I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 16365

Why Has Science Become an Old Man's Game?

Christian Fons-Rosen
Patrick Gaulé
Taras Hrendash

AUGUST 2023

DISCUSSION PAPER SERIES

IZA DP No. 16365

Why Has Science Become an Old Man's Game?

Christian Fons-Rosen
UC Merced, MPI Munich and CEPR

Patrick Gaulé
University of Bristol and IZA

Taras Hrendash
CERGE-EI

AUGUST 2023

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Why Has Science Become an Old Man's Game?*

We investigate the causes and consequences of the aging of the scientific workforce. Using novel data on the population of US chemistry faculty members over fifty years, we find that the secular increase in the age of the academic workforce has been mainly driven by the slowdown in faculty hiring combined with later retirements. By contrast, changes in the age at which scientists start their careers only contribute to about 20% of aging. Hiring more new faculty members could rejuvenate the scientific workforce and boost scientific productivity.

JEL Classification: O31, J24, J26

Keywords: aging, science, universities, knowledge production

Corresponding author:
Patrick Gaulé
School of Economics
University of Bristol
12A Priory Rd
Bristol BS8 1TU
United Kingdom

E-mail: patrickgaule@gmail.com

* We appreciate helpful comments from Gaetan de Rassenfosse, Ina Ganguli, Margaret Kyle, Fabbio Montobbio, Paula Stephan and Fabian Waldinger, as well as from seminar and workshop participants at LMU Munich, Mines ParisTech, Universita Catolica, University College London, the University of Kent and the University of Sussex. All errors are our own. Gaule and Hrendash acknowledge financial support from the Czech Science Foundation (GACR grant no 16-05082S).

1 Introduction

Scientific knowledge is increasingly produced by older people. John Goodenough was 97 when he received the 2019 Nobel Prize for Chemistry. Between 2019 and 2023, he held a faculty position at the University of Texas at Austin, published ten articles in highly respected peer-reviewed journals, and filed four US patent applications. Beyond this (arguably extreme) example, the average age of NIH grant recipients has increased from 39 to 51 between 1981 and 2008 (Daniels 2015). Among faculty members in US chemistry departments, the mean age has increased from 37 in 1960 to 53 in 2015. Moreover, the age at which Nobel-prize-winning discoveries are made has risen steadily over the course of the 20th century (Jones 2009, Jones & Weinberg 2011). Observers ranging from Paul Romer to former NIH director Zerhouni have pointed to the aging of the scientific workforce as their major concern for the future of US innovation (Packalen & Bhattacharya 2019, Kaiser 2008).

Scientific discoveries can be made by scientists in middle age - Wilhelm Röntgen discovered X-rays at the age of 50 - or indeed later. However, scientific output tends to decline as scientists age. In our data, faculty member productivity declines markedly after the age of 60. Moreover, as we document in the text, scientists from more recent vintages tend to be more productive than scientists from previous vintages. When an older faculty gets replaced by a younger one, scientific output would tend to go up not just because scientists tend to be more productive when young (or middle-aged) but also because the younger scientist would be from a more recent cohort and have either higher human capital or higher intrinsic talent.

Understanding the causes of the aging of the scientific workforce is therefore important in forming appropriate policy responses. An influential explanation for the aging of the scientific workforce is the "burden of knowledge" hypothesis (Jones 2009). As the stock of human knowledge accumulates over time, new entrants need to spend more time training to reach the knowledge frontier. This leads to a secular increase in the age at which scientists start their careers and make key discoveries, empirical patterns that are observed among Nobel Prize winners (Jones 2009, 2010). From this perspective, policymakers should seek to improve the quality of training and preserve incentives to start scientific careers, but may want to refrain from giving large grants to young scientists (Jones 2011). An alternative explanation for the aging of the scientific workforce is the decline in the retirement rate of older scientists, a process facilitated by the elimination of mandatory retirement in US universities (Ashenfelter & Card 2002, Blau & Weinberg 2017). If declines in retirement rates drive the aging of the scientific workforce, policymakers might - depending on the productivity of older scientists - reconsider end-of-career incentives and policies.

It may seem that changing age at entry and retirement dynamics are the only factors changing the age composition of the scientific workforce, but this is not so. Entry age and exit dynamics clearly matter in determining the age composition of the scientific workforce, but a potentially important third factor is the number of people hired over time. Consider what would happen if the US government made a large investment in science, perhaps a new Apollo or Manhattan Project. As the total demand for scientific labor goes up, the extra positions would be disproportionately filled by younger people graduating from universities (as opposed to older individuals moving from non-research jobs to research jobs). As a consequence, the scientific workforce would immediately become younger. In the more distant future, however, if the hiring spree does not continue, the disproportionally large cohort of new entrants will age while staying active over decades, generating aging for the whole population. Such a scenario is not hypothetical: Indeed, there was a strong faculty hiring spree in the 1960s in the US that did not continue.[1]

In this paper, we build a demographic model of the US academic workforce to shed light on the causes of its aging. The model leverages novel data on the population of US chemistry faculty members between 1960 and 2010. Having set up the model to mimic observed empirical patterns in the data, we can then use it to quantify the importance of various channels - changes in entry age; retirement dynamics; and ratio of entrants to exits - to the aging of the workforce. For instance, we can ask what would happen to the age composition of faculty members if entry ages did not increase (stayed at their initial, as of 1960, level) but retirement dynamics and hiring patterns had evolved as they did.

We find that changes in the ratio of entrants to exits are the main driver in the aging of our sample. In the data, the mean age of chemistry faculty members rose from 39 in 1960 to 53 years in 2010. Although the age at which individuals become faculty members has indeed increased (as predicted by the "burden of knowledge" hypothesis), this only accounts for about a quarter of the increase in mean faculty age. In contrast, changes in the number of people hired over time appear to be a major factor in the aging of faculty members: they explain half of the observed aging, and close to 80% when combined with changes in retirement dynamics.

We provide several pieces of institutional context to shed further light on what happened. In the 1960s, federal R&D spending was booming due to the Cold War, whereas there was also a surge in demand for higher education from the baby boom generation. As a result, each retiring faculty member was replaced by three new hires, whereas subsequently, the ratio of entrants to new hires declined significantly. Meanwhile, in the 1990s, the abolishment of mandatory retirement for faculty members led to decreases in retirements, just at the same time as the disproportionally large cohort of faculty members born in the 1930s and hired in the 1960s approached retirement ages.

[1] In an even starker case, an unusually large number of faculty members were appointed in Italy in the year 1980. For details, see Lissoni et al. (2011).

The combination of these historical events and policies appears to have played a major role in the aging of the scientific workforce.

Our results have a number of implications for understanding the role of policy in shaping the demographic composition of the academic workforce. First, the age composition of the academic workforce need not simply reflect fundamental trends in the nature of knowledge production (or of overall societal aging). Instead, it may (and perhaps should) be seen as the result of past and present policy choices, particularly in terms of the number of people hired. Second, hiring more new faculty could generate disproportionate returns. As new hires tend to be young, they generate a sort of 'demographic dividend' whereby their productivity is higher than the average faculty, at least in the beginning. Given that currently, less than 10% percent of graduate students become faculty members and that faculty positions are generally considered attractive (Ganguli, Gaule & Vuletic 2022), there is no shortage of talented young scholars interested in taking up faculty positions.

This paper contributes to the literature on the causes of the aging of the scientific workforce (Jones 2009, Jone 2010, Blau & Weinberg 2017), highlighting a new cause (a slowdown in hiring) as quantitatively important in the post-World War II era. More generally, this paper also contributes to the literature on the composition of the population of scientists and inventors (Aghion et al. 2017, Bell et al. 2019, Agarwal & Gaule 2020, Koning, Samila & Ferguson 2021, Celik 2023). This literature highlights that increasing the representation of underrepresented groups (women, individuals from developing countries, or disadvantaged socioeconomic backgrounds) in science and innovation could increase the pace of innovation. Similarly, our paper suggests that increasing the number of young faculty members (who are not just at the peak of their abilities, but also well trained and selected by a highly competitive academic labor market) could boost the productivity of academic research. This is particularly important in a world where ideas may be getting harder to find (Bloom et al. 2000) and further gains in research intensity and educational attainment may be difficult to obtain (Jones 2002).

The remainder of the paper is organized as follows. Section 2 discusses the institutional context, including US science funding after World War II, faculty retirement policies, and chemistry as a scientific discipline. Section 3 describes the data, Section 4 presents the methodology of the simulation, and Section 5 presents the results. Finally, Section 6 concludes.

2 Institutional context

In this section, we provide four pieces of institutional context relevant to understanding the causes and consequences of the aging of the scientific workforce. First, we provide some descriptive statistics on the aging of the US scientific workforce from 1960 to 2010. Second, we describe how funding for basic research has evolved in the US from World War II to the present, as this will be relevant to understanding how hiring patterns have changed over time. Third, we describe how retirement policies for university faculty evolved over time. Finally, because our analysis will focus on chemistry, we provide some background information on how knowledge production is organized in this discipline.

2.1 Age structure of the US scientific workforce

The US scientific workforce is aging. Many observers have noted changes in the age composition of NIH grantees (see, e.g., Daniels 2015). For example, from 1980 to 2010, the median age of NIH R01[2] grant recipients has increased from 40 to 50 years of age. Among chemistry faculty members, the median age has increased from 41 years in 1960 to 51 in 2010.[3] Even in mathematics, traditionally seen as the preserve of the young, the median age of authors has increased from 35 in 1960 to 44 in 2010.[4]

(Insert Figure 1 about here.)

To put the aging of the US scientific workforce in context, it is useful to consider the aging of the US labor force as a whole. It is true that the US labor has aged from 1980 to 2010, with the median age rising from 35 to slightly above 40. However, when considering a longer time period, the trends are less clear: the US workforce had a median age slightly above 40 in 1960 already. Moreover, even within 1980-2010, the rate of aging of the US workforce (with the median age rising by less than 1.5 years by decade) is clearly lower than for the NIH grantees (with the median age rising by more than 3 years by decade).

[2]Research Project Grants (R01) is the original and historically oldest grant mechanism used by NIH. They are meant to support a specified project to be performed by a principal investigator in an area representing the investigator's specific interest and competencies, based on the mission of the NIH. NIH R01 grants constitute the bulk of NIH external grant giving.

[3]This figure is based on our dataset of US chemistry faculty members that we describe in the next session.

[4]The figure on authors of mathematics articles is based on our own calculations using data from the Mathematics Genealogy Project.

2.2 US Science Funding after World War II

World War II had been a powerful demonstration of the practical utility of science and innovation. For example, the development of the radar gave the United Kingdom a key advantage in the Battle of Britain (a series of aerial battles between German and British air forces over the skies of Britain). Most ominously, World War II ended shortly after the deployment of the atomic bomb at Hiroshima and Nagasaki. The atomic bomb had been developed by the Manhattan Project, a massive research project mobilizing most of the best physicists residing in America.

After World War II, a consensus emerged that basic research was important both for national prosperity and in the ongoing geopolitical rivalry with the USSR. The argument was best encapsulated in Vannebar Bush's influential report 'Science, The Endless Frontier' (Bush 1945). In 1957, the USSR launched the first satellite into space, sparking fears that the United States was falling behind in technology. This 'Sputnik moment' provided further impetus for investment in US science and education.

After 1945, the US federal government started supporting basic research on a systematic basis and public R&D investments became very substantial. In 1950, a brand new federal agency, the National Science Foundation, was funded. In addition, the National Institutes of Health were instituted, building on prewar institutions but with a growing mandate and budget. Federal funding for R&D grew tenfold from 1949 to 1962. As a share of GDP, federal funding for R&D reached 4% in the 1960s before steadily decreasing to 1% in 1980 and below 0.5% currently (see Figure A-1).

The influx of federal funding resulted in a bonanza for US universities. In an anecdote recounted by Paula Stephan (Stephan 2018), federal grant agencies sent representatives to universities to encourage faculty members to apply for funding. Moreover, US universities also benefited from a surge in demand for education from the 'baby boom' generation going to college with the number of high school graduates enrolled in college doubling between 1960 and 1970 (Ashenfelter & Card 2002). In response to increased demand for both research and teaching, US universities hired large numbers of new faculty members, a pattern we clearly see in our data as we will later discuss.

Federal R&D funding never recovered after falling from its peak in the 1960s. It was briefly interrupted by the doubling of the NIH budget between 1998 and 2003 but this increase was not permanent. Additionally, the rapid increase in NIH spending and subsequent deceleration created substantial adjustment problems in the market for research (Freeman & Van Reenen 2009).

2.3 Faculty retirement policies

Faculty retirement policies changed substantially in the postwar period.[5] As of the 1970s, mandatory retirement was nearly universal among postsecondary faculty in the United States. However, the United States put in place bans on age discrimination in the workplace, which eventually resulted in mandatory faculty retirement policies being considered unlawful.[6] Mandatory retirement was effectively eliminated in 1994 with the US from then on offering -uniquely among developed countries- lifetime employment security for tenured faculty members (Ashenfelter & Card (2002).

Although the abolition of compulsory retirement at 70 was a notable policy change, its practical implications were somewhat blunted by a number of factors. First, some universities did not have compulsory retirement policies for their faculty members, whilst others could let particular individuals continue beyond 70 on a case-by-case basis. Second and more importantly, individuals may choose to voluntarily retire before 70 irrespective of compulsory retirement policy. In the data from Ashenfelter and Card (2002), only half of faculty members employed at age 65 are still employed at 69, both before and after the age cap was removed. Still, there were large spikes in retirement rates at 70 and 71 that disappeared after the cap was removed (ibid.).

2.4 Chemistry

Chemistry is the scientific study of the properties and behavior of matter. While chemistry is a physical science, large parts of it relate to living organisms, so chemistry also relates closely to life science. Apart from some smaller subdisciplines, such as theoretical chemistry, chemistry is largely a lab-based science. Besides being a physical space with instruments and research materials, the lab is also an organizational structure where a faculty member (principal investigator) obtains funding for the lab, directs research projects, and appears as a coauthor on all publications. Although faculty members in elite institutions are usually supported by a relatively large number of graduate students, postdoctoral researchers, and technical staff working in their lab, this varies to an extent between time and over institutions.

Traditionally, faculty members/lab directors appear as the last authors of scientific publications. The graduate student or postdoc who has done most of the day-to-day work on the research project

[5]This subsection is largely based on Ashenfelter & Card (2002)

[6]The Age Discrimination in Employment Act prohibiting age discrimination for individuals above 1960 was introduced in 1967. Initially, university professors had been an exempt category, but this changed with subsequent amendments: the compulsory retirement age in higher education was raised to 70 in 1982, and completely abolished effective January 1994.

typically receives first authorship. Authors in the middle of the authorship list have normally made relatively minor contributions to the projects, though this clearly varies across papers.

Research by chemistry faculty members is supported by a mix of federal, state, and industry sources. Depending on their specialties, chemistry faculty may apply to the National Institutes of Health, the National Science Foundation, or other federal agencies. Given that research in chemistry often has practical applications, industry funding through R&D contracts is common.

3 Data

In order to investigate the causes and consequences of the aging of the scientific workforce, we assembled an original data set that combines multiple sources. The core and most original component is the longitudinal database of academic scientists derived from the ACS directory (described below). We complement this longitudinal database with information on publications.

