EPJ.org

REGULAR ARTICLE

Open Access

Reciprocity and impact in academic careers

Weihua Li[1,2], Tomaso Aste[1,2], Fabio Caccioli[1,2,3] and Giacomo Livan[1,2*]

*Correspondence: g.livan@ucl.ac.uk
[1] Department of Computer Science, University College London, London, UK
[2] Systemic Risk Centre, London School of Economics and Political Science, London, UK
Full list of author information is available at the end of the article

Abstract

The growing importance of citation-based bibliometric indicators in shaping the prospects of academic careers incentivizes scientists to boost the numbers of citations they receive. Whereas the exploitation of self-citations has been extensively documented, the impact of reciprocated citations has not yet been studied. We study reciprocity in a citation network of authors, and compare it with the average reciprocity computed in an ensemble of null network models. We show that obtaining citations through reciprocity correlates negatively with an impactful career in the long term. Nevertheless, at the aggregate level we show evidence of a steady increase in reciprocity over the years, largely fuelled by the citations reciprocated between coauthors. Our results characterize the structure of author networks in a time of increasing emphasis on citation-based indicators, and we discuss their implications towards a fairer assessment of academic impact.

Keywords: Reciprocity; Academic career; Network science; Statistical validation

1 Introduction

The majority of measures that evaluate academic impact are based on citations of publications. Ranging from crude citation counts to the well known journal impact factor [1], the Eigenfactor metrics for journal rankings [2] and the *h*-index [3] for authors, such measures are increasingly relied upon to inform all aspects of academic decision-making, including faculty recruiting, grant attribution, and the formation of collaborations [4]. The onset of such a trend dates back to the 1950s, when the earliest citation-based indices to assess academic impact were put forward [5]. The following decades saw the proliferation and ever-increasing adoption of such indicators in all fields of science [6, 7], which eventually led to the systematic analysis of academic citations, to the emergence of *bibliometrics* as a research field [8], and, more recently, to the rise of a "science of science" devoted to understanding the determinants of scientific impact [9–11].

Given the importance that citations play nowadays in shaping the prospects of an academic career, it is certainly not surprising to see an increased attention to the study of citation patterns and of the publication strategies that can attract a larger number of citations. While an author's productivity and the quality of her work are obvious determinants of academic impact [12], other less tangible aspects have recently been identified as key contributors to success. For example, the social network a scientist is embedded in has recently been shown to play a relevant role in determining her future chances of success [13].

Springer

Other social and behavioural considerations also play relevant roles in modern academic reputation systems, especially concerning the proliferation of self-citations and citations between close collaborators [14–17]. Self-citations and their role in inflating bibliometric indicators have been studied extensively [18]. Several scholars have proposed revised metrics to mitigate their distorting impact [19], and indeed some of the most popular citation indexing services (e.g., Web of Science) and social networking sites for scientists (e.g., ResearchGate) provide detailed author-specific bibliometric information with and without the contribution of self-citations.

Yet, very little attention has been paid to the scientific community's collective response to the increasing adoption of bibliometric indicators, and whether this ultimately resulted in more sophisticated citation patterns involving collaborators and colleagues [17, 20–23]. Indeed, the quest for higher citation counts generates an obvious incentive for scientists to seek citations from and give citations to their closest circle of coauthors and collaborators. Such behaviour, known as reciprocity, has long been observed and investigated in many social and economic systems [24–26]. It reflects the tendency to return helpful acts, and it is crucial for forming and maintaining cooperative relationships among individuals and groups [27].

Following the broad stream of literature that has analyzed citation patterns from a network perspective (see, e.g., [28]), we investigate the inherent reciprocity dynamics in an author citation network constructed from the citation history of 463,348 papers in the *Physical Review* (PR) corpus of journals published by the American Physical Society (APS) between 1893 and 2010 (see Methods), spanning all research fields of physical sciences, with disambiguated author names obtained from [12].

We address two main research questions: (1) how much reciprocated citations contribute to a scientist's academic reputation, and whether these can be used to predict impact and classify different career trajectories; (2) how prevalent is the tendency to reciprocate citations in the scientific community at large, and how it has evolved over time.

We tackle such questions by representing authors as nodes in a directed weighted network $\mathcal{C}$, where the weight c_{ij} denotes the number of times author i has cited author j, and use such a representation to measure the number of reciprocated citations as $c_{ij}^{\leftrightarrow} = \min(c_{ij}, c_{ji})$. We assess the statistical significance of the empirically observed patterns of reciprocated citations with the values observed under an ensemble of null network models obtained through the controlled randomization of the original networks' topology (see Methods).

Following the definition in [29], we define the *excess* reciprocity[a] of an author i as

$$\rho^{(i)} = \frac{\rho_0^{(i)} - \langle \rho_{\text{null}}^{(i)} \rangle}{1 - \langle \rho_{\text{null}}^{(i)} \rangle}, \tag{1}$$

where $\rho_0^{(i)} = \sum_{j \in \mathcal{C}} c_{ij}^{\leftrightarrow} / \sum_{j \in \mathcal{C}} c_{ji}$ is the fraction of citations received by author i in the empirical author citation network which she reciprocated (i.e., the fraction of reciprocated incoming weight), while $\langle \rho_{\text{null}}^{(i)} \rangle$ is the average of the same quantity as computed in the null model ensemble. As we shall explain more extensively later on, the rationale of the above definition lies in its ability to quantify an *author-specific* propensity to reciprocate with respect to an expected benchmark quantified by the null model. For example, an excess reciprocity $\rho^{(i)} = 0.1$ for an author with 100 received citations would correspond to 55 of

Figure 1 Empirical frequency of the number of directed (blue dots) and reciprocated (red dots) citations between pairs of authors active in the APS dataset between 1950 and 2010
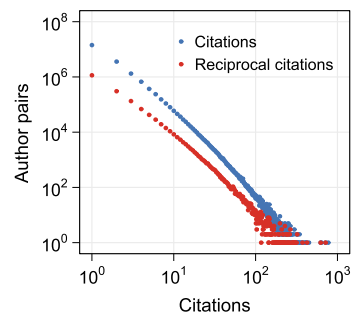
them being reciprocated under an expected reciprocity $\langle \rho_{\text{null}}^{(i)} \rangle = 1/2$, and to 28 reciprocated citations under an expected reciprocity $\langle \rho_{\text{null}}^{(i)} \rangle = 1/5$.

