https://doi.org/10.1038/s41586-018-0315-8

Hot streaks in artistic, cultural, and scientific careers

Lu Liu[1,2,3], Yang Wang[1,2], Roberta Sinatra[4,5,6], C. Lee Giles[3,7], Chaoming Song[8] & Dashun Wang[1,2,9]*

The hot streak—loosely defined as 'winning begets more winnings'—highlights a specific period during which an individual's performance is substantially better than his or her typical performance. Although hot streaks have been widely debated in sports[1,2], gambling[3–5] and financial markets[6,7] over the past several decades, little is known about whether they apply to individual careers. Here, building on rich literature on the lifecycle of creativity[8–22], we collected large-scale career histories of individual artists, film directors and scientists, tracing the artworks, films and scientific publications they produced. We find that, across all three domains, hit works within a career show a high degree of temporal regularity, with each career being characterized by bursts of high-impact works occurring in sequence. We demonstrate that these observations can be explained by a simple hot-streak model, allowing us to probe quantitatively the hot streak phenomenon governing individual careers. We find this phenomemon to be remarkably universal across diverse domains: hot streaks are ubiquitous yet usually unique across different careers. The hot streak emerges randomly within an individual's sequence of works, is temporally localized, and is not associated with any detectable change in productivity. We show that, because works produced during hot streaks garner substantially more impact, the uncovered hot streaks fundamentally drive the collective impact of an individual, and ignoring this leads us to systematically overestimate or underestimate the future impact of a career. These results not only deepen our quantitative understanding of patterns that govern individual ingenuity and success, but also may have implications for identifying and nurturing individuals whose work will have lasting impact.

According to the Matthew effect[9,23,24], victories bring reputation and recognition that can translate into tangible assets, which in turn help to bring future victories. This school of thought supports the existence of a hot streak in a career, which is also consistent with literature in the field of innovation showing that peak performance clusters in time, typically occurring around the middle of a career[8,11,21]. On the other hand, the random impact rule uncovered in the arts[10,21] and sciences[10,18] predicts the opposite: the best works occur randomly within a career, and their occurrence is primarily driven by productivity. Following this school of thought, works after a major breakthrough are not affected by what preceded them, supporting the viewpoint of regression towards the mean. The two divergent schools of thought raise a fundamental question: do hot streaks exist in creative careers?

To answer this question, we collected data sets recording the career histories of individual artists, film directors and scientists (Supplementary Information S1) and traced the impact of the artworks, films and papers they produced, approximated by auction prices[15], IMDB ratings (https://www.imdb.com/)[25] and citations garnered after 10 years of publication (C_{10})[13,16,18,26], respectively (see Methods). We started by investigating the timing of the three most impactful works produced in each career. In a sequence of N works by an individual, we denoted with N^* the position of the highest-impact work within a career, N^{**} the second highest and N^{***} the third. We found that each of the three highest-impact works was randomly distributed among all the works produced by an individual (Extended Data Fig. 1a–c), offering strong endorsement for the random impact rule[10,18,21].

However, as we show next, the randomness in individual creativity is only apparent, because the timing between creative works follows highly predictable patterns. We measured the correlation between the timing of the two biggest hits within a career, and compared it with a null hypothesis in which N^* and N^{**} each occured at random. The normalized joint probability, $\phi(N^*, N^{**}) = P(N^*, N^{**})/(P(N^*)P(N^{**}))$, is substantially overrepresented along the diagonal elements of matrices (Fig. 1a–c), demonstrating that N^* and N^{**} are much more likely to colocate with each other than would be expected from the random impact model across three domains. The diagonal pattern disappears if we shuffle the order of works within each career, thereby breaking the temporal correlations (Extended Data Fig. 1j–r).

To quantify the temporal colocation of hits observed in Fig. 1a–c, we calculated the distance between the two highest-impact works for every individual, measured by the number of works produced in between, $\Delta N = N^* - N^{**}$. We compared $P\left(\frac{\Delta N}{N}\right)$ of real careers with $P_S\left(\frac{\Delta N}{N}\right)$ of shuffled careers by defining $R\left(\frac{\Delta N}{N}\right) = P\left(\frac{\Delta N}{N}\right) / P_S\left(\frac{\Delta N}{N}\right)$. For artists, directors, and scientists, all $R\left(\frac{\Delta N}{N}\right)$ exhibit a clear peak centring around zero and decay quickly as ΔN deviates from zero (Fig. 1d–f). Notably, $R\left(\frac{\Delta N}{N}\right)$ is mostly symmetric around zero (Fig. 1d–f), indicating that the biggest hit is equally likely to arrive before or after the second biggest. The colocation patterns are not limited to the two highest-impact works within a career. We repeated our analyses for other pairs of hit works, such as N^{**} versus N^{***} and N^* versus N^{***}, and uncovered the same colocation patterns (Extended Data Fig. 1d–i).

Do high impact works come in streaks within a career? We counted the number of consecutive works whose impact exceeded the median of all works within a career (Extended Data Fig. 2d–f). We calculated the length of the longest streak L for each career. We then shuffled the order of works within each career, and measured again their longest streaks L_s. $P(L)$ was characterized by a much longer tail than $P(L_s)$ (Fig. 1g–i), indicating that real careers are characterized by long streaks of relatively high-impact works clustered together in sequence. We tested the robustness of these results by controlling for individual career length, and by varying our threshold used to calculate L, and arrived at the same conclusions (Extended Data Figs. 2–4, Supplementary Information S2). Together, these results raise an important question: what mechanisms are responsible for the temporal regularities observed across diverse career histories?

Let us first consider a null model in which the goodness of works produced in a career (that is, log(price) for artists, ratings for directors,

[1]Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL, USA. [2]Kellogg School of Management, Northwestern University, Evanston, IL, USA. [3]College of Information Sciences and Technology, Pennsylvania State University, University Park, PA, USA. [4]Department of Network and Data Science, and Department of Mathematics and its Applications, Central European University, Budapest, Hungary. [5]Center for Complex Network Research, Northeastern University, Boston, MA, USA. [6]Complexity Science Hub, Vienna, Austria. [7]Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA, USA. [8]Department of Physics, University of Miami, Coral Gables, FL, USA. [9]McCormick School of Engineering, Northwestern University, Evanston, IL, USA. *e-mail: dashun.wang@northwestern.edu

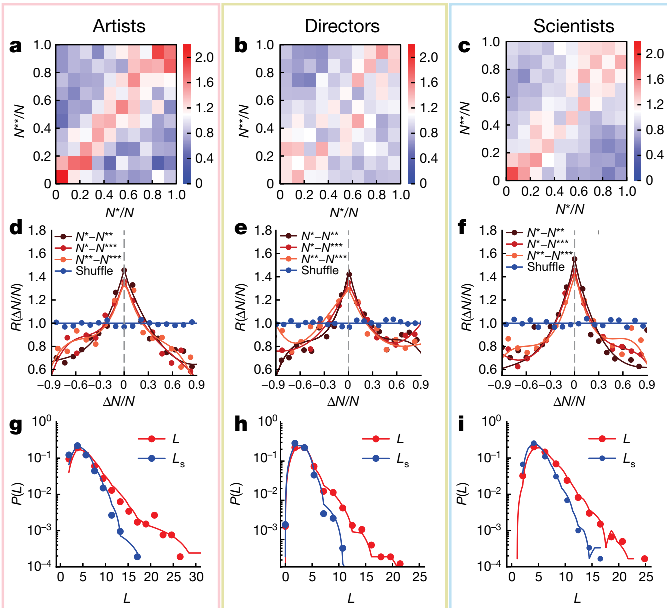

Fig. 1 | Hot streaks in artistic, cultural and scientific careers. a–c, $\phi(N^*, N^{**})$, colour coded, measures the joint probability of the two highest-impact works within a career for artists (**a**), directors (**b**), and scientists (**c**). $\phi > 1$ indicates that two hits are more likely to colocate than would be expected at random. **d–f**, $R\left(\frac{\Delta N}{N}\right)$ measures the temporal distance between highest-impact works relative to the null model's prediction. Real careers show a clear peak around 0 (red dots), which is well captured by the hot-streak model (solid lines). Different shades of red correspond to different pairs of hit works. Blue dots denote the same measurement but on shuffled careers, and blue lines are predictions from shuffled careers generated by our model. **g–i**, The distribution of the length of streaks $P(L)$ for real careers and $P(L_s)$ for shuffled careers. The hot-streak model (red lines) and its shuffled version (blue lines) closely reproduce $P(L)$ observed in real (red dots) and shuffled careers (blue dots).

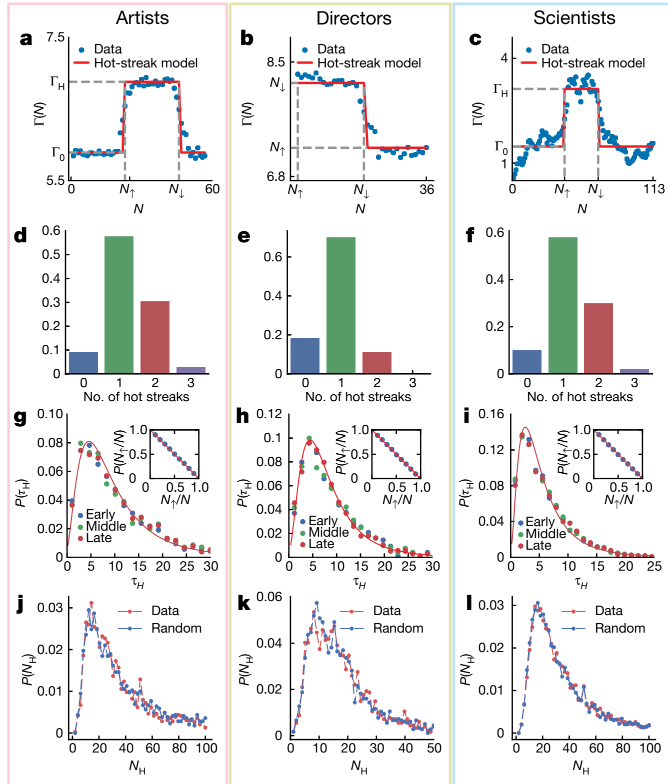

Fig. 2 | The hot-streak model. a–c, $\Gamma(N)$ for one artist (**a**), film director (**b**) and scientist (**c**), selected for illustration purposes. See Extended Data Fig. 9 for randomly selected careers. **d–f**, Histogram of the number of hot streaks in a career. We also measured several performance metrics for individuals who had one or two hot streaks, and found no detectable difference (Extended Data Fig. 10). **g–i**, The distributions of durations of hot streaks $P(\tau_H)$. Red lines are log-normal fits as a visual guide. The insets show cumulative distributions $P\left(\geq \frac{N_\uparrow}{N}\right)$, indicating that the start of a hot streak $N_\uparrow$ is distributed randomly among N works in a career. **j–l**, The distributions of the number of works produced during hot streaks $P(N_H)$, compared with a null distribution in which we randomly pick one work as the start of the hot streak. **j**, Artists ($n = 3{,}166$). **k**, Directors ($n = 5{,}098$). **l**, Scientists ($n = 18{,}121$). Two-sided Kolmogorov–Smirnov tests indicate that we cannot reject the hypothesis that the two distributions are drawn from the same distribution ($P = 0.12$ for artists, $P = 0.12$ for directors, and $P = 0.17$ for scientists).

and $\log(C_{10})$ for scientists) is drawn from a normal distribution $\mathcal{N}(\Gamma_i, \sigma_i^2)$ that is fixed for an individual. The average Γ_i characterizes the typical impact of works produced by the individual, and σ_i captures the variance. This null model can reproduce the fact that each hit occurs randomly within a career[10,18]. However, it fails to capture any of the temporal correlations observed in Fig. 1. The main reason for this failure is illustrated in Fig. 2a–c, where we selected for illustration purposes one individual from each of the three data sets and measured the dynamics of Γ_i during his or her career. We find that Γ_i is not constant throughout a career. Rather, it deviates from a baseline performance (Γ_0) at a certain point in a career ($t_\uparrow$), elevating to a higher value Γ_H ($\Gamma_H > \Gamma_0$), which is then sustained for some time before falling back to level similar to Γ_0 (Fig. 2a–c):

$$\Gamma(t) = \begin{Bmatrix} \Gamma_H & t_\uparrow \leq t \leq t_\downarrow \\ \Gamma_0 & \text{otherwise} \end{Bmatrix} \quad (1)$$

This observation, combined with the shortcomings of the null model, raises an intriguing question: could a simple model based on equation (1) explain the temporal anomalies documented in Fig. 1?