The ACS Directory. Our main data source is the 'ACS Directory of Graduate Research' (hereafter: ACS Directory). The ACS Directory was a biennial publication of the American Chemical Society that ran from 1953 to 2015, when it was discontinued. The purpose of the publication was to provide prospective graduate students with information on US chemistry departments offering Ph.D. degrees. Initially published as a book, the ACS directory was later also diffused into an electronic version (first on CD-ROM and then on a dedicated website). The publication included lists of faculty members with their names, year of birth, sex, educational history, and current affiliation among other information. Information on the year of birth is particularly useful as it is otherwise hard to find on a systematic basis.

(Insert Figure 2 about here.)

Building a dataset of US chemistry faculty members based on the ACS directory. We procured the 1961, 1971, 1981, 1991, 2001 and 2011 editions of the ACS Directory. For the 1991, 2001 and 2011 editions, we used electronic versions.[7] For the earlier versions (1961, 1971 and 1981), we digitalized the respective copies (for a total of more than 3 thousand pages) using optical character recognition (OCR) software, as well as freelancers, to correct OCR mistakes. The resulting data give us six snapshots between 1961 and 2011 of the distribution of faculty members in US chemistry departments. We then generated a longitudinal database linking individual faculty members across the different editions of the book, using names, birth years, and educational histories to

[7] A version of that database covering the years 1993 to 2009 has been used in Gaule (2014), Gaule & Piacentini (2018), Catalini, Fons-Rosen & Gaule (2020) and Ganguli, Gaule & Vuletic (2022).

create links. This longitudinal dataset also enables us to obtain a proxy for entry and exit from the profession, through the year of the first listing in the directory and the year of the last listing, respectively.

Publication data. We use publication data from Scopus, a major bibliographic database run by Elsevier. Our dataset covers scientific articles published between 1960 and 2010 in 500 scientific journals for a total of more than 1 million publications. The journals in our database include all the major chemistry journals, as well as multidisciplinary journals (*Nature*, *Science*, *PNAS*) and journals from other scientific disciplines where chemists might publish (physics, biology, etc.). For each article in our database, we know the names of the authors and their affiliations, as well as the total number of citations received (from publication to 2016).

Matching individuals to publications and constructing productivity measures. We then matched the names of faculty members with their publications in close to 500 chemistry journals using bibliographic data from Scopus. The match is based on last name, first and middle initials, and university (the last which we know from the faculty data). To generate a proxy for high achievement, we tagged papers that were in the top centile of the citation distribution for papers published in a given year. We then created two measures of scientific productivity at the individual scientist level: the number of papers written and the number of papers hitting the top centile of the citation distribution.

4 Aging and scientific output

This section discusses how the age composition of the scientific workforce may affect its productivity. We first review the prior literature on age and scientific output before studying the link between age and scientific output in our own data. Finally, we consider the differences between cohorts in productivity, which are another channel through which the age composition of faculty may affect its productivity.

4.1 Prior literature on age and scientific output

In this subsection, we briefly review previous work on aging and scientific output. Most relevant for our purposes is Levin & Stephan (1991) who analyze scientific output over the lifecycle across six science areas in models controlling for motivation and ability through individual fixed effects. They find that publishing activity initially increases but then declines somewhat in mid-career, with the exception of particle physics, where there were no life-cycle effects. Similarly, only modest life-cycle effects are found in others such as Diamond (1986) on the productivity of Berkeley mathematicians and Turner and Mairesse (2003) on French condensed matter physicists. However, Azoulay et al. (2006) find pronounced life-cycle effects when considering the patenting behavior of academic life scientists, with mid-career academics being much more likely to patent than younger and older faculty members. Yu et al. (2022) make the important point that studying the age-output relationship is subject to a selective attrition problem: the most able scientists tend to have longer careers. The existence of selective attrition suggests that traditional estimates of the age-output relationship may understate the extent to which aging decreases productivity.

Related literature investigates the age at which great achievements are made. Considering Nobel laureates during the period 1900-1992, Stephan & Levin (1993) find that while doing prize-winning work does not require extraordinary youth, the odds decrease markedly in midlife. In more recent -and highly influential- work, Jones (2010) also documents an inverse U-shaped between age and great scientific achievement, but emphasizes changes over time: Nobel Prize-winning research is performed at an average age that is 6 years older at the end of the 20th century than it was at the beginning. Jones and Weinberg (2011) further document that theorists tend to make Nobel prize-winning discoveries at an earlier age than empiricists in physics, chemistry, and medicine.

Overall, the literature suggests that while life-cycle effects tend to be modest when considering publishing, they may be larger for patenting and great achievements. A smaller and more recent literature has documented interesting patterns in how the relationship between age and output has evolved over time, with age at great achievement rising.

4.2 Age-output relationship in our data

We start by investigating the relationship between age and scientific output in our data. We have the advantage of having large and comprehensive data on individuals, albeit in a single discipline. However, as discussed above, there is a long tradition of estimating the relationship between age and scientific output, so this is not our main focus.

For simplicity, we use the number of published papers as the measure of output.[8] Figure 3 plots the mean productivity by age. Consistent with prior literature, we find an inverse U relationship between age and productivity: productivity increases in the beginning before reaching a long plateau, roughly between the ages of 40 and 60, and subsequently declining. These results, along with those of previous studies as described in the previous subsection, provide *prima facie*

[8] Note that we only count publications in a set of 500 journals that includes all major chemistry journals, as well as top multidisciplinary journals such as Science, Nature or Proceedings of the National Academy of Sciences, and top journals from neighboring disciplines (e.g. physics, biology). See the Data section for details.

evidence that the age distribution of scientists matters for how much science is produced in the aggregate.

(Insert Figure 3 about here)

4.3 Productivity differences across cohorts

We now move on to study productivity differences across cohorts of scientists. These may arise for at least two reasons. First, the quantity and quality of human capital acquired during the training (and, in particular, the doctoral degree) that scientists go through before becoming faculty may vary over time. Relatedly, even if the quantity and quality of training stayed constant, the knowledge frontier shifts over time, which affects the relevance of the human capital acquired (Jones 2009). Second, over time, the underlying talent of individuals entering academia may have changed. Both Ph.D. admission and faculty hiring are centered around assessing an individual's potential to produce new knowledge. It seems likely that the competitiveness of these processes may have increased over time. For instance, based on our own data, in the 1960s one in five graduate students became faculty in a research-intensive university. By 2015, this figure has dropped to around five percent (Ganguli, Gaule & Vuletic 2022). Similarly, the number of talented foreign students who come to the United States has increased over time. For example, there were virtually no Chinese doctoral students in the US before 1985 (Doran & Yoon 2018). By the 2000s, around 1 in 6 doctoral students in US doctoral programs were Chinese, and they tended to outperform native students (Gaule & Piacentini 2013).

(Insert Figure 4 about here)

Figure 4 displays the output over age distinguished by birth cohorts (in 10-year bins). A clear trend emerges from the figure: each cohort appears more productive than the previous one at the same age.[9] In Table 1, we report the results of formal hypothesis tests and reject the null of equal productivity between cohorts at a given age. This finding contrasts with Levin & Stephan (1991) who did not find differences across cohorts, but we have both a considerably larger sample and a longer time coverage.

One difficulty in interpreting productivity differences between cohorts is that they may reflect general trends in the difficulty of producing and publishing research over time. Advances in instrumentation and computing technology, for example, may facilitate the production of scientific

[9] Moreover, over time it seems that productivity tends to peak earlier

knowledge over time. Relatedly, the number of outlets where scientists can publish has increased over time, including among the restricted set of journals we consider. To address this difficulty, we residualize individual output by year fixed effects (Figure A-2), or by a linear time trend (Figure A-3). When we remove time trends in this way, we find essentially the same pattern.

We perform a number of additional robustness checks. One may be concerned about the productivity measure we have chosen, particularly given that research has become more collaborative over time (Wuchty, Jones & Uzzi 2007). To address this, we implement fractional counting (where each author of a 2-authored article would receive a credit of 0.5) and find similar results (Figure A-4). Finally, one might prefer an output that places a greater emphasis on quality. Although journal impact factors are not readily available for our entire time period, we do have citations. We define highly cited papers as those belonging to the top centile in terms of the number of citations received among papers published in a given year. When considering these highly cited papers as the measure of productivity, we obtain similar (albeit noisier) results on productivity differences between cohorts (see Figure A-5).

4.4 Taking stock

The well-known relationship between age and scientific output suggests that a scientific workforce with a large share of older scientists would tend to be less productive. In our data, we find that while middle-aged chemists are productive, scientific productivity decreases markedly after the age of 60. Moreover, apart from the individual age-output relationship, more recent cohorts appear to be systematically more productive than previous ones at the same age, possibly due to tighter selection to enter the scientific and academic workforce. When an older faculty gets replaced by a younger one, the scientific output would tend to go up not just because scientists tend to be more productive when young (or middle-aged), but also because the younger scientist would be from a more recent cohort and have either higher human capital or higher intrinsic talent. Overall, the evidence in this section suggests that the age composition of scientists matters for how much scientific knowledge is produced.

5 Simulation Methodology

Overview. We build a demographic model simulating the evolution of the age of the US academic workforce. We use the model to quantify the importance of three channels: changes in entry age; retirement dynamics; and ratio of new hires to exits, to the aging of the workforce.

Data inputs. The model leverages the fact that we observe (close to) the population of US academic faculty in chemistry over 50 years. The simulation starts with the sample of 3,250 scientists who were active as of 1961 in the data. Six editions of the directories are used with an interval of ten years between consequential snapshots, leading us to have a range of years between 1961 and 2011. For the purposes of the simulation, we focus only on one attribute of the individuals, their birth years.

Modeling the distribution of the entry age and the retirement probabilities. In each decade, we observe the actual distribution of the age of entrants, and we slightly coarsen (simplify) this distribution into particular percentiles (10, 50, 75, 95, 100) as described in the Appendix. Then we keep track of the simplified distribution in each decade. Similarly, for exits, we compute the share of individuals (in 10-year age bins) who no longer appear in the next edition (a decade later) of the directory. Finally, for faculty lines, we measure the ratio of new hires in the next decade to the number of exits in the current decade. We refer to these quantities (the age distribution of entrants, the exit probabilities by age, and the ratio of new hires to exits between adjacent periods) as *moments*.

Running the baseline simulation. In our baseline simulation, we aim to replicate the observed data. In each period, we retire a random subset of current scientists with the proportion of retirements varying with their age and observed coarsened retirement probabilities (as described in the previous paragraph). We then determine how many faculty members to hire based on the observed ratio of new hires to exits.[10] Finally, the age distribution of the new hires will be given by the coarsened distribution of the age of entrants.

An example. To illustrate the procedure described above, let us consider a specific example, the transition from 1955 to 1965 in our baseline simulation. We have 3,250 scientists listed in 1955 and seek to model the number and age distribution in 1965.[11] We start by retiring 95% of 70-year-olds, 76% of 60-year-olds, 31% of 50-year-olds, 29% of 40-year-olds with 95%, 76%, 31%, 29% given by observed decade-specific exit probabilities. This yields a total of 1,245 exits. In addition to the 1,245 exits that we replace, there are also 1,245*2 new faculty lines to fill, since the ratio of new hires to exits was 3 between 1955 to 1965. Thus, we introduce 1245 * 3 = 3735 new scientists to the data. We then assign an age to each of these new scientists on the basis of the coarsened distribution of the age of entrants in the actual data of the corresponding decade. For example, the youngest 10% are 24 years old, the next 30% youngest are 30 years old, etc. We then reproduce the same procedure to simulate the 1975 distribution from the 1965 distribution we just created.

[10] If the ratio of exits to new hires is above one, we fill new faculty lines in addition to replacing the faculty members exiting. If the ratio is below one, we replace only a portion of faculty members exiting.

[11] Technically, our first observed data is from the book published in 1961 but we treat it as giving us a snapshot of the faculty population in 1955. See the Appendix for a detailed discussion of this assumption.

Running a counterfactual simulation. To run counterfactual simulations, we run the same procedure as the baseline simulation, but instead of using the moments based on actual time-varying data, we use one or more moments from the 1960s. We may, for instance, assume that the ratio of new hires to exits keeps the value it had in the 1960s (i.e., a value equal to 3 whereas, in reality, it moved down to 1.4 in the 1970s and kept evolving afterward as described in Figure 7). So, whereas, in the baseline simulation, we would use the actual time-varying ratio (3, 1.4, ...); in the counterfactual simulations, we would instead hire enough scientists to replace each exit with three new hirings between 1965 and 1975, as well as between 1975 and 1985, etc.

Switching different channels on or off. In our model, aging can occur for one of three different reasons: changes in the ratio of entrants to exits, changes in retirement probabilities, and changes in the age of entrants. By using either the initial value (1960s) or time-varying moments, we can activate or deactivate each of these channels. For instance, we can activate the change in the ratio of entrants to exits while leaving the other two channels switched off, thereby shedding light on the contribution of this particular channel to aging. This would yield the counterfactual evolution of the age of faculty members in case the ratio of entrants to exits had evolved as it actually did, while the age of entrants and retirement probabilities stayed at their 1960s levels and thus did not contribute to aging. In the simulation results section, we consider all possible combinations of active and inactive channels.

6 Simulation results

6.1 Preliminaries

Our model is meant to replicate the observed increase in the mean age of faculty members from 1965 to 2005. In practice, however, the model may differ from the actual data because we make a number of simplifying assumptions. Figure 5 shows that our model very closely replicates the actual data when we activate all three channels.

(Insert Figures 5 and 6 about here)

We also check that the three channels of our model generate most of the aging observed in the data. Figure 6 shows the result of a simulation in which we turn off all three channels. Reassuringly, the mean age stays flat around 39 from 1965 to 2005.

6.2 Describing the moments

Before proceeding to the counterfactual simulation results, we briefly describe how our quantities of interest -the ratio of entrants to exits, the age of entrants, and retirement probabilities- evolve over time. During our period of interest, we observe substantial changes in the ratio of entrants to exits (see Figure 7). In the 1960s, the ratio had a value of three, meaning that each was replaced by three new hires. In the 1970s, the ratio had fallen to just 1.4, corresponding to a much slower expansion rate. In the 1980s, the ratio fell below one before modestly rebounding in the 1990s.

Regarding the age of entrants, we observe a clear upward trend with the mean age of entrants increasing from 32 to 36 years of age between 1965 and 2005 (see Figure 7).[12] While this upward trend is consistent with the burden of knowledge hypothesis, we note that in terms of magnitude, it can only explain part of the aging we observe: the mean age of faculty members increases by 14 years, while the mean age of new applicants only increases by four years. Finally, for exits, we observe decreases in the probability of exits for individuals in their 60s and 70s (see Figure A-6). For example, a 60-year-old in 1965 would have a 75% chance of exiting in the next decade, while the corresponding probability for a 60-year-old in 2005 is 60%.

(Insert Figure 7 about here.)

6.3 Simulations: switching on one channel at a time

We start the simulations by switching on one channel at a time while deactivating the other two, thus isolating the contribution of the activated channel to the aging of faculty members. Switching the entry age on explains 26.4% of the observed aging (see Table 2 or Figure A-7) while switching the retirement age on explains 6.4% (see Table 2 or Figure A-7). Clearly, either of these factors, while contributing to overall aging trends, only explains a relatively modest share of total aging.

Next, we use our model to investigate the effect of the changing ratio of entrants to exits on the aging of faculty members. Specifically, we assume that the age at entry and retirement probabilities stay at their 1960s levels while letting the ratio of entrants to exits evolve as it actually did. Figure 8 shows that the changing ratio of entrants to exits alone generates a substantial share (around half) of the aging we observe in the data over the whole period.

(Insert Figure 8 and Table 2 about here.)