Let us stress two points from the outset. First, positive values of excess reciprocity are not to be interpreted as indicators of strategic behavior aimed at deliberately manipulating bibliometric indicators. The definition in Eq. (1) simply quantifies an excess with respect to a benchmark. This, in turn, is identified by a null model, i.e., ultimately, by a null hypothesis encoding patterns of citations that can be expected within the structural constraints imposed by the activity of nodes in the network. In this respect, positive excess reciprocity is to be interpreted as the result of multiple factors contributing towards an over-expression of reciprocated citations with respect to such patterns. It is therefore important to point out that the above definition cannot identify such factors or discriminate between them. Because of this, in the following we will refrain from making judgments about the exact nature of the excess reciprocity we will detect in empirical data.
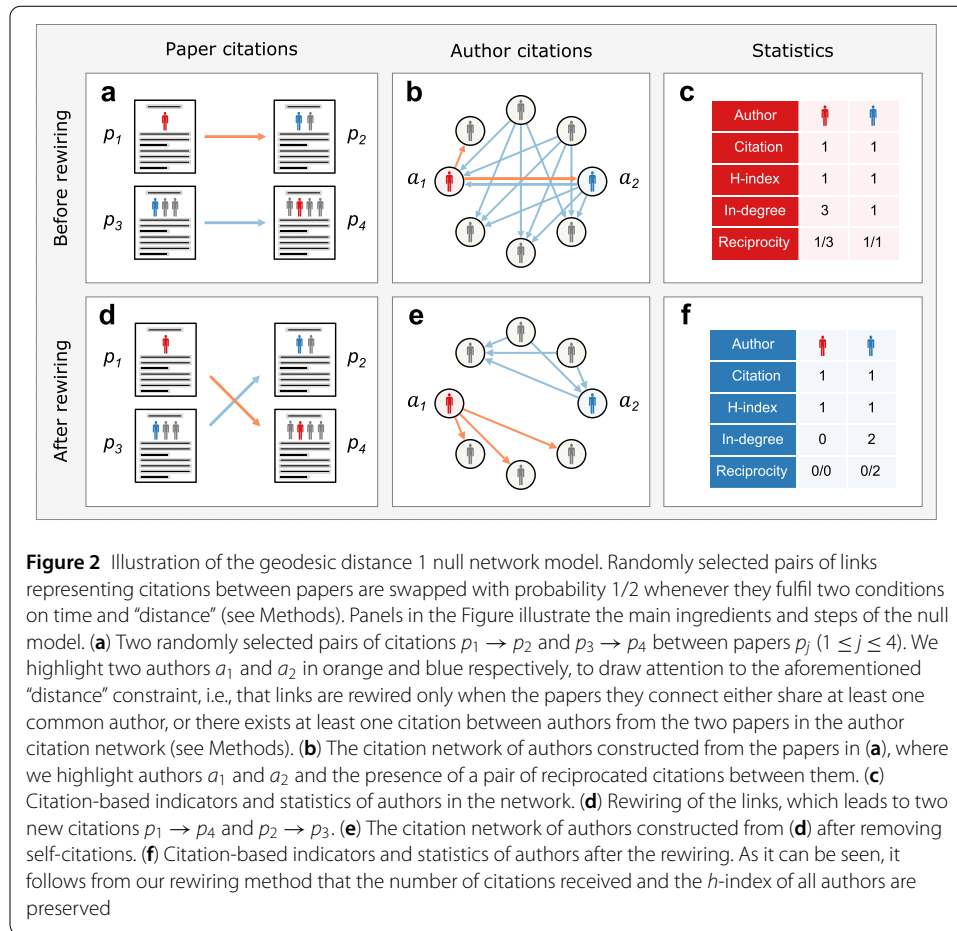
Second, let us remark here that it is not our intention to equate academic impact, which is a multifaceted concept, to mere citation counts. However, we shall take the (widely shared) position that a high citation count is an unquestionable indicator of impact for a scientist's production.

2 Results

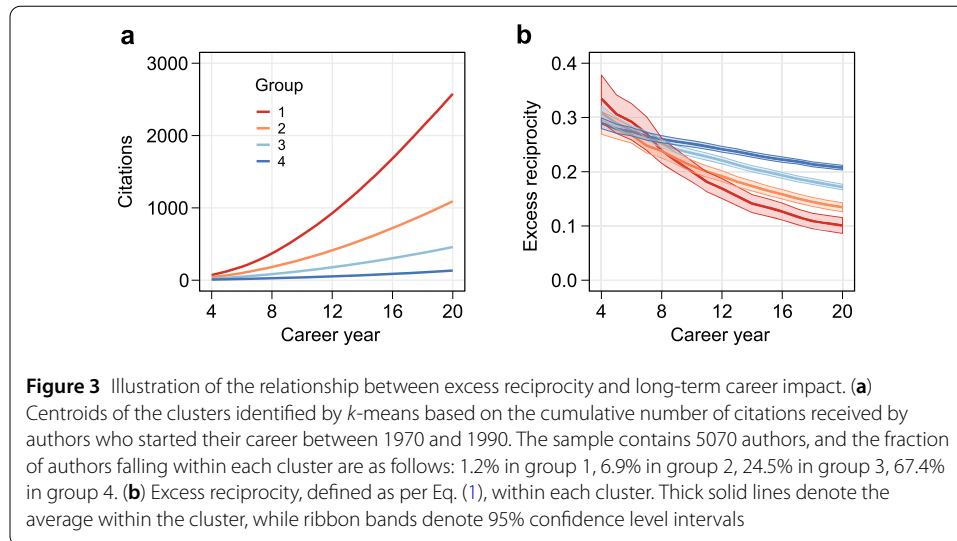### 2.1 Reciprocity and career impact

Figure 1 shows the frequency of both the number of directed citations (i.e., the number of times c_{ij} an author i has cited another author j) and the number of reciprocated citations (i.e., $c_{ij}^{\leftrightarrow} = \min(c_{ij}, c_{ji})$) across all pairs of authors active in the APS dataset from 1950 to 2010. As it can be seen, both distributions show a markedly heavy-tailed behaviour. In particular, there are more than fifty thousand pairs of authors with 10 reciprocated citations, and more than two thousand pairs with 50 or more reciprocated citations. Overall, more than 21% of the citations in the dataset are reciprocated (roughly 8.5 millions out of 40.4 millions).