To test this hypothesis, we applied equation (1) to real productivity patterns, allowing us to generatively simulate the impacts of the works produced by an individual (Supplementary Information S3.3). During the period in which Γ_H operates, the individual seemingly performs at a higher level than his or her typical performance (Γ_0), prompting us to call this model the hot-streak model (where the Γ_H period corresponds to the hot streak). We introduced to each career one hot streak that occured at random with a fixed duration and magnitude, and repeated our measurements in Fig. 1 on careers generated by the model. We find that, whereas equation (1) introduces only a simple temporal variation, the hot-streak model is sufficient to reproduce all empirical patterns observed in Fig. 1 (Fig. 1d–i and Extended Data Fig. 1s–u). Given the myriad factors that can affect career impacts[9–12,18,22,27–30], and the obvious diversity of careers we studied, the level of universality and accuracy demonstrated by the simple hot-streak model was unexpected.

The real value of the model arises, however, when we fit the model to real careers to obtain the individual specific Γ_0, Γ_H, $t_\uparrow$ and $t_\downarrow$ parameters (Supplementary Information 3.4), helping us to reveal several fundamental patterns that govern individual careers.

1. Hot streaks are ubiquitous across careers, yet at the same time usually unique within a career. The vast majority of artists (91%, Fig. 2d), film directors (82%, Fig. 2e) and scientists (90%, Fig. 2f) have at least one hot streak during their careers, documenting the practical relevance of the uncovered hot streak phenomenon. However, despite its ubiquity, the hot streak is likely to be unique within a career. Indeed, when we relaxed our fitting algorithm to allow for multiple hot streaks (up to three) with different values of Γ_H, we found that, among those who had a hot streak, 64% of artists, 80% of directors, and 68% of scientists were best captured by one hot streak only (Fig. 2d–f), documenting the precious nature of hot streaks. Second acts may occur but are less likely, particularly for film directors. Occurrences of more than two hot streaks are rare across all careers.

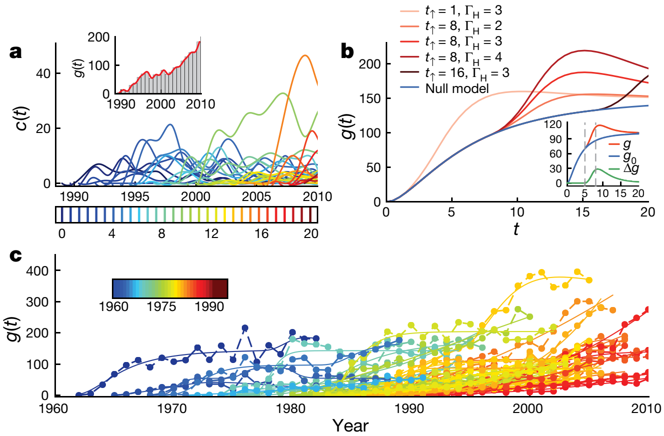

Fig. 3 | The hot streak governs the collective impact of a scientific career. a, Citation patterns of papers published by a randomly selected scientist in our data set. The publication dates are rearranged such that the individual produces a constant number of papers each year (coloured vertical lines, see Methods). Inset, collective impact of the individual, representing the sum of citation dynamics of all of his or her papers. **b**, Varying hot-streak parameters allows us to reproduce a wide variety of career dynamics (red lines) that cannot be captured by the null model (blue line). For hot-streak parameters in equation (4) (Methods), here we use $\mu=7.0$, $\sigma=1.0$, $\Gamma_0=1.0$ and $\tau_H=3$ years, but vary $t_\uparrow$ and Γ_H. Inset, $g(t)$ can be decomposed into $g_0(t)$ and $\Delta g(t)$. **c**, $g(t)$ of 50 randomly selected scientists. Colour corresponds to a career's starting year, dots denote real data, and solid lines capture the predictions from the hot-streak model.

2. The hot streak occurs randomly within a career. We estimate the beginning of hot streaks, by measuring $N_\uparrow$, the position of work produced when a hot streak starts ($t_\uparrow$). We find that, across artistic, cultural, and scientific careers, $N_\uparrow$ is randomly distributed in the sequence of N works within a career (Fig. 2g–i, insets). This finding reconciles two seemingly divergent schools of thought[9,10,18], providing a further explanation for the random impact rule: if the hot streak occurs randomly within a career, and the highest impact works are statistically more likely to appear within a hot streak, then the timing of the highest impact works is also random.

3. Across different domains, hot streaks are considerably shorter than the typical career length recorded in our database. We measure the duration distribution of hot streaks ($\tau_H=t_\uparrow-t_\downarrow$), finding $P(\tau_H)$ peaks around 5.7 years for artists, 5.2 years for directors, and 3.7 years for scientists, which is largely independent of when it occurs within a career (early, mid or late career; Fig. 2g–i).

4. Unexpectedly, individuals are not more productive during hot streaks. We measured the distribution of the total number of works produced during hot streaks $P(N_H)$. We then constructed a null distribution, by randomly picking one work in a career and designating its production year to be the start of the hot streak. We found that the two distributions aligned well with each other (Fig. 2j–l). Therefore, individuals show no detectable change in productivity during hot streaks, despite the fact that their outputs in this period are significantly better than the median, suggesting that there is an endogenous shift in individual creativity when the hot streak occurs. For additional properties of hot streaks, see Methods and Extended Data Fig. 5.

To investigate the impact of hot streaks on individual careers, we focused on scientific careers and measured the collective impact of a scientist, $g(t)$, defined as the total number of citations over time collected by all papers published by an individual (Fig. 3a). $g(t)$ can be derived analytically by combining the hot-streak model (equation (1)) and an existing model[16] for the citation patterns of papers (see Methods and Supplementary Information S5), consisting of two terms:

$$g(t)=g_0(t)+\Delta g(t) \qquad (2)$$

$g_0(t)$ captures a career's collective impact in the absence of a hot streak (that is, $\Gamma(t)=\Gamma_0$). Contributions from the hot streak are encoded in $\Delta g(t)$, driven by both the timing and magnitude of hot streaks (see Methods). Varying hot-streak parameters leads to substantial changes in the collective impact of a career (Fig. 3b). Hence the hot-streak model captures a wide range of impact trajectories that are followed by real careers (Fig. 3c), and the accuracy of the model is documented by several metrics (see Methods). Given that individuals improve substantially during hot streaks, the uncovered phenomena can be particularly crucial for understanding the long-term impact of a career (Extended Data Fig. 6).

We further tested several alternative hypotheses, each associated with possible origins of the uncovered hot streaks (see Methods and Supplementary Information S6). Of all hypotheses considered, the hot-streak model is the simplest and least flexible. However, it is the only model whose predictions are consistent with real careers (Extended Data Figs 7, 8). The identification of the true origins of hot streaks is beyond the scope of this work. As such, the hot streaks uncovered here should be treated in a metaphorical sense, highlighting an intriguing period of outstanding performance during individual careers without implying any associated drivers. Crucially, though, the findings presented here hold the same, regardless of the underlying drivers.

Online content

Any Methods, including any statements of data availability and Nature Research reporting summaries, along with any additional references and Source Data files, are available in the online version of the paper at https://doi.org/10.1038/s41586-018-0315-8.

Received: 20 October 2017; Accepted: 4 June 2018;
Published online: 11 July 2018

1. Gilovich, T., Vallone, R. & Tversky, A. The hot hand in basketball: on the misperception of random sequences. *Cognit. Psychol*. **17**, 295–314 (1985).
2. Miller, J. B. & Sanjurjo, A. Surprised by the gambler's and hot hand fallacies? A truth in the law of small numbers. *IGIER Working Paper No. 552* (2016).
3. Ayton, P. & Fischer, I. The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness? *Mem. Cognit*. **32**, 1369–1378 (2004).
4. Rabin, M. & Vayanos, D. The gambler's and hot-hand fallacies: theory and applications. *Rev. Econ. Stud*. **77**, 730–778 (2010).
5. Xu, J. & Harvey, N. Carry on winning: the gamblers' fallacy creates hot hand effects in online gambling. *Cognition* **131**, 173–180 (2014).
6. Hendricks, D., Patel, J. & Zeckhauser, R. Hot hands in mutual funds: short-run persistence of relative performance, 1974–1988. *J. Finance* **48**, 93–130 (1993).
7. Kahneman, D. & Riepe, M. W. Aspects of investor psychology. *J. Portfol. Manage*. **24**, 52–65 (1998).
8. Lehman, H. C. *Age and Achievement* (Princeton Univ. Press, Princeton, 1953).
9. Merton, R. K. The Matthew effect in science. *Science* **159**, 56–63 (1968).
10. Simonton, D. K. Age and outstanding achievement: what do we know after a century of research? *Psychol. Bull*. **104**, 251–267 (1988).
11. Jones, B. F. The burden of knowledge and the death of the renaissance man: is innovation getting harder? *Rev. Econ. Stud*. **76**, 283–317 (2009).
12. Petersen, A. M. et al. Reputation and impact in academic careers. *Proc. Natl Acad. Sci. USA* **111**, 15316–15321 (2014).
13. Stringer, M. J., Sales-Pardo, M. & Nunes Amaral, L. A. Effectiveness of journal ranking schemes as a tool for locating information. *PLoS One* **3**, e1683 (2008).
14. Radicchi, F., Fortunato, S. & Castellano, C. Universality of citation distributions: toward an objective measure of scientific impact. *Proc. Natl Acad. Sci. USA* **105**, 17268–17272 (2008).
15. Galenson, D. W. *Old Masters and Young Geniuses: The Two Life Cycles of Artistic Creativity* (Princeton Univ. Press, Princeton, 2011).
16. Wang, D., Song, C. & Barabási, A.-L. Quantifying long-term scientific impact. *Science* **342**, 127–132 (2013).
17. Way, S. F., Morgan, A. C., Clauset, A. & Larremore, D. B. The misleading narrative of the canonical faculty productivity trajectory. *Proc. Natl Acad. Sci. USA* **114**, E9216–E9223 (2017).
18. Sinatra, R., Wang, D., Deville, P., Song, C. & Barabási, A.-L. Quantifying the evolution of individual scientific impact. *Science* **354**, aaf5239 (2016).
19. Fortunato, S. et al. Science of science. *Science* **359**, eaao0185 (2018).
20. Zeng, A. et al. The science of science: From the perspective of complex systems. *Phys. Rep*. **714–715**, 1–73 (2017).
21. Simonton, D. K. *Scientific Genius: A Psychology of Science* (Cambridge Univ. Press, Cambridge, 1988).
22. Duch, J. et al. The possible role of resource requirements and academic career-choice risk on gender differences in publication rate and impact. *PLoS One* **7**, e51332 (2012).
23. Price, D. S. A general theory of bibliometric and other cumulative advantage processes. *J. Assoc. Inf. Sci. Technol*. **27**, 292–306 (1976).

24. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
25. Wasserman, M., Zeng, X. H. T. & Amaral, L. A. N. Cross-evaluation of metrics to estimate the significance of creative works. *Proc. Natl Acad. Sci. USA* **112**, 1281–1286 (2015).
26. Uzzi, B., Mukherjee, S., Stringer, M. & Jones, B. Atypical combinations and scientific impact. *Science* **342**, 468–472 (2013).
27. Wuchty, S., Jones, B. F. & Uzzi, B. The increasing dominance of teams in production of knowledge. *Science* **316**, 1036–1039 (2007).
28. Shockley, W. On the statistics of individual variations of productivity in research laboratories. *Proc. IRE* **45**, 279–290 (1957).
29. Redner, S. Citation statistics from 110 years of physical review. *Phys. Today* **58**, 49 (2005).
30. Moreira, J. A., Zeng, X. H. T. & Amaral, L. A. N. The distribution of the asymptotic number of citations to sets of publications by a researcher or from an academic department are consistent with a discrete lognormal model. *PLoS One* **10**, e0143108 (2015).