[12]Note that, for simplicity, we show the mean age but our simulations take into account various centiles of the observed age distribution of new entrants in each decade.

In summary, when considering one channel at a time, the largest effect comes from the changing ratio of exits to entrants, followed by changing age at entry, and changing retirement patterns is only a distant third.

6.4 Simulations: switching on two channels at a time

We next consider what happens when we switch two channels on while leaving the third deactivated. The effect of switching on two channels may differ from the sum of the effects of turning on each channel individually. For instance, the effect of switching on changes in retirement probabilities (i.e. having lower retirement probabilities over time) would matter more if the population of scientists is older, which in turn depends on whether the other channels are switched on. Because such interactions are possible, considering two channels at a time may be insightful.

(Insert Figure 9 and Table 3 about here.)

Figure 9 shows how the mean age of faculty members would have evolved -according to our model- if the age of entrants stayed at 1960s levels, while the ratio of entrants to exits and the exit probabilities evolved as they actually did. The simulated line follows the actual data relatively closely, although around 20% of the variation in aging remains unexplained by the conjunction of changing ratio of entrants to exits, and changing exits probabilities.

Table 3 decomposes how different pairs of channels interact with each other. Each row corresponds to a combination of channels, and based on Table 2 we report the 'main effect' of each channel (i.e. the effect of turning on just that particular channel, as per the preceding subsection) as well as the total effect of turning both channels at the same time. As discussed earlier, the sum of the main effects may differ from the total effect when the two channels interact with each other. In our case, this arises when an older population of scientists (due to the large cohort of new entrants in the 1960s) combines with changes in retirement policies. Whereas changes in retirement probabilities alone only explain about 6% of the variation in aging, the interaction with changing ratio of entrants to exits explains another 20% of the variation in aging.

6.5 Taking stock

Prior literature has focused on the age of entry and changing patterns of retirement to explain the aging of scientists. However, we find that these channels, either alone or in combination, explain less than a third of the total variation in aging of chemistry faculty members between 1965 and 2005.

Instead, a third factor, changes in total faculty employment (which we capture through the ratio of entrants to exits), seems to explain a major part of the observed aging. In the 1960s, there were a large number of new entrants to the profession, likely due to the expansion in federal R&D investment, possibly combined with greater demand for college education by the baby boomer generation. These entrants were on average young and tended to stay, while the 1960s hiring boom did not continue. Initially (in the 1970s and 1980s), the large cohorts of new entrants from the 1960s led to a decrease in the mean age; later on (from the 1990s), it led to the opposite pattern. Meanwhile, changes in retirement patterns in the 1990s, which otherwise would not have been very important, may have had a disproportionate impact because they coincided with an unusually large number of faculty members (born in the 1930s, hired in the 1960s) entering retirement age in the 1990s.

7 Discussion

The US scientific workforce has aged considerably over the last 60 years. This phenomenon has been noted by many observers, with the rising age of NIH grantees receiving particular attention (Kaiser 2008; Daniels 2015). However, the causes of the aging of the scientific workforce remain imperfectly understood.

Previous literature on the causes of aging has focused on the increase in age at entry into science (Jones 2009), as well as the tendency of scientists to retire later (Blau & Weinberg 2017). Our work highlights a third distinct reason why the scientific workforce may be aging: compositional changes arising from a slowdown in hiring over time. In a simulation based on detailed data on US chemistry faculty members between 1960 and 2010, changes in hiring over time appear to drive most of the change in the age composition of scientists. In the 1960s, a period when universities were expanding significantly, new faculty hires outnumbered retirements by a factor of three. Because new hires tend to be young, this led to a large influx of young people into the academic profession. However, as hiring slowed down in subsequent decades, the cohorts of 1960s entrants (and to a lesser extent, 1970s entrants) became disproportionally large, resulting in an aging scientific workforce.

A key theme that emerges from our analysis is that the aging of the scientific workforce is not the inevitable result of an aging society or of fundamental forces in the production of knowledge. Rather, much of it could be the reflection of a specific set of historical circumstances and policy choices, in the US post-WWII era, a massive expansion of the university sector followed by a period of no growth in faculty numbers.

Hiring larger numbers of young faculty could increase scientific productivity, as well as potentially shift the direction of research to more novel and relevant directions such as clean energy and climate change.[13] In addition, it may also advance other policy goals such as increasing gender and racial diversity (as young cohorts tend to be more diverse). Finally, the resulting increase in the share of graduate students finding positions in academia could strengthen incentives for talented individuals to pursue doctoral training and reduce excess competition. However, this type of policy is hardly discussed in policy circles.

We conclude by noting two directions for future research. This paper, and indeed much of the related literature, has focused on the US scientific workforce. However, much less is known about other countries. Of particular interest here is the case of China, which has made large investments in universities in the last two decades. Research on changes in the age of composition of the scientific workforce outside the US and in China in particular would be welcome. Another area that deserves attention is the age dynamics in industrial R&D where employment relationships are quite different from academia. To what extent has the industrial R&D workforce aged and is that aging also influenced by hiring sprees? Research along either of these lines of inquiry could further elucidate whether changes in the age composition of the scientific workforce are driven by fundamental forces or particular historic circumstances.

[13]Qualitative work by Wray (2003) suggests that scientific revolutions are more likely to be started by young and middle-age scientists. Einio, Feng & Jaravel 2022 present various pieces of evidence suggesting that entrepreneurs and innovators create products targeted to consumers like themselves in terms of various demographics, including age.

References

Agarwal, R., & Gaule, P. (2020). Invisible Geniuses: Could the Knowledge Frontier Advance Faster? *American Economic Review: Insights*, 2(4):409-424.

American Chemical Society Directory of Graduate Research (1961 and subsequent editions).

Aghion, P., Akcigit, U., Hyytinen, A., & Toivanen, O. (2017). The social origins of inventors. NBER WP 24110. National Bureau of Economic Research.

Ashenfelter, O., & Card, D. (2002). Did the elimination of mandatory retirement affect faculty retirement? *American Economic Review*, 92(4): 957-980.

Azoulay, P., Ding, W., & Stuart, T. (2007). The determinants of faculty patenting behavior: Demographics or opportunities? *Journal of Economic Behavior & Organization*, 63(4): 599-623.

Bell, A., Chetty, R., Jaravel, X., Petkova, N., & Van Reenen, J. (2019). Who becomes an inventor in America? The importance of exposure to innovation. *The Quarterly Journal of Economics*, 134(2): 647-713.

Blau, D., Weinberg, B. (2017) Why the US science and engineering workforce is aging rapidly. *Proceedings of the National Academy of Sciences*, 114:3879-3884.

Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, 110(4): 1104-1144.

Bush, V. (1945). *Science, the Endless Frontier. Report to the President*. Washington: U.S. Office of Scientific Research and Development

Catalini, C., Fons-Rosen, C., & Gaulé, P. (2020). How Do Travel Costs Shape Collaboration? *Management Science*, 66(8):3340-3360.

Celik, M. A. (2023). Does the cream always rise to the top? The misallocation of talent in innovation. *Journal of Monetary Economics*, 133: 105-128.

Daniels, R. J. (2015). A generation at risk: Young investigators and the future of the biomedical workforce. *Proceedings of the National Academy of Sciences*, 112(2):313-318.

Diamond, A. M. (1984). An economic model of the life-cycle research productivity of scientists. *Scientometrics*, 6(3): 189-196.

Einio, E., Feng, J., & Jaravel, X. (2022). Social Push and the Direction of Innovation. Available at SSRN 3383703.

Freeman, R., & Van Reenen, J. (2009). What If Congress Doubled R&D Spending on the

Physical Sciences? *Innovation Policy and the Economy, 9*: 1-38.

Ganguli, I., Gaule, P., & Cugalj, D. V. (2022). Chasing the Academic Dream: Biased Beliefs and Scientific Labor Markets. *Journal of Economic Behavior & Organization*, 202:17-33.

Gaule, P. (2014). Who comes back and when? Return Migration Decisions of Academic Scientists. *Economics Letters*, 124(3):461-464.

Gaule, P., & Piacentini, M. (2018). An advisor like me? Advisor gender and post-graduate careers in science. *Research Policy*, 47(4):805–813.

Jones, B. F. (2009). The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder? *Review of Economic Studies*, 76(1), 283-217.

Jones, B. F. (2010). Age and Great Invention. *Review of Economics and Statistics*, 92:(1):1-14.

Jones, B. F. (2011). As Science Evolves, How Can Science Policy? *Innovation Policy and the Economy* 11:103-131.

Jones, B. F., Reedy, E. J., & Weinberg, B. A. (2014). Age and Scientific Genius. The Wiley Handbook of Genius, 422-450.

Jones, B. F., & Weinberg, B. A. (2011). Age Dynamics in Scientific Creativity. *Proceedings of the National Academy of Sciences*, 108(47):18910-18914.

Jones, C. I. (2002). Sources of US economic growth in a world of ideas. *American Economic Review*, 92(1):220-239.

Kaiser, J. (2008). Zerhouni's Parting Message: Make Room for Young Scientists. *Science, 322*(5903):834–835.

Koning, R., Samila, S., & Ferguson, J. P. (2021). Who do we invent for? Patents by Women Focus more on Women's Health, but Few Women Get to Invent. *Science*, 372(6548):1345-1348.

Levin, S., Stephan, P. (1991) Research productivity over the life cycle: Evidence for academic scientists. *American Economic Review*, 81:114-132.

Lissoni, F., Mairesse, J., Montobbio, F., & Pezzoni, M. (2011). Determinants of Scientific Productivity: A study on Italian and French Physicists. *Industrial and Corporate Change*, 20(1):253-294.

Packalen, M., & Bhattacharya, J. (2019). Age and the Trying Out of New Ideas. *Journal of Human Capital*, 13(2): 341-373.

Stephan (2018) Verbal comment at NBER Summer Institute Science of Science meeting.

Stephan, P., Levin, S. (1993) Age and the Nobel Prize revisited. *Scientometrics*, 28: 387-399

Turner, L., Mairesse, J. (2003). *Individual productivity differences in scientific research: An econometric study of publishing of French physicists*. Paper presented at the Zvi Griliches Memorial Conference Paris, France. Wray, K. B. (2003). Is Science Really a Young Man's Game? *Social Studies of Science*, 33(1):137-149.

Wuchty, S., Jones, B. F., & Uzzi, B. (2007). The Increasing Dominance of Teams in Production of Knowledge. *Science*, 316(5827):1036-1039.

Yu, H., Marschke, G., Ross, M. B., Staudt, J., & Weinberg, B. A. (2022). Publish or Perish. *The Journal of Human Resources*, 5(8):4.

Figures

Figure 1: Age structure of the US scientific and labor workforce

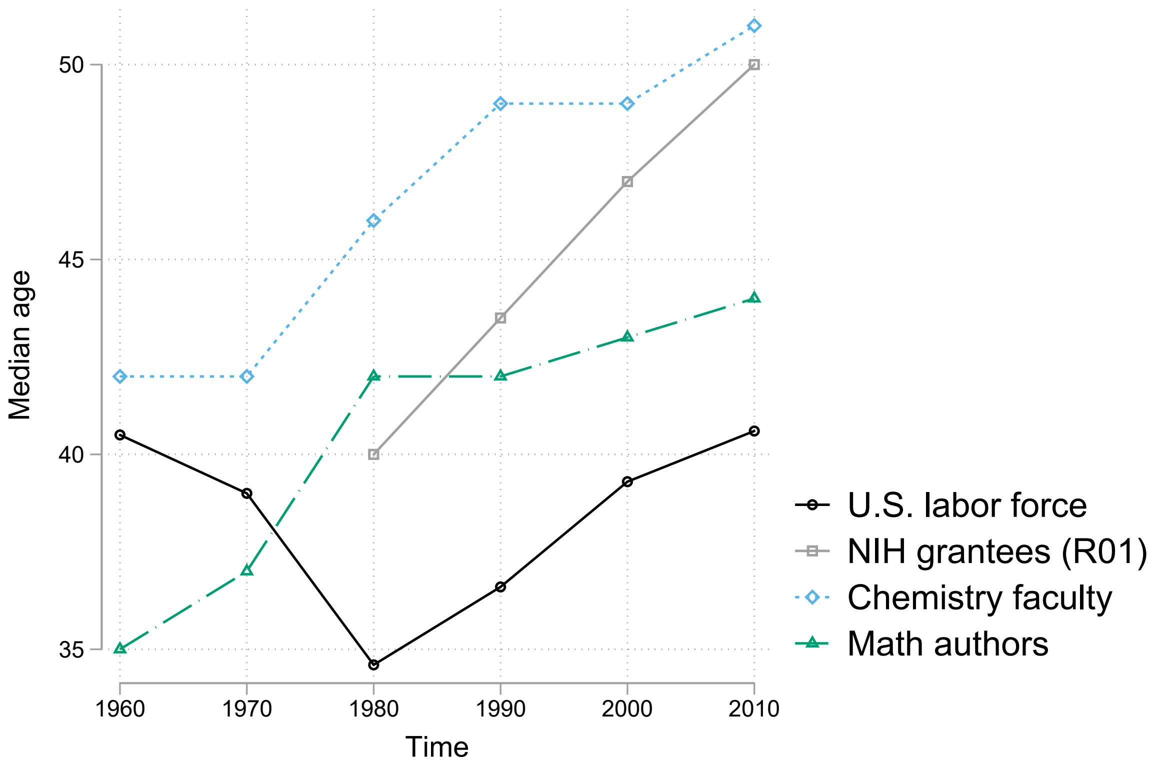

Notes: Source: authors calculations based on NIH data, and bespoke chemistry and math faculty data (based on the ACS directories and Math Genealogy Project respectively.)