The sheer scale of the aforementioned phenomenon suggests that reciprocated citations might play a key role in shaping an author's scientific impact. This, in turn, poses the question whether authors with a history of systematic reciprocity might, on average, outperform their peers. In this section we investigate citation patterns at the level of individual authors in order to answer such a question. We do so by proposing a null network model ensemble (see Fig. 2 and Methods) to estimate the average baseline level of reciprocity $\langle \rho_{\text{null}}^{(i)} \rangle$ one should expect in the author citation networks under partially random interactions, which we then use to measure excess reciprocity.

Figure 2 Illustration of the geodesic distance 1 null network model. Randomly selected pairs of links representing citations between papers are swapped with probability 1/2 whenever they fulfil two conditions on time and "distance" (see Methods). Panels in the Figure illustrate the main ingredients and steps of the null model. (**a**) Two randomly selected pairs of citations $p_1 \rightarrow p_2$ and $p_3 \rightarrow p_4$ between papers p_j ($1 \leq j \leq 4$). We highlight two authors a_1 and a_2 in orange and blue respectively, to draw attention to the aforementioned "distance" constraint, i.e., that links are rewired only when the papers they connect either share at least one common author, or there exists at least one citation between authors from the two papers in the author citation network (see Methods). (**b**) The citation network of authors constructed from the papers in (**a**), where we highlight authors a_1 and a_2 and the presence of a pair of reciprocated citations between them. (**c**) Citation-based indicators and statistics of authors in the network. (**d**) Rewiring of the links, which leads to two new citations $p_1 \rightarrow p_4$ and $p_2 \rightarrow p_3$. (**e**) The citation network of authors constructed from (**d**) after removing self-citations. (**f**) Citation-based indicators and statistics of authors after the rewiring. As it can be seen, it follows from our rewiring method that the number of citations received and the h-index of all authors are preserved

The rationale of the definition in (1) is to discount density-related effects. Indeed, simply measuring reciprocity as the fraction of reciprocated weight $\rho_0^{(i)}$ typically leads to seemingly high (low) values in dense (sparse) networks. The measure in Eq. (1) takes care of such potential spurious effects by discounting the average reciprocity observed in a null model ensemble, so that positive (negative) values of $\rho^{(i)}$ indicate authors whose citations have been received through an over-representation (under-representation) of reciprocated relationships, whereas values $\rho^{(i)} \simeq 0$ indicate levels of reciprocity compatible with the null assumption encoded in the null network model being used. In conclusion, excess reciprocity indirectly quantifies how much the academic impact of an author, as measured by her number of citations and h-index (which are both preserved by the null model, see Methods), relies on the citations from authors she cited as well.

We investigated the relationship between excess reciprocity and impact by following the career paths of authors with a traceable publication history of at least 20 years in the APS dataset. We first employed a variety of methodologies to predict a scientist's future impact (in terms of citations) based on her previous history of excess reciprocity. In all cases we found very weak to no evidence of any predictive power (see Additional file 1, Supplementary Note 1), which strongly suggests that citation strategies based on the mere exchange of citations do not contribute to attracting higher numbers of citations in the future.

Figure 3 Illustration of the relationship between excess reciprocity and long-term career impact. (**a**) Centroids of the clusters identified by *k*-means based on the cumulative number of citations received by authors who started their career between 1970 and 1990. The sample contains 5070 authors, and the fraction of authors falling within each cluster are as follows: 1.2% in group 1, 6.9% in group 2, 24.5% in group 3, 67.4% in group 4. (**b**) Excess reciprocity, defined as per Eq. (1), within each cluster. Thick solid lines denote the average within the cluster, while ribbon bands denote 95% confidence level intervals

We then applied the *k*-means clustering algorithm [30] to categorize authors in terms of career impact. Following [12], we performed this analysis considering the career trajectories of all authors with 10 or more papers published over the course of at least 20 years (with at least one paper published every 5 years) who published their first paper either in 1950–1970 or in 1970–1990. We chose to group authors whose careers started over two decades in order to assemble two large enough samples. The downside of this is that we pool together authors whose careers started and developed during rather different historical periods in terms of scientific publication standards and practices. However, we verified that the results presented in the following do not change qualitatively when pooling authors based on shorter time spans.

We used these two samples to perform a *k*-means clustering analysis based on the cumulative number of citations received over time. Since several authors did not receive citations early in their career, we performed our analysis starting from the 4th career year. In Fig. 3 we present the results for 1970–1990 (see Additional file 1, Supplementary Note 2 for the results obtained for 1950–1970), which were obtained on a pool of 5070 scientists. We identified 4 distinct groups with very different levels of career impact, ranging from a small minority of authors (1.2% of the sample) who managed to attract several thousand citations over the time period considered, to the relative majority of authors (67.4% of the sample) who only enjoyed moderate to low impact (see Fig. 3 caption for more details).

We find the above groups to be associated with markedly different behaviours. Namely, we find long-term career impact to be associated with progressively lower levels of excess reciprocity. Indeed, the two most impactful groups are associated with the lowest long-run excess reciprocity, with the small cluster of elite scientists (group 1 in Fig. 3) displaying an average excess reciprocity around 0.1 towards year 20 of their career. Conversely, the two least impactful groups are associated with consistently higher levels of long-term excess reciprocity, higher than 0.2 in the case of the single least impactful group. We further corroborated the progressive development of differences between the four groups by running two-sided Kolmogorov-Smirnov tests between the distributions of excess reciprocity in each group at career years 4 and 20. The results are reported in Additional file 1, Supplementary Note 3, and show that at year 4 the null hypothesis of excess reciprocities

