocessing decisions and design choices. Still, we let the literature guide our modeling: we distilled a theory of careers from a broad and multidisciplinary set of studies that gives the ME the central role of an inequality-creating mechanism. Multivariate regression models with variable selection then served to shed new light on findings from the first two modeling steps by focusing on early-career factors. This fleshed out the inequality-creating mechanism but also uncovered the gender impact effect even though gender inequality in impact had not been diagnosed. The strong message is that the choice of methods matters and that a systematic mix of methods can produce more robust as well as surprising results.

6. MATERIALS AND METHODS

6.1. Data

We use DBLP (DBLP Team, 2017; Ley, 2009), a comprehensive collection of computer science publications from major and minor journals and conference proceedings. From this dump, we remove *arXiv* preprints. The coverage of DBLP ranges from 55% in the 1980s to over 85% in 2011 (Way et al., 2016). Our data set consists of 1.9 million publications from 1970 to 2014 that are authored by 1.1 million authors. Of those, 292,443 started their career between 1970 and 2000. We have added citations among publications by combining DBLP with the AMiner data set (AMiner, 2017; Tang, Zhang et al., 2008) via publication titles and years. There are 6.6 million citations among publications. Author names in DBLP are disambiguated (Reitz & Hoffmann, 2013).

To infer the gender of authors, we have used a method that combines the results of name-based (genderize.io) and image-based (Face++) gender detection services. The accuracy of this method is above 90% for most nationalities. As the accuracy is very low for Chinese and Korean names, we label their gender as unknown to reduce noise in our analysis (Karimi, Wagner et al., 2016). As authors are free to choose the name under which they publish, the inferred variable is a true, socially constructed gender attribute.

6.2. Cohorts and Career Ages

Our main units of analysis are cohorts of computer scientists from 1970 to 2000. We consider a career to begin with an author's first publication in the database. As DBLP covers publication years back to 1960, this ensures that authors of the earliest cohort have been at least absent for 10 years. Imbalances in coverage over publication years cause earlier cohorts to be less homogeneous as we tend to miss more first publications. Given start years, we follow cohort members over career ages $t \in [1, 15]$.

6.3. Publication and Citation Counts

Our unit of observation is the individual author i in a cohort. For each author and career age, we measure the number of publications $p_i(t)$ authored in a career age, the cumulative number of publications $P_i(t)$ authored until and in a career age, the number of citations $c_i(t)$ received by $P_i(t)$ in a career age, and the cumulative number of citations $C_i(t)$ received by $P_i(t)$ until and in a career age. Citations are always counted coming from the whole field of computer science, not just from the same cohort.

6.4. Individual Inequality

To quantify individual inequality we use the Gini coefficient $G(t)$ of the publication and citation distributions of authors in the same cohort at the same career age:

$$G(t) = \frac{\sum_{i=1}^n \sum_{j=1}^n |x_i(t) - x_j(t)|}{2n \sum_{i=1}^n x_i(t)} \quad (1)$$

The numerator is the absolute difference of all pairs (i, j) of authors in a cohort. x is a placeholder for publication or citation counts. In the Supplementary materials, Figures S1 and S2, we use backward-looking 3-year windows, that is, to quantify inequality in productivity, $x(t) = p_{3yr}(t) = \sum_{\tau=0}^2 p(t - \tau)$, and, to quantify inequality in impact, $x(t) = c_{3yr}(t) = \sum_{\tau=0}^2 c_{t-\tau}(t - \tau)$, where $t \geq 3$ and the index $t - \tau$ of c defines the career age for the publications of which citations are counted. In Figures 2 and 3, we use cumulative counting, that is, to quantify inequality in cumulative productivity, $x(t) = P(t)$, and, to quantify inequality in cumulative impact, $x(t) = C(t)$. A Gini coefficient of zero expresses perfect equality, where all authors in one cohort have produced an equal number of papers or received an equal number of citations. A Gini of one indicates maximal inequality among authors.

