Scientific productivity as a random walk

Sam Zhang,[1,2,3,*] Nicholas LaBerge,[4] Samuel F. Way,[4] Daniel B. Larremore,[4,5,3] and Aaron Clauset[4,5,3,†]

[1] *Department of Applied Mathematics, University of Colorado, Boulder CO 80309, USA*
[2] *Department of Mathematics and Statistics, University of Vermont, Burlington VT 05405, USA*
[3] *Santa Fe Institute, Santa Fe, NM 87501, USA*
[4] *Department of Computer Science, University of Colorado, Boulder CO 80309, USA*
[5] *BioFrontiers Institute, University of Colorado, Boulder CO 80309, USA*

The expectation that scientific productivity follows regular patterns over a career underpins many scholarly evaluations, including hiring, promotion and tenure, awards, and grant funding. However, recent studies of individual productivity patterns reveal a puzzle: on the one hand, the average number of papers published per year robustly follows the "canonical trajectory" of a rapid rise to an early peak followed by a gradual decline, but on the other hand, only about 20% of individual productivity trajectories follow this pattern. We resolve this puzzle by modeling scientific productivity as a parameterized random walk, showing that the canonical pattern can be explained as a decrease in the variance in changes to productivity in the early-to-mid career. By empirically characterizing the variable structure of 2,085 productivity trajectories of computer science faculty at 205 PhD-granting institutions, spanning 29,119 publications over 1980–2016, we (i) discover remarkably simple patterns in both early-career and year-to-year changes to productivity, and (ii) show that a random walk model of productivity both reproduces the canonical trajectory in the average productivity and captures much of the diversity of individual-level trajectories. These results highlight the fundamental role of a panoply of contingent factors in shaping individual scientific productivity, opening up new avenues for characterizing how systemic incentives and opportunities can be directed for aggregate effect.

I. INTRODUCTION

Scientific productivity, which is typically measured by the number of papers that a scholar publishes, underpins many evaluative processes over the course of an academic career, including hiring decisions, tenure and promotions, grant funding, and scientific prizes [1, 2]. Due to its broad importance, scientific productivity has been studied from a variety of angles, such as productivity over time, averaging over scholars [3–5]; productivity over scholars, averaging over time [6, 7]; and extremal statistics of the most productive or impactful papers or years within careers [8, 9]. While useful, these approaches leave unanswered key questions about scientific careers that depend on knowledge about the full distribution of scholarship.

For example, a substantial literature, spanning many decades and fields, documents a "canonical trajectory" in scientific productivity over a career. The canonical trajectory is when a researcher's productivity tends to rise rapidly to a peak in the early career followed by a gradual decline, a pattern which is robustly captured when many scientists' trajectories are averaged [3, 5, 10–12]. However, recent work has revealed that this canonical trajectory is not representative of most individual scientists, who instead exhibit a rich diversity of productivity trajectories [13], even as their average productivity reliably follows the canonical trajectory.

The discovery that the canonical trajectory is a misleading description of individual productivity patterns presents a puzzle: what mechanisms lead to both dramatic variability in individual productivity trajectories and simultaneously the canonical pattern in aggregate? Past explanations of a canonical pattern at the individual level have invoked ideas ranging from cognitive mechanisms [14] to psychological development [15] and economic mechanisms [11, 16]. Other explanations focus on the scientific reward mechanisms, in which scholars tend to become more stratified over the course of a career [12, 17, 18]. However, these ideas ignore the broad heterogeneities across scientists and institutions, and do not readily explain the empirical diversity of faculty productivity patterns [13]. As a result, little is known about mechanisms that generate realistic individual productivity trajectories.

Here, we propose and investigate a parsimonious explanation which links several simple observations by modeling scientific productivity as a discrete-time Markov chain, which we refer to as a random walk. First, individual faculty productivity fluctuates from year to year due to individually contingent factors and events, including the beginning of a new collaboration [19, 20], an experiment that fails [21], parenthood [22], or changing institutions [23, 24]. While individually unpredictable, these fluctuations combine to form recognizable statistical patterns in the aggregate. Second, these factors change over a career, such that the variability of fluctuations also changes across different career stages, with higher productivity fluctuations in the early career than in the later career. In fact, we will show that a random walk with a change in variance is sufficient to produce both the canonical trajectory and much of the observed variability around it. This change in variance explana-

* sam.zhang@uvm.edu
† aaron.clauset@colorado.edu

tion builds on past work that highlights the relationship between institutional forces and systemic incentives on the one hand and global patterns of productivity on the other [12, 18], and on work that emphasizes the central role of randomness and luck in scientific careers [25], e.g., the unpredictability of when faculty tend to publish their most highly cited papers [8, 9].

We formalize this explanation as a probabilistic generative model that can simulate the evolution of individual faculty productivities, which we validate against empirical data on the productivities of 2,085 computer scientists at PhD-granting universities in the US and Canada. We produce two models—a simplified model, and a full one. The simplified model shows that a change of variance in faculty careers is sufficient to produce the canonical trajectory while preserving individual variability. It crystallizes a set of sufficient conditions for producing canonical patterns, and allows us to explore the space of possible average trajectories. The full model shows that modeling productivity as a random walk captures many of the details of both individual productivities, and aggregate patterns like the canonical trajectory, while simultaneously revealing noteworthy certain non-Markovian patterns in real faculty productivity.

The full model fits two sets of parameters: the change points between career stages, which parameterizes the change of structural influences across a scientific career, and the parameters describing the distribution of productivity fluctuations within each career stage, which parameterizes the role of contingency and luck. Together, these assumptions model an individual researcher's productivity over time as a truncated random walk that cannot become negative, where individual step sizes are drawn from a distribution whose parameters depend on the individual's career stage.

We first show that the simplified model is sufficient for generating a diverse range of trajectories that reproduce the canonical trajectory in aggregate. We then fit the full model to the empirical data on computer scientists and obtain estimates of the model's change points, which represent the timings of major career transitions for faculty researchers, and the parameters for the random walk within each career stage. We directly validate the timing of the inferred career change points by comparing them to the typical timing of faculty promotions for this population of researchers. We then check the fitted model by generating an ensemble of simulated productivity trajectories, which we contrast with the empirical trajectories across a variety of statistical measures. The full model successfully explains a substantial portion of the variability of individual careers as well as the canonical trajectory pattern, while also revealing important discrepancies between the model and the data that indicate higher-order mechanisms and other contingent forces that shape scientific productivity.

II. DATA

We combine two comprehensive datasets to perform our analysis. First, we use a hand-curated census of all tenured or tenure-track faculty employed at all 205 US and Canadian computer science departments documented in the Computing Research Association (CRA)'s Forsythe List of PhD-granting departments in computing-related disciplines [26] in the academic year 2011–2012. This dataset includes 5,032 faculty, whose PhD-granting institutions and employment histories were manually gathered from public materials such as CVs and academic websites.

