

THE SCIENCE OF SCIENCE



Dashun Wang
Albert-László Barabási

CAMBRIDGE UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314-321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre,
New Delhi – 110025, India

79 Anson Road, #06-04/06, Singapore 079906

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org

Information on this title: www.cambridge.org/9781108492669

DOI: 10.1017/9781108610834

© Cambridge University Press 2021

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2021

Printed in Singapore by Markono Print Media Pte Ltd

A catalogue record for this publication is available from the British Library.

ISBN 978-1-108-49266-9 Hardback

ISBN 978-1-108-71695-6 Paperback

Cambridge University Press has no responsibility for the persistence or accuracy of URLs for external or third-party internet websites referred to in this publication and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

CONTENTS

Acknowledgements viii

Introduction 1

Part I The Science of Career 5

- 1 [Productivity of a Scientist](#) 7
- 2 [The *h*-Index](#) 17
- 3 [The Matthew Effect](#) 28
- 4 [Age and Scientific Achievement](#) 39
- 5 [Random Impact Rule](#) 51
- 6 [The *Q*-Factor](#) 60
- 7 [Hot Streaks](#) 71

Part II The Science of Collaboration 81

- 8 [The Increasing Dominance of Teams in Science](#) 85
- 9 [The Invisible College](#) 95
- 10 [Coauthorship Networks](#) 102
- 11 [Team Assembly](#) 110
- 12 [Small and Large Teams](#) 124
- 13 [Scientific Credit](#) 134
- 14 [Credit Allocation](#) 147

Part III The Science of Impact 159

- 15 [Big Science](#) 161
- 16 [Citation Disparity](#) 174
- 17 [High-Impact Papers](#) 184
- 18 [Scientific Impact](#) 189
- 19 [The Time Dimension of Science](#) 197
- 20 [Ultimate Impact](#) 209

Part IV Outlook 221

- 21 [Can Science Be Accelerated?](#) 223
- 22 [Artificial Intelligence](#) 231
- 23 [Bias and Causality in Science](#) 241
- Last Thought: *All the Science of Science* 252

Appendix A1 Modeling Team Assembly 254

Appendix A2 Modeling Citations 257

References 270

Index 296

ACKNOWLEDGEMENTS

As teams are the fundamental engine for today's innovation, we must thank our lab members first and foremost. Indeed, our lab members contributed to this project in some of the most fundamental ways – not only did their work compose a large part of the initial book proposal, it also contributed fresh insights which grew into several new chapters, broadening and deepening our understanding of the subjects we set out to explore. These team members include Roberta Sinatra, Chaoming Song, Lu Liu, Yian Yin, Yang Wang, Ching Jin, Hua-Wei Shen, Pierre Deville, Michael Szell, Tao Jia, Lingfei Wu, Zhongyang He, Jichao Li, Binglu Wang, Suman Kalyan Maity, Woo Seong Jo, Jian Gao, Nima Dehmamy, Yi Bu, David Moser, Alex Gates, Junming Huang, Qing Ke, and Xindi Wang. Our science of science journey has been infinitely more fun because we had the opportunity to work and learn in such inspiring company.

Chapter 9 on the invisible college discusses the “bright ambience” of brilliant minds, a concept that we experienced vividly while writing this book. Indeed, so many friends and colleagues have been generous with their time and expertise along the way, that listing some of those risks excluding many others. We feel compelled to express our sincere thanks to Luis Amaral, Sinan Aral, Pierre Azoulay, Federico Battiston, Jeanne Brett, Elizabeth Caley, Manuel Cebrian, Damon Centola, Jillian Chown, Noshir Contractor, Ying Ding, Yuxiao Dong, Tina Eliassi-Rad, Eli Finkel, Santo Fortunato, Morgan Frank, Lee Giles, Danny Goroff, Shlomo Havlin, Cesar Hidalgo, Travis Hoppe, Ian B. Hutchins, Ben Jones, Brayden King, Rebecca Meseroll, Sam Molyneux,

Karim Lakhani, David Lazer, Jessie Li, Zhen Lei, Jess Love, Stasa Milojevic, Federico Musciotto, Willie Ocasio, Sandy Pentland, Alex Petersen, Filippo Radicchi, Iyad Rahwan, Lauren Rivera, Matt Salganik, George Santangelo, Iulia Georgescu, Ned Smith, Paula Stephan, Toby Stuart, Boleslaw Szymanski, Arnout van de Rijt, Alessandro Vespignani, John Walsh, Ludo Waltman, Kuansan Wang, Ting Wang, Adam Waytz, Klaus Weber, Stefan Wuchty, Yu Xie, Hyejin Youn.

Among the many people who made this journey possible, our special thanks go to two of them in particular, whose “bright ambience” is especially omnipresent throughout the pages and over the years. Brian Uzzi is not only always generous with his time, he also has remarkable insight. Time and again he manages to take ideas we’re struggling to present and effortlessly elevates them. We are also grateful to Brian for championing the somewhat radical idea that physicists can contribute to the social sciences. Our colleague James Evans has been a close friend and collaborator, and several ideas discussed in this book would have not been born without him, including but not limited to the concluding remarks on engaging “all the science of science” for the future development of the field.

We are extremely grateful for the generous research support we have received over the years. In particular Riq Parra from AFOSR has been a true believer from the very beginning, when few people knew what we meant by the “science of science.” Many concepts discussed in this book would not have been possible without his strong and continued support. Dashun also wishes to express his special thanks to the Kellogg School of Management, an institution that offered a level of support and trust that most researchers only dream of.

Many unsung heroes contributed to this text, collectively logging countless hours to help guide the book along. We have benefited tremendously from the excellent and dedicated editorial assistance of Carrie Braman, Jake Smith, and James Stanfill, as well as Enikő Jankó, Hannah Kiefer, Alanna Lazarowich, Sheri Gilbert, Michelle Guo, and Krisztina Eleki. Special thanks to Alice Grishchenko who took on the brave redesign effort, diligently redrawing all the figures and giving a visual identity to the book. Yian Yin and Lu Liu have been our go-to people behind the scenes, springing into action whenever we needed help, and contributing to a variety of essential tasks ranging from data analysis to managing references.

Special thanks to our very professional publishing team at Cambridge, from our enthusiastic champion, Simon Capelin, to the editorial team who worked tirelessly to move the project forward and across the finish line: Roisin Munnely, Nicholas Gibbons, Henry Cockburn, and Marion Moffatt. Thank you for all your help, and for putting up with numerous missed deadlines.

And finally, Dashun would like to thank his wife Tian Shen, and dedicate this book to her. Thank you, for everything.