6.5. Gender Inequality

To quantify gender inequality, we look at the differences between the cumulative distributions of productivity $x(t) = P(t)$ and impact $x(t) = C(t)$ of male and female scientists in the same career

age. For both $x(t)$, we rank all observations ascendingly (with adjusted ranks for ties) and perform the Mann–Whitney U test,

$$U(t) = R_m(t) - \frac{n_m(t)(n_m(t) + 1)}{2}, \quad (2)$$

where $R_m(t)$ is the sum of the ranks and $n_m(t)$ is the number of male scientists. The U test allows us to assess the statistical significance of the difference between the distributions of male and female scientists (Mann & Whitney, 1947). To quantify the size of the difference, we compute Cliff's d ,

$$d(t) = \frac{2U(t)}{n_m(t)n_f(t)} - 1, \quad (3)$$

where $n_f(t)$ is the number of female scientists (Cliff, 1993). The value of d ranges from -1 (when all observations for women are greater than those for men) to 1 (when all observations for men are greater than those for women). For example, if $d = 0.8$ for the cumulative publication distribution, a randomly picked man has an 80% chance to have more publications than a randomly chosen woman. If $d = -0.8$ then a randomly picked woman has an 80% chance to have more publications than a randomly picked man.

6.6. Reproductive Feedback

We quantify the ME as the extent to which authors reproduce their productivity and impact over time via positive feedback. For each cohort and career age, we measure to what extent scholars author new publications or receive new citations in a career age proportional to their productivity or impact in the previous career age. This relationship is quantified by the scaling law

$$x(t) \propto x(t-1)^{\beta(t)}, \quad x(t-1) \geq x_{\min}(t-1), \quad (4)$$

where the exponent β and the lower cutoff $x_{\min}$ are the model parameters. If the scaling law is a plausible fit and the estimated exponent $\hat{\beta} > 0$, past productivity or impact is advantageous to (because correlated with) present productivity or impact. If this advantage accumulates over subsequent career ages, we speak of the ME that is then quantified by the sequence of $\hat{\beta}$ s. To quantify the ME in productivity we predict the number of publications $x(t) = p(t)$ by the cumulative number of publications $x(t-1) = P(t-1)$, and to quantify the ME in impact we predict the number of citations $x(t) = c(t)$ by the cumulative number of citations $x(t-1) = C(t-1)$. Predicting by the number of publications $p(t-1)$ and citations $c(t-1)$ yields less variance in $x(t-1)$, shorter time series, and marginally smaller exponents, but similar trends.

In Figure 5A, we demonstrate the fitting procedure for the 2000 cohort, career age 15, and productivity. The pale points are the observations for authors with $p_i(15) \geq 1$ and $P_i(14) \geq 1$. The full points result from putting these observations into 20 bins of exponentially increasing size. The model is fitted to the binned data using the method of ordinary least squares, and the coefficient of determination R^2 quantifies how well the model fits the corresponding unbinned data. The lower cutoff is estimated by choosing $x_{\min}$ such that $R^2(x_{\min})$ has its first maximum. Such a method that includes the identification of the lower cutoff is not discussed in the literature (Perc, 2014). Ours is a simple heuristic that, in our particular application scenario, underestimates both model parameters but mitigates statistical errors on the scaling exponent as well as biases from finite-size effects.

6.7. Independent Variables

There are substantive and methodological reasons to not mix data from different cohorts. Substantively, we are interested in changes that may have occurred as computer science grew and became a team-based science. Methodologically, it prevents accounting for variations in the productivity and impact functions of authors across career ages (Penner et al., 2013). To account for this, the models contain a *cohort* variable.

To test for a gender effect, we include a *gender* category. Because gender could not be detected for all authors, we use *male*, *female*, and *undetected* as dummy variables.

We are further interested in the factors that affect an author's career. According to the ME, advantages accumulate over time. The earlier in a career an advantage sets in, the more it can accumulate. Hence, all our independent variables are computed for the early career $[1, t_e]$. We choose $t_e = 3$ to delimit the early career.

The first construct category contains the early-career *achievements* of authors. *Productivity* is the cumulative number of publications $P_i(t_e)$. *Productivity (1st author)* $P_{i(1st)}(t_e)$ is the number of publications written as a first author. *Impact* is the cumulative number of citations $C_i(t_e)$. Another way to quantify achievement is to use the reputation of the sources (journals and conference proceedings) an author publishes in. We operationalize symbolic capital based on the *h5-index* (Google, 2020) of sources. $h5_s(y)$ is the maximum cumulative number of publications *h5* published in source *s* in the years $[y - 4, y]$ that have accumulated at least *h5* citations in those years. The binary *top source* variable is then 1 if an author has at least one publication in a source that belongs to the top 25% of the distribution in a given year.