Second, we use the November 2016 snapshot of the Digital Bibliography and Library Project (DBLP, [27]), a large-scale bibliographic dataset for journals and conference proceedings relevant to computing research, although with limited coverage of interdisciplinary computing. The employment data is joined with the DBLP both algorithmically and manually, excluding preprints on the arXiv. By using publication data linked to definitive employment records, rather than inferring the start of careers from publications, as is common in the bibliometrics literature [28–30], we are able to isolate and analyze the dynamics of scholarly productivity under a relatively consistent and stable set of influences and incentives around productivity.

To account for DBLP's degraded coverage of publication records further back in time and non-stationarity in average productivities over time, we use the linear scaling developed by Way et al. [13] that adjusts the average productivity in DBLP to match the average productivity estimated from a random sample of CVs from the same population of researchers. This adjustment allows us to include researchers from different career stages into a single analysis, and to compare faculty at a similar career stage across cohorts. This adjustment results in a real-valued non-negative number for each faculty in each year t that we will denote as the *adjusted productivity* q_t. We denote the change in adjusted productivities as $\delta_t = q_{t+1} - q_t$.

We focus our analysis on the most productive years of a career, and where the population pattern of the canonical trajectory is strongest, by analyzing years 0–20 of the careers for all faculty who received their PhD on or after 1980. We refer to the number of years since the start of a professor's first assistant professorship as their *career age*, with their first year as career age 0.

To be included in our analysis, we require that faculty publish three or more papers indexed by DBLP before career age 5. These inclusion criteria result in a dataset of 2,085 faculty across 204 departments, and $128,816$ author-publication pairs. For a subset of our analyses, we select faculty whose careers span the full 21 years, which yields 510 careers. We designate these careers the *full trajectories*.

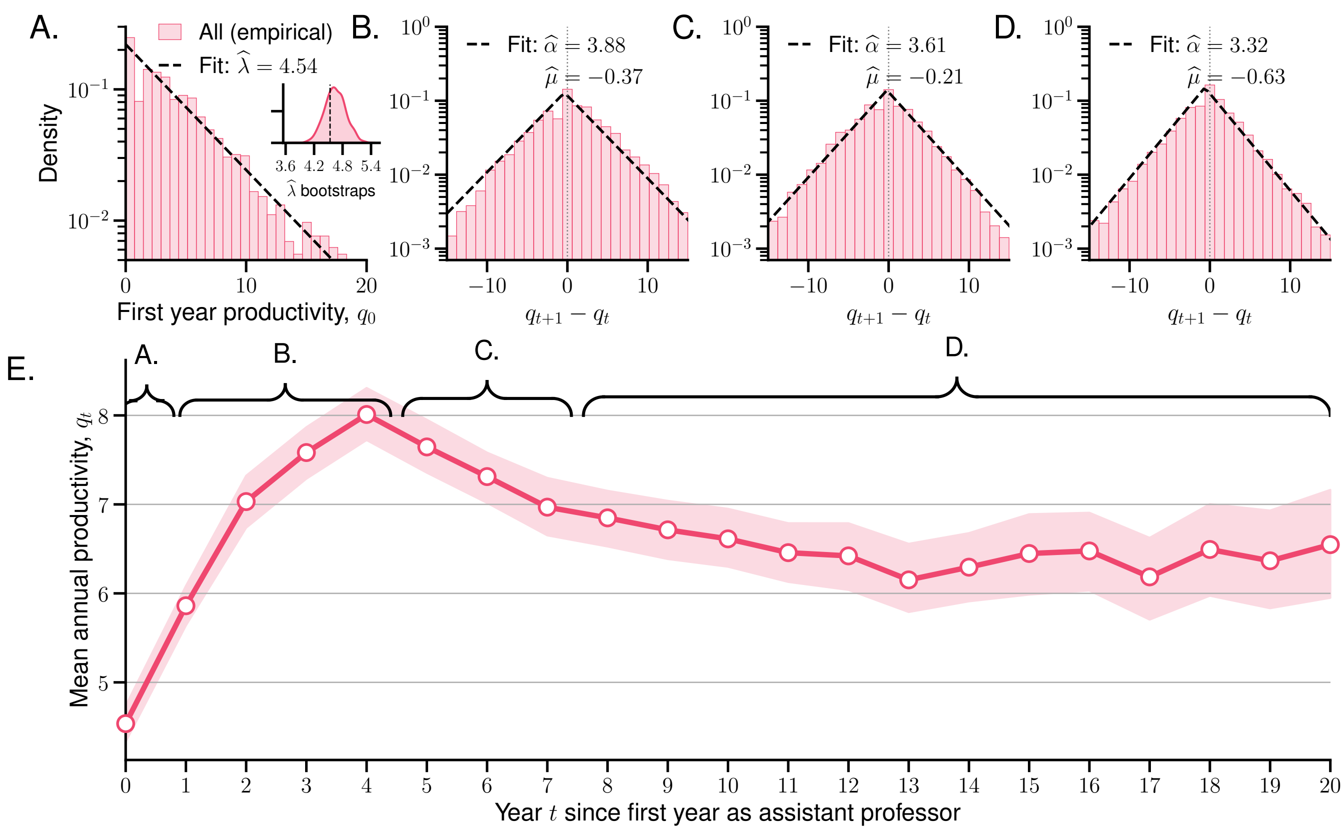

FIG. 1. **Empirical productivity data**. (A) An exponential distribution (dashed black line) accurately fits the empirical first-year productivity (pink histogram). The inset displays the estimated rate parameter against the density of estimated rates in 1,000 bootstrap replicas. (B-D) The empirical distributions of productivity changes (pink histograms) are semi-log plots, for ranges of career age, along with fitted Laplace distributions (dashed black line). (E) The average productivity for the same set of researchers, showing the "canonical trajectory" of a rapid rise followed by a gradual decline or leveling off, depicted as means of time-adjusted productivity for each career age and 95% bootstrap confidence intervals. Brackets indicate the range of career ages that were grouped together for the density plots: (A) productivity in year zero, and then changes of productivity in (B) years 1–4, (C) years 5–7, and (D) years 8–20.

III. RESULTS

A. Distribution of productivity changes

To study faculty careers from a perspective beyond average or extreme values, we characterize the stochasticity and variation within and across individuals by examining how productivity varies at the start of a career, and how it evolves empirically over time. We examine the distribution of first-year productivity q_0, and the distributions of changes in productivity $\delta_t = q_{t+1} - q_t$, and find surprising statistical regularity in both distributions: first-year productivity closely follows an exponential distribution (Fig. 1A), and the productivity changes follow a Laplace distribution regardless of career stage (Fig. 1B-D). The simple form of these empirical distributions is provocative, and suggests that the variability of initial productivity q_0 and subsequent changes to productivity δ_t may reflect relatively simple underlying stochastic processes.