INTRODUCTION

Scientific revolutions are often driven by the invention of new instruments – the microscope, the telescope, genome sequencing – each of which have radically changed our ability to sense, measure, and reason about the world. The latest instrument at our disposal? A windfall of digital data that traces the entirety of the scientific enterprise, helping us capture its inner workings at a remarkable level of detail and scale. Indeed, scientists today produce millions of research articles, preprints, grant proposals, and patents each year, leaving detailed fingerprints of the work we admire and how they come about. Access to this data is catalyzing the emergence of a new multidisciplinary field, called *science of science*, which, by helping us to understand in a quantitative fashion the evolution of science, has the potential to unlock enormous scientific, technological, and educational value.

The increasing availability of all this data has created an unprecedented opportunity to explore scientific production and reward. Parallel developments in data science, network science, and artificial intelligence offer us powerful tools and techniques to make sense of these millions of data points. Together, they tell a complex yet insightful story about how scientific careers unfold, how collaborations contribute to discovery, and how scientific progress emerges through a combination of multiple interconnected factors. These opportunities – and the challenges that come with them – have fueled the emergence of a new multidisciplinary community of scientists that are united by their goals of understanding science. These practitioners of science of science use the scientific methods to study themselves, examine projects that work as well as those that fail,

quantify the patterns that characterize discovery and invention, and offer lessons to improve science as a whole. In this book, we aim to introduce this burgeoning field – its rich historical context, exciting recent developments, and promising future applications.

We had three core audiences in mind as we wrote this book. The primary audience includes any scientist or student curious about the mechanisms that govern our passion, science. One of the founding fathers of the science of science, Thomas Kuhn, a physicist turned philosopher, triggered worldwide interest in the study of science back to 1962 with the publication of *The Structure of Scientific Revolutions*. Kuhn's notion of "paradigm shift" today is used in almost every creative activity, and continues to dominate the way we think about the emergence and acceptance of new ideas in science. In many ways, the science of science represents the next major milestone in this line of thinking, addressing a series of questions that are dear to the heart of every scientist but may well lay outside of the Kuhnian worldview: When do scientists do their best work? What is the life cycle of scientific creativity? Are there signals for when a scientific hit will occur in a career? Which kinds of collaboration triumph and which are destined to for disaster? How can young researchers maximize their odds of success? For any working scientist, this book can be a tool, providing data-driven insight into the inner workings of science, and helping them navigate the institutional and scholarly landscape in order to better their career.

A broader impact of the science of science lies in its implications for policy. Hence, this book may be beneficial to academic administrators, who can use science of science to inform evidence-based decision-making. From department chairs to deans to vice presidents of research, university administrators face important personnel and investment decisions as they try to implement and direct strategic research. While they are often aware of a profusion of empirical evidence on this subject, they lack cohesive summaries that would allow them to extract signals from potential noise. As such, this book may offer the knowledge and the data to help them better take advantage of useful insights the science of science community has to offer. What does an h -index of 25 tell us about a physics faculty member seeking tenure? What would the department most benefit from: a junior vs. a senior hire? When should we invest in hiring a superstar, and what can we expect their impact will be?

Part I

THE SCIENCE OF CAREER

Albert Einstein published 248 papers in his lifetime, Charles Darwin 119, Louis Pasteur 172, Michael Faraday 161, Siméon Denis Poisson 158, and Sigmund Freud 330 [1]. Contrast these numbers with the body of work of Peter Higgs, who had published only 25 papers by the age of 84, when he received the Nobel Prize for predicting the Higgs boson. Or think of Gregor Mendel, who secured an enduring legacy with only seven scientific publications to his name [2].

These differences show that in the long run what matters to a career is not productivity, but impact. Indeed, there are remarkable differences among the impact of the publications. Even for star scientists, of all papers they publish, at most a few may be remembered by a later generation of scientists. Indeed, we tend to associate Einstein's name with relativity and Marie Curie with radioactivity, while lacking general awareness of the many other discoveries made by each. In other words, one or at most a few discoveries – the outliers – seem to be what define a scientist's career. So, do these outliers accurately represent a scientific career? Or did these superstar scientists just get lucky in one or a few occasions along their careers?

And, if only one or at most a few papers are remembered, when do scientists make that defining discovery? Einstein once quipped, “A person who has not made his great contribution to science before the age of 30 will never do so” [3]. Indeed, Einstein was merely 26 years old when he published his *Annus Mirabilis* papers. Yet, his observation about the link between youth and discovery was not merely autobiographical. Many of the physicists of his generation too made their

defining discoveries very early in their career – Heisenberg and Dirac at 24; Pauli, Fermi, and Wigner at 25; Rutherford and Bohr at 28. But is youth a necessity for making an outstanding contribution to science? Clearly not. Alexander Fleming was 47 when he discovered penicillin. Luc Montagnier was 51 when he discovered HIV. And John Fenn was 67 when he first began to pursue the research that would later win him the Nobel Prize in chemistry. So, how is creativity, as captured by scientific breakthroughs, distributed across the lifespan of a career?

The first part of this book will dive into these sets of fascinating questions regarding scientific careers. Indeed, as we survey our young and not so young colleagues doing groundbreaking work, we are prompted to ask: Are there quantitative patterns underlying when breakthrough work happens in a scientific career? What mechanisms drive the productivity and impact of a scientist? The chapters in this part will provide quantitative answers to these questions, offering insights that affect both the way we train scientists and the way we acknowledge and reward scientific excellence.

1

PRODUCTIVITY OF A SCIENTIST

Paul Erdős, arguably the most prolific mathematician in the twentieth century, was, by all accounts, rather eccentric. The Hungarian-born mathematician – who moved to the US before the start of WWII – lived out of a ragged suitcase that he famously dragged with him to scientific conferences, universities, and the homes of colleagues all over the world. He would show up unannounced on a colleague’s doorstep, proclaim gleefully, “My mind is open.” He then spent a few days working with his host, before moving on to surprise some other colleague at some other university. His meandering was so constant that it eventually earned him undue attention from the FBI. To his fellow mathematicians, he was an eccentric but lovable scientist. But to law enforcement officers during the Cold War, it was suspicious that he crossed the Iron Curtain with such ease. Indeed, Erdős was once arrested in 1941 for poking around a secret radio tower. “You see, I was thinking about mathematical theorems,” he explained to the authorities in his thick Hungarian accent. It took decades of tracking for the Bureau to finally believe him, concluding that his rambling was indeed just for the sake of math.

His whole *life* was, too. He had no wife, no children, no job, not even a home to tie him down. He earned enough in guest lecturer stipends from universities and from various mathematics awards to fund his travels and basic needs. He meticulously avoided any commitment that might stand in the way of his work. Before he died in 1996 at the age of 83, Erdős had written or coauthored a stunning 1,475 academic papers in collaboration

with 511 colleagues. If total publication counts as a measure of productivity, how does Erdős' number compare to the productivity of an ordinary scientist? It surely seems exceptional. But how exceptional?

1.1 How Much Do We Publish?

Scholarly publications are the primary mode of communication in science, helping disseminate knowledge. The productivity of a scientist captures the rate at which she adds units of knowledge to the field. Over the past century, the number of publications has grown exponentially. An important question is whether the growth in our body of knowledge is simply because there are now more scientists, or because each scientist produces more on average than their colleagues in the past.