being drawn from the same distributions can be rejected only when comparing the least impactful group of authors (group 4) with the other ones. In contrast, at year 20 the null hypothesis can be rejected for all pairs of groups.
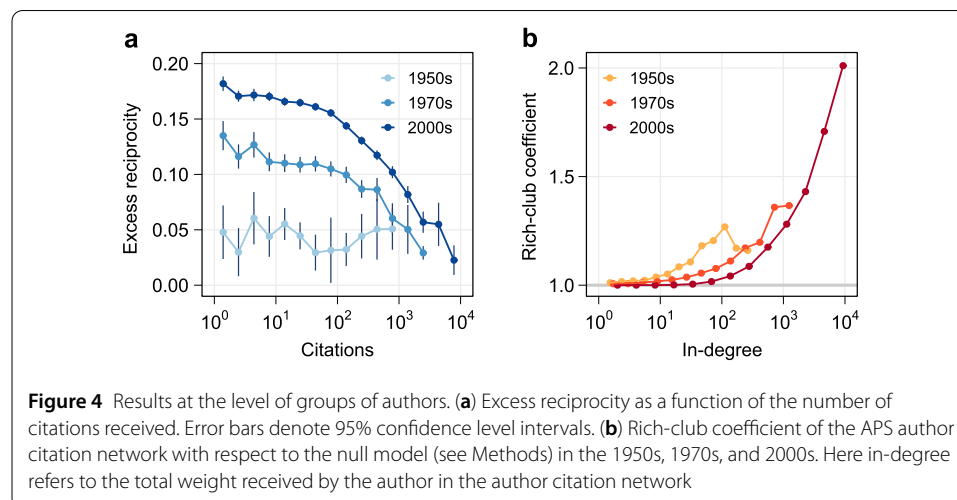
In addition, we checked whether authors belonging to a certain group tend to publish more frequently in some APS journals rather than others. The results of this analysis are presented in Additional file 1, Supplementary Note 3, and show that authors in the most impactful groups (group 1 and 2) have a higher publication rate in *Physical Review Letters* (PRL), which is somewhat unsurprising since PRL is by far the most impactful venue among those considered here. Yet, a more nuanced picture emerges when looking at the remaining journals, as the most impactful clusters do not necessarily account for the relative majority of publications in the most impactful journals and vice versa. Moreover, while the ranking and behaviours in terms of excess reciprocity are similar across the two time periods we consider, it is interesting to notice that publication rates of the different groups across journals are rather different (see Additional file 1, Supplementary Note 2).

2.2 Shifts in citation patterns

In the previous section we analyzed the relationship between excess reciprocity and long-term impact from a cross-sectional point of view by "collapsing" together the career trajectories of several authors whose actual careers developed asynchronously over the span of a few decades. We now seek to further unpack this relationship by investigating temporal snapshots of the APS citation network, testing how an author's propensity to reciprocate citations correlated with her impact during different historical periods.

We do so by performing analyses at the decade level. For each decade from the 1950s to the 1990s, we consider all authors whose career started before the end of such decade and did not end before the first year of that decade. We then pool all the papers published by such authors before the end of the decade, and their corresponding citations, to build the author citation network for the decade of interest.

Figure 4(a) shows, for three different decades, the average excess reciprocity of authors as a function of their accrued citations (see Additional file 1, Supplementary Note 7 for all six decades). As it can be seen, over time we observe the emergence of a clear negative correlation between an author's impact and her excess reciprocity. In the 1950s the entire

Figure 4 Results at the level of groups of authors. (**a**) Excess reciprocity as a function of the number of citations received. Error bars denote 95% confidence level intervals. (**b**) Rich-club coefficient of the APS author citation network with respect to the null model (see Methods) in the 1950s, 1970s, and 2000s. Here in-degree refers to the total weight received by the author in the author citation network

APS scientific community was essentially compatible with the null model, with average excess reciprocity lower than 0.05 for all groups of authors. This changes considerably from the 1970s onwards, and it becomes quite pronounced in the 2000s, with a very apparent negative relationship between an author's impact and her tendency to reciprocate.

One might intuitively expect high impact authors to display, as a group, the lowest tendency to reciprocate. Indeed, in network terms, highly impactful academics simply do not have enough outgoing links to reciprocate their incoming links, i.e., they cannot provide enough citations to match the high number of citations they receive. While this is certainly true, as shown consistently for all decades in Fig. 4(a), there are subtler aspects to this observation.

First, let us recall that the definition in Eq. (1) measures the excess of reciprocity with respect to an *expected baseline*, which in our case is computed from a null model which preserves the heterogeneity (in terms of number of publications) and level of impact (both in terms of accrued citations and h-index) of each author. In this respect, the above result shows that high impact authors simply do not reciprocate much more than one could reasonably expect. Yet, a deeper analysis of the citations received by high impact authors reveals more substantial differences with respect to our null model. Indeed, while our null model naturally incorporates the low levels of excess reciprocity of high impact authors, it does not prescribe *who* the recipients of citations from them should be.

To investigate who the recipients are, we examine the level of interconnectedness among the leading authors with the highest citation counts in each decade by measuring the rich-club coefficient [31, 32] in the author citation networks (see Methods). The rich-club coefficient quantifies the tendency to preferentially establish relationships within a group with respect to the expected tendency based on a null hypothesis. In the present case, we measure the rich club coefficient as $\phi(c) = \phi_0(c)/\langle\phi_{\mathrm{null}}(c)\rangle$, where $\phi_0(c)$ is the fraction of the total number of citations flowing between authors that have received at least c citations in the empirical network (i.e., authors with an incoming weight equal to or larger than c), and $\langle\phi_{\mathrm{null}}(c)\rangle$ is the corresponding quantity computed over our null network model ensemble.

We observe an increasingly pronounced rich-club effect among leading academics, with the effect being up to twice as strong with respect to the null model for authors with an incoming weight around 10^4 in the 2000s (Fig. 4(b)). Conversely, in earlier decades we find the effect to be much less strong, although still present (see also Additional file 1, Supplementary Note 7). This result indicates that, although the overall tendency of high impact authors to reciprocate is close to the one predicted by our null model, they overwhelmingly tend to cite their peers. The presence of such an interconnected rich core of successful scientists suggests that homophily with respect to career excellence has increasingly become one of the driving forces behind the attribution of citations.

2.3 Reciprocity, coauthorship and self-citations

We now shift our attention to the evolution of reciprocity at the aggregate level of the entire APS community. We straightforwardly generalize Eq. (1) to define a measure of network-wide excess reciprocity as $\rho = (\rho_0 - \langle\rho_{\mathrm{null}}\rangle)/(1 - \langle\rho_{\mathrm{null}}\rangle)$, where ρ_0 denotes the overall fraction of reciprocated weight in the empirical networks, whereas $\langle\rho_{\mathrm{null}}\rangle$ denotes the corresponding average quantity in the null network ensemble. We track such quantity over time by considering annual networks constructed by including all papers published by active authors up to the year under analysis. We consider an author to be active whenever the year under consideration is between the first and last of her career.