Acknowledgements We thank A.-L. Barabasi, R. Burt, J. Chown, J. Evans, C. Jin, L. Nordgren, W. Ocasio, B. Uzzi, Y. Yin, the Kellogg Insights and all members of NICO for invaluable comments. This work is supported by the Air Force Office of Scientific Research (AFOSR) under award number FA9550-15-1-0162 and FA9550-17-1-0089, and Northwestern University's Data Science Initiative. R.S. acknowledges support from AFOSR grant FA9550-15-1-0364 and from the Central European University Intellectual Themes Initiative 'Just Data'.

Reviewer information *Nature* thanks J. Walsh, J. West and the other anonymous reviewer(s) for their contribution to the peer review of this work.

Author contributions D.W. conceived the project and designed the experiments; L.L. and Y.W. collected data and performed empirical analyses with help from R.S., C.L.G., C.S. and D.W.; L.L., Y.W., C.S. and D.W. did theoretical calculations; all authors discussed and interpreted results; D.W. and L.L. wrote the manuscript; all authors edited the manuscript.

Competing interests The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41586-018-0315-8.
Supplementary information is available for this paper at https://doi.org/10.1038/s41586-018-0315-8.
Reprints and permissions information is available at http://www.nature.com/reprints.
Correspondence and requests for materials should be addressed to D.W.
Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

METHODS

Data description. We compiled three large-scale data sets of individual careers across three major domains involving human creativity. The first data set (D_1) consists of auction records curated from online auction databases, allowing us to reconstruct the career histories of 3,480 artists through the sequence of works they each produced, together with the impacts of the artworks, approximated by hammer prices in auctions[15]. D_2 contains profiles of 6,233 film directors recorded in the IMDB database, each career being represented by the sequence of films directed by the individual. As metrics that quantify the impact of a film correlate closely with each other[25], here we use the IMDB rating to measure the goodness of a film. Finally, our third data set (D_3) includes publication records of 20,040 individual scientists through a large-scale name disambiguation effort that combined the Web of Science (https://clarivate.com/products/web-of-science/) and Google Scholar (https://scholar.google.co.uk/) data sets. The impact of each paper is measured by citations garnered after 10 years of its publication[13,16,18,26] (C_{10}). Further details on data collection and curation are provided in Supplementary Information S1.

To study the impact of works across the three domains, we measured the distributions of hammer prices, IMDB ratings and paper citation counts in our data sets. Both hammer price (D_1) and C_{10} (D_3) follow fat-tailed distributions, well approximated by a log-normal function (Extended Data Fig. 2a, c), and the IMDB rating follows a normal distribution ranging between 1 and 10 (Extended Data Fig. 2b). To make sure C_{10} is not affected by citation inflation[14,18,31], we also measured a rescaled C_{10} (see Supplementary Information S1.3) and found that it also followed a fat-tailed distribution (Extended Data Fig. 2c, inset). Therefore, we take the logarithmic of hammer price and C_{10} (log(price) and log(C_{10})) to approximate the goodness of an artwork and scientific publication. Note that the choice of logarithmic for hammer prices and C_{10} is meant to be consistent with prior studies[14,18], and does not affect any of the conclusions of the paper. Indeed, the logarithmic function is a monotonically increasing function, hence it does not change the rank ordering of top hits in a career. Note that while the data sets we used in this paper cover a large collection of career histories across a wide range of domains, the data-driven nature of our study indicates that the scope of our data is limited to individuals who have had sufficiently long careers to provide enough data points for statistical analyses (Supplementary Information S1).

Additional properties of hot streaks. How much does an individual deviate from his or her typical performance during a hot streak? Do people with higher Γ_0 also experience more performance gain from hot streaks? We explored correlations between Γ_0 and Γ_H, finding them to be well approximated by a linear relationship (Extended Data Fig. 5a–c). Hence, individuals with better typical performance also perform better during their hot streaks. It is interesting to note that the coefficients are slightly less than 1 (Extended Data Fig. 5a–c). Hence $\Delta\Gamma \equiv \Gamma_H - \Gamma_0$ decreases with Γ_0 (Extended Data Fig. 5a–c, insets), suggesting that individuals with smaller Γ_0 benefit more from hot streaks. These results are again independent of when the hot streak occurs along a career (Extended Data Fig. 5a–c).

The temporally localized nature of the hot streak is also captured by its proportion over career length τ_H/T (Extended Data Fig. 5d–f). We compared the duration of hot streaks with typical career length, finding that the median hovers around 20% (0.17 for artists, 0.23 for directors, and 0.20 for scientists).

Analytical solutions for the collective impact of a scientific career, $g(t)$. Brought into the spotlight by popular websites such as Google Scholar, $g(t)$ is playing an increasingly important role in driving many critical decisions, from hiring, promotion and tenure to awarding of grants and rewards. Many factors are known to influence it, ranging from productivity[17,28] to citation disparity and dynamics[13,14,16,22,23,29] and temporal inhomogeneities along a career[11,17,18,21,22,30]. As our goal is to understand impact, here we bypass the need to evaluate the inhomogeneous nature of productivity[17,18] by rearranging the publication time of each paper, such that an individual produces a constant number of papers each year, denoted by n (Fig. 3a). To calculate $g(t)$, we need to incorporate the citations patterns of papers into our hot-streak model (equation (1)). A recent study[16] suggested that the citation dynamics of a paper published at time t_0 can be approximated by

$$C(t,t_0) = m\left(e^{\lambda\Phi\left(\frac{\ln(t-t_0)-\mu}{\sigma}\right)}-1\right) \equiv m\left(e^{\Gamma(t_0)\Phi\left(\frac{\ln(t-t_0)-\mu}{\sigma}\right)}-1\right), \qquad (3)$$

where m is a global parameter describing the typical number of references a paper contains, and $\Phi(\cdot)$ is the cumulative normal function, characterized by μ and σ, which capture the typical citation life cycle of a paper. The paper's ultimate impact is determined by its fitness[16], λ. To adapt equation (3) into our framework, we replace λ with $\Gamma(t_0)$, and for simplicity assume that μ and σ are fixed for different papers published by an individual. The resulting model, combining equations (1) and (3), can be solved analytically (Supplementary Information S5), allowing us to express $g(t)$ in terms of hot-streak parameters:

$$g(t) = \underbrace{nm\left(e^{\Gamma_0\Phi\left(\frac{\ln(t)-\mu}{\sigma}\right)}-1\right)}_{g_0(t)} + \underbrace{\begin{cases} 0 & t<t_\uparrow \\ nm(\Gamma_H-\Gamma_0)\Phi\left(\frac{\ln(t-t_\uparrow)-\mu}{\sigma}\right)C(t,t_\uparrow) & t_\uparrow \le t < t_\downarrow \\ nm(\Gamma_H-\Gamma_0)\left[\Phi\left(\frac{\ln(t-t_\uparrow)-\mu}{\sigma}\right)C(t,t_\uparrow)- \right. & \\ \left. \qquad \Phi\left(\frac{\ln(t-t_\downarrow)-\mu}{\sigma}\right)C(t,t_\downarrow)\right] & t \ge t_\downarrow \end{cases}}_{\Delta g(t)} \qquad (4)$$

Equation (4) consists of two terms. $g_0(t)$ captures a career's collective impact in the absence of hot streaks (that is, $\Gamma(t) = \Gamma_0$). Contributions from the hot streak are encoded in $\Delta g(t)$, driven by both the timing and magnitude of hot streaks ($t_\uparrow$, $t_\downarrow$, Γ_H, and $\Gamma_H - \Gamma_0$).

Evaluating the accuracy of the hot-streak model. We quantify the accuracy of our model in equation (4) using three metrics.

To account for the inherently noisy career trajectories, we first assign an impact envelope to each career, explicitly quantifying the uncertainty of model predictions (Extended Data Fig. 6g). We simulated $g(t)$ for each individual by assigning a Gaussian noise $\mathcal{N}(0, \sigma_s^2)$ to the fitted Γ_0 and Γ_H. For each paper i we randomly draw its Γ_i from a normal distribution, depending on whether the paper was published within the hot streak (Γ_H during hot streak, Γ_0 otherwise). The standard deviation σ_s represents the inherent noise of the goodness parameter defined in Supplementary Information S3.5. For each individual, we simulated $g(t)$ for 1,000 realizations, allowing us to obtain a distribution of $g(t)$, with one standard deviation offering an uncertainty envelope. We repeated the same procedures for the null model. We measure the fraction of $g(t)$ that falls within the envelope, finding that the distribution peaks close to 1 (Extended Data Fig. 6h), which indicates that most career trajectories are well encapsulated within the predicted envelopes.

The superior accuracy of our model is also captured by the mean absolute percentage error (MAPE). We compared the distribution of MAPE between the data and the predictions of the model (Extended Data Fig. 6f), finding again that the hot-streak model outperformed the null model. The improvement was most pronounced for an early onset of hot streaks (Extended Data Fig. 6i), which is also consistent with our model's predictions.

To account for model complexity, we also calculated the Bayesian information criterion (BIC) measure, which penalizes the number of parameters in the model. Compared with the null model, the hot-streak model has systematically smaller BIC (Extended Data Figs 6e), documenting that the hot-streak model better captures the collective impact of a career than the null model.

Implications of hot streaks for long-term career impact. The analytical framework presented here not only offers a new theoretical basis for our quantitative understanding of dynamical patterns governing individual career impact, but also may have implications for comparing and evaluating scientists (Extended Data Fig. 6). Indeed, for individuals whose hot streaks are yet to come, ignoring the hot streak may lead to underestimation of their impacts (Extended Data Fig. 6a, b), especially given the ubiquitous nature of hot streaks (Fig. 2f). On the other hand, an early onset of a hot streak leads to a high impact that peaks early but may not be sustained unless a second hot streak occurs (Extended Data Fig. 6c).

Testing alternative hypotheses. To explore the possibility that alternative hypotheses might explain the observed patterns, we tested several models that capture different dynamics of hot streaks (Supplementary Information S6.3), each associated with possible origins of the uncovered hot streaks. (A) A right trapezoid (Extended Data Fig. 7b) captures a sudden onset of a hot streak with a more gradual decline, as innovators may stumble upon a groundbreaking idea, which manifests itself in the forms of multiple artworks, films, or publications. Hence from an evolutionary perspective, the duration of a hot streak may characterize the time it takes for the temporary competitive advantage to dissipate. (B) An isosceles trapezoid model (Extended Data Fig. 7c) captures hot streaks that evolve and dissolve gradually over time, which may approximate social tie dynamics, as one individual's hot streak could be the result of a fruitful, repeated collaboration[27,32]. (C) Furthermore, individual performance may peak at a certain point in a career, prompting us to test inverted-U shape (Extended Data Fig. 7d) and tent functions (Extended Data Fig. 7e). (D) Last, a left trapezoid function (Extended Data Fig. 7f) captures a gradual startup period with a sharp cutoff, corresponding to career opportunities that can augment impact but last for a fixed duration.

We tested the validity of the four alternative hypotheses (A–D) by comparing each model's prediction with empirical observations on the relative order of the top six hits within a career. The symmetric patterns of ϕ and $R\left(\frac{\Delta N}{N}\right)$ observed in real careers suggest that the biggest hit is equally likely to appear before or after the second-biggest hit. The randomness of the relative ordering among hits is not limited to the two biggest hits. Indeed, we measured the position of the top three hits ($\tilde{N}$: mover) relative to the top six hits of the career, and compute $P(\tilde{N})$ for each of the three hits for artists, directors and scientists. We found a lack of predictive patterns for $P(\tilde{N})$ across the three domains, suggesting that the relative orders among the top six hits in real careers are random (Extended Data Fig. 7g, o, w). We tested hypotheses A–D systematically to describe real careers (Supplementary Information S6.3), and found that the hot-streak model was the only model whose predictions were consistent with real careers (Extended Data Fig. 7h–m, p–u, x–ac). As such, the hot-streak model also offers a superior fit to the data than the other models (Extended Data Fig. 7n, v, ad).