An analysis of over 53 million authors and close to 90 million papers published across all branches of science shows that both the number of papers and scientists grew exponentially over the past century [4]. Yet, while the former grew slightly faster than the latter (Fig. 1.1a), meaning that the number of publications per capita has been decreasing over time, for each scientist, individual productivity has stayed quite stable over the past century. For example, the number of papers a scientist produces each year has hovered at around two for the entire twentieth century (Fig. 1.1b, blue curve), and has even increased slightly during the past 15 years. As of 2015, the typical scientist authors or coauthors about 2.5 papers per year. This growth in individual productivity has its origins in collaborations: Individual productivity is boosted as scientists end up on many more papers as coauthors (Fig. 1.1b, red curve). In other words, while in terms of how many scientists it takes to produce a paper, that number has been trending downwards over the past century, thanks to collaborative work individual productivity has increased during the past decade.

1.2 Productivity: Disciplinary Ambiguities

But, when it comes to a scientist's productivity, it's not easy to compare across disciplines. First, each publication may represent a unit of knowledge, but that unit comes in different sizes. A sociologist may

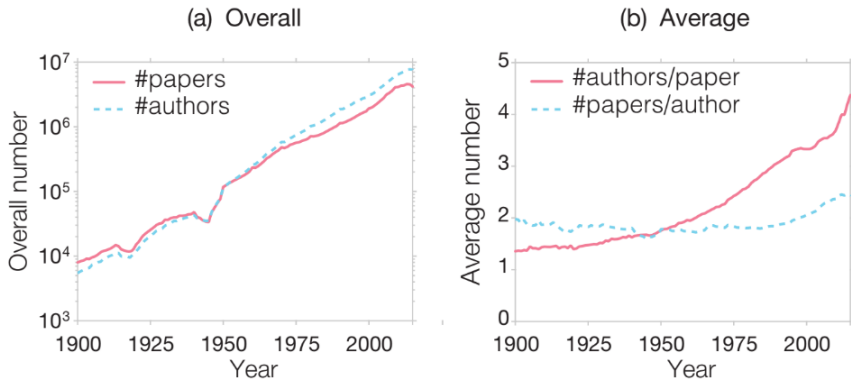


Figure 1.1 The growing number of scientists. (a) During the past century, both the number of scientists and the number of papers has increased at an exponential rate. (b) The number of papers coauthored by each scientist has been hovering around two during the past 100 years, and increased gradually in the past 15 years. This growth is a direct consequence of collaborative effects: Individual productivity is boosted as scientists end up on many more papers as coauthors. Similar trends were reported using data within a single field [5]. For physics, for example, the number of papers coauthored by each physicist has been less than one during the past 100 years, but increased sharply in the past 15 years. After Dong et al. [4] and Sinatra et al. [5].

not feel their theory is fully articulated unless the introduction of the paper spans a dozen pages. Meanwhile, a paper published in *Physical Review Letters*, one of the most respected physics journals, has a strict four-page limit, including figures, tables, and references. Also, when we talk about individual productivity, we tend to count publications in scientific journals. But in some branches of the social sciences and humanities, books are the primary form of scholarship. While each book is counted as one unit of publication, that unit is admittedly much more time-consuming to produce.

And then there is computer science (CS). As one of the youngest scientific disciplines (the first CS department was formed at Purdue University in 1962), computer science has adopted a rather unique publication tradition. Due to the rapidly developing nature of the field, computer scientists choose conference proceedings rather than journals as their primary venue to communicate their advances. This approach has served the discipline well, given everything that has been accomplished in the field – from the Internet to artificial intelligence – but it can be quite confusing to those outside the discipline.




Figure 1.2 Productivity distribution. The blue symbols show the number of papers published by all authors listed in the INSPECT database of scientific and technical literature, in the period 1969–2004 (over 3 million authors). The red line corresponds to the lognormal fit to the data (1.1). After Fronczak et al. [10].

Box 1.1 The study of productivity has a long history [9–15]

In 1926, Alfred J. Lotka [11] observed that the number of papers produced by scientists follows a fat-tailed distribution. In other words, he found that a small fraction of scientists are responsible for the bulk of scientific literature. Lotka studied 6,891 authors listed in Chemical Abstracts publishing between 1907 and 1916, concluding that the number of authors making N contributions follows a power law

$$P(N) \sim N^{-\alpha}, \quad (1.2)$$

where the exponent $\alpha \approx 2$. A power law predicts that productivity has a long tail, capturing major variations among individuals. Note that it often requires a large amount of data to reliably distinguish a power law from a lognormal distribution [9], which Lotka did not have in 1926.

This lognormal distribution of productivity is rather odd, as Shockley quickly noticed. Indeed, in most competitive arenas, individual performance metrics almost always follow a narrow distribution. Think about running. At the Rio Olympics in 2016, Usain Bolt finished the 100-meter final in just 9.81 seconds. Justin Gatlin came in second and

Andre De Grasse in third, with running times 9.89 s and 9.91 s, respectively. These numbers are awfully close, reflecting a well-known fact that performance differences between individuals are typically bounded [16]. Similarly, Tiger Woods, even on his best day, only took down his closest contenders by a few strokes, and the fastest typist may only type a few words more per minute than a merely good one. The bounded nature of performance reminds us that it is difficult, if not impossible, to significantly outperform the competition in any domain. Yet, according to Fig. 1.2, this boundedness does not hold for scientific performance. Apparently, it *is* possible to be much better than your competitors when it comes to churning out papers. Why is that?

1.4 Why So Productive?

Shockley proposed a simple model to explain the lognormal productivity distribution he observed (Eq. 1.1) [9]. He suggested that in order to publish a paper, a scientist must juggle multiple factors, like:

- F₁. Identify a good problem.
- F₂. Make progress with it.
- F₃. Recognize a worthwhile result.
- F₄. Make a decision as to when to stop the research and start writing up the results.
- F₅. Write adequately.
- F₆. Profit constructively from criticism.
- F₇. Show determination to submit the paper for publication.
- F₈. Make changes if required by the journal or the referees.

If any of these steps fail, there will be no publication. Let us assume that the odds of a person clearing hurdle F_i from the list above is p_i . Then, the publication rate of a scientist is proportional to the odds of clearing each of the subsequent hurdles, that is $N \sim p_1 p_2 p_3 p_4 p_5 p_6 p_7 p_8$. If each of these odds are independent random variables, then the multiplicative nature of the process predicts that $P(N)$ follows a lognormal distribution of the form (1.1).

To understand where the outliers come from, imagine, that Scientist A has the same capabilities as Scientist B in all factors, except that A is twice as good at solving a problem (F₂), knowing when to stop (F₄), and determination (F₇). As a result, A's productivity will be eight times higher than B's. In other words, for each paper published by

Scientist B, Scientist A will publish eight. Hence small differences in scientists' ability to clear individual hurdles can together lead to large variations in overall productivity.