Careers are affected by being able to reap benefits from embedding into social networks. Hence, our second construct category is *social support*. *Collaboration network* is the size of the social support network, measured in terms of the number of distinct coauthors in the early career. The maturing of the computer science field is marked by the emergence of team science (Wuchty et al., 2007). Therefore, we study the effect of *team size*, defined as the median number of authors of all publications produced in the early career. *Senior support* quantifies the extent to which an author enjoys mentorship from a senior scientist. Our proxy is the largest *h-index* (Hirsch, 2005) of all coauthors *j* in the social support network: $\max(h_j(y))$. $h_j(y)$ is the maximum cumulative number of publications *h* that each has accumulated at least *h* citations until *y*, where *y* is the year in which author *i* is in career age t_e .

All independent variables are standardized by subtracting the median and dividing the result by the range between the 1st and 3rd quartiles.

6.8. Dependent Variables

Authors can leave academia for a certain number of years in a row. We label each author in our corpus as a *dropout* if she or he has not published for 10 consecutive years in the first 15 career ages. Some 59% of the authors are labeled as dropouts. This label is used as a binary variable in dropout predictions.

To quantify the *success* of authors, we define $C_i^+(15) = C_i(15) - C_i(t_e)$. This variable measures the increase in the cumulative number of *citations* received by all publications published until and in career age 15 after the early career period. This measure avoids autocorrelation with the independent predictor $C_i(t_e)$ and an inflated coefficient of determination (Penner et al., 2013).

Dependent variables are standardized like the independent ones.

6.9. Prediction Models

In *dropout prediction*, we regress dropout against the independent variables using a logistic model. In *success prediction*, we regress citation increase against the independent variables using a linear model. We use the elastic net variant because it contains regularization techniques to ensure that the model generalizes well (to avoid overfitting). These techniques estimate weights that penalize regression coefficients. This is useful when multiple independent variables are correlated with each other (Zou & Hastie, 2005). There are two parameters. The mixing parameter λ controls the extent to which overfitting is avoided by L1 regularization (which makes some weights zero, i.e., selects variables to remove) as opposed to L2 regularization (which makes weights small but not zero). When $\lambda = 1$ only L1 penalties are applied; when $\lambda = 0$ only L2 penalties are applied. We use the default $\lambda = 0.5$, that is, the elastic net will perform variable selection but will keep highly correlated variables in the model. The regularization parameter α is a constant that multiplies the penalty weights. When $\alpha = 0$, the model becomes an ordinary least squares regression (without any regularization). The optimal value for α is learned from the data.

In all prediction models, there are 292,443 observations. Regression coefficients and their weights are learned in 10-fold cross-validation. That is, the data is randomly divided into 10 folds of 29,244 observations, and in 10 iterations the model is trained on nine folds and tested on the remaining one (Hox, 2017). Regression coefficients are reported as averages across the 10 folds. When means are far from zero, effects are sizable; when standard deviations are low, coefficients are robust.

For the binary prediction model (dropout prediction), we use two scores as evaluation metrics. The *F1* score is the weighted average of the precision (proportion of predicted positives that are correct) and recall (proportion of known positives that are predicted correctly). The *average precision* summarizes a precision-recall curve as the weighted mean of precisions achieved for every highest value of recall. Both range from 0 to 1. For the linear models (success prediction), we use two other goodness-of-fit measures. The *mean squared error* quantifies the mean squared distance of all observations to the regression line. The *adjusted R^2* coefficient of determination measures the proportion of the variance in the dependent variable that is predictable from the independent variables. It corrects for the number of independent variables that the models use. It increases only if the new term improves the model more than would be expected by chance. Both measures range between 0 and 1, where higher values are better. For all four evaluation metrics, we report the average value across 10 folds.

ACKNOWLEDGMENTS

This research was made possible through the generous support of the Volkswagen Foundation (<https://www.volkswagenstiftung.de/>) under Grant Ref. 92173. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

AUTHOR CONTRIBUTIONS

Haiko Lietz: Conceptualization, Data curation, Formal analysis, Methodology, Writing—original draft, Writing—review, and editing. Mohsen Jadidi: Formal analysis, Methodology, Writing—original draft. Daniel Kostic: Data curation, Formal analysis. Milena Tsvetkova: Funding acquisition, Methodology, Writing—original draft. Claudia Wagner: Conceptualization, Funding acquisition, Methodology, Writing—original draft.

COMPETING INTERESTS

The authors have no competing interests.