We also tested whether Markov models could account for our observations (Supplementary Information S6.2). We explored multiple variants of Markov models by introducing short-range correlations between the impacts of adjacent works, correlations between the volatility of their impacts, and hidden Markov model with two states, finding again that the hot-streak model stood out in its ability to describe the observed patterns (Extended Data Fig. 8 and Supplementary Information S6.2). Together, these results demonstrate that none of these alternative hypotheses alone can account for the empirical observations in real careers.

Code availability. Code is available at https://lu-liu.github.io/hotstreaks/.

Data availability. The data are available at https://lu-liu.github.io/hotstreaks/.

31. Pan, R. K., Petersen, A. M., Pammolli, F. & Fortunato, S. The memory of science: inflation, myopia, and the knowledge network. Preprint at https://arxiv.org/abs/1607.05606 (2016).
32. Palla, G., Barabási, A.-L. & Vicsek, T. Quantifying social group evolution. *Nature* **446**, 664–667 (2007).

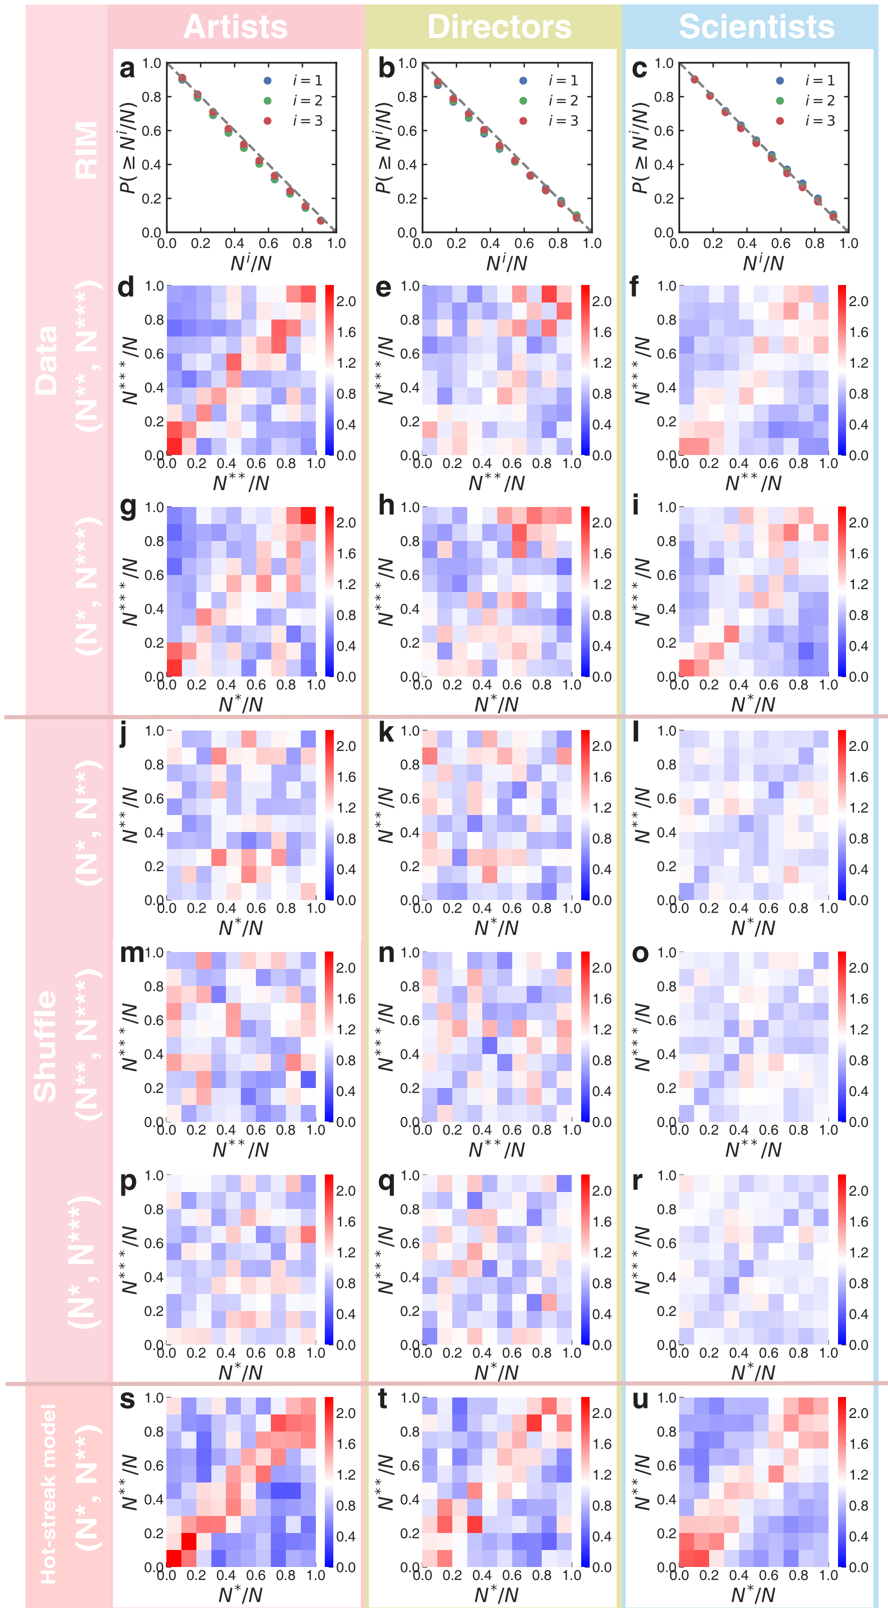

Extended Data Fig. 1 | Additional results on hot streaks in artistic, cultural, and scientific careers. a–c, The cumulative distribution $P(\geq N^i/N)$ for the order of the top three highest impact works within a career for artists (**a**), directors (**b**) and scientists (**c**). N^i denotes the order of the i^{th} highest-impact work within a career. The colours denote different hit works, and the dashed grey line denotes $P(\geq N^i/N)$ for a uniform distribution. **d–f**, $\phi(N^{**}, N^{***})$ for the second- and third-highest-impact works within a career. $\phi(N^{**}, N^{***})$ is also overrepresented along the diagonal. **g–i**, $\phi(N^{*}, N^{***})$ for the first- and third-highest-impact works within a career. **j–r**, We shuffled the order of each work in a career while keeping their impact intact. The diagonal patterns in **d–i** and Fig. 1a–c disappeared for shuffled careers. **s–u**, $\phi(N^{*}, N^{**})$ predicted by the hot-streak model successfully recovered the diagonal patterns observed in **a–c**. For **d–u** and Fig. 1a–c, we applied the same binning procedure to data, using bins that ranged from 0 to 1 with increments of 0.1.