Shockley's model not only explains why productivity follows lognormal distribution, but it also offers a framework to improve our own productivity. Indeed, the model reminds us that publishing a paper does not hinge on a single factor, like having a great idea. Rather, it requires scientists to excel at multiple factors. When we see someone who is hyper-productive, we tend to attribute it to a single exceptional factor. Professor X is really good at coming up with new problems (F_1), or conveying her ideas in writing (F_5). The model suggests, however, that the outliers are unlikely to be explained by a single factor; rather, a researcher is most productive when she excels across many factors and fails in none.

The hurdle model indicates that a single weak point can choke an individual's productivity, even if he or she has many strengths. It also tells us that Erdős may have not been as super-human as we often think he was, or that his productivity might be attainable with careful honing of various skills. Indeed, if we could improve at every step of writing a paper, and even if it's just a tiny bit in each step, these improvements can combine to exponentially enhance productivity. Admittedly, this is easier said than done. But you can use this list to diagnose yourself: What step handicaps your productivity the most?

The remarkable variations in productivity have implications for reward. Indeed, Shockley made another key observation: while the productivity of a scientist is multiplicative, his salary – a form of reward often tied to performance – is additive. The highest paid employees earn at best about 50–100 percent more than their peers. There are many reasons why this is the case – it certainly seems fairer, and it helps ensure a collaborative environment. Yet, from a paper-per-dollar perspective, Shockley's findings raise some interesting questions about whether the discrepancy between additive salaries and multiplicative productivities could be exploited. Indeed, an institution may be better off employing a few star scientists, even if that means paying them a great deal more than their peers. Shockley's arguments are often used as a rationale for why top individuals at research-intensive institutions are offered much higher salaries and special

perks, and why top departments within a university get disproportionately more funding and resources.

To be sure, gauging a career based on publication count alone grossly misrepresents how science works. Yet, individual productivity has been shown to closely correlate with the eminence of a scientist as well as her perceived contributions to the field. This pattern was documented by Wayne Dennis, dating back at least to 1954 [1], when he studied 71 members of the US National Academy of Sciences and eminent European scientists. He found that, almost without exception, highly productive individuals have also achieved scientific eminence, as demonstrated by their listing in the *Encyclopedia Britannica* or in histories of important developments they have contributed to the sciences. Higher productivity has been shown to increase the odds of receiving tenure [17], and of securing funding for future research [18]. At the institutional level, the publication rates of the faculty are not only a reliable predictor of a program's reputation, they also influence the placement of graduates into faculty jobs [19].

In sum, sustained high productivity is rare, but it correlates with scientific impact and eminence. Given this evidence, it may appear that productivity is the key indicator for a meaningful career in science. Yet, as we show in the following chapters, among the many metrics used to quantify scientific excellence, productivity is the least predictive. The reason is simple: While great scientists tend to be very productive, not all scientists who are productive make long-lasting contributions. In fact, most of them do not. Multiple paths can lead to achieving high productivity. For example, lab technicians in certain fields may find their names on more than a hundred – or sometimes as many as a thousand – papers. Hence, they appear to be exceptionally prolific based on their publication counts, but are rarely credited as the intellectual owner of the research. The way people publish is also changing [20]. Coauthorship is on the rise, as are multiple publications on the same data. There have also been more discussions about LPUs, which stands for least publishable unit [20] or the “salami publishing” approach, which could further contribute to inflated productivity counts.

So, if productivity is not the defining factor of a successful career, what is?

Box 1.2 Name disambiguation

Our ability to accurately track individual productivity relies on our skill to identify the individual(s) who wrote a paper and all other work that belongs to that individual [21, 22]. This seemingly simple task represents a major unsolved problem [21–23], limited by four challenges. First, a single individual may appear in print under multiple names because of orthographic and spelling variants, misspellings, name changes due to marriage, religious conversion, gender reassignment, or the use of pen names. Second, some common names can be shared by multiple individuals. Third, the necessary metadata is often incomplete or missing. This includes cases where publishers and bibliographic databases failed to record authors' first names, their geographical locations, or other identifying information. Fourth, an increasing percentage of papers is not only multi-authored, but also represents multidisciplinary and multi-institutional efforts. In such cases, disambiguating some of the authors does not necessarily help assign the remaining authors.

While multiple efforts are underway to solve the name disambiguation problem, we need to be somewhat mindful about the results presented in this and following chapters, as some conclusions may be affected by the limitations in disambiguation. In general, it is easier to disambiguate productive scientists, who have a long track record of papers, compared with those who have authored only a few publications. Therefore, many studies focus on highly productive scientists with unusually long careers instead of “normal” scientists.

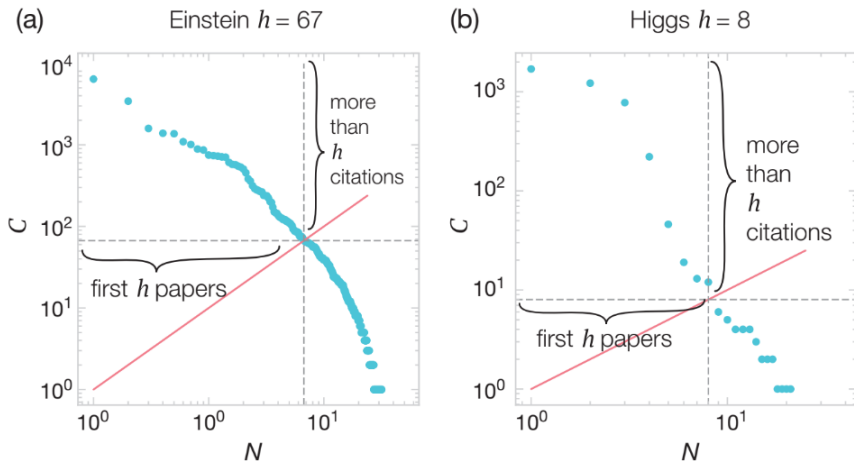


Figure 2.1 The h -index of Albert Einstein (a) and Peter Higgs (b). To calculate the h -index, we plot the number of citations versus paper number, with papers listed in order of decreasing citations. The intersection of the 45° line with the curve gives h . The total number of citations is the area under the curve [26]. According to Microsoft Academic Graph, Einstein has an h -index of 67, and Higgs 8. The top three most cited papers by Einstein are: (1) Can quantum mechanical description of physical reality be considered complete, *Physical Review*, 1935; (2) Investigations on the theory of Brownian movement, *Annalen der Physik*, 1905; and (3) On the electrodynamics of moving bodies, *Annalen der Physik*, 1905. The top three for Higgs are: (1) Broken symmetries and the masses of gauge bosons, *Physical Review Letters*, 1964; (2) Broken symmetries, massless particles and gauge fields, *Physics Letters*, 1964; (3) Spontaneous symmetry breakdown without massless bosons, *Physical Review*, 1966.