FUNDING INFORMATION

This research was made possible through the generous support of the Volkswagen Foundation (<https://www.volkswagenstiftung.de/>) under Grant Ref. 92173. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

DATA AVAILABILITY

To reproduce this study, download the data (Lietz, Jadidi et al., 2023) and follow the instructions given in the paper's GitHub repository (<https://github.com/gesiscss/inequality/>).

REFERENCES

- Abbate, J. (2017). *Recoding gender: Women's changing participation in computing*. Cambridge, MA: MIT Press.
- Acuna, D. E., Allesina, S., & Kording, K. P. (2012). Future impact: Predicting scientific success. *Nature*, 489(7415), 201–202. <https://doi.org/10.1038/489201a>, PubMed: 22972278
- Aksnes, D. W., Langfeldt, L., & Wouters, P. (2019). Citations, citation indicators, and research quality: An overview of basic concepts and theories. *SAGE Open*, 9(1), 2158244019829575. <https://doi.org/10.1177/2158244019829575>
- Aksnes, D. W., Rorstad, K., Piro, F., & Sivertsen, G. (2011). Are female researchers less cited? A large-scale study of Norwegian scientists. *Journal of the American Society for Information Science and Technology*, 62(4), 628–636. <https://doi.org/10.1002/asi.21486>
- Albarrán, P., Crespo, J. A., Ortuño, I., & Ruiz-Castillo, J. (2011). The skewness of science in 219 sub-fields and a number of aggregates. *Scientometrics*, 88(2), 385–397. <https://doi.org/10.1007/s11192-011-0407-9>
- Albert, R., & Barabási, A.-L. (2002). Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1), 47–97. <https://doi.org/10.1103/RevModPhys.74.47>
- Allison, P. D., & Long, J. S. (1990). Departmental effects on scientific productivity. *American Sociological Review*, 55(4), 469–478. <https://doi.org/10.2307/2095801>
- Allison, P. D., Long, J. S., & Krauze, T. K. (1982). Cumulative advantage and inequality in science. *American Sociological Review*, 47(5), 615–625. <https://doi.org/10.2307/2095162>
- AMiner. (2017). *DBLP-Citation-network V10*. <https://lfs.aminer.cn/lab-datasets/citation/dblp.v10.zip>
- Azoulay, P., & Lynn, F. B. (2020). Self-citation, cumulative advantage, and gender inequality in science. *Sociological Science*, 7, 152–186. <https://doi.org/10.15195/v7.a7>
- Bol, T., de Vaan, M., & van de Rijdt, A. (2018). The Matthew effect in science funding. *Proceedings of the National Academy of Sciences*, 115(19), 4887–4890. <https://doi.org/10.1073/pnas.1719557115>, PubMed: 29686094
- Bornholdt, S., Jensen, M. H., & Sneppen, K. (2011). Emergence and decline of scientific paradigms. *Physical Review Letters*, 106(5), 058701. <https://doi.org/10.1103/PhysRevLett.106.058701>, PubMed: 21405444
- Bourdieu, P. (1988). *Homo academicus*. Stanford, CA: Stanford University Press.
- Bradford, S. C. (1985 [1934]). Sources of information on specific subjects. *Journal of Information Science*, 10(4), 176–180. <https://doi.org/10.1177/016555158501000407>
- Burris, V. (2004). The academic caste system: Prestige hierarchies in PhD exchange networks. *American Sociological Review*, 69(2), 239–264. <https://doi.org/10.1177/000312240406900205>
- Chase, I. D. (1980). Social process and hierarchy formation in small groups: A comparative perspective. *American Sociological Review*, 45(6), 905–924. <https://doi.org/10.2307/2094909>
- Clauset, A., Arbesman, S., & Larremore, D. B. (2015). Systematic inequality and hierarchy in faculty hiring networks. *Science Advances*, 1(1), e1400005. <https://doi.org/10.1126/sciadv.1400005>, PubMed: 26601125
- Clauset, A., Larremore, D. B., & Sinatra, R. (2017). Data-driven predictions in the science of science. *Science*, 355(6324), 477–480. <https://doi.org/10.1126/science.aal4217>, PubMed: 28154048
- Cliff, N. (1993). Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological Bulletin*, 114(3), 494–509. <https://doi.org/10.1037/0033-2909.114.3.494>
- Cole, J. R. (1979). *Fair science: Women in the scientific community*. New York, NY: Free Press.
- Cole, J. R., & Singer, B. (1991). A theory of limited differences: Explaining the productivity puzzle in science. In H. Zuckerman, J. R. Cole, & J. T. Bruer (Eds.), *The outer circle: Women in the scientific community* (pp. 277–323). New York, NY: Norton.
- Cole, J. R., & Zuckerman, H. (1984). The productivity puzzle: Persistence and changes in patterns of publication of men and women scientists. In P. Maehr & M. W. Steinkamp (Eds.), *Advances in motivation and achievement* (Vol. 2, pp. 217–258). Greenwich, CT: JAI Press.
- Cole, R. J., & Cole, S. (1973). *Social stratification in science*. Chicago, IL: University of Chicago Press.
- Costas, R., Bordons, M., van Leeuwen, T. N., & van Raan, A. F. J. (2009). Scaling rules in the science system: Influence of field-specific citation characteristics on the impact of individual researchers. *Journal of the American Society for Information Science and Technology*, 60(4), 740–753. <https://doi.org/10.1002/asi.21017>
- DBLP Team. (2017). *DBLP computer science bibliography. Monthly snapshot release of June 2017*. <https://dblp.org/xml/release/dblp-2017-06-01.xml.gz>
- DiPrete, T. A., & Eirich, G. M. (2006). Cumulative advantage as a mechanism for inequality: A review of theoretical and empirical developments. *Annual Review of Sociology*, 32, 271–297. <https://doi.org/10.1146/annurev.soc.32.061604.123127>
- Dong, Y., Johnson, R. A., & Chawla, N. V. (2015). Will this paper increase your h-index? Scientific impact prediction. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining* (pp. 149–158). New York, NY: Association for