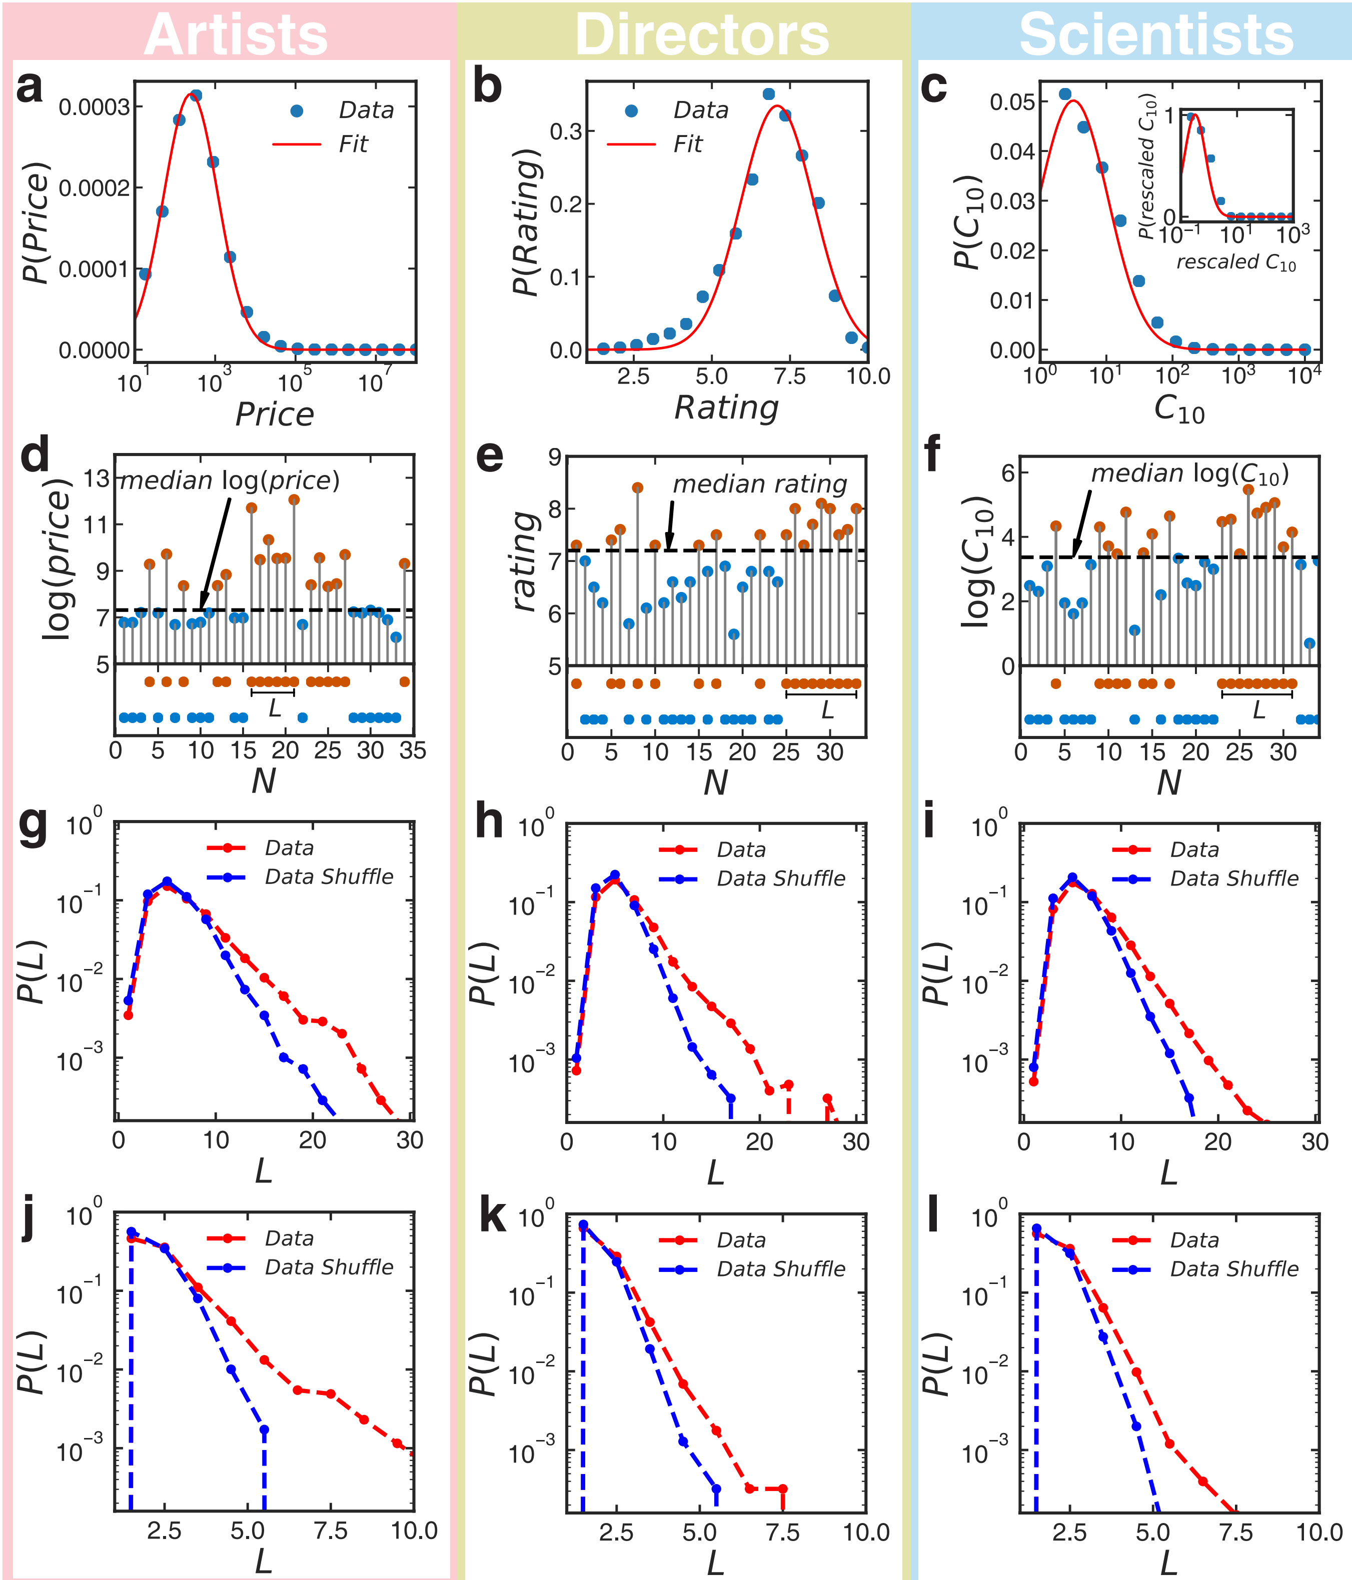

Extended Data Fig. 2 | Measuring the length of streaks using different thresholds. a, The distribution of auction price $P(\text{Price})$ for artists. Blue dots denote data, and the red line is a log-normal distribution with average $\mu = 7.9$ and standard deviation $\sigma = 1.5$. **b**, The distribution of film rating $P(\text{Rating})$ for directors. The red line is a normal distribution with average $\mu = 7.1$ and standard deviation $\sigma = 1.2$. **c**, The distribution of raw and rescaled C_{10} (inset) for scientists. The red line is a log-normal distribution, with $\mu = 2.3$ and $\sigma = 1.3$ for **c** and $\mu = -0.4$ and $\sigma = 0.8$ for the inset. **d–f**, Definitions of the longest streak L for artists (**d**), directors (**e**) and scientists (**f**). Dots are coloured orange above the threshold, blue otherwise. The lower panel highlights the longest streak in a career. **g–i**, $P(L)$ for real careers and $P(L_s)$ for shuffled careers using the mean impact within a career as the threshold. **j–l**, As in **g–i**, but using the top 10% impact as the threshold to calculate L and L_s. In all cases, $P(L)$ has a wider tail than $P(L_s)$, indicating that high-impact works in real careers tend to cluster together.

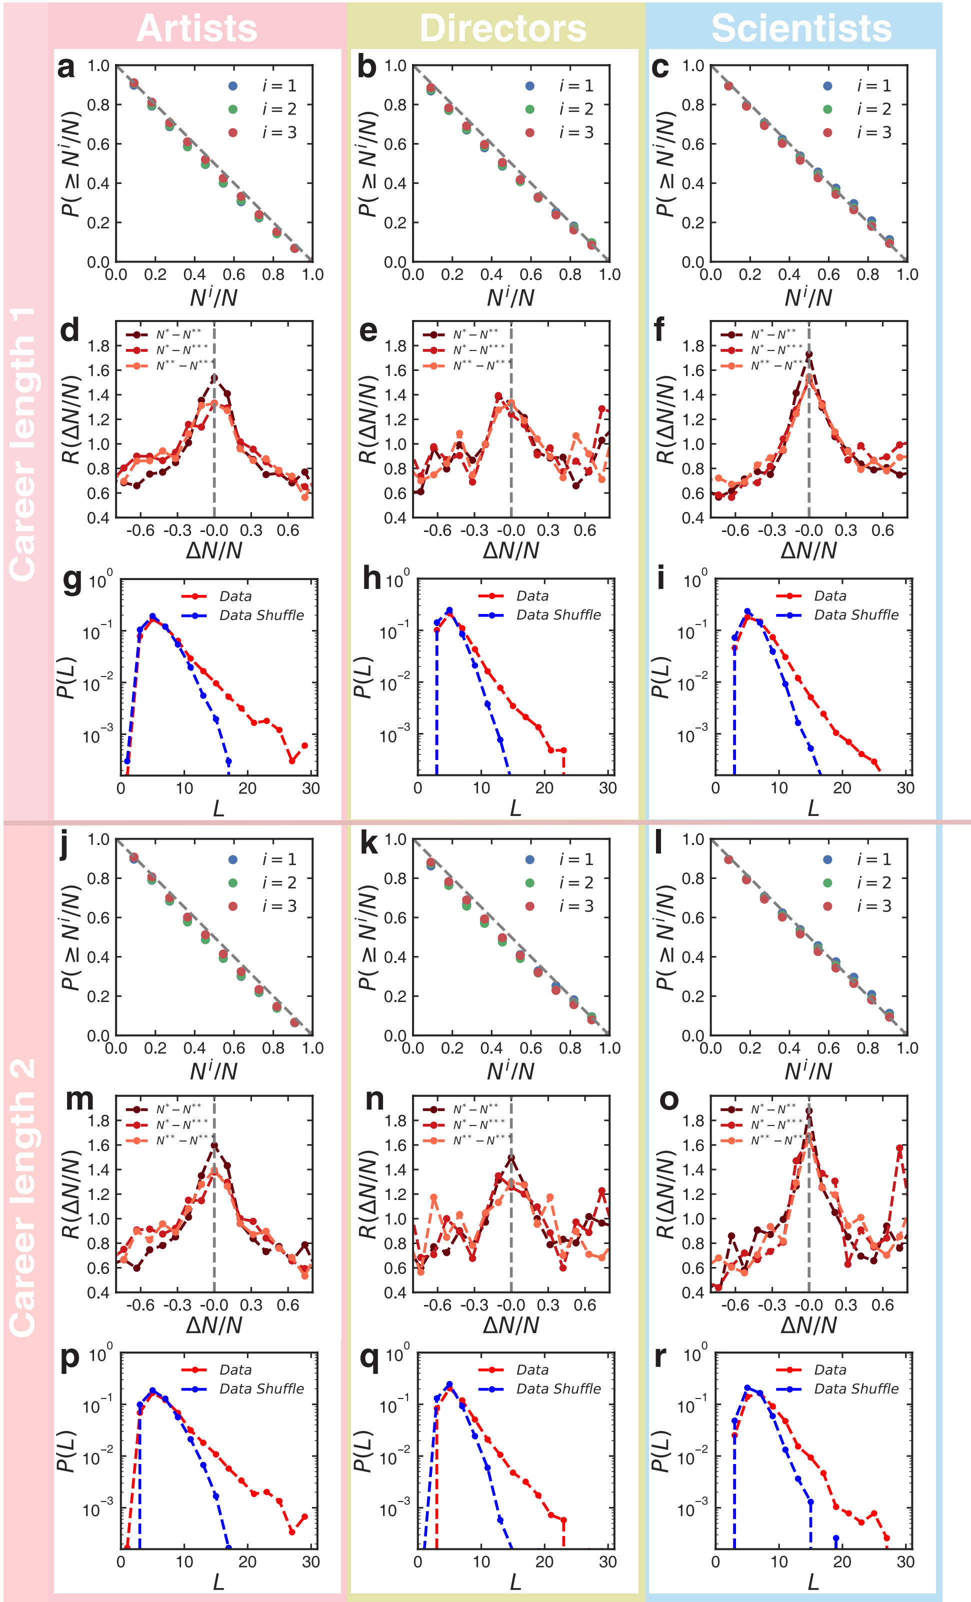

Extended Data Fig. 3 | Varying career length. To test the robustness of our results, we repeated our measurements by controlling for the career length of individuals. **a–i**, Artists and directors with careers of at least 20 years and scientists with careers of at least 30 years. **a–c**, $P(\geq N^i/N)$ of the top three highest-impact works within a career. **d–f**, $R\left(\frac{\Delta N}{N}\right)$ among the top three highest-impact works in a career. **g–i**, $P(L)$ for real careers and $P(L_s)$ for shuffled careers. **j–r**, As in **a–i** but for artists and directors with careers of at least 30 years and scientists with careers of at least 40 years. These results demonstrate that the patterns observed in Fig. 1 hold for individuals with different career lengths.

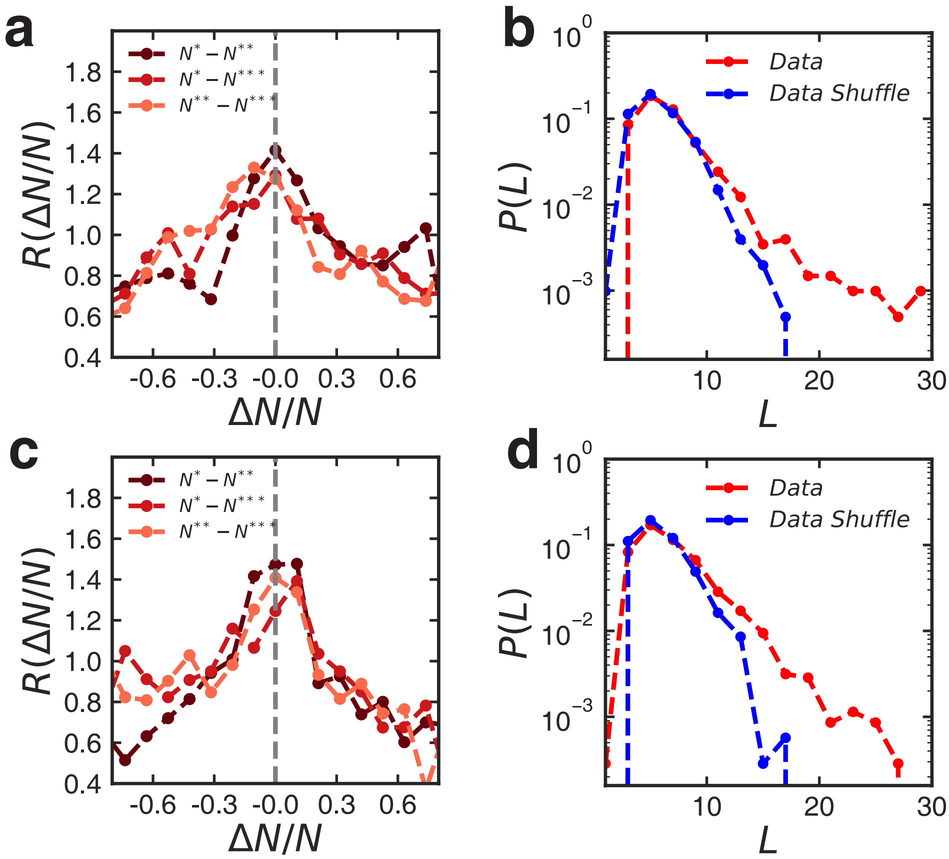

Extended Data Fig. 4 | Artistic careers from different eras. a, $R\left(\frac{\Delta N}{N}\right)$ for artists who started their careers before 1850. **b**, $P(L)$ for real careers and $P(L_s)$ for shuffled careers for artists who started their careers before 1850. **c**, $R\left(\frac{\Delta N}{N}\right)$ for artists who started careers between 1850 and 1900. **d**, $P(L)$ for real careers and $P(L_s)$ for shuffled careers for artists who started their careers between 1850 and 1900. These results demonstrate that the patterns observed in Fig. 1 hold for artists from different eras.