Chapter 19). Yet, despite the model's simplicity, the linear relationship predicted by (2.3) holds up generally well for scientists with long scientific careers [26].

This linear relationship (2.3) has two important implications:

- (1) If a scientist's h -index increases roughly linearly with time, then its speed of growth is an important indicator of her eminence. In other words, the differences between individuals can be characterized by the slope, m . As (2.2) shows, m is a function of both n and c . So, if a scientist has higher productivity (a larger n), or if her papers collect more citations (higher c), she has a higher m . And the higher the m , the more eminent is the scientist.
- (2) Based on typical values of m , the linear relationship (2.3) also offers a guideline for how a typical career should evolve. For

example, Hirsch suggested in 2005 that for a physicist at major research universities, $h \approx 12$ might be a typical value for achieving tenure (i.e., the advancement to associate professor) and that $h \approx 18$ might put a faculty member into consideration for a full professorship. Fellowship in the American Physical Society might typically occur around $h \approx 15\text{--}20$, and membership in the US National Academy of Sciences may require $h \approx 45$ or higher.

Since its introduction, the h -index has catalyzed a profusion of metrics and greatly popularized the idea of using objective indicators to quantify nebulous notions of scientific quality, impact or prestige [27]. As a testament to its impact, Hirsh's paper, published in 2005, had been cited more than 8,000 times as of the beginning of 2019, according to Google Scholar. It even prompted behavioral changes – some ethically questionable – with scientists adding self-citations for papers on the edge of their h -index, in hopes of boosting it [28–30]. Given its prevalence, we must ask: can the h -index predict the future impact of a career?

Box 2.1 The Eddington number

The h -index for scientists is analogous to the Eddington number for cyclists, named after Sir Arthur Eddington (1882–1944), an English astronomer, physicist, and mathematician, famous for his work on the theory of relativity. As a cycling enthusiast, Eddington devised a measure of a cyclist's long-distance riding achievements. The Eddington number, E , is the number of days in your life when you have cycled more than E miles. Hence an Eddington number of 70 would mean that the person in question has cycled at least 70 miles a day on 70 occasions. Achieving a high Eddington number is difficult, since jumping from, say, 70 to 75 may require more than 5 new long-distance rides. That's because any rides shorter than 75 miles will no longer be included. Those hoping to increase their Eddington number are forced to plan ahead. It might be easy to achieve an E of 15 by doing 15 trips of 15 miles – but turning that $E = 15$ into an $E = 16$ could force a cyclist to start over, since an E number of 16 only counts trips of 16 miles or more. Arthur Eddington, who reached an $E = 87$ by the time he died in 1944, clearly understood that if he wanted to achieve a high E number, he had to start banking long rides early on.

2.2 The Predictive Power of the *h*-Index

To understand the value of the *h*-index, let's take a look at the “usual suspects” – metrics that are commonly used to evaluate a scientist's performance, and review their strengths and limitations [26].

- (1) Total number of publications (N).
 Advantage: Measures the productivity of an individual.
 Disadvantage: Ignores the impact of papers.
- (2) Total number of citations (C).
 Advantage: Measures a scientist's total impact.
 Disadvantage: It can be affected by a small number of big hits, which may not be representative of the individual's overall career, especially when these big hits were coauthored with others. It also gives undue weight to highly cited reviews as opposed to original research contributions.
- (3) Citations per paper (C/N).
 Advantage: Allows us to compare scientists of different ages.
 Disadvantage: Outcomes can be skewed by highly cited papers.
- (4) The number of “significant papers,” with more than c citations.
 Advantage: Eliminates the disadvantages of (1), (2), (3), and measure broad and sustained impact.
 Disadvantage: The definition of “significant” introduces an arbitrary parameter, which favors some scientists or disfavors others.
- (5) The number of citations acquired by each of the q most-cited papers (for example, $q = 5$).
 Advantage: Overcomes many of the disadvantages discussed above.
 Disadvantage: Does not provide a single number to characterize a given career, making it more difficult to compare scientists to each other. Further, the choice of q is arbitrary, favoring some scientists while handicapping others.

The key advantage of the *h*-index is that it *sidesteps all of the disadvantages* of the metrics listed above. But, is it more effective at gauging the impact of an individual's work? When it comes to evaluating the predictive power of metrics, two questions are often the most relevant.

Q1: Given the value of a metric at a certain time t_1 , how well does it predict the value of itself or of another metric at a future time t_2 ?

This question is especially interesting for hiring decisions. For example, if one consideration regarding a faculty hire is the likelihood of the candidate to become a member of the National Academy of Sciences 20 years down the line, then it would be useful to rank the candidates by their projected *cumulative* achievement after 20 years. Hirsch tested Q_1 by selecting a sample of condensed matter physicists and looked at their publication records during the first 12 years of their career and in the subsequent 12 years [31]. More specifically, he calculated four different metrics for each individual based on their career records in the first 12 years, including the *h*-index (Fig. 2.2a), the total number of citations




Figure 2.2 Quantifying predictive power of the *h*-index. Scatter plots compare the total number of citations, C , after $t_2 = 24$ years vs. the value of the various indicators at $t_1 = 12$ year for each individual within the sample. Hirsch hypothesized C may grow quadratically with time, and hence used its square root when calculating the total number of citations. By calculating the correlation coefficient, he found that the *h*-index (a) and the number of citations at t_1 (b) are the best predictors of the future cumulative citations at t_2 . The number of papers correlates less (c), and the number of citations per paper performs the worst (d). After Hirsch [31].

(Fig. 2.2b), the total number of publications (Fig. 2.2c), and the average number of citations per paper (Fig. 2.2d). He then asked if we want to select candidates that have the most total citations by year 24, which one of the four indicators gives us the best chance? By measuring the correlation coefficient between future cumulative citations at time t_2 and four different metrics calculated at time t_1 , he found that the *h*-index and the number of citations at time t_1 turn out to be the best predictors (Fig. 2.2).

While Fig. 2.2 shows that the *h*-index predicts cumulative impact, in many cases it's the future scientific output that matters the most. For example, if we're deciding who should get a grant, how many more citations an applicant's *earlier* papers are expected to collect in the next few years is largely irrelevant. We're concerned, instead, with papers that the potential grantee has not yet written and the impact of those papers. Which brings us to Q2:

Q2: How well do the different metrics predict *future* scientific output?

To answer Q2, we need to use indicators obtained at t_1 to predict scientific achievement occurring only in the subsequent period, thereby omitting all citations to work performed prior to t_1 . Hirsch repeated the similar prediction task for the four metrics, but this time used each of them to predict total citations accrued by papers published only in the next 12 years. Naturally, this is a more difficult task, but an important one for allocating research resources. Hirsch found that the *h*-index again emerges as the best predictor for achievement incurred purely in future time frame [31].