- Computing Machinery. <https://doi.org/10.1145/2684822.2685314>
- Flack, J. C. (2017). Coarse-graining as a downward causation mechanism. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2109), 20160338. <https://doi.org/10.1098/rsta.2016.0338>, PubMed: 29133440
- Fuchs, S. (2001). *Against essentialism: A theory of culture and society*. Cambridge, MA: Harvard University Press <https://doi.org/10.4159/9780674037410>
- Garfield, E. (1996). What is the primordial reference for the phrase “publish or perish”? *The Scientist*, 10(12), 11–12.
- Gibbons, M. (1994). *The new production of knowledge: The dynamics of science and research in contemporary societies*. Thousand Oaks, CA: Sage. <https://doi.org/10.4135/9781446221853>
- Google. (2020). *Google Scholar metrics*. <https://scholar.google.com/intl/en/scholar/metrics.html>
- Guimerà, R., Uzzi, B., Spiro, J., & Amaral, L. A. N. (2005). Team assembly mechanisms determine collaboration network structure and team performance. *Science*, 308(5722), 697–702. <https://doi.org/10.1126/science.1106340>, PubMed: 15860629
- Hirsch, J. E. (2005). An index to quantify an individual’s scientific research output. *Proceedings of the National Academy of Sciences*, 102(46), 16569–16572. <https://doi.org/10.1073/pnas.0507655102>, PubMed: 16275915
- Hofman, J. M., Watts, D. J., Athey, S., Garip, F., Griffiths, T. L., ... Yarkoni, T. (2021). Integrating explanation and prediction in computational social science. *Nature*, 595(7866), 181–188. <https://doi.org/10.1038/s41586-021-03659-0>, PubMed: 34194044
- Holman, L., Stuart-Fox, D., & Hauser, C. E. (2018). The gender gap in science: How long until women are equally represented? *PLOS Biology*, 16(4), e2004956. <https://doi.org/10.1371/journal.pbio.2004956>, PubMed: 29672508
- Hox, J. J. (2017). Computational social science methodology, anyone? *Methodology*, 13(Supplement), 3–12. <https://doi.org/10.1027/1614-2241/a000127>
- Huang, J., Gates, A. J., Sinatra, R., & Barabási, A.-L. (2020). Historical comparison of gender inequality in scientific careers across countries and disciplines. *Proceedings of the National Academy of Sciences*, 117(9), 4609–4616. <https://doi.org/10.1073/pnas.1914221117>, PubMed: 32071248
- Huber, J. C. (2002). A new model that generates Lotka’s law. *Journal of the American Society for Information Science and Technology*, 53(3), 209–219. <https://doi.org/10.1002/asi.10025>
- Jadidi, M., Karimi, F., Lietz, H., & Wagner, C. (2018). Gender disparities in science? Dropout, productivity, collaborations and success of male and female computer scientists. *Advances in Complex Systems*, 21(3–4), 1750011. <https://doi.org/10.1142/S0219525917500114>
- Jeong, H., Neda, Z., & Barabási, A. L. (2003). Measuring preferential attachment in evolving networks. *Europhysics Letters*, 61(4), 567–572. <https://doi.org/10.1209/epl/i2003-00166-9>
- Karimi, F., Wagner, C., Lemmerich, F., Jadidi, M., & Strohmaier, M. (2016). Inferring gender from names on the web: A comparative evaluation of gender detection methods. In *Proceedings of the 25th International Conference Companion on World Wide Web* (pp. 53–54). International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/2872518.2889385>
- Larivière, V., Gingras, Y., & Archambault, E. (2009). The decline in the concentration of citations, 1900–2007. *Journal of the American Society for Information Science and Technology*, 60(4), 858–862. <https://doi.org/10.1002/asi.21011>