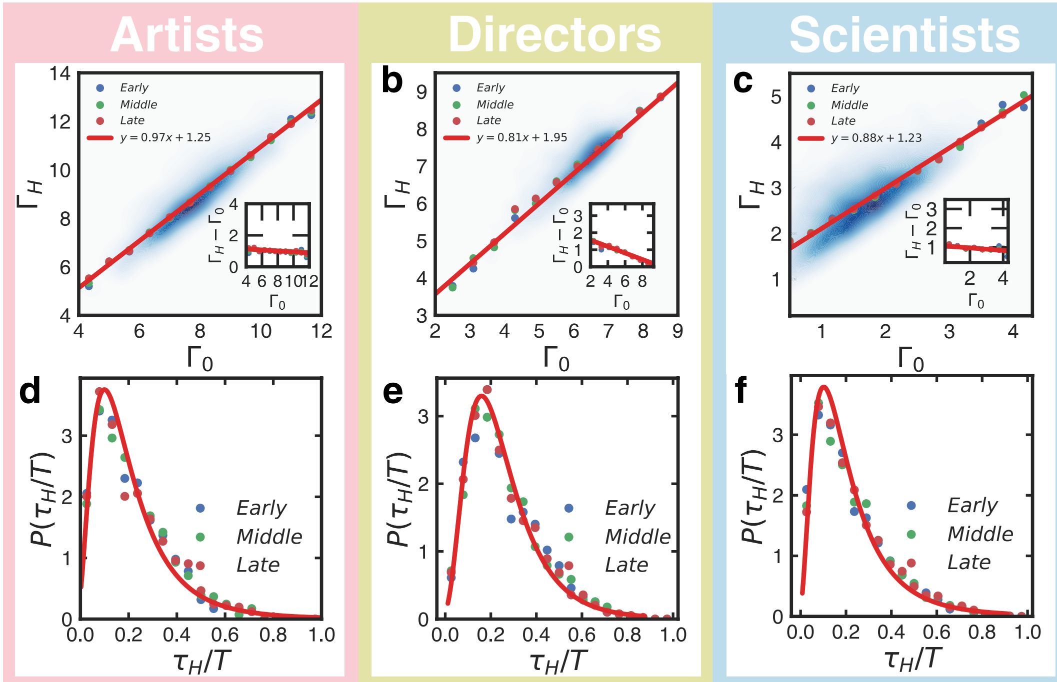

Extended Data Fig. 5 | Additional properties of hot streaks. **a–c**, Correlations between Γ_H and Γ_0 for artists (**a**; $n = 3{,}166$), directors (**b**; $n = 5{,}098$) and scientists (**c**; $n = 18{,}121$). The blue background denotes the kernel density of data, dots represent binning results of data, and the red lines depict a linear fit. Inset, the relationship between $\Delta\Gamma$ ($= \Gamma_H - \Gamma_0$) and Γ_0. **d–f**, The distribution of τ_H/T, representing the duration of hot streaks over total career lengths. The temporally localized nature of a hot streak is also captured by its proportion over career length τ_H/T.

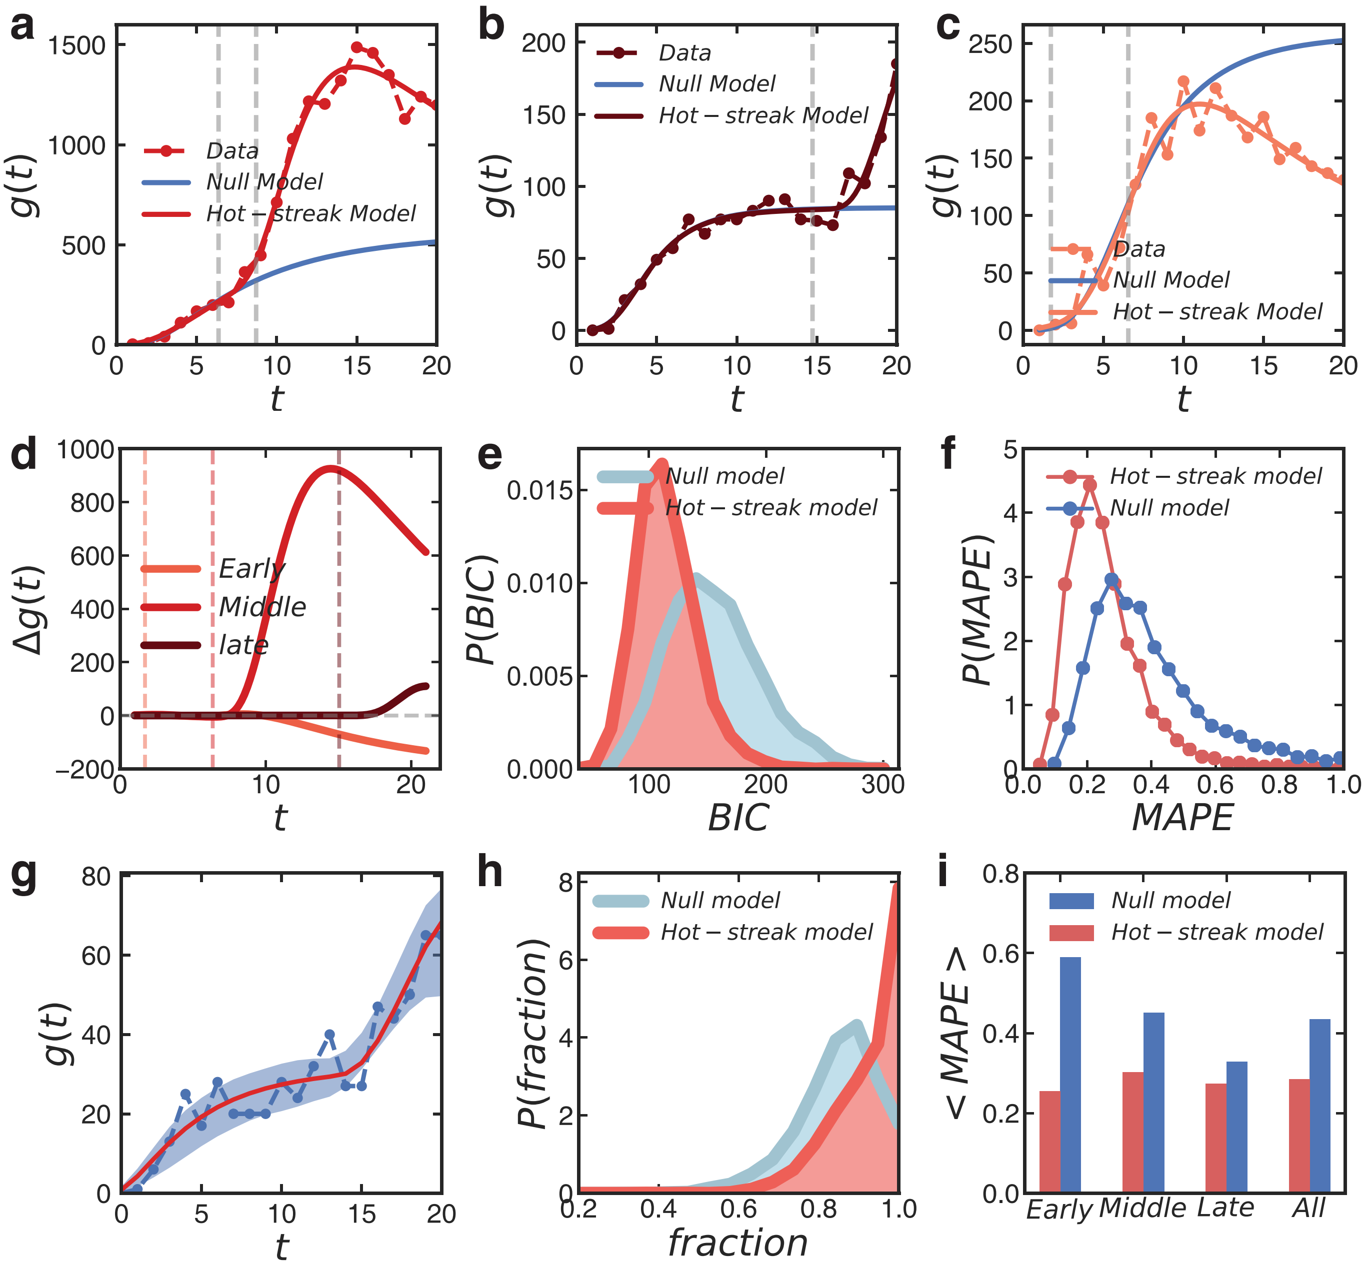

Extended Data Fig. 6 | Comparison of $g(t)$ between the null model and the hot-streak model. **a**–**c**, $g(t)$ of three scientists in our data set with mid-career (**a**), late-career (**b**) and early-career (**c**) onset of hot streaks. Red dots denote data, the blue line is the null model's prediction based on early performance, and the red line captures the predictions from the hot-streak model, with dashed grey lines denoting the start and end of hot streaks. **d**, The difference between our hot-streak model and the null model for each individual, $\Delta g(t)$. Dashed lines with corresponding colours denote the start of the hot streak. **d** illustrates the discrepancies in estimating an individual's future impact if we ignore the uncovered hot streaks. **e**, The distribution of the BIC measure, $P(\text{BIC})$, showing that the hot-streak model outperforms the null model in describing $g(t)$ after accounting for model complexity. **f**, The distribution of the MAPE measure, $P(\text{MAPE})$, showing that the hot-streak model outperforms the null model in describing $g(t)$. **g**, The uncertainty envelope of $g(t)$ for an individual in our data set. Blue dots denote data, and the red line is the fitting result of equation (4). Shaded area illustrates predicted uncertainty (one standard deviation). **h**, The fraction of $g(t)$ falling within the envelope for the null model (blue) and our hot-streak model (red). Fraction = 1.0 indicates that the entire $g(t)$ trajectory falls within the envelope. **i**, Average MAPE of our hot-streak model and the null model for individuals with early-career, mid-career and late-career onset of hot streaks. The difference is largest for individuals with early-onset hot streaks and smallest for those with late-onset ones.

Extended Data Fig. 7 | Testing alternative hot-streak dynamics. **a**–**f**, Illustrative examples of $\Gamma(N)$ for the hot-streak model (**a**), right trapezoid function (**b**), isosceles trapezoid function (**c**), quadratic function (**d**), tent function (**e**) and left trapezoid function (**f**). **g**, The distribution of the relative position $P(\tilde{N})$ of the three highest-impact works among the six highest-impact works within a career for artists, where $\tilde{N}$ denotes the relative order among the top six hits. **h**–**m**, $P(\tilde{N})$ predicted by corresponding models shown in **a**–**f**, respectively, according to artists' real productivity profiles. To test whether data agree with model predictions, we measured their statistical difference using the P value of the Kolmogorov–Smirnov test for discrete distributions. We colour the distributions green if we cannot reject the hypothesis that the data and the model predictions come from the same distributions, and red otherwise. Among the six models, the hot-streak model is the only model whose predictions are consistent with the data in terms of the relative ordering among the six highest-impact works observed in real careers. **n**, The proportion of real careers that are captured by the model with the smallest BIC among different hypotheses. The hot-streak model again stands out as the best model to describe real careers. We repeated the analyses for directors (**p**–**v**) and scientists (**x**–**ad**), the conclusions remained the same across all three domains.