Fitting exponential and Laplace distributions to the data, we notice that the estimated variances decrease from $\widehat{\alpha} = 3.88$ (95% CI: [3.78, 3.97]; all CIs are individual-level block bootstraps with 10,000 bootstrap replicas) to $\widehat{\alpha} = 3.64$ (95% CI: [3.54, 3.77]) and $\widehat{\alpha} = 3.32$ (95% CI: [3.19, 3.39]) over the course of a career (Fig. 1). On the other hand, the location parameters exhibit much more inferential uncertainty as well as the lack of any clear pattern, where between career years 1–4 and 5–7, the mode increases from $\widehat{\mu} = -0.37$ (95% CI: [-0.48, 0.35]) to $\widehat{\mu} = -0.21$ (95% CI: [-1.77, 0.32]), despite a change in the average trajectory from increasing to decreasing. This pattern suggests that the variance, rather than the location, of these distributions plays the key role in shaping the appearance of the canonical trajectory. The fact that across all career stages $\hat{\mu} < 0$ is intriguing, as it suggests a downward pressure on productivity over time, i.e., the mode of next year's productivity will be slightly lower than this year's.

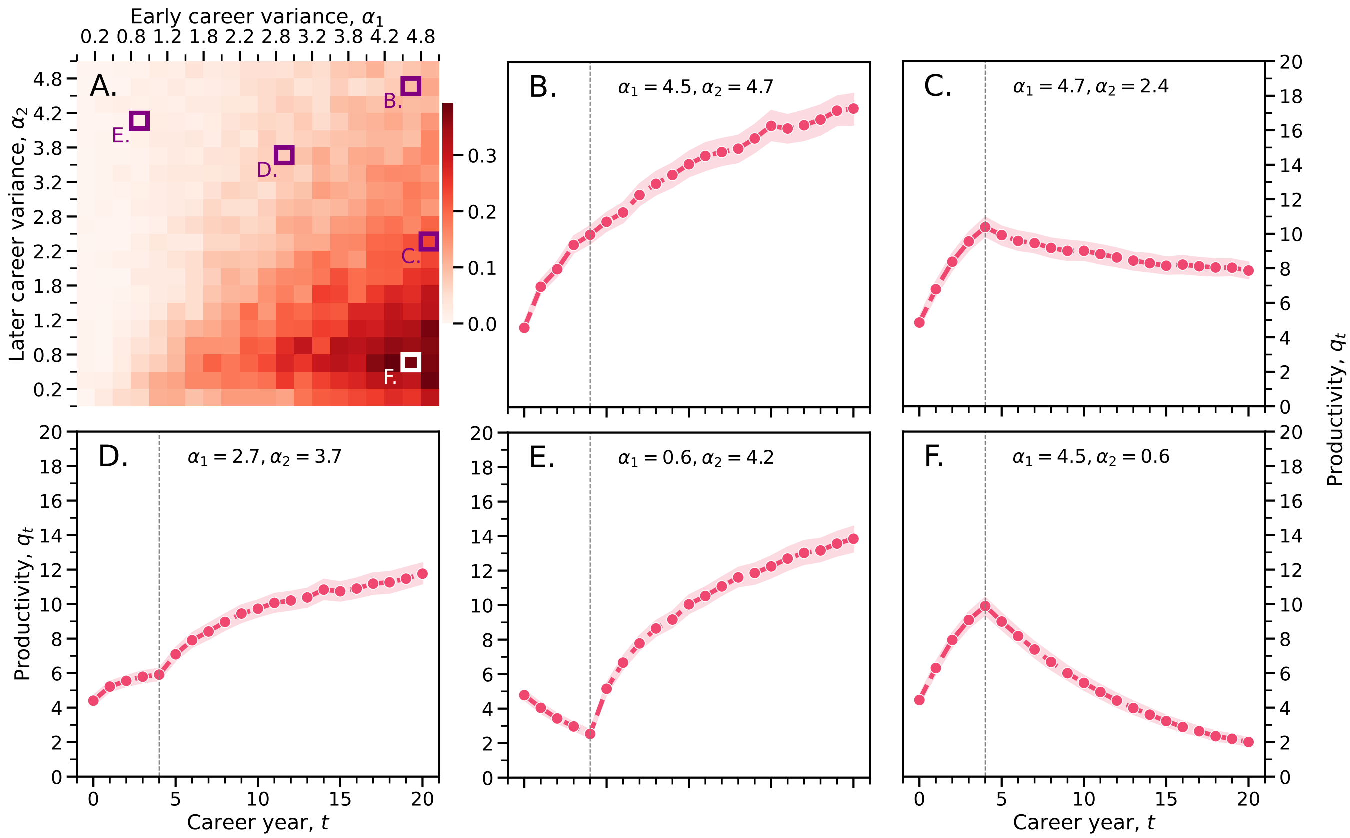

FIG. 2. Reproducing canonical trajectories with a simplified model. (A) Simulating $N = 400$ trajectories for each pair of α_1 and α_2 with $\mu = -1$ fixed, we display the fraction of those trajectories that are canonical. Some regions of the parameter space generate non-canonical trajectories (B, D, E), while others generate more canonical trajectories on average (C, F). Shaded intervals denote pointwise 95% confidence intervals for $N = 1000$ simulations at those parameters.

B. Modeling the canonical trajectory

Given the statistical regularity of the q_0 and δ_t distributions, we test whether changes in variance could drive the shape of the canonical trajectory by building a simple model. To do so, we build on the literature suggesting simple two-stage careers—that faculty productivity experiences a qualitative transformation around tenure, with rapid rise before and gradual decline after—to construct a simplified model with separate variance parameters for either stage [3, 5, 10–12].

Our simplified random walk model of the productivity of a faculty career is a discrete-time Markov chain with two free parameters: the variance in the early career α_1 (before year 5), and the variance in the later career α_2 (after year 5). Following our empirical observation that the mode is typically negative, we fix the mode of the distribution at $\mu = -1$. By simulating career trajectories at each pair of possible variances (α_1, α_2), we examine whether there exist necessary criteria on the variances of faculty productivity for producing canonical trajectories at the individual level.