- Larivière, V., Ni, C., Gingras, Y., Cronin, B., & Sugimoto, C. R. (2013). Bibliometrics: Global gender disparities in science. *Nature*, 504(7479), 211–213. <https://doi.org/10.1038/504211a>, PubMed: 24350369
- Lee, E., Karimi, F., Wagner, C., Jo, H.-H., Strohmaier, M., & Galesic, M. (2019). Homophily and minority-group size explain perception biases in social networks. *Nature Human Behaviour*, 3(10), 1078–1087. <https://doi.org/10.1038/s41562-019-0677-4>, PubMed: 31406337
- Ley, M. (2009). DBLP: Some lessons learned. *Proceedings of the VLDB Endowment*, 2(2), 1493–1500. <https://doi.org/10.14778/1687553.1687577>
- Lietz, H., Jadidi, M., Kostic, D., & Wagner, C. (2023). *Computer science (1970–2014)*. GESIS, Köln. Data file version 1.0.0. <https://doi.org/10.7802/2642>
- Lincoln, A. E., Pincus, S., Koster, J. B., & Leboy, P. S. (2012). The Matilda effect in science: Awards and prizes in the US, 1990s and 2000s. *Social Studies of Science*, 42(2), 307–320. <https://doi.org/10.1177/0306312711435830>, PubMed: 22849001
- Long, J. S. (1992). Measures of sex differences in scientific productivity. *Social Forces*, 71(1), 159–178. <https://doi.org/10.1093/sf/71.1.159>
- Long, J. S., & Fox, M. F. (1995). Scientific careers: Universalism and particularism. *Annual Review of Sociology*, 21, 45–71. <https://doi.org/10.1146/annurev.soc.21.080195.000401>
- Lotka, A. J. (1926). The frequency distribution of scientific productivity. *Journal of the Washington Academy of Sciences*, 16(12), 317–323.
- Lutter, M., & Schröder, M. (2016). Who becomes a tenured professor, and why? Panel data evidence from German sociology, 1980–2013. *Research Policy*, 45(5), 999–1013. <https://doi.org/10.1016/j.respol.2016.01.019>
- Mann, H. B., & Whitney, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. *Annals of Mathematical Statistics*, 18(1), 50–60. <https://doi.org/10.1214/aoms/1177730491>
- Mazloumian, A. (2012). Predicting scholars’ scientific impact. *PLOS ONE*, 7(11), e49246. <https://doi.org/10.1371/journal.pone.0049246>, PubMed: 23185311
- Mazloumian, A., Eom, Y.-H., Helbing, D., Lozano, S., & Fortunato, S. (2011). How citation boosts promote scientific paradigm shifts and Nobel prizes. *PLOS ONE*, 6(5), e18975. <https://doi.org/10.1371/journal.pone.0018975>, PubMed: 21573229
- Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63. <https://doi.org/10.1126/science.159.3810.56>, PubMed: 5634379
- Merton, R. K. (1973). The normative structure of science. In *The sociology of science: Theoretical and empirical investigations* (pp. 267–278). Chicago, IL: University of Chicago Press.
- Merton, R. K. (1988). The Matthew effect in science, II: Cumulative advantage and the symbolism of intellectual property. *Isis*, 79(4), 606–623. <https://doi.org/10.1086/354848>
- Newman, M. E. J. (2009). The first-mover advantage in scientific publication. *Europhysics Letters*, 86(6), 68001. <https://doi.org/10.1209/0295-5075/86/68001>
- Padgett, J. F., & Powell, W. W. (2012). *The emergence of organizations and markets*. Princeton, NJ: Princeton University Press. <https://doi.org/10.23943/princeton/9780691148670.001.0001>
- Page, S. E. (2015). What sociologists should know about complexity. *Annual Review of Sociology*, 41, 21–41. <https://doi.org/10.1146/annurev-soc-073014-112230>
- Pan, R. K., Petersen, A. M., Pammolli, F., & Fortunato, S. (2018). The memory of science: Inflation, myopia, and the knowledge