Figure 5 Network-wide reciprocity in the empirical networks and in the null model. (**a**) Annual excess reciprocity values from 1950 to 2009. (**b**) Fraction of reciprocity ρ_0 contributed by past coauthors in the empirical networks (dark blue) and in the null model (light blue). (**c**) Fraction of reciprocated citations between two authors versus the overlap in their research interests, proxied by the Jaccard index between the list of references cited by the two authors, with 95% confidence level intervals. The grey dots represent 10^4 author pairs sampled from real data. When controlling for the overlap, on average coauthors are found to reciprocate more than non-coauthors. (**d**) Fraction of self-citations in the empirical networks (dark red) and in the null model (orange). The standard error bars for the null model results are small and not graphically visible in (**b**) and (**d**)

During the entire period of study we systematically observe positive values of network-wide excess reciprocity, indicating a stronger propensity of the APS community to reciprocate citations than the one expected in our null model ensemble. Furthermore, we find reciprocity to increase steadily (and roughly linearly) up to the early 1990s, after which it stabilizes around 0.15 (Fig. 5(a) and Additional file 1, Supplementary Note 8).

A closer look reveals that, over the entire period of observation, a substantial proportion of the overall reciprocity ρ_0 is accounted for by citations between coauthors. Such proportion grows from about 40% in the 1950s to about 50% in the 1990s. This is in contrast with the expected proportion computed in the null model, which instead shows a steady decline over time (Fig. 5(b)).

In order to better understand the impact of citations from coauthors on a scientist's career, we pool all authors over the entire period of observation and compare the tendency to reciprocate between coauthors and non-coauthors. Namely, we define the reciprocity between a *pair* of authors *i* and *j* as the number of reciprocated citations between them, divided by the total number of citations received by both authors. In Fig. 5(c) we show such quantity as a function of the distance between research interests, quantified by the Jaccard similarity index between the sets of papers cited by a pair of authors over their career [33]. Higher Jaccard indices indicate higher proportions of common references used by both authors, which we interpret as a proxy for a substantial overlap of research inter-
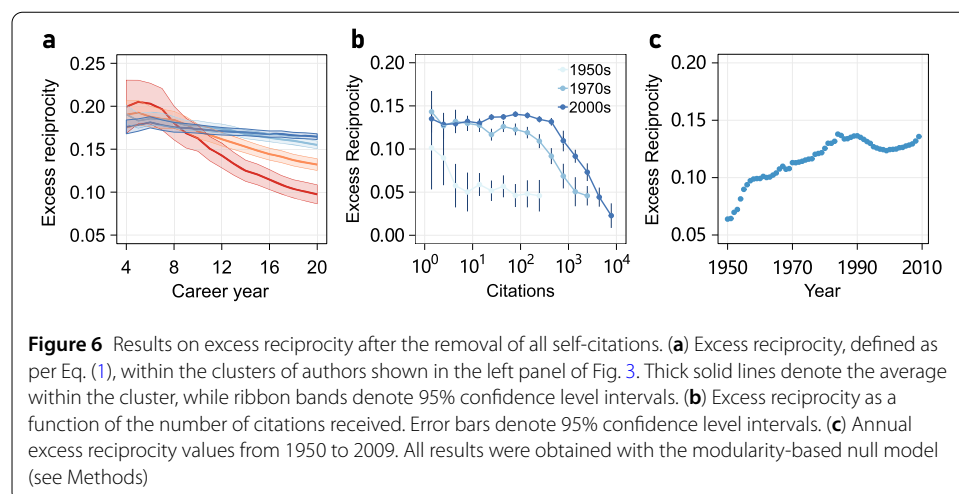
ests. As one would intuitively expect, we observe an overall positive correlation between research interests and the tendency to reciprocate citations. However, on average we find this relationship to be stronger in the case of coauthors, regardless of the specific level of proximity between research interests.

Let us conclude this section with a short digression devoted to the investigation of self-citations through the lens of our null model. Figure 5(d) compares the observed fraction of self-citations with the corresponding expected proportion in our null model ensemble over time. As it can be seen, the empirical rate of self-citation has remained around a fairly stable level of around 20% (which decreases to roughly 18% when including authors no longer active in the time frame under consideration). Yet, the null model predicts a sharp downward trend, which, as in the case of reciprocity between coauthors, highlights a growing gap between empirical citation patterns and those expected under our null hypothesis. Interestingly, the aforementioned rate of self-citation is much larger than those observed in citation datasets from other disciplines (e.g., Law, Political Science, Mathematics), which in most cases are between 5% and 12% [34].

2.4 Robustness checks

We tested the robustness of our main findings in a number of ways, in order to rule out spurious effects due to possible confounding factors. First, following [35] we modified our null model in order to account for modularity-related effects in the author network, i.e., that authors belonging to the same scientific sub-communities can be naturally expected to cite each other at an above average rate. To this end, we used two popular community detection algorithms (InfoMap [36] and the modularity-based algorithm published in [37]) in order to extract the community structure of the author network at different granularity levels (see Additional file 1, Supplementary Note 4), and constrained our null model to partially preserve it (see Methods).

Second, we proceeded to discount self-citation as a potential confounding factor in our analyses. Indeed, pairs of past coauthors both citing their own work naturally give rise to reciprocated citations. We therefore ran our analyses after removing all self-citations in the paper network (which, as shown in the bottom right panel of Fig. 5, roughly amount to 20% of all citations in the dataset). The main results are shown in Fig. 6, and, as it can be seen, are very much in line with those obtained from the full dataset. Indeed, we still

Figure 6 Results on excess reciprocity after the removal of all self-citations. (**a**) Excess reciprocity, defined as per Eq. (1), within the clusters of authors shown in the left panel of Fig. 3. Thick solid lines denote the average within the cluster, while ribbon bands denote 95% confidence level intervals. (**b**) Excess reciprocity as a function of the number of citations received. Error bars denote 95% confidence level intervals. (**c**) Annual excess reciprocity values from 1950 to 2009. All results were obtained with the modularity-based null model (see Methods)

observe a clear negative relationship between excess reciprocity and impact (left panel), whose strength becomes clearer in more recent decades (middle panel), in analogy with the results reported in the right panel of Fig. 3 and the left panel of Fig. 4. Also, in the right panel we see a pattern in the evolution of network-wide excess reciprocity similar to the one observed in the full dataset (top-left panel of Fig. 5), albeit with systematically lower annual values. Based on these results, we can safely conclude that self-citations do produce some extra excess reciprocity, but do not represent its main driver.