These findings indicate that two individuals with similar *h* are comparable in terms of their overall scientific achievement, even if their total number of papers or citations are quite different. Conversely, two individuals of the same scientific age can have a similar number of total papers or citation counts but very different *h* values. In this case, the researcher with the higher *h* is typically viewed by the community as the more accomplished. Together, these results highlight the key strength of the *h*-index: When evaluating scientists, it gives an easy but relatively accurate estimate of an individual's overall scientific achievements. Yet at the same time, we must also ask: What are the limitations of the *h*-index?

must account for the field-dependent nature of citations [43]. This can be achieved by the h_g -index, which rescales the rank of each paper n by the average number of papers written by author in the same year and discipline, n_o [43] or the h_s -index, which normalizes the h -index by the average h of the authors in the same discipline [44].

- **Time dependence.** As we discussed in Chapter 2.2, the h -index is time dependent. When comparing scientists in different career stages, one can use the m quotient (2.2) [26], or contemporary h -index [45].
- **Collaboration effects.** Perhaps the greatest shortcoming of the h -index is its inability to discriminate between authors that have very different coauthorship patterns [46–48]. Consider two scientists with similar h indices. The first one is usually the intellectual leader of his/her papers, mostly coauthored with junior researchers, whereas the second one is mostly a junior author on papers coauthored with eminent scientists. Or consider the case where one author always publishes alone whereas the other one routinely publishes with a large number of coauthors. As far as the h -index is concerned, all these scientists are indistinguishable. Several attempts have been proposed to account for the collaboration effect, including fractionally allocating credit in multi-authored papers [48–50], and counting different roles played by each coauthor [51–54] by for example differentiating the first and last authorships. Hirsch himself has also repeatedly acknowledged this issue [46, 47], and proposed the h_α -index to quantify an individual's scientific leadership for their collaborative outcomes [47]. Among all the papers that contribute to the h -index of a scientist, only those where he or she was the most senior author (the highest h -index among all the coauthors) are counted toward the h_α -index. This suggests that a high h -index in conjunction with a high h_α/h ratio is a hallmark of scientific leadership [47].

In addition to these variations of the h -index, there are other metrics to quantify the overall achievement of individual scientists, including the i_{10} -index, used exclusively by Google Scholar [55], which computes the number of articles with at least 10 citations each; or the SARA method [56], which uses a diffusion algorithm that mimics the spreading of scientific credits on the citation network to quantify an individual's scientific eminence. Despite the multitude of metrics attempting to correct the shortcomings of the h -index, to date no other bibliometric

index has emerged as preferable to the *h*-index, cementing the status of the *h*-index as a widely used indicator of scientific achievement.

As we dug deeper into *h*-index and the voluminous body of work motivated by it, it was easy to forget a perhaps more important point: No scientist's career can be summarized by a single number. Any metric, no matter how good it is at achieving its stated goal, has limitations that must be recognized before it is used to draw conclusions about a person's productivity, the quality of her research, or her scientific impact. More importantly, a scientific career is not just about discoveries and citations. Rather, scientists are involved in much broader sets of activities including teaching, mentoring, organizing scientific meetings, reviewing, and serving on editorial boards, to name a few. As we encounter more metrics for scientific eminence, it's important to keep in mind that, while they may help us understand certain aspects of scientific output, none of them alone can capture the diverse contributions scientists make to our community and society [57, 58]. Just as Einstein cautioned: "Many of the things you can count, don't count. Many of the things you can't count, do count."

Therefore, we must keep in mind that the *h*-index is merely a proxy to quantify scientific eminence and achievement. But the problem is, in science, status truly matters, influencing the perception of quality and importance of one's work. That's what we will focus on in the next chapter, asking if and when status matters, and by how much.

3

THE MATTHEW EFFECT

Lord Rayleigh is a giant of physics, with several laws of nature carrying his name. He is also known beyond the profession thanks to Rayleigh scattering, which answers the proverbial question, “Why is the sky blue?” Rayleigh was already a respected scientist when, in 1886, he submitted a new paper to the *British Association for the Advancement of Science* to discuss some paradoxes of electrodynamics. The paper was promptly rejected on the grounds that it did not meet the journal’s expectation of relevance and quality. Yet, shortly after the decision, the editors reversed course. Not because anything changed about the paper itself. Rather, it turns out that Rayleigh’s name had been inadvertently omitted from the paper when it was first submitted. Once the editors realized it was Rayleigh’s work, it was immediately accepted with profuse apologies [59, 60]. In other words, what was initially viewed as the scribblings of some “paradoxe,” suddenly became worth publishing once it became clear that it was the work of a world-renowned scientist.

This anecdote highlights a signaling mechanism critical in science: the role of scientific reputation. Robert K. Merton in 1968 [60] called this the Matthew effect after a verse in the biblical Gospel of Matthew pertaining to Jesus’ parable of the talents: “For to everyone who has will more be given, and he will have an abundance. But from the one who has not, even what he has will be taken away.” The Matthew effect as a concept has been independently discovered in multiple disciplines over the last century, and we will encounter it again in Chapter 17, when we discuss citations. In the context of careers, the

Matthew effect implies that a scientist's status and reputation alone can bring additional attention and recognition. This means that status not only influences the community's perception of the scientist's credibility, playing an important role in how her work is evaluated, but it also translates into tangible assets – from research funding to access to outstanding students and collaborators – which in turn further improve her reputation. The goal of this chapter is to unpack the role of the Matthew effect in careers. When does it matter? And to what extent?

3.1 What's in a Name?

The Internet Engineering Task Force (IETF) is a community of engineers and computer scientists who develop the protocols that run the Internet. To ensure quality and functionality, engineers must submit all new protocols as manuscripts that undergo rigorous peer review. For a while, each manuscript included the name of every author. However, beginning in 1999, some manuscripts replaced the full author list with a generic “et al.,” concealing the name of some authors from the review committee.

By comparing cases where well-known authors were hidden by the et al. label with those where the hidden names were little-known, researchers effectively conducted a real-world Lord Rayleigh experiment [61]. They found that when an eminent name was present on a submission, like the chair of a working group, which signals professional standing, the submission was 9.4 percent more likely to be published. However, the “chair effect” declined by 7.2 percent when the senior author's name was masked by the et al. label. In other words, name-based signaling accounts for roughly 77 percent of the benefits of having an experienced author as a coauthor on the manuscript.

Interestingly, when the analysis was restricted to a small pool of manuscripts that were “pre-screened,” or closely scrutinized, the author name premium disappeared. This suggests that the status effect only existed when the referees were dealing with high submission rates. In other words, when the reviewers do actually read the manuscript, carefully judging their content, status signals tend to disappear.

Given the exponential growth of science, we frequently encounter the “too many to read” situations. Yet, typically, peer review is a rather involved process, with multiple rounds of communication

between authors and expert reviewers, suggesting that the status signaling may be less of a concern for scientific manuscripts. Indeed, through those rebuttals and revisions, an objective assessment of the work is expected to prevail. Yet, as we see next, the status effect is rarely eliminated.