Figure 2: A sample entry from the ACS directory

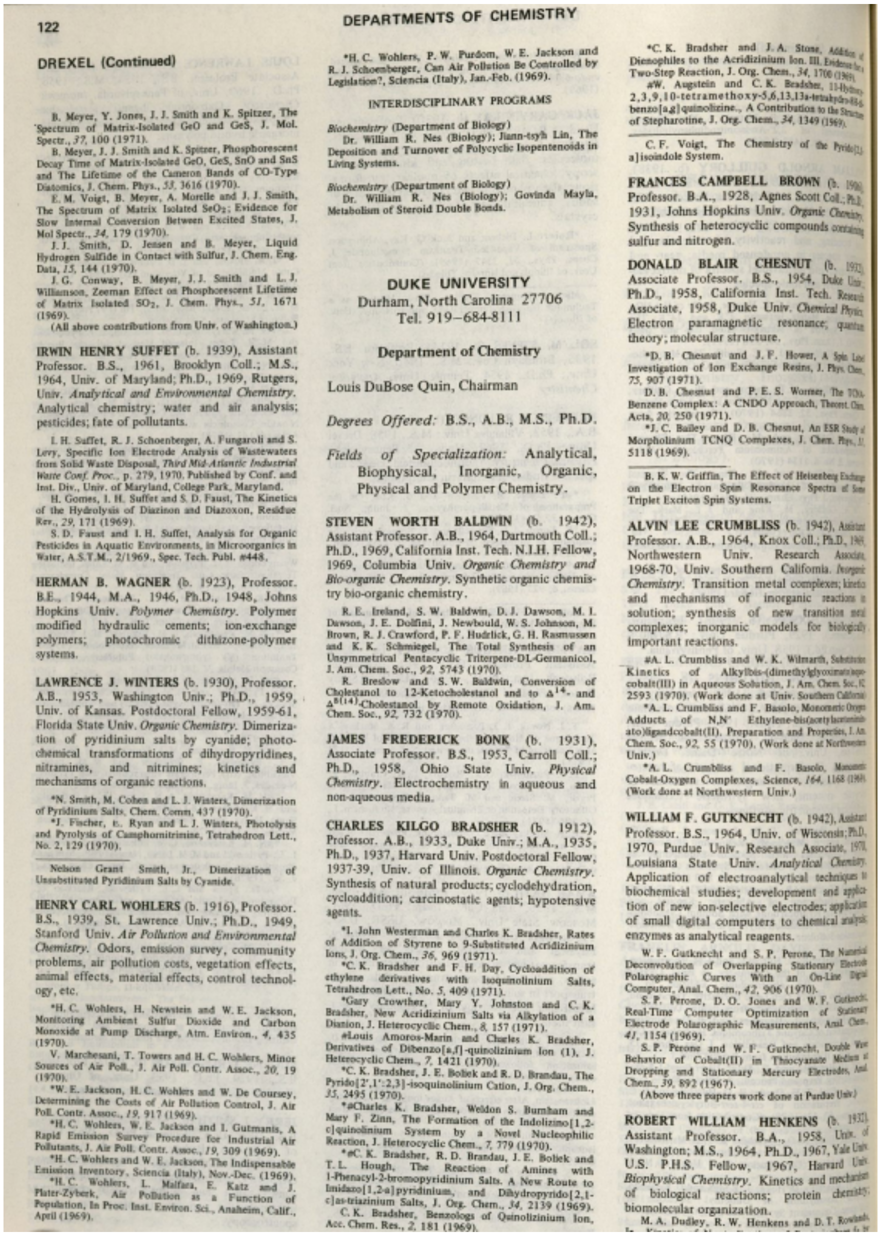

122 DEPARTMENTS OF CHEMISTRY

DREXEL (Continued)

B. Meyer, Y. Jones, J. J. Smith and K. Spitzer, The Spectrum of Matrix-Isolated GeO and GeS, J. Mol. Spectr., 37, 100 (1971).
B. Meyer, J. J. Smith and K. Spitzer, Phosphorescent Decay Time of Matrix-Isolated GeO, GeS, SnO and SnS and The Lifetime of the Cameron Bands of CO-Type Diatomics, J. Chem. Phys., 55, 3616 (1970).
E. M. Voigt, B. Meyer, A. Morelle and J. J. Smith, The Spectrum of Matrix Isolated SeO_2; Evidence for Slow Internal Conversion Between Excited States, J. Mol Spectr., 34, 179 (1970).
J. J. Smith, D. Jensen and B. Meyer, Liquid Hydrogen Sulfide in Contact with Sulfur, J. Chem. Eng. Data, 15, 144 (1970).
J. G. Conway, B. Meyer, J. J. Smith and L. J. Williamson, Zeeman Effect on Phosphorescent Lifetime of Matrix Isolated SO_2, J. Chem. Phys., 51, 1671 (1969).
(All above contributions from Univ. of Washington.)

IRWIN HENRY SUFFET (b. 1939), Assistant Professor. B.S., 1961, Brooklyn Coll.; M.S., 1964, Univ. of Maryland; Ph.D., 1969, Rutgers, Univ. *Analytical and Environmental Chemistry.* Analytical chemistry; water and air analysis; pesticides; fate of pollutants.

I. H. Suffet, R. J. Schoenberger, A. Fungaroli and S. Levy, Specific Ion Electrode Analysis of Wastewaters from Solid Waste Disposal, *Third Mid-Atlantic Industrial Waste Conf. Proc.*, p. 279, 1970. Published by Conf. and Inst. Div., Univ. of Maryland, College Park, Maryland.
H. Gomes, I. H. Suffet and S. D. Faust, The Kinetics of the Hydrolysis of Diazinon and Diazoxon, Residue Rev., 29, 171 (1969).
S. D. Faust and I. H. Suffet, Analysis for Organic Pesticides in Aquatic Environments, in Microorganics in Water, A.S.T.M., 2/1969., Spec. Tech. Publ. #448.

HERMAN B. WAGNER (b. 1923), Professor. B.E., 1944, M.A., 1946, Ph.D., 1948, Johns Hopkins Univ. *Polymer Chemistry.* Polymer modified hydraulic cements; ion-exchange polymers; photochromic dithizone-polymer systems.

LAWRENCE J. WINTERS (b. 1930), Professor. A.B., 1953, Washington Univ.; Ph.D., 1959, Univ. of Kansas. Postdoctoral Fellow, 1959-61, Florida State Univ. *Organic Chemistry.* Dimerization of pyridinium salts by cyanide; photochemical transformations of dihydropyridines, nitramines, and nitrimines; kinetics and mechanisms of organic reactions.

*N. Smith, M. Cohen and L. J. Winters, Dimerization of Pyridinium Salts, Chem. Comm. 437 (1970).
*J. Fischer, E. Ryan and L. J. Winters, Photolysis and Pyrolysis of Camphornitrimine, Tetrahedron Lett., No. 2, 129 (1970).

Nelson Grant Smith, Jr., Dimerization of Unsubstituted Pyridinium Salts by Cyanide.

HENRY CARL WOHLERS (b. 1916), Professor. B.S., 1939, St. Lawrence Univ.; Ph.D., 1949, Stanford Univ. *Air Pollution and Environmental Chemistry.* Odors, emission survey, community problems, air pollution costs, vegetation effects, animal effects, material effects, control technology, etc.

*H. C. Wohlers, H. Newstein and W. E. Jackson, Monitoring Ambient Sulfur Dioxide and Carbon Monoxide at Pump Discharge, Atm. Environ., 4, 435 (1970).
V. Marchesani, T. Towers and H. C. Wohlers, Minor Sources of Air Poll., J. Air Poll. Contr. Assoc., 20, 19 (1970).
*W. E. Jackson, H. C. Wohlers and W. De Coursey, Determining the Costs of Air Pollution Control, J. Air Poll. Contr. Assoc., 19, 917 (1969).
*H. C. Wohlers, W. E. Jackson and I. Gutmanis, A Rapid Emission Survey Procedure for Industrial Air Pollutants, J. Air Poll. Contr. Assoc., 19, 309 (1969).
*H. C. Wohlers and W. E. Jackson, The Indispensable Emission Inventory, Sciencia (Italy), Nov.-Dec. (1969).
*H. C. Wohlers, L. Malfara, E. Katz and J. Plater-Zyberk, Air Pollution as a Function of Population, In Proc. Inst. Environ. Sci., Anaheim, Calif., April (1969).
*H. C. Wohlers, P. W. Purdom, W. E. Jackson and R. J. Schoenberger, Can Air Pollution Be Controlled by Legislation?, Sciencia (Italy), Jan.-Feb. (1969).

INTERDISCIPLINARY PROGRAMS

Biochemistry (Department of Biology)
Dr. William R. Nes (Biology); Jiann-tsyh Lin, The Deposition and Turnover of Polycyclic Isopentenoids in Living Systems.

Biochemistry (Department of Biology)
Dr. William R. Nes (Biology); Govinda Mayla, Metabolism of Steroid Double Bonds.

DUKE UNIVERSITY
Durham, North Carolina 27706
Tel. 919–684-8111

Department of Chemistry

Louis DuBose Quin, Chairman

Degrees Offered: B.S., A.B., M.S., Ph.D.

Fields of Specialization: Analytical, Biophysical, Inorganic, Organic, Physical and Polymer Chemistry.

STEVEN WORTH BALDWIN (b. 1942), Assistant Professor. A.B., 1964, Dartmouth Coll.; Ph.D., 1969, California Inst. Tech. N.I.H. Fellow, 1969, Columbia Univ. *Organic Chemistry and Bio-organic Chemistry.* Synthetic organic chemistry bio-organic chemistry.

R. E. Ireland, S. W. Baldwin, D. J. Dawson, M. I. Dawson, J. E. Dolfini, J. Newbould, W. S. Johnson, M. Brown, R. J. Crawford, P. F. Hudrlick, G. H. Rasmussen and K. K. Schmiegel, The Total Synthesis of an Unsymmetrical Pentacyclic Triterpene-DL-Germanicol, J. Am. Chem. Soc., 92, 5743 (1970).
R. Breslow and S. W. Baldwin, Conversion of Cholestanol to 12-Ketocholestanol and to Δ^{14}- and $\Delta^{8(14)}$-Cholestanol by Remote Oxidation, J. Am. Chem. Soc., 92, 732 (1970).

JAMES FREDERICK BONK (b. 1931), Associate Professor. B.S., 1953, Carroll Coll.; Ph.D., 1958, Ohio State Univ. *Physical Chemistry.* Electrochemistry in aqueous and non-aqueous media.

CHARLES KILGO BRADSHER (b. 1912), Professor. A.B., 1933, Duke Univ.; M.A., 1935, Ph.D., 1937, Harvard Univ. Postdoctoral Fellow, 1937-39, Univ. of Illinois. *Organic Chemistry.* Synthesis of natural products; cyclodehydration, cycloaddition; carcinostatic agents; hypotensive agents.

*I. John Westerman and Charles K. Bradsher, Rates of Addition of Styrene to 9-Substituted Acridizinium Ions, J. Org. Chem., 36, 969 (1971).
*C. K. Bradsher and F. H. Day, Cycloaddition of ethylene derivatives with Isoquinolinium Salts, Tetrahedron Lett., No. 5, 409 (1971).
*Gary Crowther, Mary Y. Johnston and C. K. Bradsher, New Acridizinium Salts via Alkylation of a Diazion, J. Heterocyclic Chem., 8, 157 (1971).
#Louis Amoros-Marin and Charles K. Bradsher, Derivatives of Dibenzo[a,f]-quinolizinium Ion (1), J. Heterocyclic Chem., 7, 1421 (1970).
*C. K. Bradsher, J. E. Boliek and R. D. Brandau, The Pyrido[2',1':2,3]-isoquinolinium Cation, J. Org. Chem., 35, 2495 (1970).
*#Charles K. Bradsher, Weldon S. Burnham and Mary F. Zinn, The Formation of the Indolizino[1,2-c]quinolinium System by a Novel Nucleophilic Reaction, J. Heterocyclic Chem., 7, 779 (1970).
*#C. K. Bradsher, R. D. Brandau, J. E. Boliek and T. L. Hough, The Reaction of Amines with 1-Phenacyl-2-bromopyridinium Salts. A New Route to Imidazo[1,2-a]pyridinium, and Dihydropyrido[2,1-c]as-triazinium Salts, J. Org. Chem., 34, 2139 (1969).
C. K. Bradsher, Benzologs of Quinolizinium Ion, Acc. Chem. Res., 2, 181 (1969).
*C. K. Bradsher and J. A. Stone, Addition of Dienophiles to the Acridizinium Ion. III. Evidence for a Two-Step Reaction, J. Org. Chem., 34, 1700 (1969).
#W. Augstein and C. K. Bradsher, 11-Hydroxy-2,3,9,10-tetramethoxy-5,6,13,13a-tetrahydro-8H-benzo[a,g]quinolizine., A Contribution to the Structure of Stepharotine, J. Org. Chem., 34, 1349 (1969).

C. F. Voigt, The Chemistry of the Pyrido[2,1-a]isoindole System.

FRANCES CAMPBELL BROWN (b. 1908), Professor. B.A., 1928, Agnes Scott Coll.; Ph.D., 1931, Johns Hopkins Univ. *Organic Chemistry.* Synthesis of heterocyclic compounds containing sulfur and nitrogen.

DONALD BLAIR CHESNUT (b. 1932), Associate Professor. B.S., 1954, Duke Univ.; Ph.D., 1958, California Inst. Tech. Research Associate, 1958, Duke Univ. *Chemical Physics.* Electron paramagnetic resonance; quantum theory; molecular structure.

*D. B. Chesnut and J. F. Hower, A Spin Label Investigation of Ion Exchange Resins, J. Phys. Chem., 75, 907 (1971).
D. B. Chesnut and P. E. S. Wormer, The TCNQ Benzene Complex: A CNDO Approach, Theoret. Chim. Acta, 20, 250 (1971).
*J. C. Bailey and D. B. Chesnut, An ESR Study of Morpholinium TCNQ Complexes, J. Chem. Phys., 51, 5118 (1969).

B. K. W. Griffin, The Effect of Heisenberg Exchange on the Electron Spin Resonance Spectra of Some Triplet Exciton Spin Systems.

ALVIN LEE CRUMBLISS (b. 1942), Assistant Professor. A.B., 1964, Knox Coll.; Ph.D., 1968, Northwestern Univ. Research Associate, 1968-70, Univ. Southern California. *Inorganic Chemistry.* Transition metal complexes; kinetics and mechanisms of inorganic reactions in solution; synthesis of new transition metal complexes; inorganic models for biologically important reactions.

#A. L. Crumbliss and W. K. Wilmarth, Substitution Kinetics of Alkylbis-(dimethylglyoximato)aquocobalt(III) in Aqueous Solution, J. Am. Chem. Soc., 92, 2593 (1970). (Work done at Univ. Southern California.)
*A. L. Crumbliss and F. Basolo, Monomeric Oxygen Adducts of N,N' Ethylene-bis(acetylacetoniminato)ligandcobalt(II). Preparation and Properties, J. Am. Chem. Soc., 92, 55 (1970). (Work done at Northwestern Univ.)
*A. L. Crumbliss and F. Basolo, Monomeric Cobalt-Oxygen Complexes, Science, 164, 1168 (1969). (Work done at Northwestern Univ.)

WILLIAM F. GUTKNECHT (b. 1942), Assistant Professor. B.S., 1964, Univ. of Wisconsin; Ph.D., 1970, Purdue Univ. Research Associate, 1970, Louisiana State Univ. *Analytical Chemistry.* Application of electroanalytical techniques to biochemical studies; development and application of new ion-selective electrodes; application of small digital computers to chemical analysis; enzymes as analytical reagents.

W. F. Gutknecht and S. P. Perone, The Numerical Deconvolution of Overlapping Stationary Electrode Polarographic Curves With an On-Line Digital Computer, Anal. Chem., 42, 906 (1970).
S. P. Perone, D. O. Jones and W. F. Gutknecht, Real-Time Computer Optimization of Stationary Electrode Polarographic Measurements, Anal. Chem., 41, 1154 (1969).
S. P. Perone and W. F. Gutknecht, Double Wave Behavior of Cobalt(II) in Thiocyanate Medium at Dropping and Stationary Mercury Electrodes, Anal. Chem., 39, 892 (1967).
(Above three papers work done at Purdue Univ.)

ROBERT WILLIAM HENKENS (b. 1937), Assistant Professor. B.A., 1958, Univ. of Washington; M.S., 1964, Ph.D., 1967, Yale Univ. U.S. P.H.S. Fellow, 1967, Harvard Univ. *Biophysical Chemistry.* Kinetics and mechanisms of biological reactions; protein chemistry; biomolecular organization.

M. A. Dudley, R. W. Henkens and D. T. Rowlands

Notes: The Figure displays a sample page from the 1971 ACS directory. In the directory, departments list their faculty members in alphabetical order along with brief standardized biographical information. Among other information, each entry includes the name, birth year, and education history of the person. Whereas biographical information about particular famous scientists is often available from many sources, the strength of the directory is that it covers all faculty members employed in the reporting departments irrespective of their status in their profession.