Third, we controlled for the effect of productivity in a number of ways. We did so by repeating our analysis after clustering the authors based on the number of citations received per paper (rather than on absolute citation counts), and after removing authors with low productivity and impact (i.e., those with a total of less than 10 citations over their first 20 career years, see Methods). In the same spirit, we also controlled for the presence of large collaborations by repeating our analyses after further restricting our sample to papers published by three or less authors.

Fourth, we restricted our analyses to US-based authors only in order control for potential geographical biases, and to assess the robustness of our results with respect to the name disambiguation procedure we used to identify authors (see Additional file 1, Supplementary Note 4).

In all the above cases we still detected the same negative relationship between excess reciprocity and impact based on the clusters of authors identified with k-means. As a further robustness check, we tested such relationship when separating authors based on different clustering criteria, i.e., we also grouped authors based on quartiles and with the Expectation-Maximization clustering algorithm [38]. In both cases, we still detected the same negative relationship (see Additional file 1, Supplementary Note 4).

Lastly, we tested such relationship from the opposite perspective, i.e., by grouping authors based on excess reciprocity and then measuring the impact of different groups. We resorted to matched pair analysis, and divided the authors whose careers started in 1970-1990 into "high reciprocators" (treatment) and "low reciprocators" (control) groups based on their excess reciprocity pattern over the first 10 career years, and performed a t-test on the average number of citations attracted by authors in the two groups over the following 10 career years after pairing them based on productivity (i.e., on the number of papers published in the first 10 years). Consistently with our results based on clustering, we found the treatment group to attract substantially less citations per author (272.2) than the control group (331.6), with $p < 0.001$ (see Additional file 1, Supplementary Note 6).

3 Discussion

This study addressed two main research questions, i.e., whether the constantly increasing attention to citation counts and bibliometric indicators has led to an incentive for academics to boost such metrics through the exchange of citations, and, if so, whether this behaviour is rewarding in terms of career impact. Let us stress from the outset that our results, being based on a comparison between empirical data and null network models, cannot provide direct evidence of strategic behaviour explicty aimed at manipulating bibliometric indicators. Yet, they inform us on how the academic community has collectively organized in response to the increased emphasis on such metrics, revealing a nuanced picture which we discuss in the following.

Starting from the 1950s, we observe two main eras in the APS citation landscape. From 1950 to 1990 we observe a steady increase in network-wide excess reciprocity, which then

remains roughly constant around 0.15 for the following twenty years. It is tempting to relate the former trend with the onset and the ensuing rise of Scientometrics as a research field. Indeed, the concept of bibliometric indicators was first proposed in 1955 by Garfield [5], and put in practice in 1964 with the launch of the Science Citation Index, with an ensuing proliferation of indicators [6] and the establishment of the first academic journal entirely devoted to Scientometrics in 1978 [39].

It is somewhat surprising to observe the above trend plateauing over the last two decades, when the average number of citations made by papers in the APS increased [17], and the incorporation of citations and bibliometric indicators in academic decision-making witnessed further increase [40]. Yet, this is accompanied by a widening gap between the observed reciprocity taking place between coauthors (which, after the 1990s accounts for almost *half* of the whole network's reciprocity), and the expected level of reciprocity between coauthors in our null model. These two seemingly at odds trends might be related to the onset of the Internet, which dramatically reduced barriers to access to published research. We speculate that the expansion of the Internet might have, on average, widened a scientist's potential pool of authors to cite [17], therefore diluting the overall reciprocity in the scientific community while at the same time resulting in a more prominent role of the reciprocated citations between collaborators and colleagues.

The above change is accompanied by a structural evolution in the author-author citation network. Over the years we witness the emergence of a very clear rich club of top scientists citing each other at a rate much higher than the one predicted by our null model. When looked at from the perspective of the whole network, however, the impact of such a rich club gets diluted, and top scientists end up being below average "reciprocators". Starting from the 1970s the author network organizes in such a way that most of the reciprocity takes place at the level of low to medium impact authors.

The latter result is corroborated by our clustering analysis (see Fig. 3 and Additional file 1, Supplementary Note 4), which shows that only low-impact career trajectories tend to be associated with high excess reciprocity. All in all, these results suggest that even potentially sophisticated shortcuts to artificially boost bibliometric indicators based on the exchange of citations, rather than on mere self-citations, cannot circumvent the fact that consistently high-quality publications are, by far, the main determinant of academic impact.

Yet, our results show that a remarkable proportion of most scientists' citations come, on average, from their immediate "neighbourhood" (Fig. 5(b)), and that the gap between such proportion and the one expected under our null hypothesis has been constantly increasing for more than 40 years. We interpret this as an echo of the academic community's *collective* incentive to boost individual metrics of academic reputation. In this respect, it is worth stressing once more that our results are of a statistical nature. As such, they cannot provide insight about the countless reasons that might induce *individual* authors to reciprocate citations. A proper investigation of potentially malicious practices deliberately aimed at boosting bibliometric indicators is well beyond the scope of the present work, and most likely would entail collecting data by interviewing authors.

We believe our results should caution against the current practice of condensing the entirety of a scientist's production into a single bibliometric indicator, i.e., into a single number. We argue that, in analogy to what some academic platforms do to display information about a scientist's impact with and without self-citations, such a profiling should

be complemented by providing a more detailed breakdown about the origin of a scientist's citations.

4 Materials and methods

4.1 APS citation data and network construction

The APS corpus of journals publish articles spanning all research fields in Physics since 1893. The dataset contains all papers published between 1893 and 2010, which we complemented with the work done to disambiguate authors in [12], where 236,884 unique authors were identified. Following [12], we only retained papers in the dataset with no more than 10 authors, which left us with 415,342 papers and 4,125,843 citations.

We used the above data to construct networks of citations both between papers and between authors. Although our analysis starts from 1950, we still employed all papers published from the beginning of the APS dataset in order to account for the fact that reciprocity in citations is an inherently cumulative phenomenon. Paper networks are directed and unweighted, while author networks are directed and weighted, with a link of weight c_{ij} from node i to node j denoting that author i has cited c_{ij} times author j. We used this general structure to extract annual networks, by selecting all the active authors who had started their careers before and ended after the year of interest, and choosing all papers published by such authors before that same year.

In analogy, in order to perform our analyses at the decade level, we selected all authors active in at least one year of the decade of interest. We then retained all papers published by the selected authors before the end of the decade.

4.2 Geodesic distance 1 null network model

We propose a null network model ensemble to estimate an average baseline level of reciprocity $\langle \rho_{\text{null}}^{(i)} \rangle$ one should expect to observe from a certain author i in the citation networks under partially random interactions. This quantity can be then used to compute the author's excess reciprocity as per Eq. (1).