Across the parameter space, we find that high variance in the early career paired with low variance in the later career $\alpha_2 < \alpha_1$, reliably produces a canonical trajectory at the individual level (Fig. 2C,F), while other choices of variances typically do not (Fig. 2B,D,E). In contrast, low variance in the early career followed by a higher variance later $\alpha_1 < \alpha_2$ tends to produce an aggregate trajectory with a "bounce", in which the average productivity falls to an early nadir, and then gradually rises over time. When the variances are equal or nearly so, the average productivity instead tends to rise to a level that is proportional to the variance's magnitude. Finally, regardless of the parameterization, most individual trajectories do not follow the corresponding aggregate trajectory, and instead individual trajectories exhibit the broad diversity of shapes observed in empirical data [13].

The appearance of the canonical trajectory when $\alpha_1 > \alpha_2$ occurs for a straightforward mathematical reason: because the random walk tends to drift toward zero ($\mu = -1$), but productivity cannot be negative ($q_t \geq 0$), the Markov chain's expected value will tend to relax onto a value that is roughly proportional to the variance. (We derive this behavior analytically in the Supporting Information.) The canonical pattern appears because initial productivity q_0 is close to zero, causing the average productivity to rise initially. But, because $\alpha_1 > \alpha_2$, the

Markov chain overshoots the expected productivity of the later career period, and at the beginning of that period, when the variance shifts to its lower value, the expected productivity then gradually falls. Hence, the canonical pattern can be explained as a natural consequence of a reduction in the variance of annual productivity over a career.

C. Modeling empirical productivity trajectories

While the simple model confirms that a change in variance is sufficient to produce a canonical trajectory in a two-stage career, real productivity trajectories may exhibit more than two stages. We therefore introduce a full model that decides on the number of career stages from the data, as well as the years spanned by each stage. To prevent overfitting to the data by adding overly many career stages, we regularize this model by fitting a productivity-dependent mode that allows greater shrinkage from high productivity values (see Supporting Information).

In this model, initial productivity is drawn from an exponential distribution with rate $\widehat{\lambda}_0$, and we estimate the number and location of breakpoints between career stages. In each career stage i, we further fit both scale $\widehat{\alpha}_i$ and location slope $\widehat{\beta}_i$ for the Laplace distribution governing the change in productivity, such that the conditional probability of observing a change in productivity δ following a year with productivity x is given by:

$$f(x, \delta) = \mathbb{I}\{\delta > -x\}(2 - e^{-\beta_i x/\alpha_i})^{-1}(1/\alpha_i)e^{-|\delta-\beta_i x|/\alpha_i}$$

These parameters can be accurately and efficiently estimated from data, and we confirm this fact by recovering known parameters, including various career breakpoints, from simulated data (see Supporting Information).

Fitted parameters. Despite the full model's increased complexity relative to the simplified model, its estimated parameters remain fully interpretable. The estimated career stages denote regimes with similar productivity dynamics, meaning a relatively stable set of factors, both systematic and contingent, that influence a scientist's productivity.

After fitting the full model to the set of 2,085 productivity time series in our data, we perform an initial check of the model's fit by examining the estimated parameters. The maximum likelihood fit yields four career stages: years 0–4, 5–7, 8–13, and 14–20 (Fig. 3A). These inferred career stages align well with common transitions that correspond to promotions in faculty careers, such as tenure evaluation which typically occurs in career years 5–7, and promotion to full professor, which often occurs about 12–15 years into a faculty career [31]. We note that the inferred change points varied across bootstrap replicas, with no set of maximum likelihood change points occurring in over 13% of replicates. The change points our procedure infers from the empirical data (4, 7, and 13) were the third most common set of change points in the bootstraps, occurring in 6.3% of replicas, behind (2, 4, 10) (12.9%) and (4, 5, 10) (6.4%) (Fig. 3A). Fitting the model to each of 1000 block bootstrapped resamples using individual faculty as the unit of resampling provides uncertainty estimates for all of the model's parameters. The relative instability of the inferred change point at year 13 is largely due to the fact that longer careers are less common in the data (full trajectories comprise only 510 (24.4%) of total trajectories, see Supporting Information); and only in the resamples with more of the full trajectories would the later career ages be detected as a change point. As a robustness check, we also fitted the full model to only the full trajectories, and find that the change point sets (4, 7, 11) and (4, 7, 13) are much more common across bootstrap replicas (23% in total).

Within the maximum likelihood career stages estimated from the full model (4, 7, 13), the estimated variances in the pre-tenure early career $\widehat{\alpha_1} = 4.5$ (95% CI: $[4.3, 4.6]$), $\widehat{\alpha_2} = 4.3$ (95% CI: $[4.1, 4.4]$) were higher than the variances in the later career $\widehat{\alpha_3} = 3.8$ (95% CI: $[3.7, 3.9]$), $\widehat{\alpha_4} = 3.5$ (95% CI: $[3.4, 3.7]$). Meanwhile, the estimated β_i parameter, which determines the mode of the career-stage Laplace distribution in conjunction with the productivity, fluctuated in an uncorrelated way with the average productivity, as did the mode μ_i when we fit it directly as a robustness check. This finding confirms the insights from the simplified model: the fitted full model produces the canonical trajectory through changes in variance, rather than changes in the typical productivity. Hence, counter-intuitively, the distribution of the number of papers that a researcher is likely to produce in the next year (given their current year's productivity) does not need to shift across a career in order to produce the aggregate pattern observed in the canonical trajectory. Rather, the canonical pattern can emerge merely from mid-career reductions in the variance in annual productivity.

Canonical trajectory. If the fitted full model includes the most salient aspects of individual productivity dynamics, then we expect simulations from the model to be statistically similar to the empirical trajectories.

First, we examine whether the model simulations display a canonical trajectory in aggregate. Indeed, our simulated trajectories evolve similarly to empirical productivity trajectories on average, successfully recovering the rapid rise and gradual decline (Fig. 3A). In fact, the average productivity is closely aligned between simulated and empirical trajectories, such that the largest average within-year difference between the two is less than one unit of productivity across an entire faculty career. This level of agreement is particularly notable because the model was fitted to individual level data, and yet it produces synthetic time series that yield the same aggregate pattern as the empirical data.