Figure 3: Relationship between age and scientific output

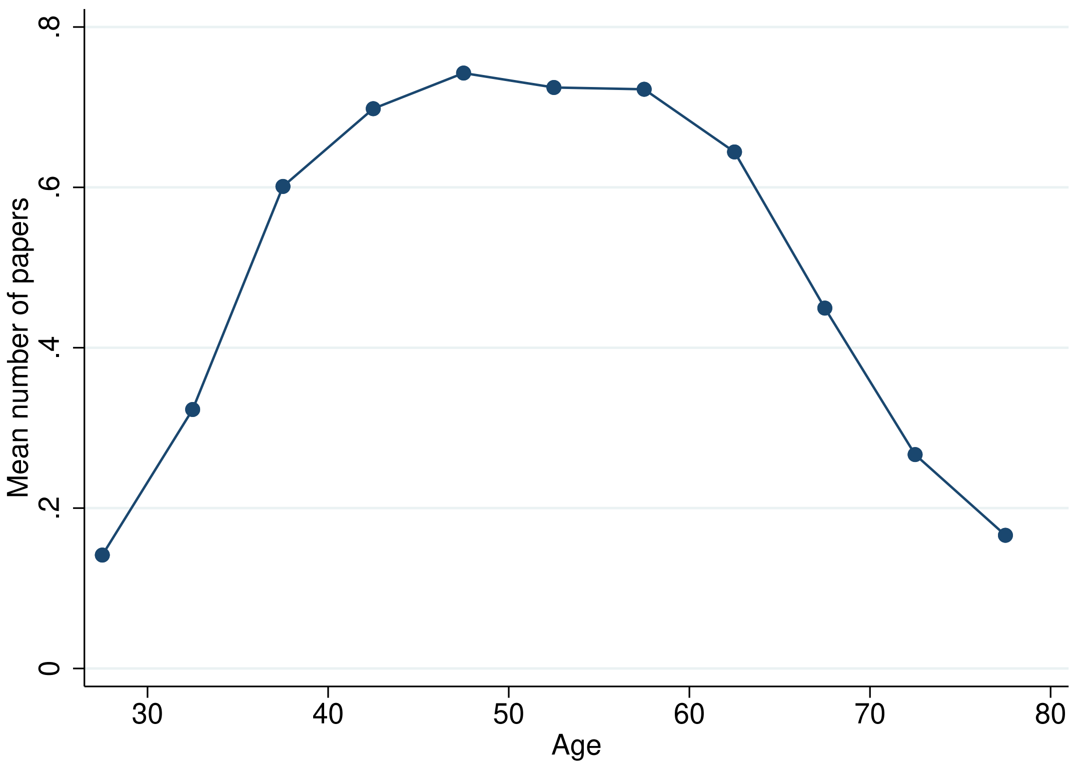

Notes: This Figure plots mean scientific output by age. We use the number of published papers measure of productivity.

Figure 4: Productivity differences across cohorts

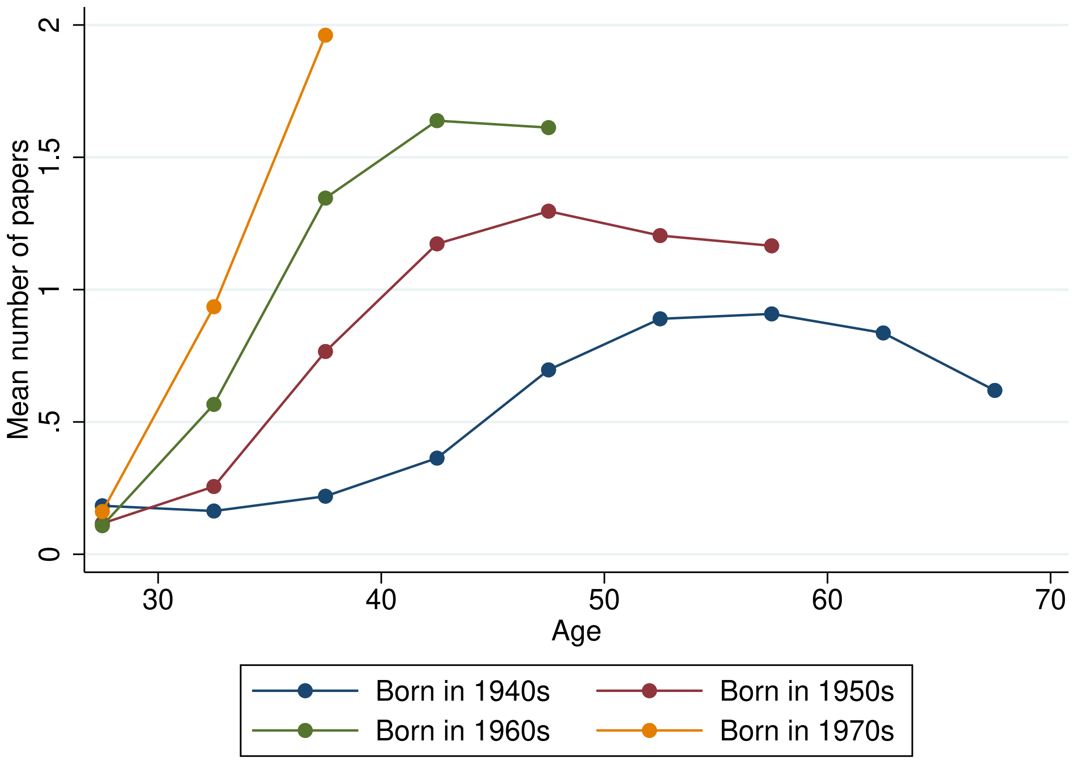

Notes: This Figure plots mean scientific output by birth cohort (10-year bins) and age (also in 10-year bins). The underlying data is a longitudinal panel of 11,854 faculty members running from PhD graduation year to age 80. We use the number of published papers measure of productivity.

Figure 5: Aging - actual data and model simulation (all aging channels switched on)

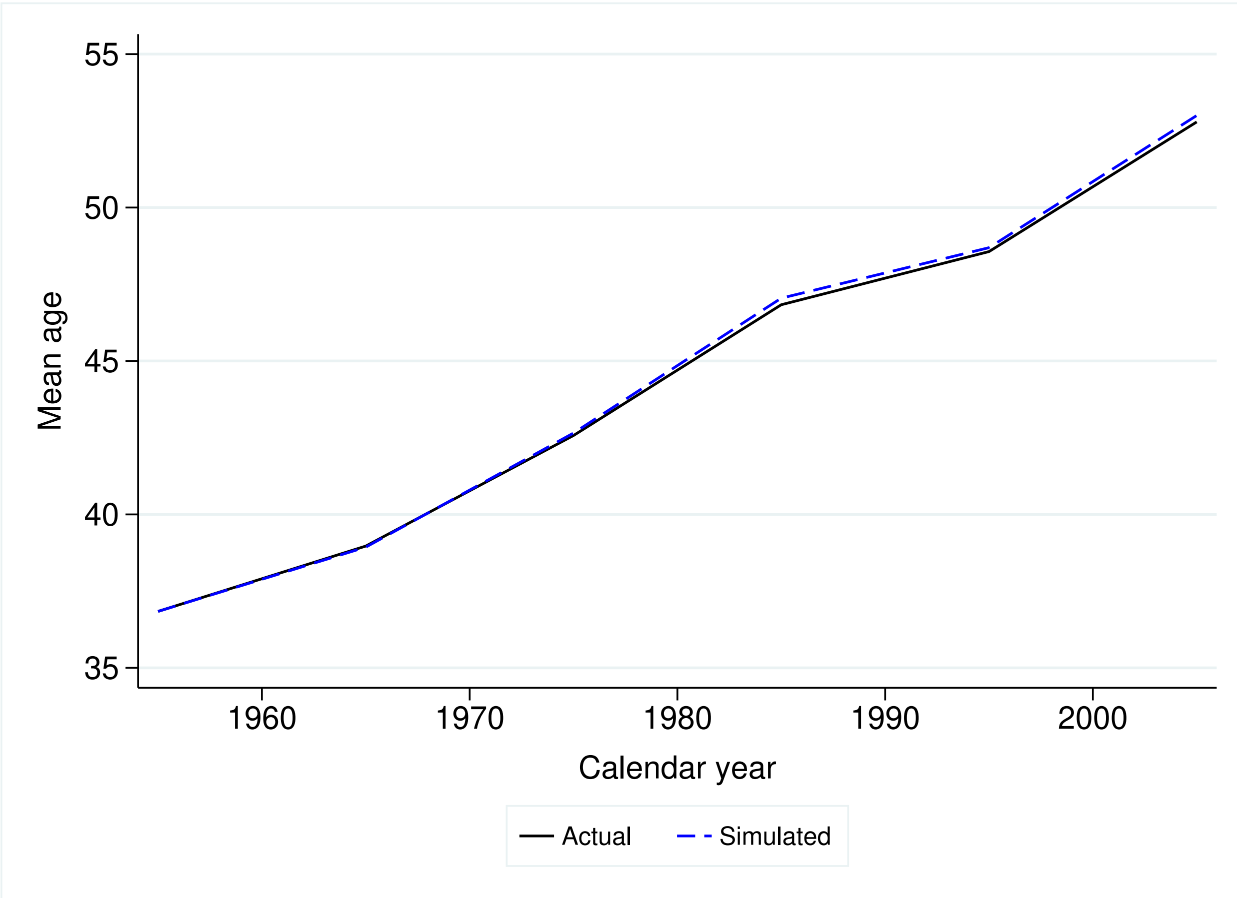

Notes: This Figure assesses the accuracy of the model in simulating the evolution of the aging of faculty members over time. The dashed blue line denotes the outcome of our simulated model, and the black line denotes the actual data. In this version of the simulated model, we switch on all aging channels: we assume that retirements, age at entry, and the number of faculty lines all evolved as they did. However, because our model makes simplifying assumptions on the distribution of age at entry and retirement probabilities, the two lines may differ.

Figure 6: Aging - actual data and model simulation (all aging channels switched off, i.e. stay at 1960s values)

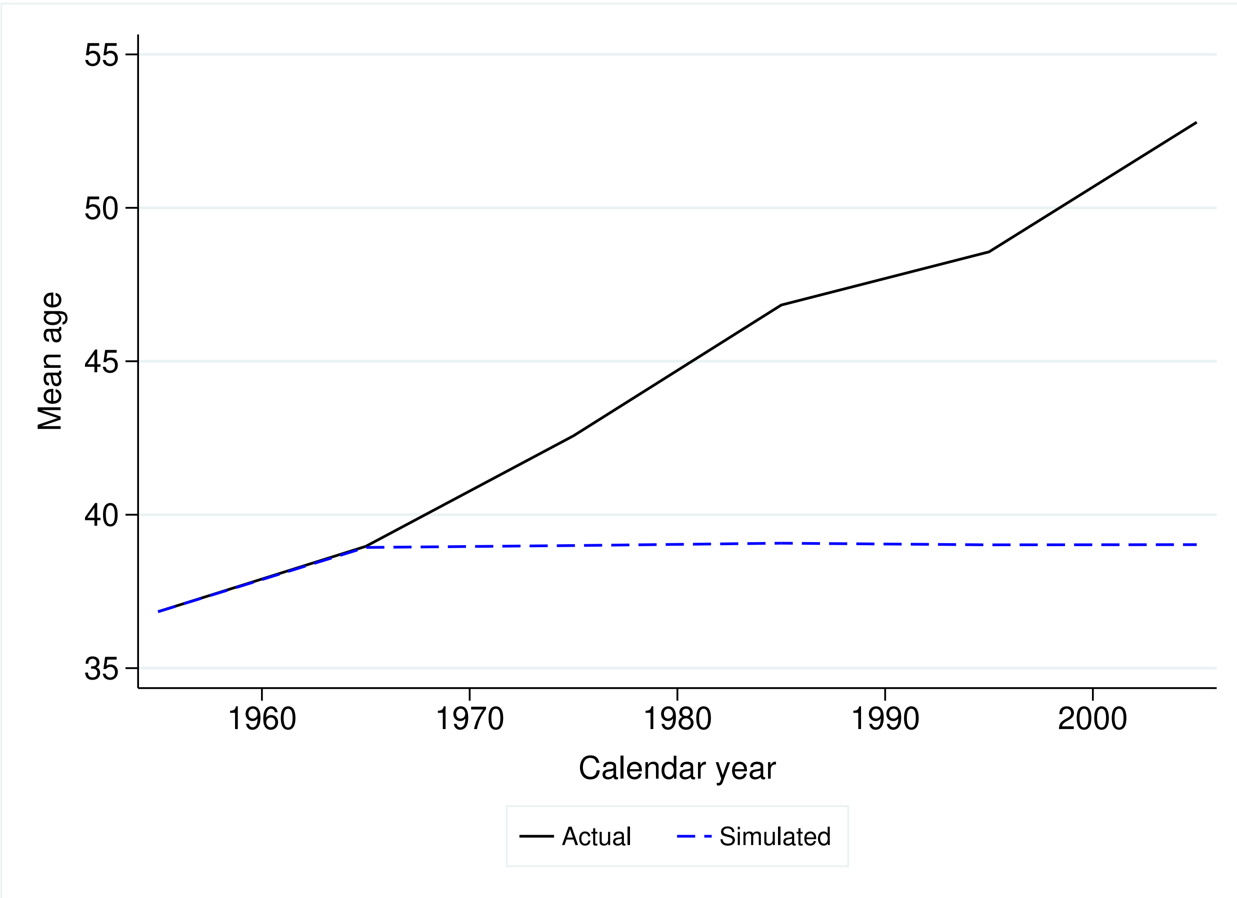

Notes: This Figure assesses the accuracy of the model in simulating the evolution of the aging of faculty members over time. The red line denotes the outcome of our simulated model, and the black line denotes the actual data. In this version of the simulated model, we switch off all aging channels: we assume that exits, age at entry, and the ratio of entrants to exits all stay as they were in the 1960s.

Figure 7: Evolution of the age of new entrants and of the ratio of new entrants to exits over time .

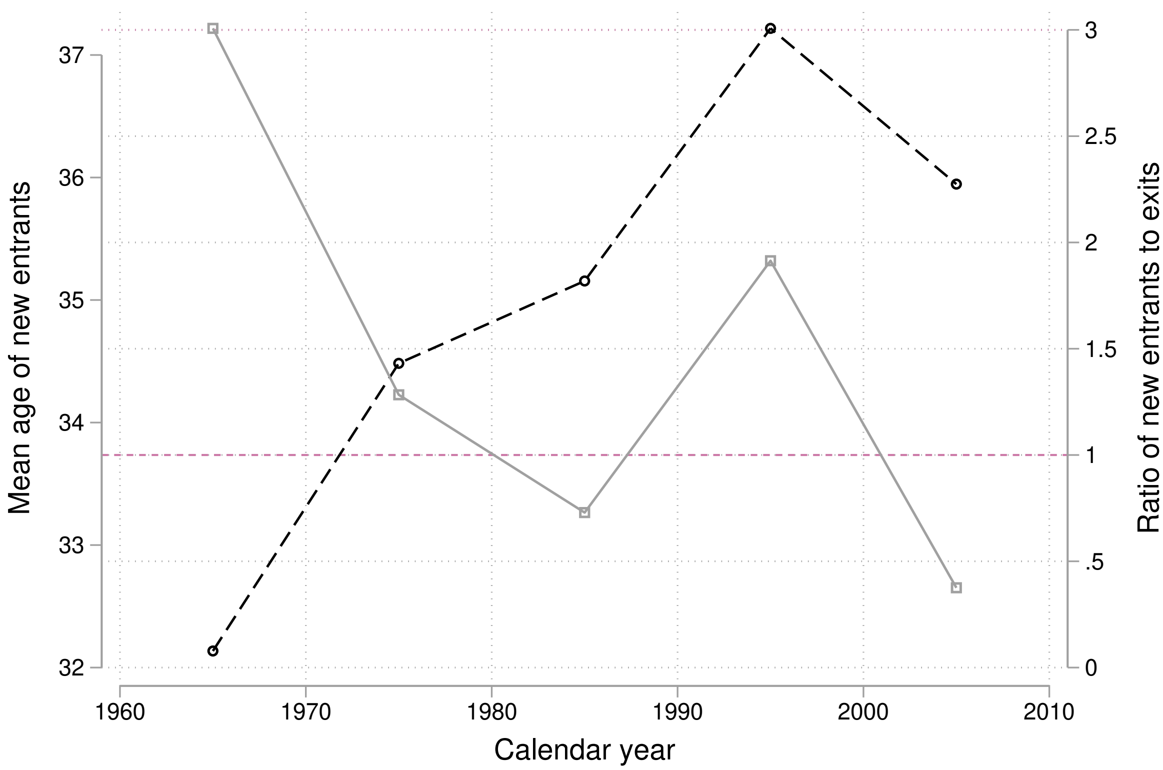

Notes: This Figure shows how the mean age of new entrants (dashed line) and the ratio of new entrants to exits (solid line) evolved over time. The horizontal dashed line corresponds to a ratio of 1 with the number of new entrants equaling exits (and the total number of faculty members staying the same). The horizontal dotted line corresponds to a ratio of 3 with each faculty exiting replaced by three new hires, as was the case in the 1960s.

Figure 8: Aging - counterfactual where the retirement probabilities and age of entrants stay at 1960s levels, but the ratio of entrants to exits evolves as it actually did.

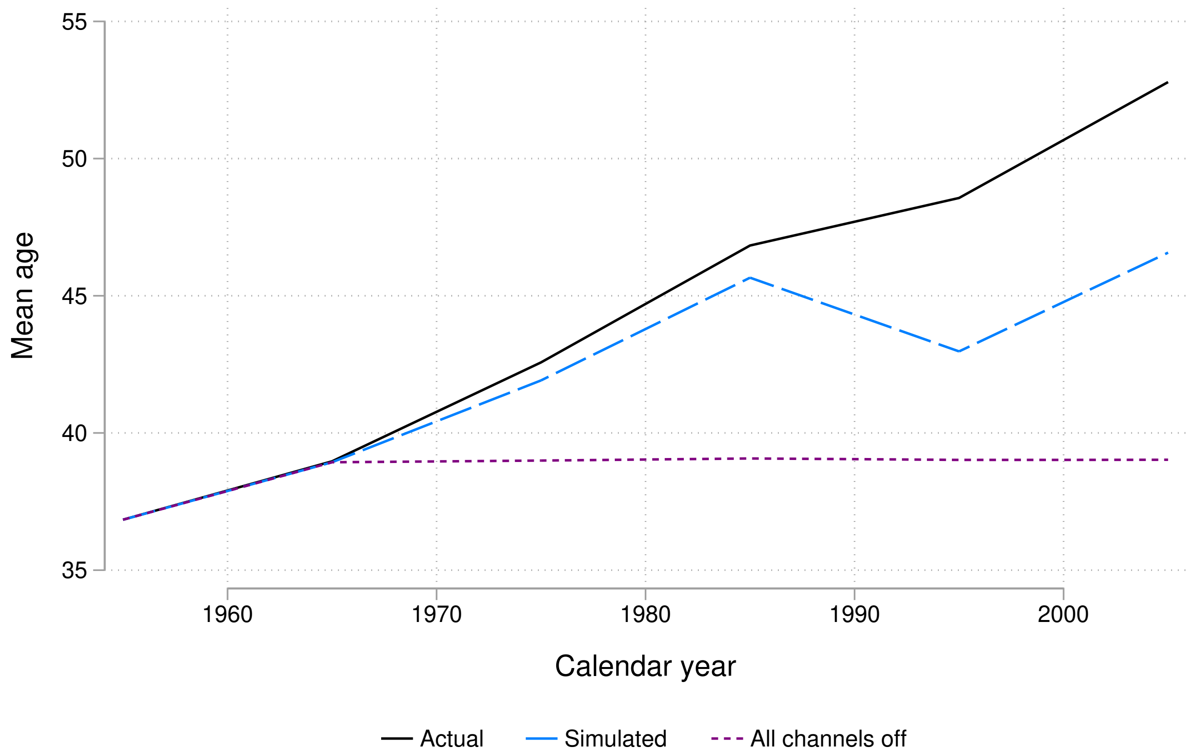

Notes: The blue line shows how the mean age of faculty members would have evolved -according to our model- if the exit probabilities and the age of entrants stayed at 1960s levels, but the ratio of entrants to exits changed as it actually did. Therefore, the blue line represents the change in mean age coming exclusively from the changing ratios of entrants to exits, with the other channels that may drive aging being turned off. The black line denotes the actual data. The black line denotes the actual data, and the short purple short-dashed line indicates simulation results with all aging channels turned off.

Figure 9: Aging - counterfactual where the age of entrants stays at 1960s levels, but the ratio of entrants to exits and retirement probabilities evolve as they actually did.

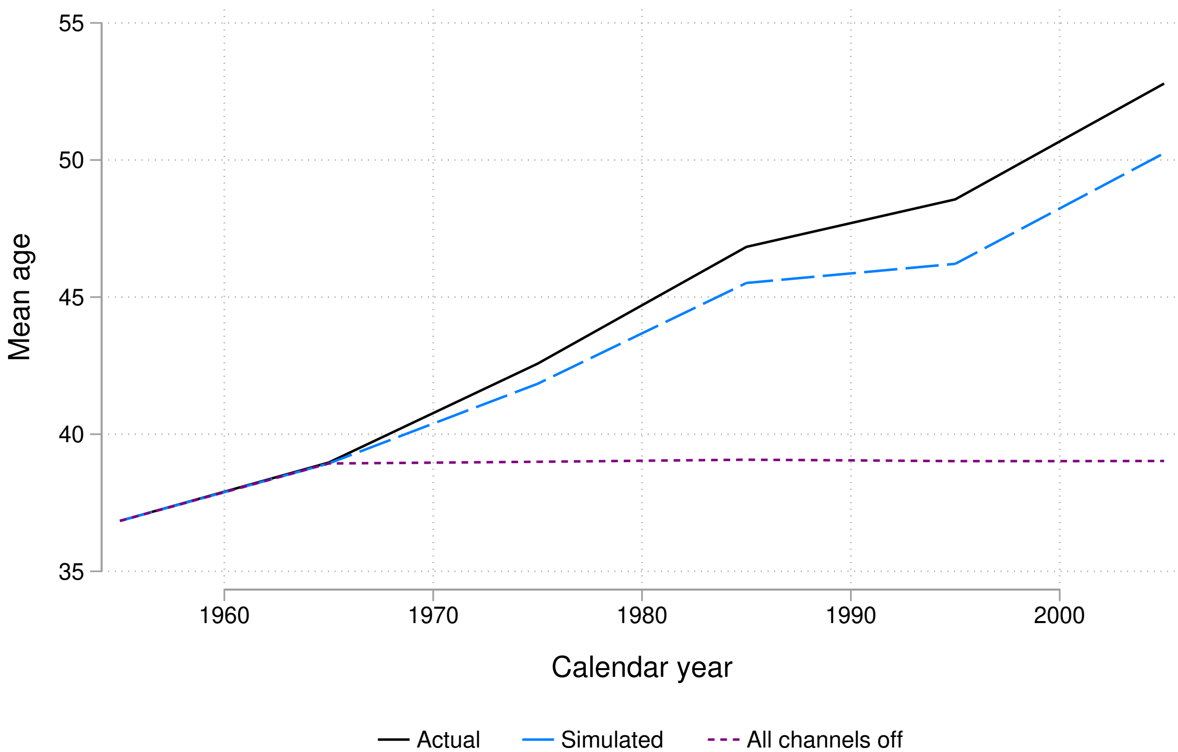

Notes: The blue line shows how the mean age of faculty members would have evolved -according to our model- if the age of entrants stayed at 1960s levels, but the ratio of entrants to exits changed, as well as exit probabilities, evolved as they actually did. Therefore, the blue line represents the change in mean age coming from the conjunction of (1) the changing ratio of entrants to exits and (2) changing exit probabilities, with the remaining channel that may drive aging (changing age at entry) being turned off. The black line denotes the actual data, and the short purple short-dashed line indicates simulation results with all aging channels turned off.

Tables

Table 1: Productivity differences across cohorts -formal tests

	(1)	(2)
	D.V.=Nr of papers	
Age in decade x Decade of birth FE	Yes	Yes
Hypotheses (reporting p-values):		
(1) born in 1930s=born in 1940s	< 0.001	< 0.001
(2) born in 1940s= born in 1950s	< 0.001	< 0.001
(3) born in 1950s=born in 1960s	< 0.001	< 0.001
(4) born in 1960s=born in 1970s	0.001	< 0.001
Joint test of (1), (2), (3) and (4)	< 0.001	< 0.001
Year fixed effects	Yes	No
Linear time trend	No	Yes
Estimation	Poisson	Poisson
Obs.	356,640	356,640
Mean of D.V.	0.596	0.596
Pseudo R2	0.1069	0.0950

Notes: This Table investigates productivity differences across cohorts in the spirit of Figure 4. The underlying data is a longitudinal panel of 11,854 faculty members running from Ph.D. graduation year to age 80. We run a Poisson regression with the number of papers on the left-hand side and a full set of indicator variables for all possible age (10-year bins) x decade of birth interactions on the right-hand side. To adjust for the fact that producing research or publishing may become easier over time, we control for calendar-year fixed effects. We then test whether two adjacent cohorts had equal productivity at all observed ages (in 10-year bins). For instance, when we test the null hypothesis that the cohort born in the 1930s has equal productivity to the cohort born in the 1940s, we run a composite test for the 1930s cohort having equal productivity in their thirties to the 1940s cohort in their thirties, the 1930s cohort having equal productivity in their forties to the 1940s cohort in their forties, etc. See Figure Figure 4 for a visual representation.

Table 2: Contribution of different channels to aging-summary of simulation results

Channels (yes=time-varying data, no=keep 1960s level)			Contribution
Change in the ratio of entrants to exits	Retirement probabilities	Entry age	
Yes	No	No	54.8%
No	Yes	No	6.4%
No	No	Yes	26.4%
Yes	Yes	No	81.8%
Yes	No	Yes	79.8%
No	Yes	Yes	30.1%

Notes: This Table summarizes simulation results regarding the contribution of different channels to faculty aging between 1965 and 2005 (between these years the mean age increased from 38.9 to 52.8). Each row corresponds to a different scenario where we switch off (keep at 1960s level) one or more channels.

Table 3: Understanding interactions between two channels

Channels	Main effect 1	Main effect 2	Interaction	Total
Ratio of entrants to exits + Retirement prob.	54.8%	6.4%	20.6%	81.8%
Retirement prob. + Entry age	6.4%	26.4%	-2.7%	30.1%
Ratio of entrants to exits + Entry age	54.8%	26.4%	-1.4%	79.8%

Notes: This Table highlights the role of interactions between different channels in generating aging. Typically, the effect of switching on two channels would not be the same as the sum of turning on each channel individually. For instance, the effect of switching on changes in retirement probabilities (i.e. have lower retirement probabilities over time) would matter more if the population of scientists is older, which in turn depends on whether the other channels are switched on.

Appendix: Details on the simulation methodology

Assumptions. We assume that someone who appeared in a given snapshot s for the first time actually entered the data set in $\underline{s} = s-6$. Due to the left truncation, we cannot make any assumptions about the actual entrance time of people who were active as of the initial snapshot $s = 1961$, that is, $\underline{s} = 1955$. Therefore, after excluding 1955, our range of entry years is: $\underline{s} \in 1965 + 10k_{k=0}^{4}$. Similarly, we assume that someone who was active in a given snapshot s for the last time actually exited the dataset in $\overline{s} = s+4$. Due to the right truncation, we cannot make this assumption in the last snapshot $s = 2011$, that is, $\overline{s} = 2015$. Therefore, we end up with the range: $\overline{s} \in 1965 + 10k_{k=0}^{4}$.[14]

We refer to the mid-decade points throughout the simulation as the simulated periods t. In other words, whenever we discuss time periods, we will be referring to these mid-decade points.

Transitions between simulated periods. The dataset, starting from the initial data we have available for 1961, is subject to hirings and exits over time that lead to changes in the *number of people*. Furthermore, we identify three channels that lead to changes in the *mean age*: (1) Exits; (2) Hirings; (3) Age composition of hirings. In the following, we will sequentially describe these three channels.

Exits. For each person who is active in the current simulated period t, we assign an updated status – "active" or "exited" – for the next period $t+10$ based on age-specific and time-specific exit probabilities e_{at} calculated from the actual data, where a is the age decade (ie, people in their 20s, 30s, etc.) and t is time. This implies that we account for the fact that older scientists are more likely to exit than younger ones, and that these likelihoods also evolve over time.

For every middle-decade point t and 10-year age group a, we calculate the probabilities $e_{at} = N_{at}^{e}/N_{at}^{a}$. The numerator, denoted as N_{at}^{e}, is the number of people in the age group a who are active at time t but are no longer active at time $t+10$. The denominator, denoted as N_{at}^{a}, is the total number of currently active people in the age group a at time t.[15]

In our simulations, suppose that a group of individuals (defined by an age range a at time t) has a likelihood of exit defined by e_{at}, say, equal to 20%. In this case, we randomly pick a fifth of these individuals and assign them a "exited" status in the next time period.

Hirings. In each simulated period t, we simulate hiring $\hat{N}_t^{h}$ people, which we refer to as the "extensive" margin of hiring, meaning that we do not account for scientist age.[16] It is determined

[14]In other words, we assume that the hiring of people who appear for the first time in snapshot s and the exit of people who were active for the last time a decade earlier in snapshot $s-10$, actually happen at the middle-decade point.

[15]For clarify, the subindex e stands for "exits" and the subindex a stands for "active".

[16]The subindex h stands for "hiring" and the hat notation stands for "simulated".

by ${\hat{N}_t}^h = {\hat{N}_t}^e \cdot h_t$, where ${\hat{N}_t}^e$ is the number of simulated exits including all age groups, and h_t is an expansion rate, i.e. the number of simulated hirings for each simulated exit in the same period. We define this expansion rate as $h_t = N_t^h / N_t^e$, that is, calculating the ratio from actual data on hirings and exits in a given time period (N_t^h and N_t^e).

The exit channel directly affects both the intensive and extensive margins, since each exiting individual is assigned an age that we track. But the hirings channel only affects the extensive margin, since no age is assigned to entrants yet. To allocate an age to these hirings, we need to add a third channel, which will be described below. In other words, we distinguish between the "extensive" (number of people) and "intensive" (age composition) margins of *hiring* represented by two separate channels, while the *exit* channel is not split into two parts.

Age composition of hirings. In each simulated period t, years of birth of newly hired people (${\hat{N}_t}^h$) are assigned so that the resulting age composition of simulated hirings reconstructs the actual age composition of newly hired people in the data.[17] All parameters calculated from actual data (i.e., N_t^h and N_t^e, e_{at}, and all $P_t^{x\%}$) are time-specific data moments.

Extensive vs intensive margins. From these three definitions it follows that exits and hirings affect *not only* the total number of scientists *but also* the age composition of active scientists in each simulated period t. We justify the asymmetry by which the exit channel does not incorporate an intensive margin channel as follows.

First, we will show that the results of a simulation relying on time-specific data moments will be equivalent, independently of whether the exit activity is collapsed into a single channel or rather decomposed into two different channels – the number and the age composition of exits.

Consider an alternative definition of exits in which we explicitly account for the two channels: in each simulated period t, we first impose a fixed number of total exits N_t^e and then divide them into age groups according to age-specific shares of the exits – E_{at}. We then randomly assign an "exited" status to a number of people from a given age group according to $\hat{N}_{at}^e = N_t^e \cdot E_{at}$. In this way, we would explicitly distinguish between the "extensive" and "intensive" margins of exits.

Although this alternative share of exits E_{at} only relies on one age-specific component, N_{at}^e, in our preferred share of exits e_{at} we additionally incorporate a second age-specific component, N_{at}^a. Thus, without loss of generality, one can replicate the actual data by relying on *either* a broader set of data moments in our initial definition *or* a narrower set of data moments with the alternative definition.

[17] In particular, we assign $(P_t^{min} + P_t^{10\%})/2$ to 10% of randomly chosen simulated hires, $(P_t^{10\%} + P_t^{50\%})/2$ to 40%, $(P_t^{50\%} + P_t^{75\%})/2$ to 25%, $(P_t^{75\%} + P_t^{90\%})/2$ to 15% and $(P_t^{90\%} + P_t^{95\%})/2$ to 10% of hires, where $P_t^{x\%}$ is a corresponding percentile of the actual age composition of hires at the time t.

Second, it follows from defining the simulated exits as $e_{at} = N^e_{at}/N^a_{at}$ that we simulate the number of exits for the age range a in the following way: $\hat{N}^e_{at} = N^a_{at} \cdot e_{at}$. Summing up across all age ranges, the total number of exits is therefore $\hat{N}^e_t = \sum_a \hat{N}^e_{at} = \sum_a N^a_{at} e_{at}$. Since hirings is a function of exits, $\hat{N_t}^h = \hat{N_t}^e \cdot h_t$, we can rewrite this expression as $\hat{N_t}^h = h_t(\sum_a N^a_{at} e_{at})$. Thus, our approach allows us to account for the age-specific stock of active scientists in a given period t, N^a_{at}, not only in the exits, but also in the hirings channel.

Counterfactual scenarios. Starting with the initial sample of scientists described above, we generate transitions to the next periods through the three channels described above by: (1) adding new people to the dataset; (2) randomly dropping people conditional on each given age group; (3) randomly assigning age to each newly added people based on statistics coming from actual data.

For the three channels, we differentiate between two modes: (1) "ON" – time-specific data moments (N^h_t and N^e_t, r_{at}, and $P^{x\%}_t$) are used in each time period; (2) "OFF" – only the initial 1965 period data moments (N^h_{1965} and N^e_{1965}, r_{a1965}, and $P^{x\%}_{1965}$) are used for all periods.

For clarification, in our baseline scenario "ALL ON", all three channels are simulated in the "ON" mode. Simulated data moments ($h_t = N^h_t/N^e_t$, r_{at}, and $P^{x\%}_t$) for a given time period are generated based on the statistics of the actual data in that same time period. Instead, in the baseline scenario "ALL OFF" only actual data from 1965 is used to generate the simulated data moments of all time periods.

We will evaluate alternative counterfactual scenarios by switching on only one channel at a time. For example, suppose that the time-varying expansion rate is "ON", while the exits and age distribution of newly hired are "OFF". In this scenario, actual data is used for the parameter $h_t = N^h_t/N^e_t$ across all decades. Instead, data from 1965 (r_{a1965}, $P^{x\%}_{1965}$) is used for the exits and moments of the age distribution of newly hired across all decades.

The other two counterfactual scenarios will be switching on either the exits or the age distribution of the newly hired channel. That is, we will use actual data from all decades only for the parameters r_{at} or $P^{x\%}_t$, respectively, while the remaining channels will be in the "OFF" mode, and only data from 1965 will be used.

Appendix Figures

Figure A-1: US Federal R&D funding over time

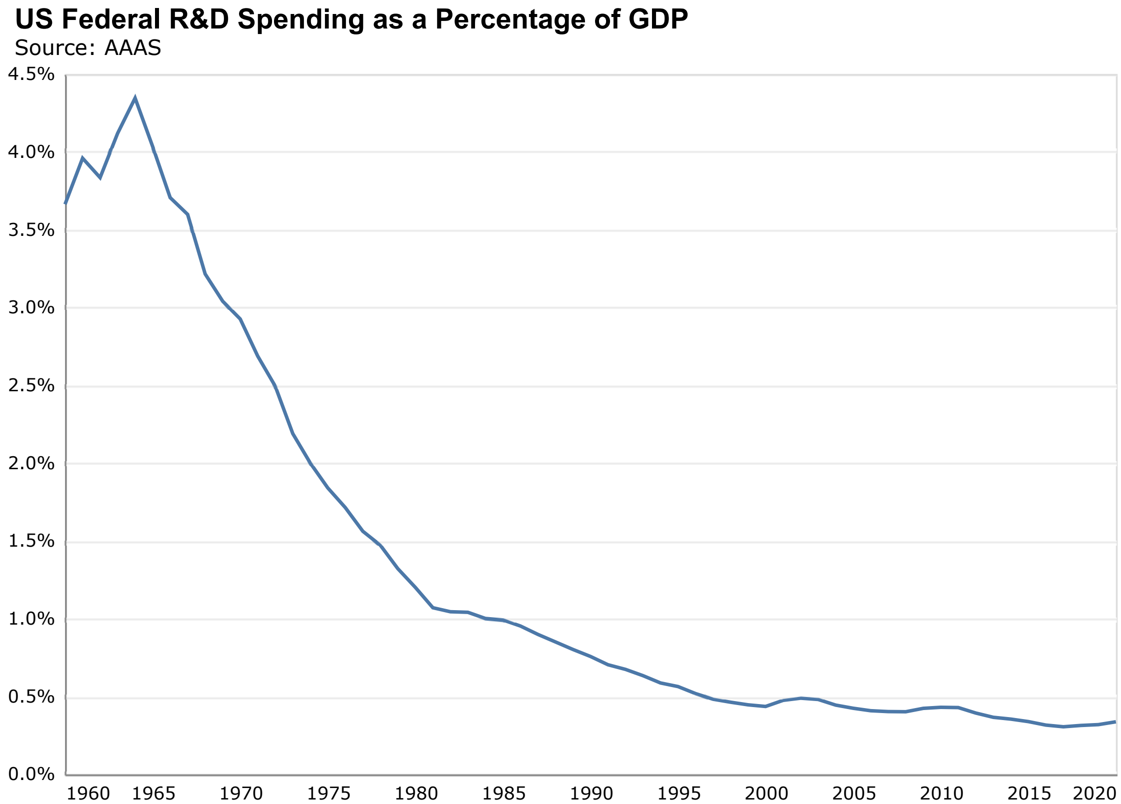

Notes: This Figure is based on data from American Asssociation for the Advancement of Science "Historical Trends in Federal R&D" https://www.aaas.org/programs/r-d-budget-and-policy/historical-trends-federal-rd.

Figure A-2: Productivity differences across cohorts -accounting for calendar year effects

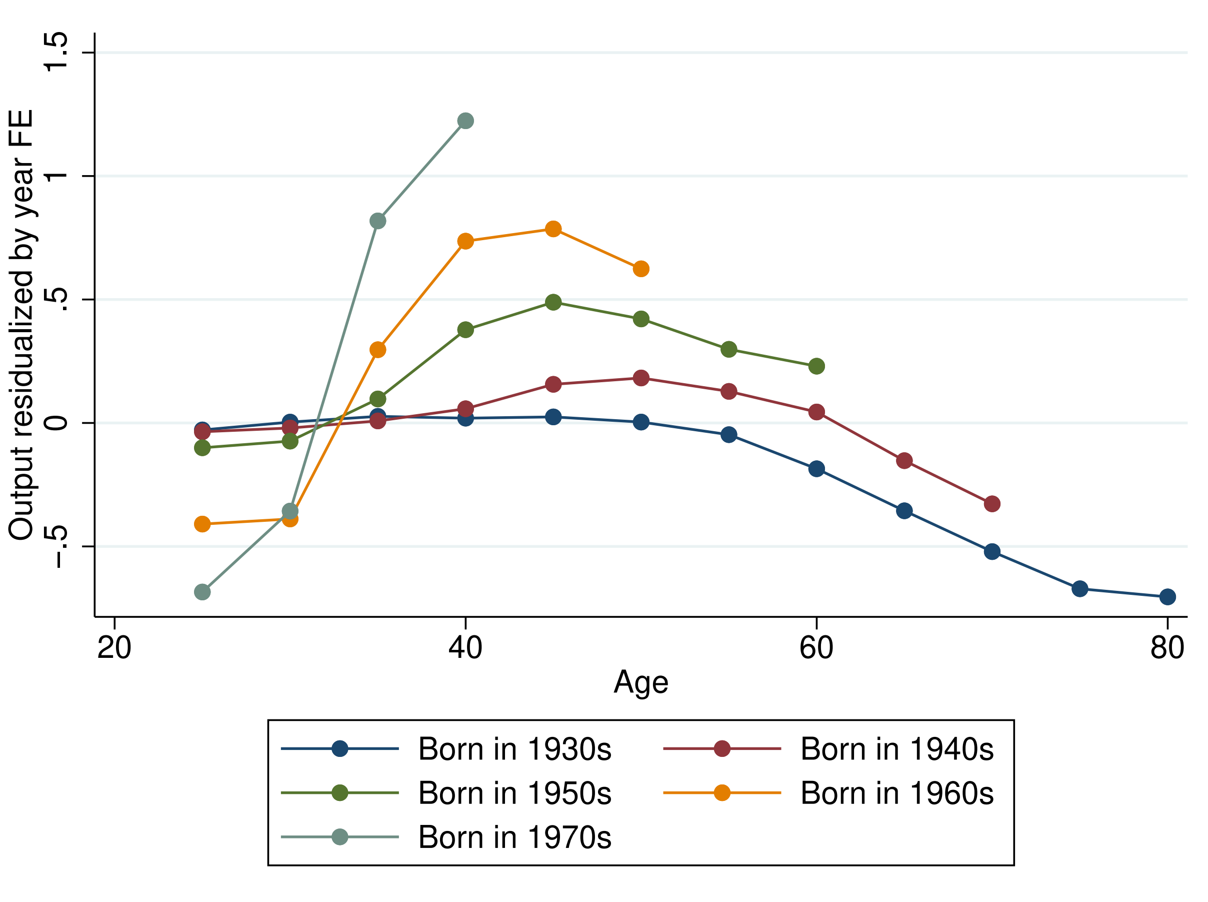

Notes: This Figure shows average output (residualized by year fixed effects) by birth cohort and age. We first regress individual output (the number of published papers) on year fixed effects and generate the residuals from that regression. We then compute the average of these residuals by birth decade and age.

Figure A-3: Productivity differences across cohorts -accounting for a linear time trend

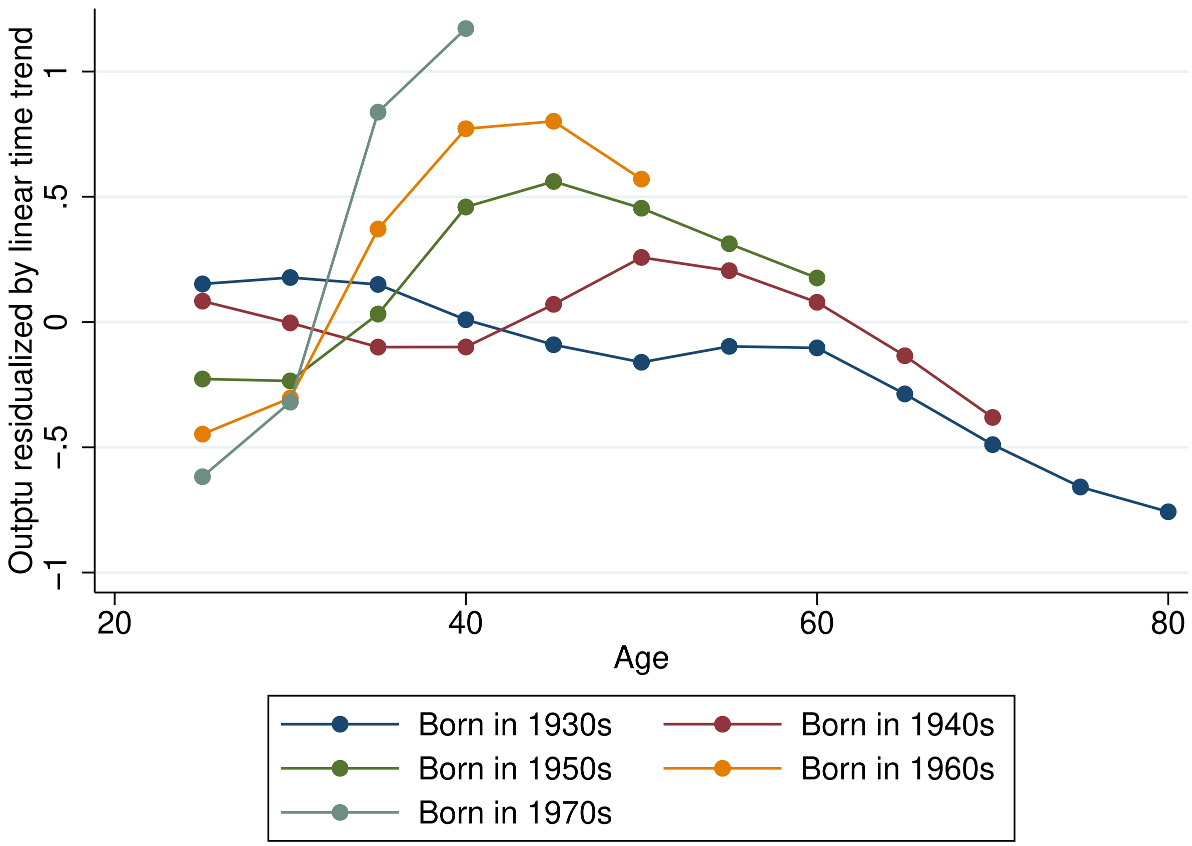

Notes: This Figure shows average output (residualized by a linear time trend) by birth cohort and age. We first regress individual output (the number of published papers) on a linear time trend and generate the residuals from that regression. We then compute the average of these residuals by birth decade and age.

Figure A-4: Productivity differences across cohorts- fractional counting of output

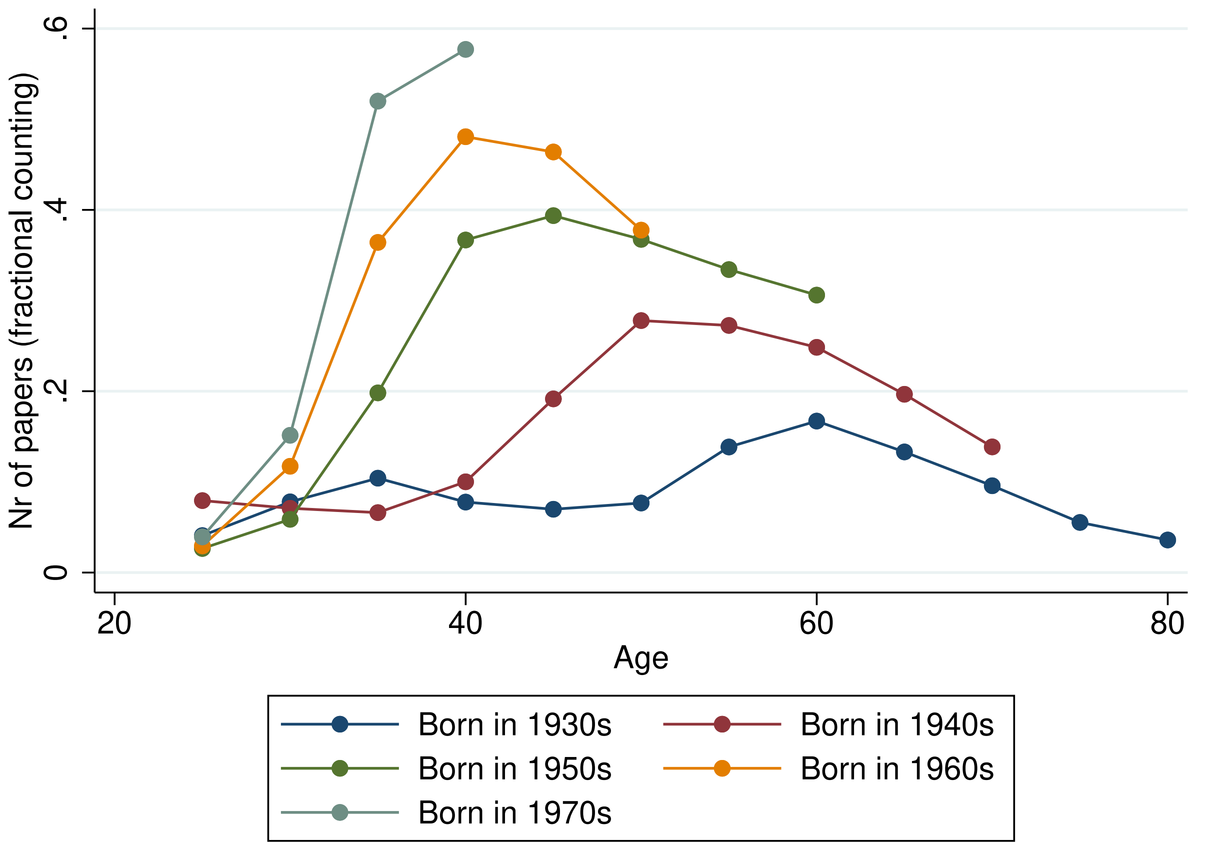

Notes: This Figure shows average scientific output by birth decade and age using fractional counting to adjust for the number of coauthors. Under fractional counting, the author of a paper with n coauthors receives 1/n of the credit for the paper.

Figure A-5: Productivity differences across cohorts- highly cited papers

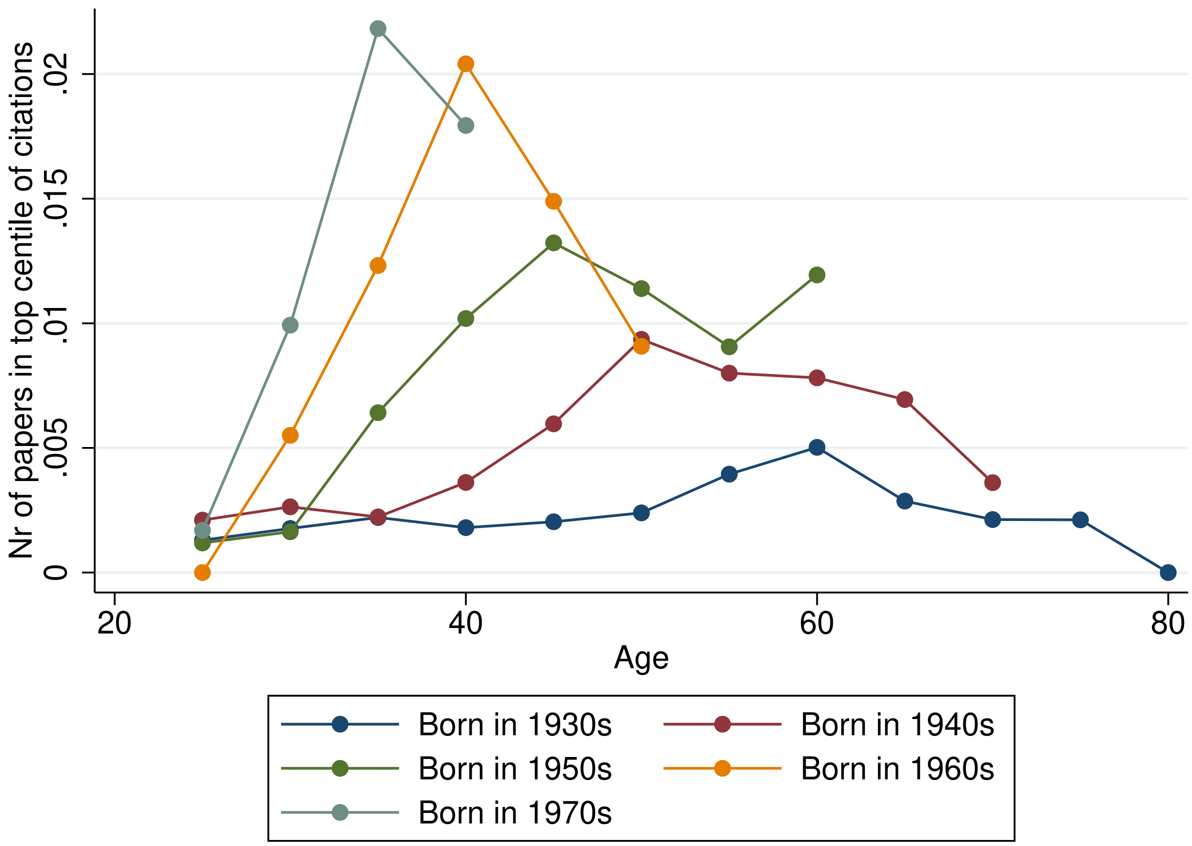

Notes: This Figure shows average scientific output by birth decade and age using an alternative measure of productivity: highly cited papers. We refer to a paper as highly cited if it is in the top centile of the citation distributions of papers published in any given year.

Figure A-6: Age-specific probabilities of exits across time

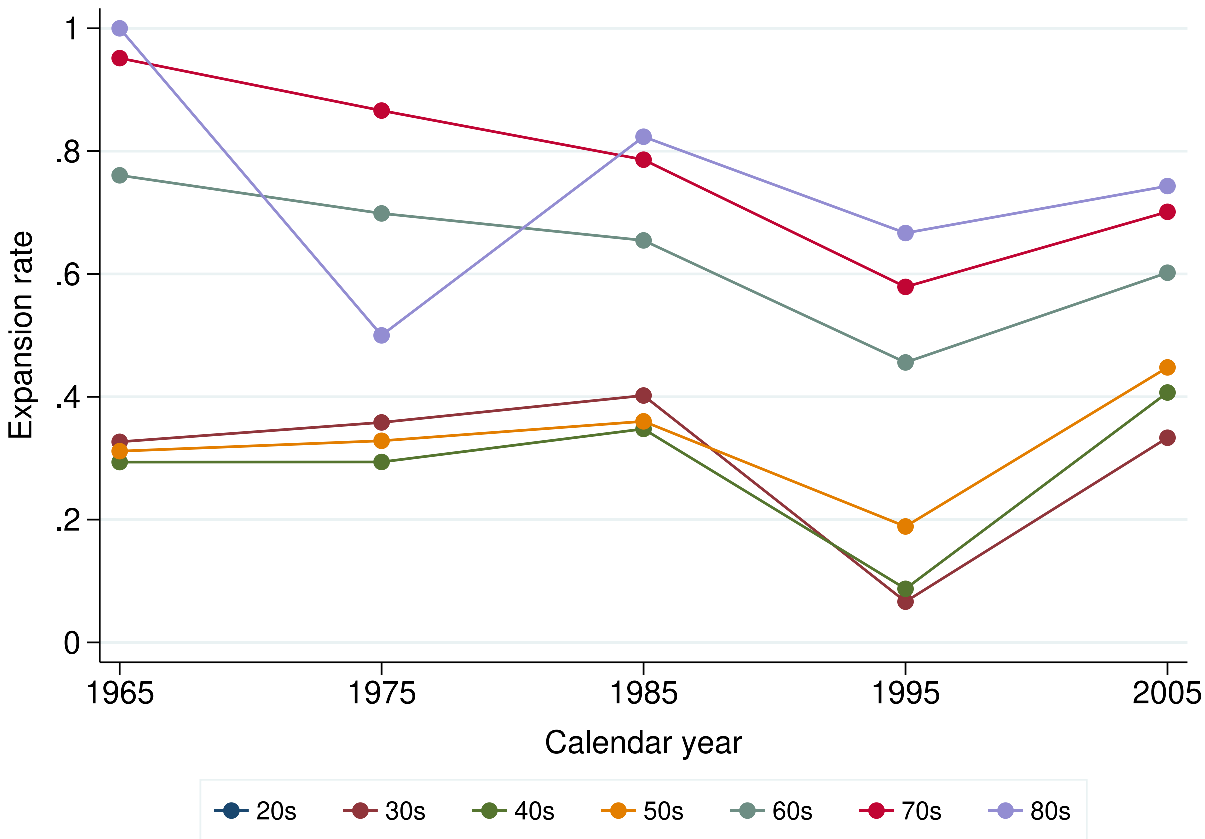

Notes: This Figure shows how the probabilities of exits in each age group evolved over time. There is a clear trend only in groups of 60s and 70s.

Figure A-7: Time-varying age distribution of newly hired is “ON”, expansion rate and exits are “OFF”

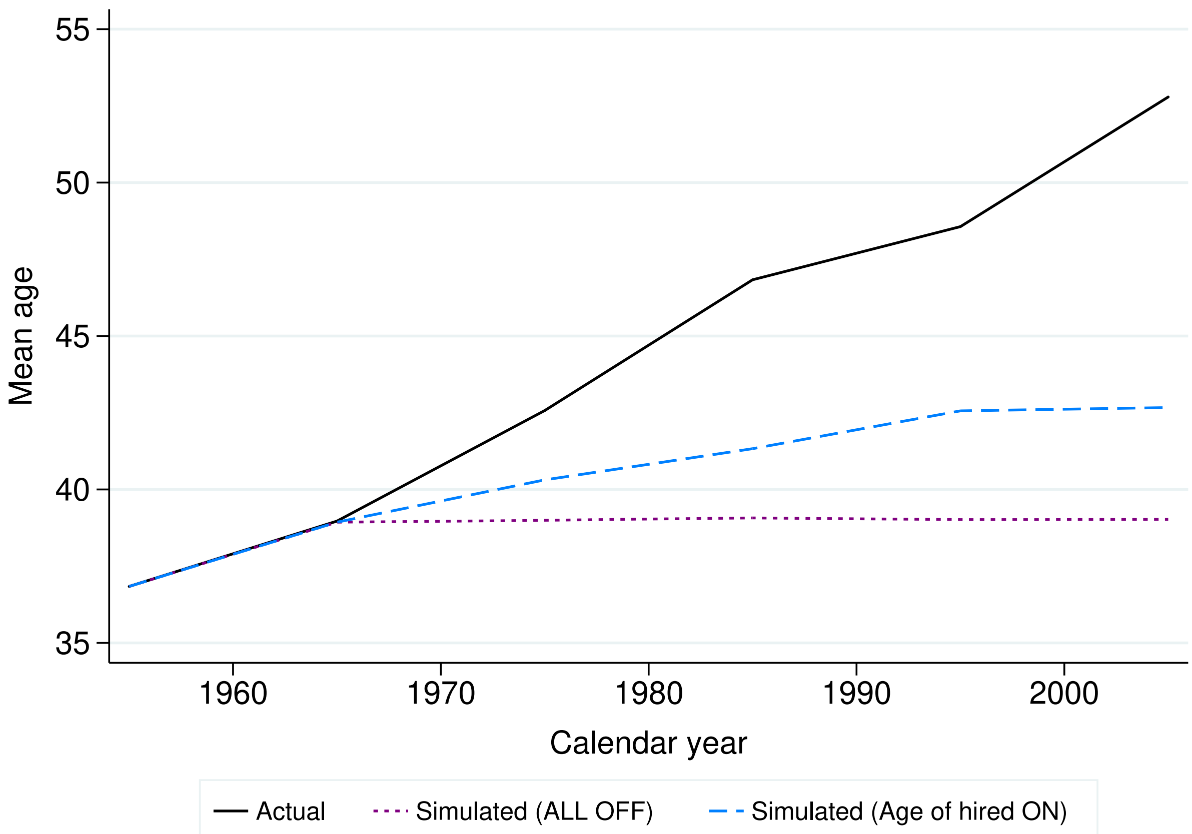

Notes: This Figure shows how the changes observed during the last 50 years in the age distribution of newly hired (the difference between the blue and red dashed lines) contributed to the increasing mean age of scientists (solid black line).

Figure A-8: Time-varying exits are "ON", age distribution of newly hired and expansion rate are "OFF"

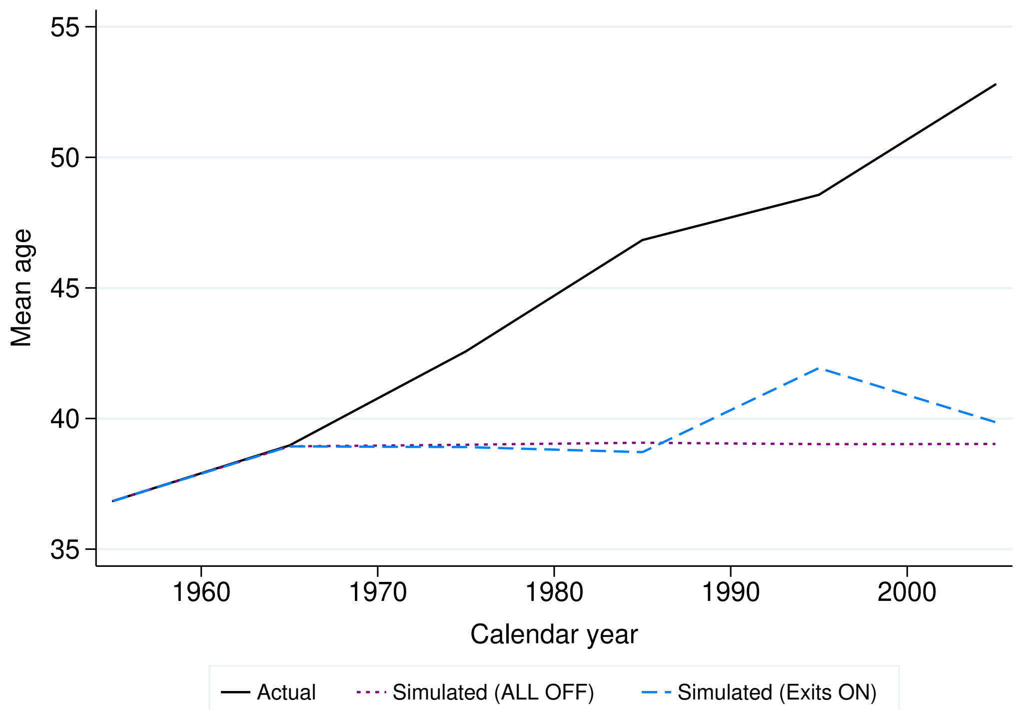

Notes: This figure shows how the changes observed during the last 50 years in the probability of exits (the difference between the blue and red dashed lines) contributed to the increasing mean age of scientists (solid black line).