We build on and generalize the null network model put forward in [41] to measure reciprocity in directed unweighted networks. Given a set of papers, we construct our null network model ensemble according to the following principles. First, the ensemble should reflect the fact that citations are attributed via papers, i.e., that papers are the fundamental units of interaction. Second, it should preserve the inherent directionality of citations between papers due to time ordering, i.e., that more recent papers cite older ones, and not the other way around. Third, the ensemble should account for the fact that most citations occur within well defined scientific communities, which can be proxied in terms of a "distance" constraint between authors, consistently with the vast amount of literature which shows that homophily and peer influence typically do not extend beyond a few degrees of separation in a variety of social networks [42–44].

Following the above prescriptions, we define our null network model ensemble as follows. We start from the real citation network structure, and iteratively select random pairs of citations $p_1 \rightarrow p_2$ and $p_3 \rightarrow p_4$ (see Fig. 2), where p_i ($1 \leq i \leq 4$) denotes a paper and $p_i \rightarrow p_j$ denotes a citation from p_i to p_j. With probability 1/2 we swap the two links representing the citations if the two following conditions are both met: (i) both new citations $p_1 \rightarrow p_4$ and $p_3 \rightarrow p_2$ preserve the time ordering of publication dates (i.e., p_1 was published before p_4, and p_3 was published before p_2), and (ii) both new citations $p_i \rightarrow p_j$ (with

$(i, j) \in \{(1, 4), (3, 2)\})$ are such that either papers p_i and p_j share at least one common author, or there exist authors a_i from p_i and a_j from p_j such that a_i has cited a_j or a_j has cited a_i at least once in the empirical citation network. The latter condition encodes the aforementioned distance constraint, limiting the set of allowed link rewiring moves to papers whose author lists feature at least one pair of authors at geodesic distance 1 in the empirical author network (see Additional file 1, Supplementary Note 8 for the results obtained with a less constrained null model). When controlling for modularity-related effects, we retain the same rewiring procedure, but only accept rewiring moves when the new potential citations $p_i \rightarrow p_j$ (with $(i, j) \in \{(1, 4), (3, 2)\}$) are such that there exist authors a_i from p_i and a_j from p_j who both belong to the same community (see Additional file 1, Supplementary Note 4).

The above operations are repeated until the system has reached an equilibrium state. This is identified by tracking the autocorrelation of the total reciprocated weight during the rewiring procedure, which is stopped after such quantity has reached values compatible with those of a random process. At that point, we collect independent samples of the paper-paper citation network, which are then used to compute the average fraction of reciprocated weight for each author $\langle \rho_{\text{null}}^{(i)} \rangle$ in the null model ensemble and excess reciprocity as per Eq. (1).

Let us remark that our rewiring procedure does not change the number of citations cited and received by a paper, thus the most relevant bibliometric indicators (i.e., citation counts and the h-index) are preserved for each author. Hence, the rewiring preserves the bibliometric impact of each individual author, and effectively amounts to probing alternative citation patterns through which it might have been produced. Indeed, as shown in Fig. 2, the null model changes the authors' in-degree in their network, i.e., the number of different peers that have cited their work. Overall, the changes induced by the rewiring alter the in-degree distribution in a substantial manner. Still, we ran two-sided Kolmogorov–Smirnov tests on the in-degree distributions before and after the rewiring, and found that the null hypothesis of the two being the same distribution could not be rejected at a 1% significance level.

4.3 Rich-club coefficient

The presence of the rich-club phenomenon in a weighted network can be detected by first ranking the authors in terms of a "richness" parameter r [32]. Here r equals the weighted in-degree of a scientist in the author citation network, i.e. the number of received citation at a certain level of temporal aggregation. For each value of r, we select the group of all authors whose total number of received citations is larger than r. We thus obtain a series of increasingly selective clubs. For each of such clubs, we count the number $E_{>r}$ of links connecting the members, and measure the sum $W_{>r}$ of the weights attached to such links. We then measure the ratio $\phi^w(r)$ between $W_{>r}$ and the sum of the weights attached to the top $E_{>r}$ strongest citations within the whole network. We thus measure the fraction of weights shared by the most successful authors compared with the total amount they could share if they were connected through the strongest links of the network:

$$\phi_0^w(r) = \frac{W_{>r}}{\sum_{l=1}^{E_{>r}} w_l^{\text{rank}}}, \tag{2}$$

where $w_l^{\text{rank}} \geq w_{l+1}^{\text{rank}}$ with $l = 1, 2, \ldots, E$ are the ranked weights on the citations of the network, and E is the total number of links. The weighted rich-club effect can be detected by measuring the ratio

$$\phi^w(r) = \frac{\phi_0^w(r)}{\phi_{\text{null}}^w(r)}, \tag{3}$$

where $\phi_{\text{null}}^w(r)$ refers to the average weighted rich-club effect assessed on the null model. When ϕ^w is larger than 1, it means that leading authors are concentrating most of their citations towards other successful authors compared with what happens in the random null model. Conversely, if it is smaller than 1, the citations among club members are weaker than randomly expected.

Additional material

Additional file 1: Supplementary information (PDF 904 kB)

Abbreviations
APS, American Physical Society; PRL, Physical Review Letters.

Availability of data and materials
The datasets generated and/or analysed during the current study are available at American Physical Society, https://journals.aps.org/datasets.

Competing interests
The authors declare that they have no competing interests.

Authors' contributions
WL, TA, FC, and GL designed and performed research, contributed new analytic tools, and wrote the paper. WL analyzed data. All authors read and approved the final manuscript.

Author details
[1]Department of Computer Science, University College London, London, UK. [2]Systemic Risk Centre, London School of Economics and Political Science, London, UK. [3]London Mathematical Laboratory, London, UK.

Endnote
[a] The quantity in Eq. (1) was simply referred to as reciprocity in [29], where it was first introduced. We shall instead refer to it as excess reciprocity in order to highlight that such a definition always entails a comparison with a benchmark defined by a null model.

Publisher's Note
Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References
1. Garfield E (2006) The history and meaning of the journal impact factor. JAMA 295(1):90–93
2. Bergstrom CT, West JD, Wiseman MA (2008) The eigenfactor™ metrics. J Neurosci 28(45):11433–11434
3. Hirsch JE (2005) An index to quantify an individual's scientific research output. Proc Natl Acad Sci USA 102(46):16569
4. Clauset A, Larremore DB, Sinatra R (2017) Data-driven predictions in the science of science. Science 355(6324):477–480
5. Garfield E (1955) Citation indexes for science. A new dimension in documentation through association of ideas. Science 122:108–111
6. Garfield E et al (1970) Citation indexing for studying science. Nature 227(5259):669–671
7. Bornmann L, Mutz R, Neuhaus C, Daniel H-D (2008) Citation counts for research evaluation: standards of good practice for analyzing bibliometric data and presenting and interpreting results. Ethics Sci Environ Polit 8(1):93–102

8. Hood W, Wilson C (2001) The literature of bibliometrics, scientometrics, and informetrics. Scientometrics 52(2):291–314
9. Petersen AM, Fortunato S, Pan RK, Kaski K, Penner O, Rungi A, Riccaboni M, Stanley HE, Pammolli F (2014) Reputation and impact in academic careers. Proc Natl Acad Sci USA 111(43):15316–15321
10. Weiss CH, Poncela-Casasnovas J, Glaser JI, Pah AR, Persell SD, Baker DW, Wunderink RG, Amaral LAN (2014) Adoption of a high-impact innovation in a homogeneous population. Phys Rev X 4(4):041008
11. Fortunato S, Bergstrom CT, Börner K, Evans JA, Helbing D, Milojević S, Petersen AM, Radicchi F, Sinatra R, Uzzi B, Vespignani A, Waltman L, Wang D, Barabási A-L (2018) Science of science. Science 359(6379):0185
12. Sinatra R, Wang D, Deville P, Song C, Barabási A-L (2016) Quantifying the evolution of individual scientific impact. Science 354(6312):5239
13. Sarigöl E, Pfitzner R, Scholtes I, Garas A, Schweitzer F (2014) Predicting scientific success based on coauthorship networks. EPJ Data Sci 3(1):9
14. Hellsten I, Lambiotte R, Scharnhorst A, Ausloos M (2007) Self-citations, co-authorships and keywords: a new approach to scientists' field mobility? Scientometrics 72(3):469–486
15. Fowler JH, Aksnes DW (2007) Does self-citation pay? Scientometrics 72(3):427–437
16. Radicchi F, Fortunato S, Markines B, Vespignani A (2009) Diffusion of scientific credits and the ranking of scientists. Phys Rev E 80(5):056103
17. Martin T, Ball B, Karrer B, Newman M (2013) Coauthorship and citation patterns in the physical review. Phys Rev E 88(1):012814
18. Seeber M, Cattaneo M, Meoli M, Malighetti P (2017) Self-citations as strategic response to the use of metrics for career decisions. Research Policy
19. Bartneck C, Kokkelmans S (2010) Detecting h-index manipulation through self-citation analysis. Scientometrics 87(1):85–98
20. Guimera R, Uzzi B, Spiro J, Amaral LAN (2005) Team assembly mechanisms determine collaboration network structure and team performance. Science 308(5722):697–702
21. Haines VA, Godley J, Hawe P (2011) Understanding interdisciplinary collaborations as social networks. Am J Community Psychol 47(1–2):1–11
22. Zaggl MA (2017) Manipulation of explicit reputation in innovation and knowledge exchange communities: the example of referencing in science. Res Policy 46(5):970–983
23. Petersen A (2018) Megajournal mismanagement: manuscript decision bias and anomalous editor activity at plos one. SSRN: e-print 2901272
24. Trivers RL (1971) The evolution of reciprocal altruism. Q Rev Biol 46(1):35–57
25. Berg J, Dickhaut J, McCabe K (1995) Trust, reciprocity, and social history. Games Econ Behav 10(1):122–142
26. Hilbe C, Chatterjee K, Nowak MA (2018) Partners and rivals in direct reciprocity. Nat Hum Behav 2:469–477
27. Mahmoodi A, Bahrami B, Mehring C (2018) Reciprocity of social influence. Nat Commun 9(1):2474
28. Newman ME (2004) Coauthorship networks and patterns of scientific collaboration. Proc Natl Acad Sci USA 101(suppl 1):5200–5205
29. Garlaschelli D, Loffredo MI (2004) Patterns of link reciprocity in directed networks. Phys Rev Lett 93(26):268701
30. MacQueen J et al (1967) Some methods for classification and analysis of multivariate observations. In: Proceedings of the fifth Berkeley symposium on mathematical statistics and probability, Oakland, CA, USA, vol 1, pp 281–297
31. Zhou S, Mondragón RJ (2004) The rich-club phenomenon in the Internet topology. IEEE Commun Lett 8(3):180–182
32. Opsahl T, Colizza V, Panzarasa P, Ramasco JJ (2008) Prominence and control: the weighted rich-club effect. Phys Rev Lett 101(16):168702
33. Ciotti V, Bonaventura M, Nicosia V, Panzarasa P, Latora V (2016) Homophily and missing links in citation networks. EPJ Data Sci 5(1):7
34. King MM, Bergstrom CT, Correll SJ, Jacquet J, West JD (2017) Men set their own cites high: gender and self-citation across fields and over time. Socius 3:2378023117738903
35. Guimera R, Sales-Pardo M, Amaral LA (2007) Classes of complex networks defined by role-to-role connectivity profiles. Nat Phys 3(1):63
36. Rosvall M, Bergstrom CT (2008) Maps of random walks on complex networks reveal community structure. Proc Natl Acad Sci 105(4):1118–1123
37. Clauset A, Newman ME, Moore C (2004) Finding community structure in very large networks. Phys Rev E 70(6):066111
38. Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the em algorithm. J R Stat Soc, Ser B, Stat Methodol 39:1–38
39. de Solla Price D (1978) Editorial statements. Scientometrics 1(1):3–8
40. OECD, SCImago Research Group et al (2016) Compendium of bibliometric science indicators. OECD, Paris
41. Livan G, Caccioli F, Aste T (2017) Excess reciprocity distorts reputation in online social networks. Sci Rep 7(1):3551
42. Christakis NA, Fowler JH (2007) The spread of obesity in a large social network over 32 years. N Engl J Med 357(4):370–379
43. Li W, Bradshaw AE, Clary CB, Cranmer SJ (2017) A three-degree horizon of peace in the military alliance network. Sci Adv 3(3):1601895
44. Centola D (2010) The spread of behavior in an online social network experiment. Science 329(5996):1194–1197