Career year of greatest productivity. The year of greatest productivity is not directly parameterized by

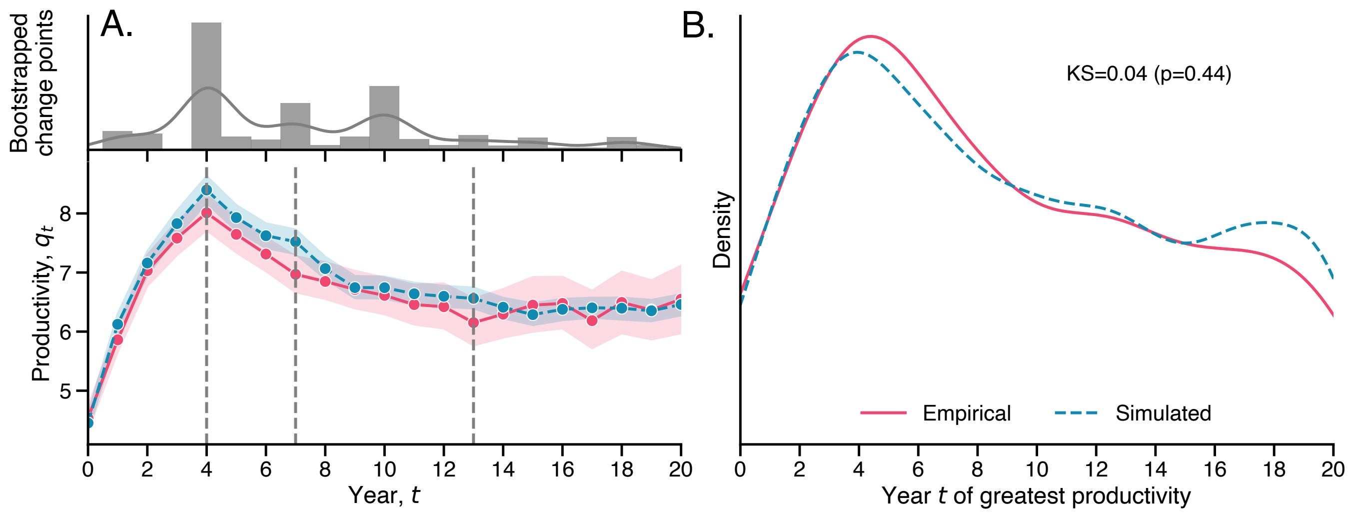

FIG. 3. **Fitting the empirical data** (A) Average productivity by career year for real and simulated trajectories, where shaded ribbons denote 95% confidence intervals. Dashed gray lines denote estimated career change points (at years 4, 7, and 13). Above, the bootstrap distribution of change points across 1000 bootstrap iterations, where bootstrap is conducted at the individual level. (B) Distribution of the years with greatest productivity among the full empirical and simulated trajectories. Distributions are similar across the entire career (KS = 0.04; $p = 0.44$).

the random walk model. To evaluate the model's accuracy on this pattern of productivity, when fitted to the full trajectories only, we examine the distribution of the year in which a trajectory reaches its maximum productivity for the full trajectories and for 10,000 trajectories simulated from the fitted model. We find that these two distributions (Fig. 3B) are statistically indistinguishable ($KS = 0.03, p = 0.75$), indicating that the model naturally explains this pattern in the data.

Variance within and across careers. Focusing on the full trajectories and computing the variance and standard deviations of productivity within each empirical and simulated trajectory, we find that the empirical trajectories tend to exhibit slightly lower variance than simulated trajectories ($KS = 0.21, p < 0.001$, Fig. 4A). The prevalence of years with zero publications in empirical trajectories, however, is not sufficient to explain this difference (Fig. 4A).

Empirically, faculty produce more cumulative papers by career year 5 than do simulated trajectories ($t = 9.16, p < 0.001$, Fig. 4C). This discrepancy is driven by a longer tail of cumulatively productive individuals in the empirical data who are not reproduced by the model: since researchers' productivity is lower variance than our model predicts (Fig. 4A), researchers with higher productivity are more consistently highly productive as well.

Years with zero publications. Comparing the empirical and simulated productivity distributions of the full trajectories, we observe that years with zero publications are substantially more common in the empirical data (15% vs 9%, Fig. 4B). Across empirical and simulated trajectories, the proportion of careers with exactly zero or one year of zero publications is similar, but empirical trajectories tend to have more zeros per trajectory than simulated ones (Fig. 4D). We note that the prevalence of years of zeros cannot be explained due to data quality issues within DBLP (see Supporting Information), and hence this discrepancy suggests that the dynamics that occur around a non-publishing state are not currently captured in our random walk model.

IV. DISCUSSION

Scientific understanding about large-scale patterns in faculty productivity has been overly focused average phenomena, such as the canonical trajectory, rather than on the dramatic variability of individuals. This focus has drawn the field to incomplete theories of scientific productivity, such as individual-level theories that posit an increase and decline of individual capabilities (e.g., scientific creativity and energy) over the course of a career, that attempt to explain the canonical average without accounting for the environmental determinants of productivity [24, 32] or the broad diversity of real scientific trajectories. This empirical diversity of real productivity patterns poses a major challenge to all individual-level theories of scientific productivity, because it requires a successful theory to explain both the average pattern as well as the large variations across and within individuals.

In this work, we discover two previously unknown statistical regularities: one in the distribution of early-career productivity and one in the distribution of year-to-year fluctuations in productivity (Fig. 1). We leverage these regularities to create a parsimonious explanation of productivity as a random walk where the variance in step size itself changes in a specific way across career stages.