- network. *Journal of Informetrics*, 12(3), 656–678. <https://doi.org/10.1016/j.joi.2018.06.005>
- Penner, O., Pan, R. K., Petersen, A. M., Kaski, K., & Fortunato, S. (2013). On the predictability of future impact in science. *Scientific Reports*, 3(1), 3052. <https://doi.org/10.1038/srep03052>, PubMed: 24165898
- Perc, M. (2014). The Matthew effect in empirical data. *Journal of the Royal Society Interface*, 11(98), 20140378. <https://doi.org/10.1098/rsif.2014.0378>, PubMed: 24990288
- Petersen, A. M., Fortunato, S., Pan, R. K., Kaski, K., Penner, O., ... Pammolli, F. (2014a). Reputation and impact in academic careers. *Proceedings of the National Academy of Sciences*, 111(43), 15316–15321. <https://doi.org/10.1073/pnas.1323111111>, PubMed: 25288774
- Petersen, A. M., Jung, W.-S., Yang, J.-S., & Stanley, H. E. (2011). Quantitative and empirical demonstration of the Matthew effect in a study of career longevity. *Proceedings of the National Academy of Sciences*, 108(1), 18–23. <https://doi.org/10.1073/pnas.1016733108>, PubMed: 21173276
- Petersen, A. M., Pavlidis, I., & Semendeferi, I. (2014b). A quantitative perspective on ethics in large team science. *Science and Engineering Ethics*, 20(4), 923–945. <https://doi.org/10.1007/s11948-014-9562-8>, PubMed: 24919946
- Petersen, A. M., & Penner, O. (2014). Inequality and cumulative advantage in science careers: A case study of high-impact journals. *EPJ Data Science*, 3(1), 24. <https://doi.org/10.1140/epjds/s13688-014-0024-y>
- Petersen, A. M., Riccaboni, M., Stanley, H. E., & Pammolli, F. (2012). Persistence and uncertainty in the academic career. *Proceedings of the National Academy of Sciences*, 109(14), 5213–5218. <https://doi.org/10.1073/pnas.1121429109>, PubMed: 22431620
- Price, D. J. d. S. (1986 [1963]). *Little science, big science ... and beyond*. New York, NY: Columbia University Press. <https://doi.org/10.7312/pric91844>
- Reitz, F., & Hoffmann, O. (2013). Learning from the past: An analysis of person name corrections in the DBLP collection and social network properties of affected entities. In *The influence of technology on social network analysis and mining* (pp. 427–453). Vienna: Springer. https://doi.org/10.1007/978-3-7091-1346-2_19
- Reskin, B. F., & Hargens, L. L. (1979). Scientific advancement of male and female chemists. In R. Alvarez & K. G. Lutterman (Eds.), *Discrimination in organizations* (pp. 100–122). San Francisco, CA: Jossey-Bass.
- Ronda-Pupo, G. A., & Pham, T. (2018). The evolutions of the rich get richer and the fit get richer phenomena in scholarly networks: The case of the strategic management journal. *Scientometrics*, 116(1), 363–383. <https://doi.org/10.1007/s11192-018-2761-3>
- Ruiz-Castillo, J., & Costas, R. (2014). The skewness of scientific productivity. *Journal of Informetrics*, 8(4), 917–934. <https://doi.org/10.1016/j.joi.2014.09.006>
- Sarigöl, E., Pfitzner, R., Scholtes, I., Garas, A., & Schweitzer, F. (2014). Predicting scientific success based on coauthorship networks. *EPJ Data Science*, 3(1), 9. <https://doi.org/10.1140/epjds/s13688-014-0009-x>
- Sinatra, R., Wang, D., Deville, P., Song, C., & Barabási, A.-L. (2016). Quantifying the evolution of individual scientific impact. *Science*, 354(6312), aaf5239. <https://doi.org/10.1126/science.aaf5239>, PubMed: 27811240
- Starmans, C., Sheskin, M., & Bloom, P. (2017). Why people prefer unequal societies. *Nature Human Behaviour*, 1(4), 0082. <https://doi.org/10.1038/s41562-017-0082>
- Stewart, F. (2005). Horizontal inequalities: A neglected dimension of development. In A. B. Sørensen, F. E. Weinert, & L. R. Sherrod (Eds.), *Wider perspectives on global development* (pp. 101–135). Springer. https://doi.org/10.1057/9780230501850_5
- Sørensen, A. B. (1986). Social structure and mechanisms of life course processes. In A. B. Sørensen, F. E. Weinert, & L. R. Sherrod (Eds.), *Human development and the life course: Multidisciplinary perspectives* (pp. 177–197). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Tang, J., Zhang, J., Yao, L., Li, J., Zhang, L., & Su, Z. (2008). Arnet-Miner: Extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 990–998). Las Vegas, NV: Association for Computing Machinery. <https://doi.org/10.1145/1401890.1402008>
- Uzzi, B., & Spiro, J. (2005). Collaboration and creativity: The small world problem. *American Journal of Sociology*, 111(2), 447–504. <https://doi.org/10.1086/432782>
- Wallace, M. L., Larivière, V., & Gingras, Y. (2009). Modeling a century of citation distributions. *Journal of Informetrics*, 3(4), 296–303. <https://doi.org/10.1016/j.joi.2009.03.010>
- Wang, D., Song, C., & Barabási, A.-L. (2013). Quantifying long-term scientific impact. *Science*, 342(6154), 127–132. <https://doi.org/10.1126/science.1237825>, PubMed: 24092745
- Way, S. F., Larremore, D. B., & Clauset, A. (2016). Gender, productivity, and prestige in computer science faculty hiring networks. In *Proceedings of the 25th International Conference on World Wide Web* (pp. 1169–1179). Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee. <https://doi.org/10.1145/2872427.2883073>
- Way, S. F., Morgan, A. C., Clauset, A., & Larremore, D. B. (2017). The misleading narrative of the canonical faculty productivity trajectory. *Proceedings of the National Academy of Sciences*, 114(44), E9216–E9223. <https://doi.org/10.1073/pnas.1702121114>, PubMed: 29042510
- Way, S. F., Morgan, A. C., Larremore, D. B., & Clauset, A. (2019). Productivity, prominence, and the effects of academic environment. *Proceedings of the National Academy of Sciences*, 116(22), 10729–10733. <https://doi.org/10.1073/pnas.1817431116>, PubMed: 31036658
- West, J. D., Jacquet, J., King, M. M., Correll, S. J., & Bergstrom, C. T. (2013). The role of gender in scholarly authorship. *PLOS ONE*, 8(7), e66212. <https://doi.org/10.1371/journal.pone.0066212>, PubMed: 23894278
- White, D. R., Owen-Smith, J., Moody, J., & Powell, W. W. (2004). Networks, fields and organizations: Micro-dynamics, scale and cohesive embeddings. *Computational & Mathematical Organization Theory*, 10(1), 95–117. <https://doi.org/10.1023/B:CMOT.0000032581.34436.7b>
- White, H. C. (1970). *Chains of opportunity: System models of mobility in organizations*. Cambridge, MA: Harvard University Press. <https://doi.org/10.4159/harvard.9780674437203>
- Wray, K. B. (2011). *Kuhn's evolutionary social epistemology*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511997990>
- Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The increasing dominance of teams in production of knowledge. *Science*, 316(5827), 1036–1039. <https://doi.org/10.1126/science.1136099>, PubMed: 17431139
- Xie, Y. (2014). “Undemocracy”: Inequalities in science. *Science*, 344(6186), 809–810. <https://doi.org/10.1126/science.1252743>, PubMed: 24855244

- Xie, Y., & Shauman, K. A. (1998). Sex differences in research productivity: New evidence about an old puzzle. *American Sociological Review*, 63(6), 847–870. <https://doi.org/10.2307/2657505>
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320. <https://doi.org/10.1111/j.1467-9868.2005.00503.x>
- Zuckerman, H., & Merton, R. K. (1972). Age, aging, and age structure in science. In M. W. Riley, M. Johnson, & A. Foner (Eds.), *A sociology of age stratification, Vol. 3 of Aging and society*. New York, NY: Russel Sage Foundation.