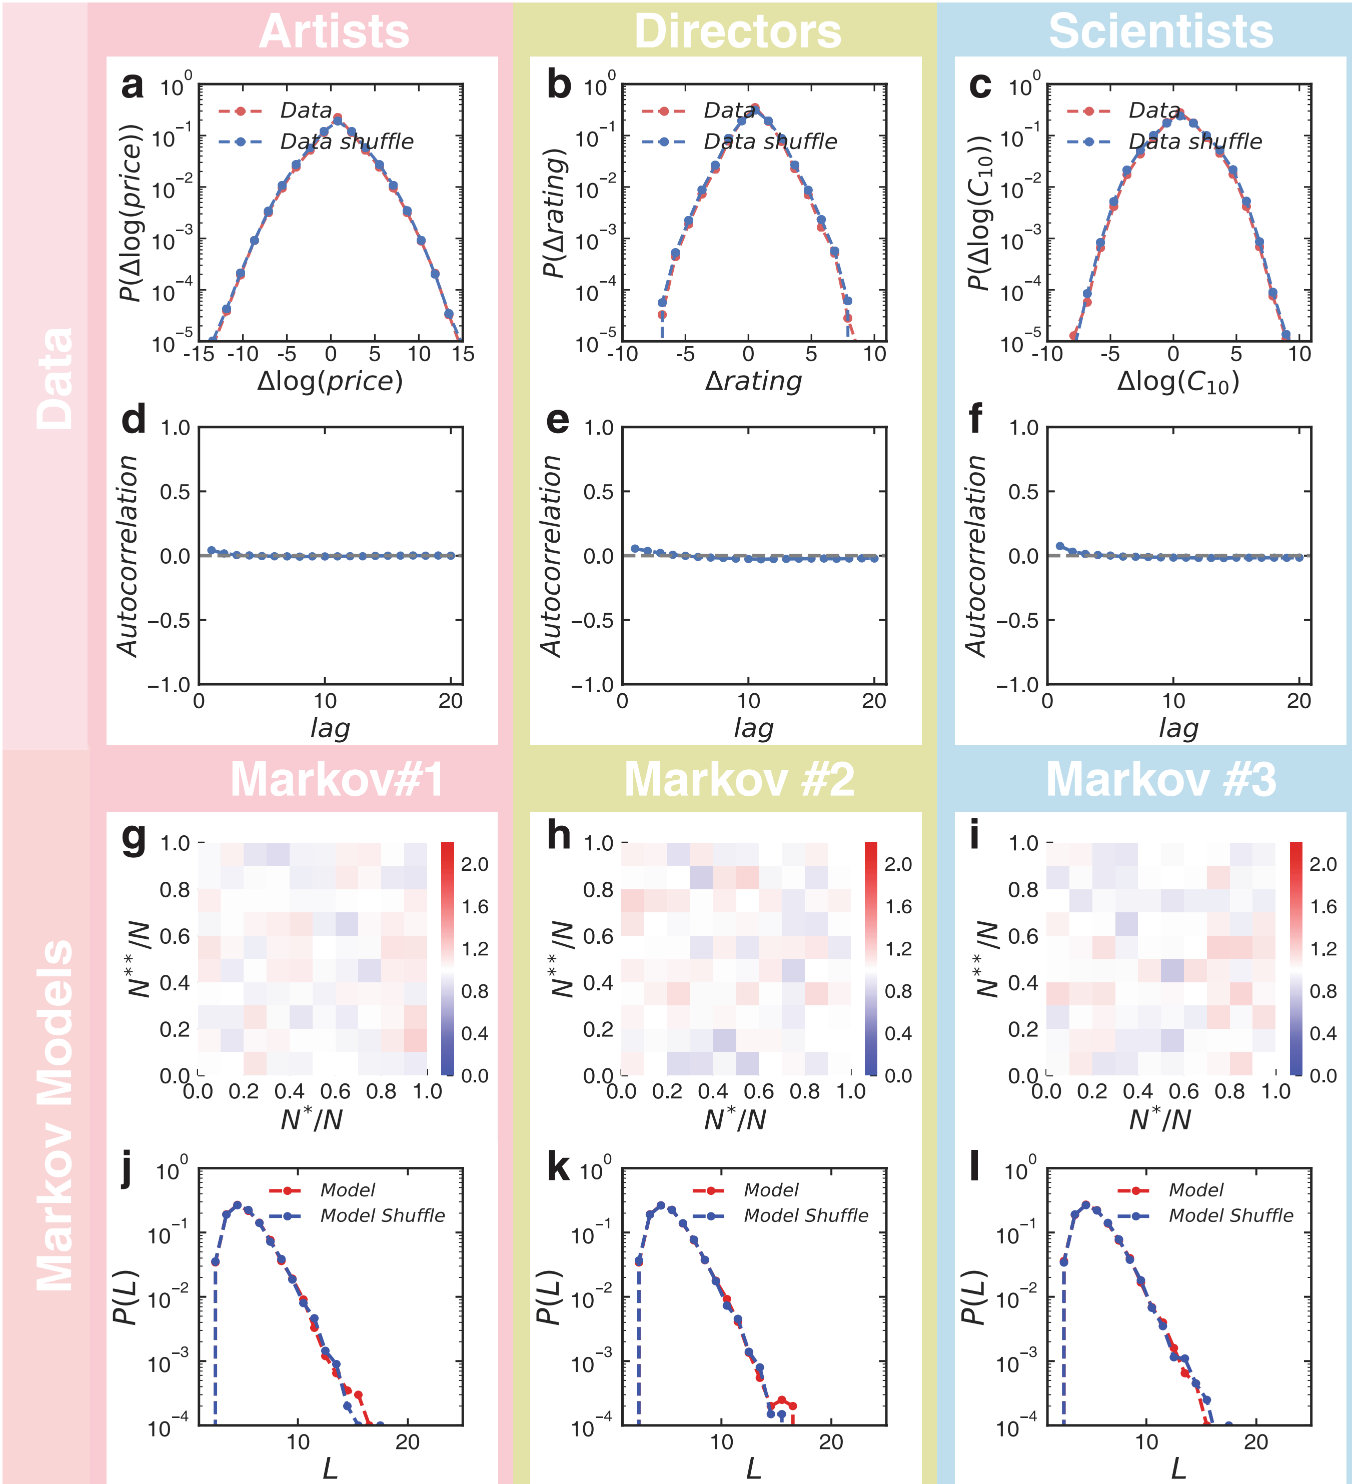

Extended Data Fig. 8 | Testing Markovian hypotheses. Here we test whether the observed patterns can be explained by Markovian dynamics that introduce correlations between neighbouring data points. We first test the assumptions of the Markovian hypothesis from the data (**a**–**f**). **a**–**c**, The distribution of N, $N+1$ differences between adjacent data points observed in real careers for artists (**a**, $n=3{,}480$), directors (**b**, $n=6{,}233$) and scientists (**c**, $n=20{,}040$). **d**–**f**, The autocorrelation measured in real careers for artists (**d**, $n=3{,}480$), directors (**e**, $n=6{,}233$), and scientists (**f**, $n=20{,}040$). **a**–**f** suggest that there is little short-range correlation in data across the three domains. We test three variants of Markovian models (**g**–**l**). The details of these models are outlined in Supplementary Information S6.2. **g**–**i**, $\phi(N^*, N^{**})$ of the top two highest-impact works within a career for three Markovian models using scientists' profiles as input. **j**–**l**, The distribution of the longest streak length $P(L)$ and $P(L_s)$ using median impact within a career as threshold for the three Markovian models. **g**–**l** demonstrate that the three Markovian models failed to capture the observed colocations among hits.

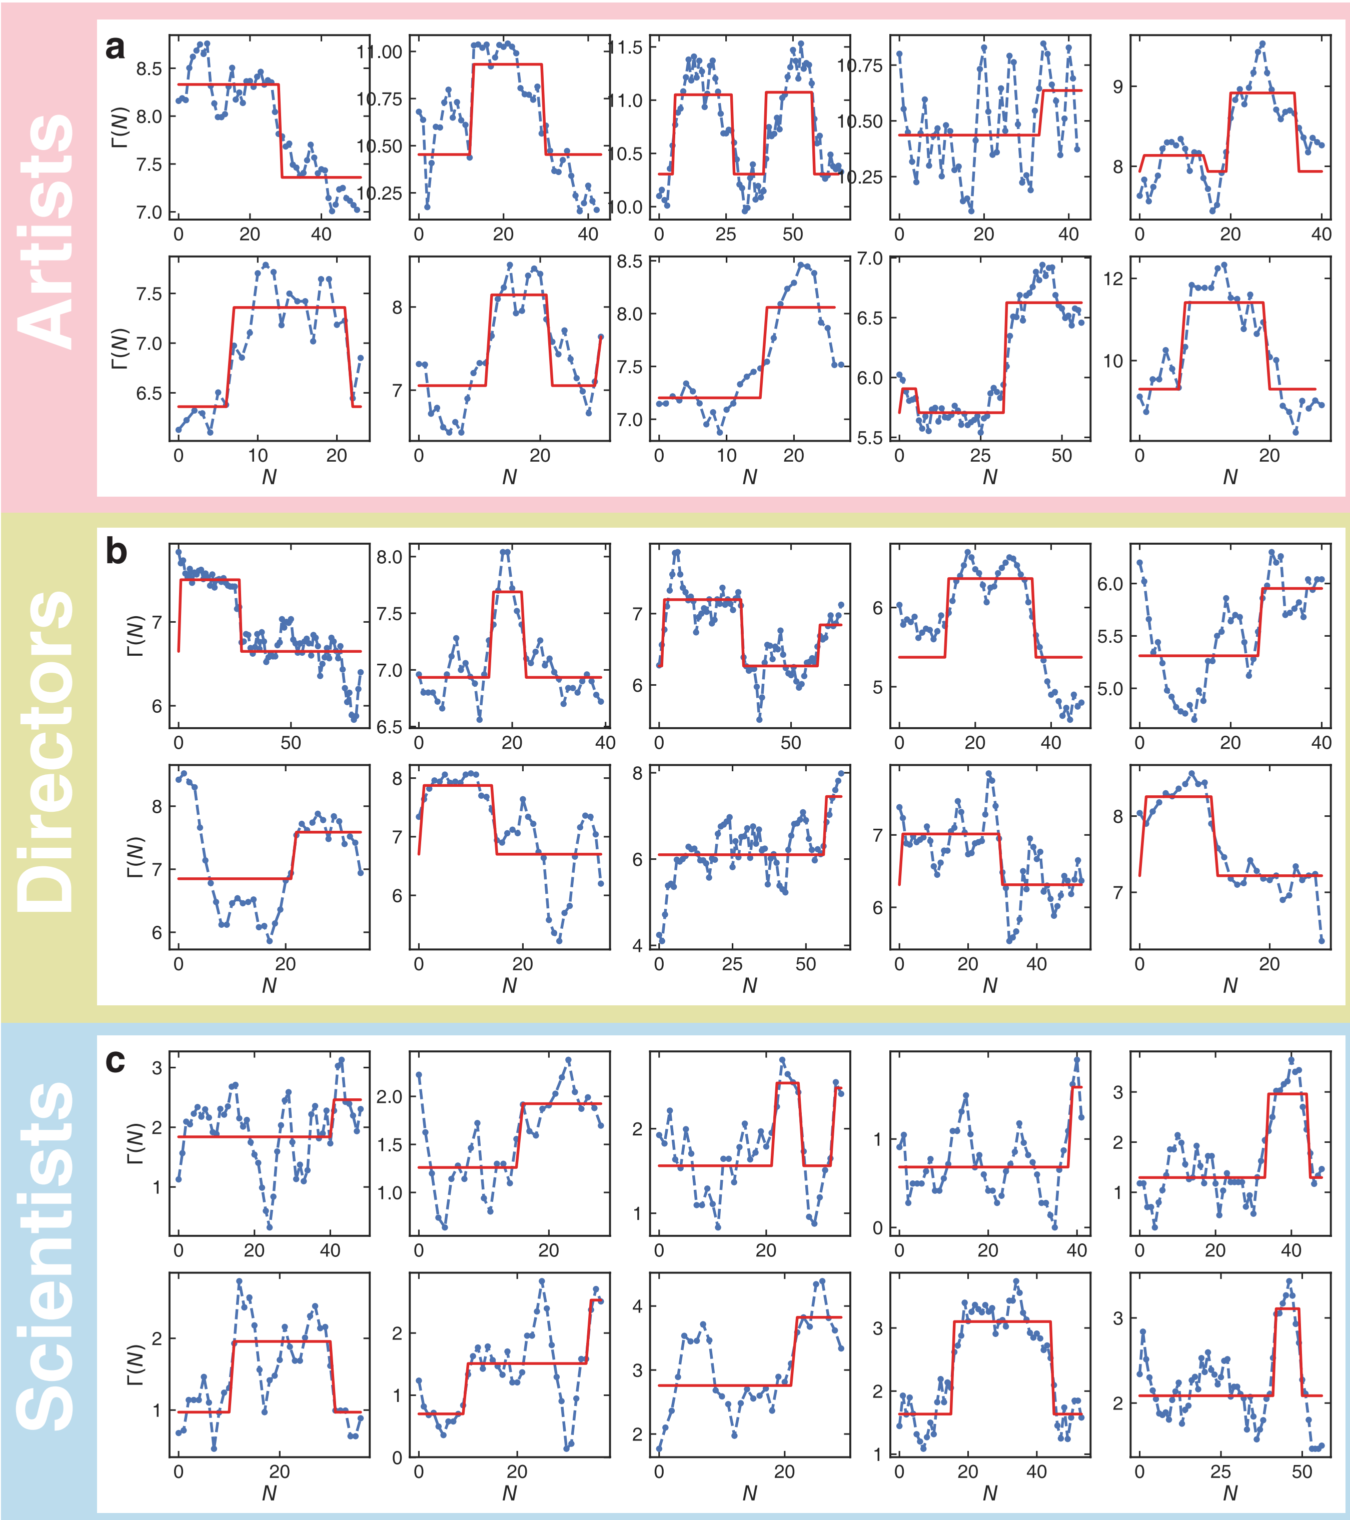

Extended Data Fig. 9 | Additional examples of Γ. a–c, Each subplot denotes the fitting result on Γ sequence for a randomly selected career for artists (**a**), directors (**b**) and scientists (**c**). Blue dots denote the moving average Γ(*N*) from data and red lines denote the best fitting result of the hot-streak model for each individual.