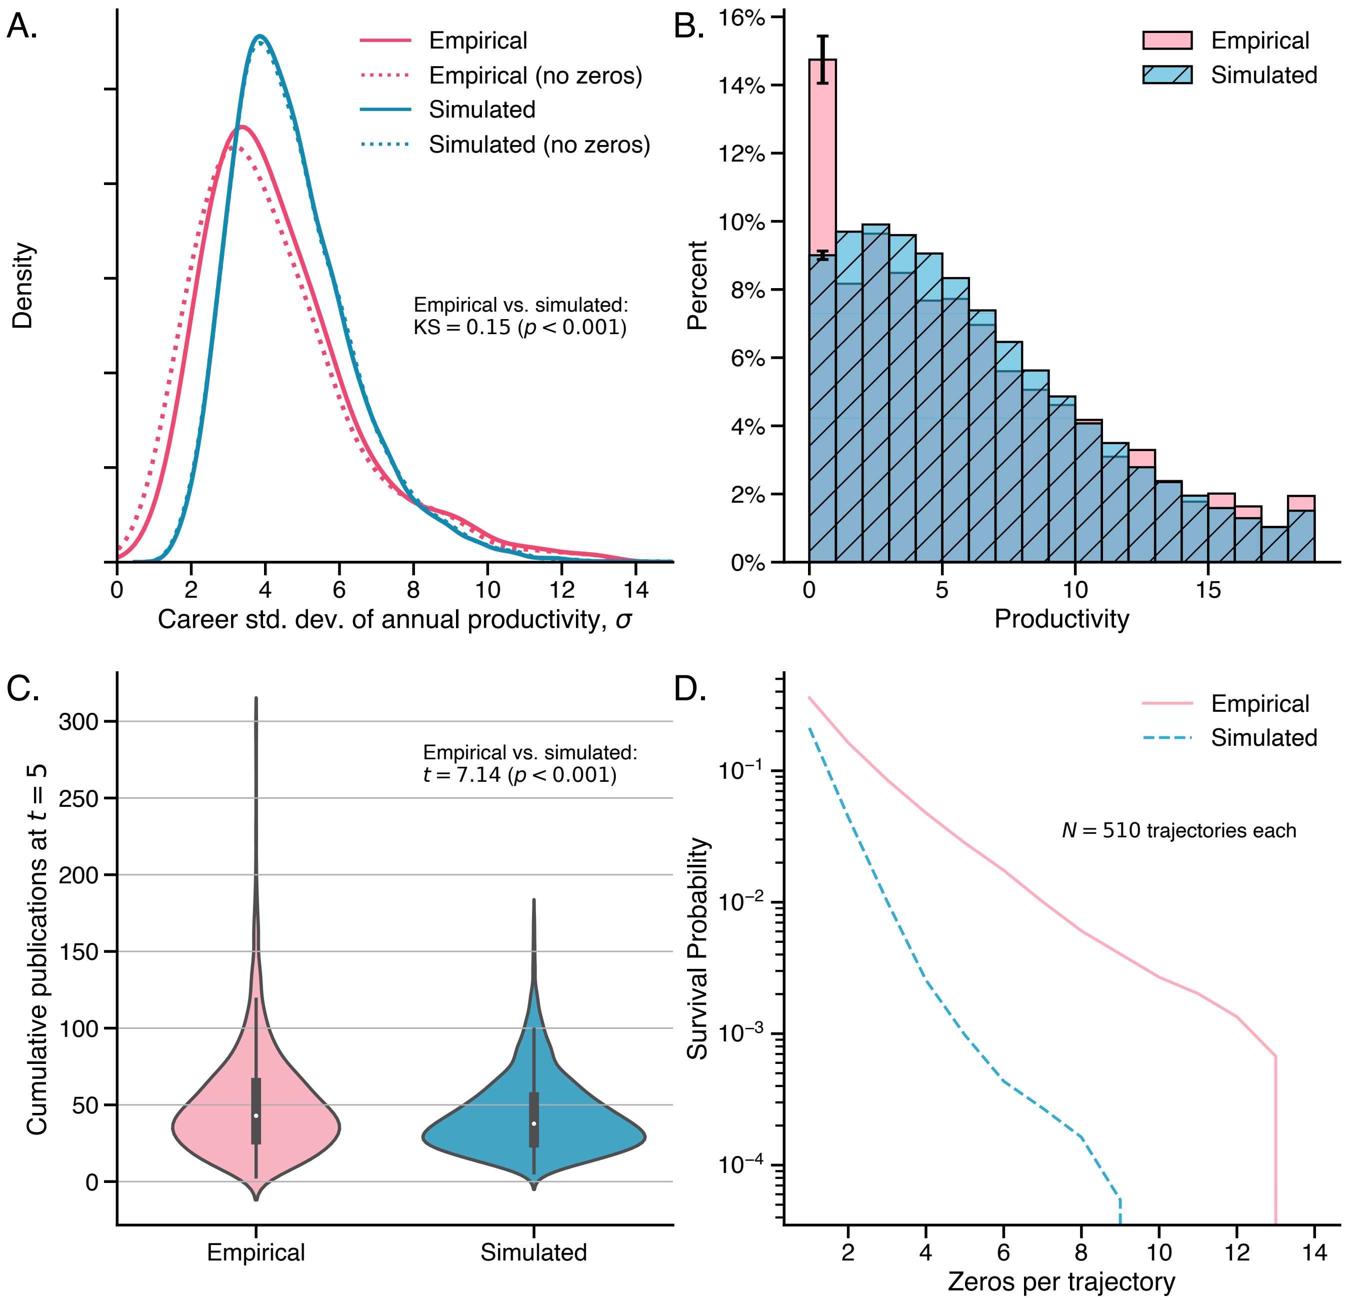

FIG. 4. **Comparing the random walk model to empirical data.** (A) Distributions of within-career standard deviations of productivity, for full empirical and simulated trajectories showing that empirical productivity variation tends to be slightly smaller (KS $= 0.14$; $p < 0.001$), even if we omit zeros. (B) The distribution of annual productivities (full trajectories), showing a close match for all values except at zero between empirical and simulated careers. Black bars indicate the binomial 95% Wald confidence intervals for the probability of zero publications. (C) Distributions of productivity of empirical and simulated trajectories at career year 5. Inside the violin plot, the white circle indicates the median, the thick bar indicates the interquartile range, and the thin bar indicates the centered 95% containment interval. By career year 5, the simulated trajectories tend to have fewer publications than the empirical trajectories on average ($t = 0.16$; $p < 0.001$), and the difference is especially pronounced among the tail of the most productive individuals. (D) Distributions of career years with zero publications within full empirical and simulated trajectories. The distribution of simulated and empirical trajectories with exactly one zero is similar, but more empirical trajectories exhibit more than one zeros than the simulated trajectories.

The model recapitulates both the canonical trajectory in average productivity and many empirical characteristics of the diversity of individual trajectories. These results, as well as the career statistics that the model does not fully reproduce, constitute a new perspective of scientific productivity and faculty careers rooted in randomness.

The key insight of this model—that a random walk with high variance in the early career followed by decreased variance in the later career can produce the canonical trajectory in aggregate while maintaining high individual diversity—highlights a critical open question: what drives this decrease in variance from the early to the later career of a scientist? A sociological explanation for the higher early career variance focuses on the structure of faculty career incentives: acquiring research grants, forming research groups, and publishing papers constitutes a critical component of tenure evaluation, so faculty are pressured in their early career to accelerate their research output in a short timespan, in a way unlike in the later career when the "start up" effects of an early career are more distant [12]. However, the institutional pressures that drive productivity patterns may differ outside of US and Canadian universities, and future work can assess the generalizability of these explanations to different global contexts.

For senior researchers, having an existing research group makes it more difficult to expand as much in relative terms—e.g., to quadruple the number of active group researchers from four to sixteen is much more challenging than to grow from one to four. Established researchers can also be more selective about grant applications to avoid the logistical difficulties of managing a rapidly expanding and contracting group. Additionally, in the later career, faculty have access to many more career paths than do early-career faculty, such as major university service roles related to curricular design and university administration, and scholarly service like editorships and professional society leadership, while the requirements for receiving tenure force all junior researchers into a narrower set of paths [12].

The existence of research groups and career roles point toward latent structure that is more complex than our model. Random walks are Markovian, or "memoryless", in that this year's productivity only depends on the prior year's productivity. In addition, faculty who enter research inactive career roles can be expected to exhibit more years with zero papers than what our simulation predicts, which is precisely what we find in the data (Fig. 4D). By contrast, graduate student, postdoctoral, and research staff contracts are generally longer than a year [24], meaning that a researcher's group size constitutes an unobserved latent variable that decreases the variance in faculty productivity. Both research groups and research inactive career roles reduce the variance in faculty productivity relative to a random walk, and indeed we observe slightly lower variances within empirical careers than what our model predicts (Fig. 4A), and higher cumulative variances across faculty (Fig. 4C). Even if individual productivity is more correlated across time than a memoryless model predicts, the discrepancy due to research inactive states is practically small relative to the remaining variance within careers (Fig. 4A). Nevertheless, future work could model latent variables such as research group size and faculty research roles directly using a hidden Markov Model to potentially capture these non-Markovian aspects of productivity trajectories [24].

The relationship between tenure evaluations (and faculty retention more broadly) and productivity is complex, and may potentially filter the data that we observe, especially in the full trajectory data. Faculty leave tenure-track positions for a variety of reasons, such as workplace climate, work-life balance, and professional reasons, and these reasons interact in complicated ways with attrition [31]. Attrition can happen among highly productive scholars who are pulled into industry positions, less productive scholars who fail to secure tenure, or average scholars who leave for non-professional reasons. For the results that we can compute using all of the trajectories, the corresponding analyses using the full trajectory data produce qualitatively similar outcomes, suggesting that the role of attrition on the findings are negligible.

The dynamical approach we construct here effectively subsumes more specific mechanistic models, and poses a further puzzle for researchers: why does faculty productivity follow such clear mathematical distributions (the exponential distribution for early-career productivity and the Laplace distribution for year-to-year changes in productivity), and why does a simple random walk model reproduce so many features of the empirical data, despite ignoring the main heterogeneities in academic careers such as prestige [32, 33], gender [19, 34, 35], parenthood [22], race [36], socioeconomic status [37], and subfield [38]?

One answer is that those heterogeneities are a subset of a panoply of contingent factors—tasks fundamental to the production of science such as delays in funding, student recruiting, peer review, coordination with collaborators including students, and regular variation due to the nature of research itself (experiments, data collection, computation, mistakes, dead ends, etc), not to mention non-academic sources of randomness, such as unexpected or variable life events—which are so numerous and unpredictable that together they constitute the bulk of the variation in productivity over time, giving rise to the appearance of dominating randomness. Indeed, the Laplace distribution can appear when heterogeneous random walks are themselves aggregated together [39].

The close agreement between the empirical data on changes in annual productivity and a Laplace distribution, which is symmetric, highlights a striking fact: the probability that a scientist's productivity increases next year by some amount very nearly equals the probability that it also decreases by the same amount in the following year. An interesting direction of future work would be to untangle the underlying factors and contingencies that make the distribution so symmetric. Ultimately, any symmetry between increases and decreases in productivity is imperfect, because scientists cannot produce fewer than zero papers any given year. This "hard" boundary plays a crucial role in explaining how an increase in productivity variance becomes an increase in average productivity. That is, when annual productivity is close to zero, the zero boundary censors the distribution of changes in productivity, and that censoring shifts the average displacement upward [40]. The higher the distribution's variance, the greater the censoring effect, and the larger the induced upward shift in the average change. In this way, the zero boundary induces a coupling between the variance in the distribution of changes to productivity with the average productivity itself.

The nuanced interplay between variance and productivity might illuminate unexplored pathways for shaping policy initiatives. Accelerating the process of obtaining extramural funding and hiring new team members could expedite the channeling of resources to innovative ideas, increasing the variance in downstream productivity. Decreased variability in later career stages could also result from adaptive learning—through planning, budgeting, research strategies, and so on—to mitigate the burdens associated with research fluctuations. If faculty had fewer unpredictable elements to manage, they might be able to devote that effort toward more research. In addition, researchers who become research inactive tend to remain so. While our results do not differentiate between researchers who choose to become research inactive and those who do not, the possibility that some researchers become stuck in inactivity simply through chance suggests potential policy interventions for recovering potential research contributions from seasoned researchers. In particular, grants or other mechanisms to assist faculty in research inactive career roles to transition back to research could prevent their subsequent contributions from becoming permanently lost to science.

The quantitative study of scientific careers and faculty productivity has been approached by many scholars, typically using techniques from a social science methodological toolkit such as descriptive data analysis and observational causal inference that aim to identify averages behind a veil of variability. Our results, based on a mechanistic model that centers this variability, show that changes in variance drive changes in the average, and that incentives and other system-level factors constrain and shape the way the fluctuations at the local level generate the aggregate trends. Our work suggests a shift in perspective: that individual-level fluctuations are an inherent part of research productivity, and that the panoply of contingent factors are an inherent part of the system to be understood rather than averaged away. This shift toward randomness and variability, away from deterministic laws, illuminates the broad diversity that characterizes real productivity patterns, within and across scientific careers.

V. ACKNOWLEDGMENTS

Funding: This work was supported in part by an Air Force Office of Scientific Research Award FA9550-19-1-0329 (NL, DBL, AC), an NSF Graduate Research Fellowship Award DGE 2040434 (SZ), and the NSF Alan T. Waterman Award SMA-2226343 (DBL). The complete data and source code for replicating the analysis and figures are available on Github (https://github.com/samzhang111/faculty-trajectories)

VI. AUTHOR CONTRIBUTIONS

SZ: Conceptualization, data curation, formal analysis, investigation, methodology, software, validation, visualization, writing; NL: Data curation, writing; SFW: Data curation; DBL: Conceptualization, data curation, funding acquisition, resources, visualization, writing; AC: Conceptualization, data curation, funding acquisition, investigation, methodology, project administration, resources, supervision, visualization, writing.

[1] Stephen Cole and Jonathan R Cole. Scientific output and recognition: A study in the operation of the reward system in science. *American Sociological Review*, 32(3):377–390, 1967.

[2] Paula Stephan. *How Economics Shapes Science*. Harvard University Press, 2012.

[3] Harvey Christian Lehman. *Age and Achievement*. Princeton University Press, 1953.

[4] Sharon G Levin and Paula E Stephan. Age and research productivity of academic scientists. *Research in Higher Education*, 30:531–549, 1989.

[5] Yves Gingras, Vincent Lariviere, Benoît Macaluso, and Jean-Pierre Robitaille. The effects of aging on researchers' publication and citation patterns. *PloS one*, 3(12):e4048, 2008.

[6] Alfred J Lotka. The frequency distribution of scientific productivity. *Journal of the Washington Academy of Sciences*, 16(12):317–323, 1926.

[7] Derek J De Solla Price. *Little science, big science*. Columbia University Press, 1963.

[8] Roberta Sinatra, Dashun Wang, Pierre Deville, Chaoming Song, and Albert-László Barabási. Quantifying the evolution of individual scientific impact. *Science*, 354(6312):aaf5239, 2016.

[9] Lu Liu, Yang Wang, Roberta Sinatra, C Lee Giles, Chaoming Song, and Dashun Wang. Hot streaks in artistic, cultural, and scientific careers. *Nature*, 559(7714):396–399, 2018.

[10] Wayne Dennis. Age and productivity among scientists. *Science*, 123(3200):724–725, 1956.

[11] Arthur M Diamond. The life-cycle research productivity of mathematicians and scientists. *J. Gerentol.*, 41(4):520–525, 1986.

[12] Stephen Cole. Age and scientific performance. *Am. J. Sociol.*, 84(4):958–977, 1979.

[13] Samuel F. Way, Allison C. Morgan, Aaron Clauset, and Daniel B. Larremore. The misleading narrative of the canonical faculty productivity trajectory. *Proceedings of the National Academy of Sciences of the USA*, 114(44):E9216–E9223, 2017.

[14] Karen L Horner, J Philippe Rushton, and Philip A Vernon. Relation between aging and research productivity of academic psychologists. *Psychol. Aging*, 1(4):319, 1986.

[15] Michael D Mumford. Age and outstanding occupational achievement: Lehman revisited. *J. Vocat. Behav.*, 25(2):225–244, 1984.

[16] Gary S Becker. *Human Capital: A theoretical and empirical analysis, with special reference to education*. University of Chicago Press, 2009.

[17] Jonathan R Cole, Stephen Cole, and Donald deB. Beaver. Social stratification in science. *American Journal of Physics*, 42(10):923–924, 1974.

[18] Barbara F Reskin. Scientific productivity and the reward structure of science. *American sociological review*, pages 491–504, 1977.

[19] Weihua Li, Sam Zhang, Zhiming Zheng, Skyler J Cranmer, and Aaron Clauset. Untangling the network effects of productivity and prominence among scientists. *Nature Communications*, 13(1):1–11, 2022.

[20] Sooho Lee and Barry Bozeman. The impact of research collaboration on scientific productivity. *Social Studies of Science*, 35(5):673–702, 2005.

[21] Yang Wang, Benjamin F Jones, and Dashun Wang. Early-career setback and future career impact. *Nature Communications*, 10(1):1–10, 2019.

[22] Allison C Morgan, Samuel F Way, Michael JD Hoefer, Daniel B Larremore, Mirta Galesic, and Aaron Clauset. The unequal impact of parenthood in academia. *Science Advances*, 7(9):eabd1996, 2021.

[23] J. Scott Long. Productivity and academic position in the scientific career. *American Sociological Review*, 43(6):889, 1978.

[24] Sam Zhang, K Hunter Wapman, Daniel B Larremore, and Aaron Clauset. Labor advantages drive the greater productivity of faculty at elite universities. *Science Advances*, 8, 2022.

[25] Aaron Clauset, Daniel B Larremore, and Roberta Sinatra. Data-driven predictions in the science of science. *Science*, 355(6324):477–480, 2017.

[26] Computing Research Association. CRA Forsythe List. https://archive.cra.org/reports/forsythe.html, 2012.

[27] The dblp team. dblp computer science bibliography. https://dblp.org/xml/release/dblp-2016-11-02.xml.gz, 2016.

[28] Wei Wang, Shuo Yu, Teshome Megersa Bekele, Xiangjie Kong, and Feng Xia. Scientific collaboration patterns vary with scholars' academic ages. *Scientometrics*, 112:329–343, 2017.

[29] Marek Kwiek and Wojciech Roszka. Academic vs. biological age in research on academic careers: A large-scale study with implications for scientifically developing systems. *Scientometrics*, 127(6):3543–3575, 2022.

[30] Balázs Győrffy, Boglárka Weltz, Gyöngyi Munkácsy, Péter Herman, and István Szabó. Evaluating individual scientific output normalized to publication age and academic field through the Scientometrics.org project. *Methodology*, 18(4):278–297, 2022.

[31] Katie Spoon, Nicholas LaBerge, K Hunter Wapman, Sam Zhang, Allison C Morgan, Mirta Galesic, Daniel B Larremore, and Aaron Clauset. Gender and retention patterns among U.S. faculty. https://osf.io/preprints/socarxiv/u26ze/ (2023).

[32] Samuel F. Way, Allison C. Morgan, Daniel B. Larremore, and Aaron Clauset. Productivity, prominence, and the effects of academic environment. *Proc. Natl. Acad. Sci. USA*, 116(22):10729–10733, 2019.

[33] K Hunter Wapman, Sam Zhang, Aaron Clauset, and Daniel B Larremore. Quantifying hierarchy and dynamics in us faculty hiring and retention. *Nature*, 610(7930):120–127, 2022.

[34] Vincent Larivière, Chaoqun Ni, Yves Gingras, Blaise Cronin, and Cassidy R Sugimoto. Bibliometrics: Global gender disparities in science. *Nature*, 504(7479):211–213, 2013.

[35] Nicholas LaBerge, Kenneth Hunter Wapman, Aaron Clauset, and Daniel B Larremore. Gendered hiring and attrition on the path to parity for academic faculty. *eLife*, 13:RP93755, jul 2024.

[36] Travis A Hoppe, Aviva Litovitz, Kristine A Willis, Rebecca A Meseroll, Matthew J Perkins, B Ian Hutchins, Alison F Davis, Michael S Lauer, Hannah A Valantine, James M Anderson, et al. Topic choice contributes to the lower rate of NIH awards to African-American/black scientists. *Science Advances*, 5(10):eaaw7238, 2019.

[37] Allison C Morgan, Nicholas LaBerge, Daniel B Larremore, Mirta Galesic, Jennie E Brand, and Aaron Clauset. Socioeconomic roots of academic faculty. *Nature Human Behaviour*, pages 1–9, 2022.

[38] Nicholas Laberge, K Hunter Wapman, Allison C Morgan, Sam Zhang, Daniel B Larremore, and Aaron Clauset. Subfield prestige and gender inequality among US computing faculty. *Communications of the ACM*, 65(12):46–55, 2022.

[39] Samuel Kotz, Tomasz Kozubowski, and Krzysztof Podgórski. *The Laplace Distribution and Generalizations: A revisit with applications to communications, economics, engineering, and finance*. Number 183. Springer Science & Business Media, 2001.

[40] William Feller. *An introduction to probability theory and its applications, Volume 1.* John Wiley & Sons, third edition, 1991.