Whether an author's status affects the *perceived* quality of his/her papers has been long debated in the scientific community. To truly assess the role of status, we need randomized control experiments, where the same manuscript undergoes two separate reviews, one in which the author identities are revealed and another in which they are hidden. For obvious ethical and logistical reasons, such an experiment is difficult to carry out. Yet, in 2017, a team of researchers at Google were asked to co-chair the program of the Tenth Association for Computing Machinery International Conference on Web Search and Data Mining (WSDM), a highly selective computer science conference with a 15.6 percent acceptance rate. The researchers decided to use the assignment as a chance to assess the importance of status for a paper's acceptance [62].

There are multiple ways to conduct peer review. The most common is the "single-blind" review, when the reviewers are fully aware of the identity of the authors and the institution where they work, but, the authors of the paper are not privy to the reviewer's identity. In contrast, in "double-blind" review, neither the authors nor the reviewers know each other's identity. For the 2017 WSDM conference the reviewers on the program committee were randomly split into a single-blind and a double-blind group. Each paper was assigned to four reviewers, two from the single-blind group and two from the double-blind group. In other words, two groups of referees were asked to independently judge the same paper, where one group was aware of who the authors were, while the other was not.

Given the Lord Rayleigh example, the results were not surprising: Well-known author – defined as having at least three papers accepted by previous WSDM conferences and at least 100 computer science papers in total – were 63 percent more likely to have the paper accepted under single-blind review than in double-blind review. The papers under review in these two processes were exactly the same, therefore, the difference in acceptance rate can only be explained by author identity. Similarly, authors from top universities had a 58 percent increase in acceptance once their affiliation was known. Further, for

cited again. This phenomenon is called preferential attachment, which we will discuss again in detail in Chapter 17. To see how an author's reputation affects the impact of her publications, we can measure the early citation premium for well-known authors [65]. For example, for a group of well-known physicists, their paper has acquired around 40 citations ($c_x \approx 40$) before preferential attachment turns on (Fig. 3.1). In contrast, for junior faculty in physics (assistant professors), c_x drops from 40 to 10. In other words, right after its publication, a senior author's paper appears four times more likely to be cited than a junior author's.

Figure 3.1 suggests that reputation plays an important role early on, when the number of citations is small (i.e., when $c < c_x$). Yet, with time, the reputation effect fades away, and the paper's long-term impact is primarily driven by mechanisms inherent to *papers* rather than their *authors*. In other words, well-known authors enjoy an early citation premium, representing better odds of their work to be noticed by the community. This leads to a leg-up in early citations. But with time, this reputation effect vanishes, and preferential attachment takes over, whose rate is driven primarily by the collective perception of the inherent value of the discovery.

The reputation boost discussed above is not limited to new papers. Eminence can spill over to earlier works as well, boosting their impact. Sudden recognitions, like receiving the Nobel Prize, allow us to quantify this effect. Consider, for example, John Fenn, who received the 2002 Nobel Prize in chemistry for the development of the




Figure 3.1 The cross-over effect of reputation on citations. The linear attachment rate breaks down for $c < c_x$, suggesting that additional forces provide a citation boost which elevates $c(t)$ to deviate from what is predicted by the pure preferential attachment mechanism. Datasets include 100 top-cited physicists, and another 100 highly prolific physicists. After Petersen et al. [65].

electrospray ionization technique. His original discovery, published in *Science* in 1989 [69], is Fenn's most cited work, collecting close to 8,500 citations by 2018 according to Google Scholar. But as his landmark paper started to collect citations at an exceptional rate following its publication, the citation rates of several of Fenn's older papers also started to grow at a higher pace. Analyses of 124 Nobel laureates show that this boost is common [70]: The publication of a major discovery increases the citation rates of papers the author published *before*. Interestingly, the older papers that enjoyed the citation boosts are *not necessarily related* to the topic of the new discovery. In other words, reputational signaling operates by bringing professional attention to the individual. Consequently, when an author becomes prominent in one area of science, her reputation may be extended to her other line of work, even in unrelated fields.

Box 3.2 From boom to bust: The reverse Matthew effect

If a major breakthrough blesses both past and future scholarship, what does a scandal do to a career? Scientists are certainly fallible, and the scientific community regularly confronts major mistakes or misconduct. These incidents lead to retractions of articles, particularly in top journals [71], where they receive enhanced scrutiny. To what degree does a retracted paper affect a scientific career? Are eminent authors affected more or less severely than their junior colleagues? While retractions are good for science, helping other researchers avoid false hypotheses, retractions are never good for the authors of the retracted paper: they experience a spillover, leading to citation losses to their prior body of work as well [72–74]. The negative impact is not distributed equally, however: Eminent scientists are more harshly penalized for their retracted papers than when retractions happen to their less-distinguished peers [74]. Importantly, this conclusion only holds when the retractions involve fraud or misconduct. In other words, when the retraction is perceived to be the consequence of an “honest mistake,” the penalty differential between high- and low-status authors disappears [74].

When a senior and junior scientists are on the same retracted paper, however, the status penalty becomes quite different [75]: Senior authors often escape mostly unscathed, whereas their junior collaborators carry the blame, sometimes even to a career-ending degree. We will return to this effect in Chapter 13, where we explore the benefits and the drawbacks of collaborative work.

3.3 Is it Really the Matthew Effect After All?

Great scientists are seldom one-hit wonders [60, 76]. Newton is a prime example: beyond the Newtonian mechanics, he developed the theory of gravitation, calculus, laws of motion, optics, and optimization. In fact, well-known scientists are often involved in multiple discoveries, another phenomenon potentially explained by the Matthew effect. Indeed, an initial success may offer a scientist legitimacy, improve peer perception, provide knowledge of how to score and win, enhance social status, and attract resources and quality collaborators, each of these payoffs further increasing her odds of scoring another win. Yet, there is an appealing alternative explanation: Great scientists have multiple hits and consistently succeed in their scientific endeavors simply because they're exceptionally talented. Therefore, future success again goes to those who have had success earlier, *not* because of advantages offered by the previous success, but because the earlier success was indicative of a hidden talent. The Matthew effect posits that success *alone* increases the future probability of success, raising the question: Does status dictate outcomes, or does it simply reflect an underlying talent or quality? In other words, is there really a Matthew effect after all?

Why should we care about which is the more likely explanation, if the outcome is the same? Indeed, independent of the mechanism, people who have previously succeeded are more likely to succeed again in the future. But, if innate differences in talent is the only reason why some people succeed while others don't, it means that the deck is simply stacked in favor of some – at the expense of others – from the outset. If, however, the Matthew effect is real, each success you experience will better your future chances. You may not be Einstein, but if you are lucky to get that early win, you may narrow the gap between yourself and someone of his eminence, as your success snowballs.

Unfortunately, it is rather difficult to distinguish these two competing theories, as they yield similar empirical observations. One test of these contrasting hypotheses was inspired by the French Academy's mythical "41st chair." The Academy decided early on to have only 40 seats, limiting its membership to 40 so-called "immortals," and would only consider nominations or applications for new members if one of the seats became vacant through the death of a member. Given this restriction, many deserving individuals were never elected into the Academy, being eternally delegated to the 41st chair. It's a crowded

seat, shared by true immortals like Descartes, Pascal, Molière, Rousseau, Saint-Simon, Diderot, Stendahl, Flaubert, Zola, and Proust [60]. At the same time, many of those who did hold a seat in the esteemed club are (unfortunately) utterly irrelevant to us today. With time, the 41st chair became a symbol of the many talented scientists who *should* have been, but were never, recognized as giants of their discipline.

But, does it actually matter if someone is formally recognized or not? Indeed, how does the post-award perception of major prizewinners compare to scientists who had comparable performance, but who were not officially recognized? In other words, how does the career of those that occupied the 41st chair differed, had they been elected to the French Academy? The answer is provided by a study, exploring the impact of a major status-conferring prize [77].

As a prestigious private funding organization for biomedical research in the United States, the Howard Hughes Medical Institute (HHMI) selects “people, not projects,” generously supporting scientists rather than awarding them grants for specific peer-reviewed research proposals. The HHMI offers about US\$1 million per investigator each year, providing long-term, flexible funding that allows awardees the freedom to follow their instincts, and if necessary, change research directions. Beyond the monetary freedom, being appointed an HHMI investigator is seen as a highly prestigious honor. To measure the impact of the HHMI award, the challenge is to create a control group of scientists who were close contenders but who were not selected for the award and compare their scientific outputs with those of the HHMI investigators.

But, let’s assume that we identify this control group of scientists, and do find evidence that HHMI investigators have more impact. How can we know that the difference is purely because of their newfound status? After all, the US\$1 million annual grant gives them the resources to do better work. To sort this out, we can focus only on articles written by the awardees *before* they received the award. Therefore, any citation differences between the two groups couldn’t be simply the result of the superior resources offered to awardees. Sure enough, the analysis uncovered a post-appointment citation boost to *earlier* works, offering evidence that in science, the haves are indeed more likely to have more than the have-nots.

This success-breeds-success effect is not limited to HHMI investigators. When a scientist moves from a laureate-to-be to a Nobel laureate, her previously published work – whether of Nobel prize-winning caliber or not – gathers far more attention [78]. Once again,

like the case of John Fenn discussed above, a person’s previous work doesn’t change when she becomes an HHMI investigator or a Nobel laureate. But with new accolades casting warm light on her contribution, attention to her work increases.

Interestingly, though, strictly controlled tests suggest that status has only a modest role on impact, and that role is limited to a short window of time. Consistent with theories of the Matthew effect, a prize has a significantly larger effect when there is uncertainty about article quality, and when prizewinners are of (relatively) low status at the time of the award. Together, these results suggest that while the standard approach to estimating the effect of status on performance is likely to overstate its true influence, prestigious scientists do garner greater recognition for their outputs, offering further support for the Matthew effect.

Box 3.3 Causal evidence for the Matthew effect: Field experiments

Randomized experiments offer the best way to untangle the role of status from individual differences such as talent. We can select two groups – a control and a treatment group – and randomly assign an advantage to some while denying it to others. If success is allocated independent of prior success or status, any discrepancy in the subsequent advantage of recipients over non-recipients can only be attributed to the exogenously allocated early success.

While we can’t assign life-altering awards or grants to randomly chosen scientists [79], we *can* explore the phenomenon using experiments carried out in real-world settings where the intervention causes minimal harm. This is what Arnout van de Rijt and his collaborators did in a series of experiments [80, 81]. They randomly selected the most productive Wikipedia contributors within a subset of the top 1 percent of editors and randomly assigned them to one of two groups. Then they gave out “barnstars” to the experimental group – an award used within the community to recognize outstanding editors, while leaving the control group unrecognized. As shown in Fig. 3.2, prior to intervention, the activities of the two groups are indistinguishable, as they were drawn randomly from the same sample of productive editors. Yet once the fake barnstars were bestowed on the experimental group, the awardees exhibited more engagement than their peers in the control group, demonstrating greater sustained productivity and less likelihood of discontinuing their editorial




Figure 4.1 The graying of science. Changes in the percentage of NIH RoI grant recipients, aged 36 and younger and aged 66 and older, 1980–2010. After Alberts et al. [82].

4.1 When Do Scientists Do their Greatest Work?

The earliest investigation into the link between a person’s age and exceptional accomplishment dates back to 1874, when George M. Beard estimated that peak performance in science and the creative arts occurred between the ages of 35 and 40 [84]. Subsequently, Harvey C. Lehman devoted around three decades to the subject, summarizing his findings in *Age and Achievement*, a book published in 1953 [85]. Since then, dozens of studies have explored the role of age in a wide range of creative domains, revealing a remarkably robust pattern: No matter what creative domain we look at or how we define achievement, one’s best work tends to occur around mid-career, or between 30 to 40 years of age [2, 66, 85–87].

Figure 4.2 shows the age distribution of signature achievements, capturing Nobel prizewinners and great technological innovators of the twentieth century [88]. The figure conveys three key messages:

- (1) There is a large variance when it comes to age. While there are many great innovations by individuals in their 30s (42%), a high fraction contributed in their 40s (30%), and some 14 percent had their breakthrough beyond the age of 50.
- (2) There are no great achievers younger than 19. While Einstein had his *annus mirabilis* at the tender age of 26, and Newton’s *annus mirabilis* came even earlier, at the age of 23, the Einsteins and Newtons of the world are actually rare, because only 7 percent of




Figure 4.2 Age distribution of great innovation. The plot shows the age of innovators at the moment of their great innovation, combining all twentieth-century observations. After Jones [88].

the sample have accomplished their great achievement at or before the age of 26.

- (3) The Nobel laureates and the inventors come from two independent data sources, with only 7 percent overlap between the two lists. Yet, the age distributions of these two samples are remarkably similar.

Thus, Fig. 4.2 demonstrates that scientific performance peaks in middle age [2, 66, 85–87]. The life cycle of a scientist often begins with a learning period, absent of major creative outputs. This is followed by a rapid rise in creative output that peaks in the late 30s or 40s and ends with a slow decline as he advances through his later years. These patterns are remarkably universal. Researchers have explored them in a variety of ways, identifying important scientists by their Nobel Prizes, by their listings in encyclopedias, and by their membership in elite groups like the Royal Society or the National Academies. No matter how you slice the data, the patterns observed in Fig. 4.2 remain essentially the same, raising two questions: Why does creativity take off during our 20s and early 30s? And why does it decline in later life?

4.2 The Life Cycle of a Scientist

4.2.1 The Early Life Cycle

A remarkable feature of a scientific career is the lack of contributions in the beginning of life [89]. Simply put, no 18 year old has managed to produce anything worthy of a Nobel. The early life cycle coincides with schooling, suggesting that the need for education may be