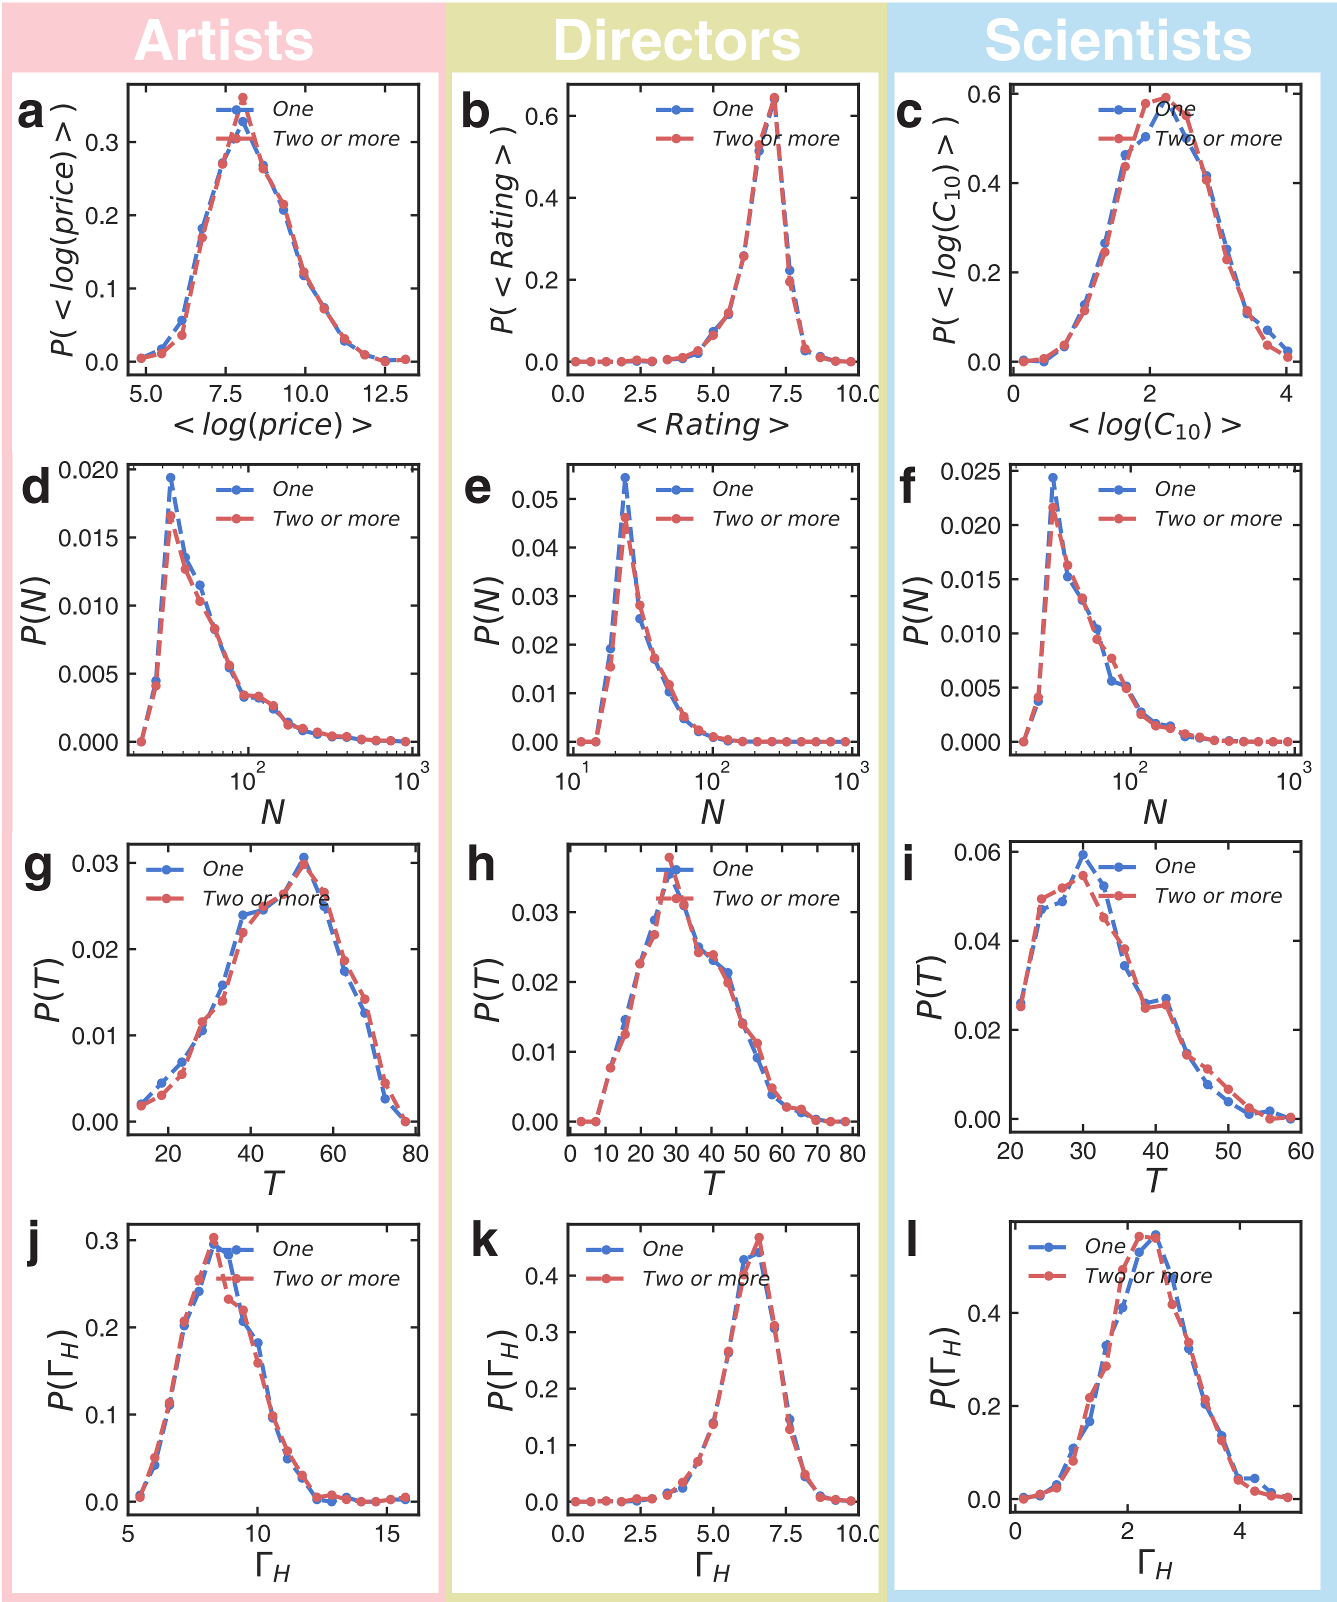

Extended Data Fig. 10 | Individuals with one or more hot streaks. **a**–**c**, The distribution of average impacts for individuals with one or more than one hot streaks for artists (**a**), directors (**b**) and scientists (**c**). Blue dots denote individuals with one hot streak, and red dots denote individuals with at least two hot streaks. **d**–**f**, The distribution of the number of works $P(N)$ within a career for individuals with one or more than one hot streak for for artists (**d**), directors (**e**) and scientists (**f**). **g**–**i**, The distribution of career length $P(\tau)$ for individuals with one or more than one hot streaks for artists (**g**), directors (**h**) and scientists (**i**). **j**–**l**, The distribution of $P(\Gamma_H)$ for individuals with one or more than one hot streaks for artists (**j**), directors (**k**) and scientists (**l**). Between those who have one or two hot streaks, there is no detectable difference in terms of typical performance metrics, including impact, productivity and career length, suggesting that the hot streak captures an orthogonal dimension to current metrics characterizing individual careers.

nature research

Corresponding author(s): Dashun Wang

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see Authors & Referees and the Editorial Policy Checklist.

Statistical parameters

When statistical analyses are reported, confirm that the following items are present in the relevant location (e.g. figure legend, table legend, main text, or Methods section).

n/a	Confirmed	
☐	☒	The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	An indication of whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistics including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. F, t, r) with confidence intervals, effect sizes, degrees of freedom and P value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☒	☐	Estimates of effect sizes (e.g. Cohen's d, Pearson's r), indicating how they were calculated
☐	☒	Clearly defined error bars *State explicitly what error bars represent (e.g. SD, SE, CI)*

Our web collection on statistics for biologists may be useful.

Software and code

Policy information about availability of computer code

Data collection	Profiles of artists were automatically downloaded from www.artprice.com and www.findartinfo.com. Profiles of movie directors were directly downloaded from www.imdb.com/interfaces. Profiles of scientists were automatically collected from Google Scholar. The Web of Science data was obtained through a purchase agreement with Thomson Reuters. Data collection was conducted using Python 2.7.
Data analysis	All data analyses were conducted using Python 2.7.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers upon request. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:
- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

Data are available at https://lu-liu.github.io/careerbursts/

Field-specific reporting

Please select the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/authors/policies/ReportingSummary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	This is a quantitative study of individual career histories based on existing datasets
Research sample	We compiled a comprehensive database consisting of three large-scale datasets of individual careers across three different domains: Dataset D1 contains profiles of artists obtained from online auction databases, Artprice (www.artprice.com) and Findartinfo (www.findartinfo.com). We analyzed 3,480 artists with at least 15 works and 10 years of career length, with careers dating back to as far as 1460. Dataset D2 contains profiles of movie directors recorded in the IMDB database (www.imdb.com/interfaces). We focused on 6,233 directors who have at least 15 movies and 10 years of career length dating back to as far as 1890. Dataset D3 contains the publication and citation records of 20,040 individual scientists, obtained by combining Google Scholar and Web of Science. We chose 20,040 scientists with at least 15 papers and 20 years of career length for our analysis. Data are available at http://personal.psu.edu/lpl5107/data/.
Sampling strategy	SI, Section S1.1-S1.3 - The sample size in this paper was chosen to have enough statistics for individual careers and to follow closely the same procedures used by existing studies in this area.
Data collection	Profiles of artists were automatically downloaded from www.artprice.com and www.findartinfo.com. Profiles of movie directors were directly downloaded from www.imdb.com/interfaces. Profiles of scientists were automatically collected from Google Scholar. The Web of Science data was collected from a local database. The experiments were not randomized. Researchers were not blinded to allocation during experiments and outcome assessment.
Timing	Data for artists and directors were collected in spring 2017. Google Scholar and Web of Science were collected in summer 2015.
Data exclusions	The analysis has no data exclusions. Selection criteria within datasets are described in SI, Section S1.1-S1.3.
Non-participation	There are no participants in this study.
Randomization	This is a data driven study, not a randomized experiment.

Reporting for specific materials, systems and methods

Materials & experimental systems

n/a	Involved in the study
☒	☐ Unique biological materials
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☒	☐ Animals and other organisms
☒	☐ Human research participants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging