INSTITUTE
OF ECONOMICS

Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

Citations are Forever: Modeling Constrained Network Formation

Pietro Battiston °

° Department of Economics, University of Milan-Bicocca, Italy and Institute of Economics, Scuola Superiore Sant'Anna, Pisa, Italy

2014/19 October 2014

ISSN (online) 2284-0400

Citations are Forever: Modeling Constrained Network Formation

Pietro Battiston*

October 17, 2014

Abstract

Determining the extent to which citation flows, and hence bibliometric indicators based on them, reflect some intrinsic value of scientific works is an important task made very difficult by endogeneity issues. This paper presents an approach which allows to go beyond the abundant anecdotal evidence by testing whether the citation behavior is free from environmental factors. The hypothesis of independence is strongly rejected, providing causal evidence of a Matthew effect at work: namely, the publication of a new work on behalf of an author increases the flow of citations to previous works. Such result is a step towards the estimation of biases affecting bibliometric indicators, at least when interpreted as measures of scientific productivity.

The study is based on a novel framework for the study of endogenous network growth subject to constraints. Constraints can be both positive and negative, and change in time depending on the actions of the agents. The framework is not limited to citation networks, and can be applied to any context in which the formation of a link inhibits or implies the formation of another one.

Keywords: Bibliometric indicators, Endogenous growth, Matthew effect, Research evaluation.

JEL Classification: D85, C10, C72, I23.

*Department of Economics, University of Milan Bicocca. Current address: Istituto di Economia, Scuola Superiore Sant'Anna, Piazza Martiri della Libertà 33, 56127 Pisa, Italy. E-mail: me@pietrobattiston.it

"Always go to other people's funerals; otherwise they won't go to yours."
Yogi Berra (confused by a typical "constrained growth" network)

1 Introduction

Evaluating scientific research is a difficult but important task. Measures of scientific productivity of researchers, research projects or research institutions often form the basis, officially or implicitly, for decisions on how to allocate available funds. This poses however a problem not just of asymmetric and costly information, but also of a conceptual nature: *what* makes good research such? Is it *originality*, *adherence to scientific standards*, or merely *impact* on the public opinion or on the academia? The right answer may well depend on the context.

The availability of *citation indexes*, such as ISI *Web of Knowledge* and Elsevier *Scopus*, has increased tremendously the popularity of *bibliometric indicators*, numeric indicators of research output which are typically computed from publications on peer-reviewed journals, and citations among them. Compared i.e. to ad-hoc peer review panels, such indicators have some obvious advantages: they are inexpensive to estimate, they can be reproduced reliably once the criteria (including the set of accepted journals) are determined, and they associate a single number to each item being evaluated, be it a researcher or group of researchers (i.e. the *h-index*) or a scientific journal (i.e. the Impact Factor), making comparisons straightforward. On the other hand, a growing literature has been pointing at the potential problems connected with the use of such measures. Namely, citations may proxy *impact*, rather than some form of *quality*; moreover, attributing an economic value to publications and citations may distort the behavior of researchers and journal editors; finally, even leaving aside the possibility of strategic behavior, several environmental factors may in principle bias bibliometric indicators, when considered as measures of research quality. Unfortunately, although the anecdotal evidence in this direction is abundant, empirically assessing and estimating such bias is a daunting task, because of endogeneity issues. In this paper, I employ a new approach to the problem, by delimiting the definition of quality of a published work as something uniquely dependent from its *content*. Although this criterion is extremely general, when coupled with an appropriate model of endogenous network formation it is restrictive enough to unambiguously detect the presence of environmental effects at work: namely, it will allow to empirically verify that citation flows to a scientific paper causally depend on the status of authors rather than uniquely on their scientific content.

The contribution I make to the empirical literature on citation behavior hence consists in providing a comprehensive lower bound to the importance of *environmental variables*, which is not based on merely anecdotal evidence, and which can be attributed a causal meaning. To the best of my knowledge, the only other successful attempt at identifying a causal effect of fame on citation patterns is due to Azoulay *et al.* (2014), who find a clear, although short-lived, increase in the inflow of citations after a scientist is appointed the title of Howard Hughes Medical Investigator. While a clear advantage of their approach is the simple and unambiguous interpretation of the results in terms of publicity and fame, it is also true that it can only be used for the study a small population of scientists in a specific field. The approach described in sections from 4 to 6 of the present study can in principle be applied to any scientific field, or even to sub-populations of researchers defined according to various criteria (nationality, other proxies of fame. . .)

My results are complementary to the literature on statistical properties of bibliometric indicators, such as the analysis of the h-index by Pratelli *et al.* (2012). Their model accounts for the strong and intrinsic non-independence of citation flows across time, but does not consider the possibility of effects such as the *fame* of the author, or even strategic citing behavior, influencing the *citation flows*. Having a lower bound on the determinants of citations which are entirely unrelated to the *content* of articles provides an additional error term to be considered on top of what they find, and can have important policy implications concerning the use of bibliometric indicators for the distribution of resources among publicly funded research facilities (and, in general, research institutions targeting scientific goals not limited to the mere *impact* on the academia).

The next section contextualizes my contribution in the existing literature; the generalized model which provides the base for my analysis is presented in Section 3 and then specialized to the case of the network of citations in Section 4; Sections 5 and 6 present respectively the data on which the analysis is ran, and the results; finally, Section 7 concludes.

2 Previous literature

In the last twenty years, the theory of networks has been recognized an important role in explaining the formation and functioning of social and economic settings in which *relationships* among agents are of fundamental importance. In particular, several models of network formation were developed targeting the mechanisms by which some characteristics of nodes (typically, the cost of creating/keeping alive a link, compared to the utility received from becoming - directly or indirectly - connected to some other nodes) endogenously determine the structure of a network. A stream of literature, starting from the seminal work of Bala & Goyal (2000), has developed focusing on a *noncooperative* approach. Under such approach, the choice of adding a link between two nodes is made independently by only one of them, which bears all the cost - although both potentially benefit from such link. Based on this framework, a definition of *stability* can be given, typically based on the concept of *pairwise Nash* equilibrium (such as in Galeotti, 2006 and Haller *et al.* , 2007), or some refinement of it (for instance Dutta & Mutuswami, 1997 consider *coalition* choices, while the concept of *"farsightedly stable networks"* formulated by Herings *et al.* , 2009 assumes that nodes have a longer horizon of reasoning).

Those studies share the implicit assumption that links can be added and destroyed freely (though at some cost). Even experimental works on endogenous network formation have usually been based on the assumption that participants can *at any point in time* - or at least *repeatedly* - decide to create/break a link (see for instance Goeree *et al.* , 2009 and Kirchsteiger *et al.* , 2011). This is a very natural setup for several reasons. Firstly, many real world networks (i.e. computer networks, social relationships. . .) are indeed characterized by links which are at least potentially volatile. Secondly, even in cases in which the network under study is characterized by exogenous restrictions, such as geographic ones, the interest of researchers, as well as the available data, is often focused on an inherently volatile *flow* of some good (such as influence, or information) over it. Thirdly, even considering networks which are typically characterized by a *stratification* of links over time (such as connections in Internet social networks, or the network of roads between cities), most databases in use are *snapshots* of networks at given points in time, often sacrificing information on their temporal evolution.

Only recently a form of *temporal hierarchy* of nodes/links has been considered in the context of such theories, by Haller (2012). His study provides interesting conclusions applying to networks growing around an exogenously fixed subset of links, which is shown to potentially change drastically the existence, numerosity, stability and efficiency of pairwise Nash equilibria. An interesting insight is that such *backbone infrastructures*, that is, sets of links which are guaranteed to exist independently from individual incentives, and which hence *restrict* the set of possible actions available to nodes, can nevertheless cause *welfare improvements*. Section 3 of the present study generalizes the approach to the analysis of *repeated* addition of nodes and/or links, under positive *and negative* constraints. Differently from the work of Haller, the set of guaranteed/forbidden links will not necessarily be exogenously given, but can come instead from the previous iteration of the network formation process. On the other hand, the proposed model differs from the literature on *growing network evolution models* (Toivonen *et al.* , 2009) not just for the presence of constraints, but also because node attributes matter crucially for the creation of links.

In Section 4, I will then adopt the new setting to model the network of citations among scientific publications, and test some of its stylized predictions. The scientometric literature, although relatively young, has developed extremely in the last decades.[1] The first bibliometric studies, such as the work by Gross & Gross (1927), were motivated by the practical aim of determining which scientific works ought to be present in a scientific library in order to satisfy the needs of faculty members. The idea that the encoding of citations between scientific papers could be of use to researchers themselves, and in particular that a *citation index* could prove of tremendous utility, is due to Garfield (1955), who went as forward as to propose the term "impact factor". Most importantly, he created the first implementation of such an index, the evolution of which is today the *ISI Web of Knowledge*. The stream of literature qualifying and quantifying the different roles that a citations can have in enhancing the quality of exposition, or trustworthiness, of a scientific paper can also be traced back to the work of Garfield *et al.* (1965), who proposed a classification, admittedly incomplete, of fifteen different motivations which can explain why a scientific paper cites another one. Such classification also includes categories for *negative* citations, aimed at *"disclaiming work or ideas of others"* or *"disputing priority claims of others"*, which can be seen today as an implicit criticism to the *normative* view, according to which instead citations are a mere way to attribute credit to previous works,

[1]The following paragraph is largely based on the review by Baccini (2010).

and even to reward their authors. The spirit of such categorization is however distinct from the alternative *constructive* view, according to which citations are only the result of *strategic* decisions, targeted at gaining *"a dominant position in their scientific community"* (Moed & Garfield, 2004). The approach taken in the present study is also, in principle, agnostic with respect to those main views, in the sense that it does not recognize citations as the mere consequence of a norm, but neither assumes they are necessarily taken strategically. For instance, the results I will expose are consistent with the hypothesis that citations are mainly an instrument to attribute credit, but that they are influenced by environmental factors. On the other hand, the empirical evidence, in particular on self-citations, is also compatible with strategic motives. In any case, the implications point towards reconsidering the value which can be attributed to bibliometric indicators as measures of scientific productivity and/or impact.

From the empirical point of view, the analysis of the *scientific network* - a broad term which can refer to the networks defined by several different *relationships* characterizing the functioning of the academia - has been the subject of many studies. Virtually all of them employ data related to scientific journals and publications on them: in particular, many focus on the *co-authorship* relation between researchers (see Goyal *et al.* , 2006, Cainelli *et al.* , 2010 and De Stefano *et al.* , 2013 for recent examples), or on the links between journals (for instance Baccini & Barabesi, 2010 considers the relation "having a non-empty intersection between members of editorial boards", while West *et al.* , 2010 discuss the *Eigenfactor*, a method for ranking journals by importance, based on the network of citations between them).

The analysis of the network of citations *between paper themselves* also has a long history, dating at least to de Solla Price (1965), who estimated some of its relevant statistics, focusing in particular on the very skewed distribution of indegree (when compared to the outdegree).[2] Some attempts have also been targeted at extracting from the characteristics of the network some insights on the behavioral aspects of the act of citing, such as in Baldi (1998). This approach is however hampered by huge endogeneity issues. The problem is worsened by the tendency to attribute citations an important role as a measure of research *impact* and/or *quality*: this increases the incentives to act *strategically* in the selection of papers to cite (i.e. citing papers from

[2]The *indegree* of a given paper is, here and in the following, the number of papers citing it, while the outdegree is the number of papers it cites.

a given journal in order to increase the chances of getting published, citing papers from a given author in order to be cited/positively referred in return).

Probably because of the historical difficulties in obtaining and analyzing *entire* bibliographic databases, most works on the network of citations have also been limited to the observation of *strictly local* properties - typically, correlations among different characteristics on a *per node* (paper) basis, or at most among characteristics of both the citing and the cited paper. Section 4 and subsequent ones of the present study, although still based on local properties, go beyond this limit, adopting an empirical strategy based on particularly defined sub-networks of diameter 3, and exploiting information about authors and journals. The aim is to provide a *causal identification* of the importance of environmental factors on the citing process (more precisely, establishing a lower bound to the importance of such effects). This approach is complementary to the literature on statistical properties of bibliometric indicators, such as the analysis of the h-index by Pratelli *et al.* (2012). Their model accounts for the strong and intrinsic non-independence of citation flows across time, but does not consider the possibility of effects such as the *fame* of the author, or even strategic citing behavior, influencing the *citation flows*. Having a lower bound on the determinants of citations which are entirely unrelated to the *content* of articles provides an additional error term to be considered on top of what they find, and can have important policy implications concerning the use of bibliometric indicators for the distribution of resources among publicly funded research facilities (and, in general, research institutions targeting scientific goals not limited to the mere *impact* on the academia).

The fame effect was already considered in several studies. As suggested by de Solla Price (1976), "success breeds success": the more a paper is known/cited, the more its fame/flow of citations will grow in the future, and the same can hence be said for an *author*. However, while this is undoubtedly true from the *predictive* point of view (a clear correlation between past and future flow of citations has been found in the empirical literature), it is particularly hard to identify and quantify a *causal effect*, which often goes under the name of *Matthew effect* (Merton, 1968). The scope of my work is hence precisely to extract from bibliometric data some evidence of *herding*: may papers gain popularity (citations) not (just) because of their content but because of some *environmental* characteristics which focus on them the attention of the scientific community? This could very well be a circular process (citations themselves could influence such environmental characteristics). Notice the term “herding” does not necessarily imply any irrationality on the behalf of agents (in the same way in which it does not when used in the context of the financial market): it can be the consequence of bounded rationality, but also of limited information (or *costly* information - the duty of keeping at pace with the existing literature typically taking up a relevant share of the work time of a researcher), and, as already mentioned, strategic behavior.

3 The model

Since the framework used for studying the network of citations does not depend from the nature of the network, but rather can be applied to any process of endogenous formation subject to constraints, it will be now presented in general terms. Its fundamental block is the model of Galeotti *et al.* (2006). A network is composed by $N = \{1, \ldots, n\}$ nodes: for each pair of nodes i, j, a cost parameter $c_{ij} > 0$ and a value parameter $v_{ij} > 0$ are given. A *directed* network g is formally a set of pairs of nodes: if a pair (i, j) is in g, we say that i *sponsors* a link to j, and we write $g_{ij} = 1$. $\bar{g}$ represents the corresponding *undirected* network, that is, the smallest network containing g and (j, i) for each (i, j) contained in g. Each node extracts from the network a benefit which depends on the values of the nodes which are *connected* to it. That is, denoting as $N_i(g)$ the set of nodes j such that the network g contains a path from i to j, the benefit extracted by i is defined as:

$$B_i(g) = B_i(\bar{g}) = \sum_{j \in N_i(\bar{g})} v_{ij};$$

i also pays a cost which is the sum of costs of sponsored (outgoing) links:

$$C_i(g) = \sum_{j : g_{ij} = 1} c_{ij},$$

and the resulting payoff deriving from the network is simply the difference between the benefit and the cost:

$$\Pi_i(g) = B_i(g) - C_i(g).$$

Some other standard graph-theoretic concepts and notations will be used. The letter e denotes the empty network. A connected set of nodes $S \subset N$

is said to be a *component* if for any $i \in S$, $j \notin S$ there is no path from i to j in $\bar{g}$; a link is said to be a *bridge* if the number of components of the network changes (increasing by 1) once it is removed; a network is said to be *minimal* if all links are bridges, and *minimally connected* if it is connected and minimal. Moreover, the notation

$$g_i = (g_{i,1}, \ldots, g_{i,n}) \in \{0,1\}^n$$

summarizes the outgoing links from a given node i in the network g (in the present work, it is always assumed that $g_{ii} = 0$). A *strategy* for a node is also an element of the set $\{0,1\}^n$, and it is *included* in another one ($g_i \subset g'_i$) if it involves sponsoring only links which are sponsored according to the other one (that is, if $g_{ij} = 0$ whenever $g'_{ij} = 0$).

3.1 Internal constraints

Haller (2012) enriches this basic model with the presence of *constraints*: in his work, a pre-existing and exogenously given network $\mathfrak{g} \in G$. The payoff function is modified by setting the cost of links in $\mathfrak{g}$ to 0, and as a consequence such links are always incentive compatible. The aim of the present section is to generalize this seminal idea with the concept of *negative* constraints: a model of network formation will be characterized not only by $\mathfrak{g}$, which will be denoted henceforth as $\mathfrak{g}^+$, but also by another network $\mathfrak{g}^-$, containing links which will be *absent* in any possible network (by assumption, $\mathfrak{g}^+$ and $\mathfrak{g}^-$ will be disjoint). Although it is possible to introduce this generalization by setting the cost of links in $\mathfrak{g}^-$ high enough, a more tractable approach is to neglect their benefits in the payoff function,[3] which is hence defined as

$$\Pi_i \left(\mathfrak{g}^+, \mathfrak{g}^-, g\right) = B_i \left(g \oplus \mathfrak{g}^+ \ominus \mathfrak{g}^-\right) - C_i (g) \qquad \text{for } g \in \mathcal{G}.$$

where $\oplus$ and $\ominus$ denote respectively the operations of union and difference between networks.[4] It can be easily verified that when $\mathfrak{g}^- = e$, this coincides with the payoff function defined by Haller (2012). With all the components of the model exposed, we can proceed to the generalization of some of his

[3] As in the approach of Haller (2012), the original cost of links in $\mathfrak{g}^+$ and $\mathfrak{g}^-$ should be taken again into consideration when doing comparative statics and welfare analysis.

[4] With a slight abuse of notation, when the network to be added/removed is composed of a single link, I will write $g \oplus (i,j)$ or $g \ominus (i,j)$, rather than $g \oplus \{(i,j)\}$ or $g \ominus \{(i,j)\}$.

results concerning Nash networks - that is, networks which are stable with respect to individual deviations.

Proposition 1. *Consider a strategic model of network formation with payoff functions* $\Pi_i(\mathfrak{g}^+, \mathfrak{g}^-, g)$, $g \in \mathcal{G}$, $i \in N$. *Suppose that costs are owner-homogeneous. Then there exists a Nash network* g^*.

Proof. See Appendix A. □

This proposition immediately generalizes of Proposition 1 by Haller.[5] What follows is instead the natural generalization to negative restrictions of his Proposition 2.

Proposition 2. *Consider a strategic model of network formation with payoff functions* $\Pi_i(\mathfrak{g}^+, \mathfrak{g}^-, g), g \in \mathcal{G}, i \in N$. *Suppose that the pre-existing network or infrastructure* $\mathfrak{g}^+ \in \mathcal{G}$ *is Pareto optimal. Then the empty network is a strict Nash network and the only Nash network.*

Proof. See Appendix A. □

The effort in generalizing the theory of network extension to the presence of negative constraints can be motivated with two main arguments:

1. considering negative constraints is important in order to understand the growth of some real world network settings,

2. from a social planner point of view, imposing negative constraints can improve the beneficial effects of an endogenously formed network, possibly at a lower cost than through positive constraints.

The first motivation has already been discussed, and will be at the core of Sections 4 to 6. I will now devote some attention to the second. Haller (2012) shows several ways in which positive exogenous constraints can impact on the equilibria of a network: examples include a *stabilizing* effect (some models of

[5]The assumption that costs are owner-homogeneous is one of the reasons why it is impractical to define negative constraints just as arbitrarily costly links: if this was the case, in order for a owner-homogeneous model of network formation to remain such after the imposition of negative constraints, such constraints could not consist in arbitrary sets of links, but rather include all outgoing links from a given set of nodes. Another reasons is that this would make the definition of *endogenous* negative constraints, as described in Section 3.2, much more complicated.

network formation exhibit non-existence of Nash network, which can instead exist for given $\mathfrak{g}^+$), a *welfare improvement* effect (constraints can raise the overall sum of payoffs in Nash equilibrium), and others. Those exogenous constraints can hence be imagined as publicly provided infrastructures which are provided by the social planner. Can some of the described effects be attained as well through *negative* constraints - i.e. with a social planner acting through *prohibition* of a set of given links? For what concerns the attainment of *efficient* Nash equilibria, negative constraints cannot simply replace positive ones. See for instance Example 3 by Haller (2012):[6] it is clear that in the absence of *positive* constraints, no link will ever connect the two sets of nodes $\{1, 2, 3, 4\}$ and $\{5, 6, 7, 8\}$ (and hence the welfare improvement represented by such a link will be lost), since for any i and k in the two different sets, we have

$$\sum_{j \neq i} v_{ij} = 8 < 16 = c_{ik}.$$

Something more can be said about the stabilizing effect. An example of it is easily found: by simply setting $\mathfrak{g}^-$ as the complementary of $\mathfrak{g}^+$ (hence e if we are assuming $g^+ = 0$), we transform g^+ itself into a trivial Nash equilibrium - analogously, the trivial but uninteresting way to stabilize any model of network formation with only positive constraints is clearly to set $\mathfrak{g}^+$ to the complete network. A more interesting case comes from Example 2 of Haller *et al.* (2007)[7], represented in Figure 1, where $v_{i,j} = 1$ for any i, j (all edges absent from the picture have a cost higher than 3 and are hence irrelevant). Haller (2012) shows that by setting $g^+ = \{[4, 2]\}$, we obtain the Nash equilibrium $g^* = \{[3, 4], [4, 2], [2, 1]\}$; the same result can be obtained however by setting $g^- = \{[3, 2]\}$.

[6] For ease of reading, the example is reported here:
Let $N = \{1, 2, 3, 4\}$, $K = \{1, 2, 3, 4\}$, $L = \{5, 6, 7, 8\}$. *Further set* V_{ij} *for all* $i \neq j$, *and*

$$c_{ij} = \begin{cases} 0.8 & \textit{if } i \neq j,\ i, j \in K, \\ 0.8 & \textit{if } i \neq j,\ i, j \in L, \\ 16 & \textit{if } i \in K, j \in L, \\ 16 & \textit{if } i \in L, j \in K. \end{cases}$$

[7] Example 1 by Haller (2012).

Figure 1: An example of the stabilizing effect of negative constraints.

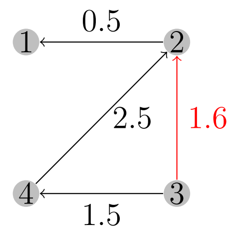

3.2 Repeated internal constraints, and non-decreasing network models

A crucial ingredient of any real world example of endogenous network formation is *time*: as will be clear later, it is particularly important for the analysis of the network of citations, which will be developed in the following sections. A study of the consequences of repeated internal constraints, going beyond the analysis of static Nash equilibria relative to exogenous constraints, is hence a natural development of the theory exposed so far. In what follows, I will assume that the formation of the network happens in a discrete time setting. For each $t = 1, 2 \ldots,$ I will define as $\mathfrak{g}^{t^+}$ and $\mathfrak{g}^{t^-}$ respectively the positive and negative constraints at that period. At each time, each node's best reply is the one maximizing $\Pi_i(\mathfrak{g}^{t^+}; \mathfrak{g}^{t^-}; \cdot)$.[8] The outcome, if any, of the step t, denoted as g^t, will hence be a Nash equilibrium for these payoffs functions. Clearly, such outcome needs not be unique, neither to exist: if it does not, the network formation process *terminates* at time t.

The result which follows considers the specific class of *non-decreasing network models*, defined as those for which $\mathfrak{g}^{t^+} = g^{t-1}$ (the positive restriction coincides with the outcome of the previous step of the process): such class naturally maps to several real world contexts, including the case of bibliometric networks. A peculiarity of non-decreasing network models is that, since the number of links present at time t is (weakly) increasing in t itself, and since it can never exceed $n^2 - n$, it must, for some t, terminate *or* stabilize in some configuration, which I will call a *limit network*. A limit network will then be defined as *strict* if there is no other limit network composed by a subset *or* a superset of its links.

[8]Clearly, the framework could also be an ideal context for the study of a less myopic type of rationality, such as the farsightedly stable networks (Herings *et al.* , 2009).

Proposition 3. *If $\mathfrak{g}^{t^-} = \mathfrak{g}^{t-1^-}$ for every t, then the set of (strict) limits of the non-decreasing network model corresponds to the set of (strict) Nash equilibria of the static model associated to the payoffs function $\Pi_i(\mathfrak{g}^{1^+}; \mathfrak{g}^{1^-}; \cdot)$.*

Proof. Consider a (strict) Nash equilibrium g^* of the model associated to payoffs functions $\Pi_i(\mathfrak{g}^{1^+}; \mathfrak{g}^{1^-}; \cdot)$. By, definition, it is also a (strict) Nash equilibrium for the first step of the non-decreasing network model. In order to prove that it is a limit network, it is hence sufficient to show that it is still a (strict) Nash equilibrium for $\Pi_i(g^*; \mathfrak{g}^{1^-}; \cdot)$. Assume it is not: that is, there is some i who (weakly) prefers some $g'_i \supset g^*_i$. But then, g^* was not a (strict) Nash network in the first place. The same applies hence for $t = 2, 3 \ldots$

Now assume g^* is a (strict) limit for the non-decreasing network model, reached at some time t^*. By construction, g^* is a union of t^* subsequent Nash extensions (some of them possibly empty), so because of Lemma 1 we know that all its elements are bridges. Given any time t and any link (i, j) in $g^t \ominus \mathfrak{g}^{t^+}$, let $\Delta^t_{i,j}$ be the profit which the link (i, j) yields to i in g^t, that is,

$$\Delta^t_{i,j} = \Pi(\mathfrak{g}^{1^+}, \mathfrak{g}^{1^-}, g^t) - \Pi(\mathfrak{g}^{1^+}, \mathfrak{g}^{1^-}, g^t \ominus (i, j)).$$

This profit is necessarily (strictly) positive, since the node is part of the best reply of i. Any $\Delta^{t'}_{i,j}$ with $t' > t$ will also be positive - all new links are bridges, and so the connected component of j can only grow, while no paths from i to j alternative to (i, j) can appear. So no node has a (weakly) positive individual incentive to simply break one or more existing links in g^t. If g^* is not a (strict) Nash equilibrium of $\Pi_i(\mathfrak{g}^{1^+}; \mathfrak{g}^{1^-}; \cdot)$, then necessarily some node has a (weakly) positive individual incentive to *add* some link, or to *replace* some link with some other. The first case is impossible: since $\mathfrak{g}^{t^-}$ is constant, this would make g^* unstable also at time t^*. But the second is also impossible: since $\Pi_i(\mathfrak{g}^{1^+}; \mathfrak{g}^{1^-}; \cdot)$ is positive, the new link should still connect i to $N_j(g^* \ominus (i, j))$. So to be incentive compatible, it should cost less than (i, j). But then, it would have been chosen at time t in its place. □

Proposition 3 in particular implies that when no Nash equilibrium exists, no limit network exists, and the iterative process necessarily terminates at time 1. Another interesting implication is that network models satisfying only the more general condition $\mathfrak{g}^{t+} \supseteq g^{t-1}$ do not exhibit richer limit structures than non-decreasing network models: imposing $(i, j) \in \mathfrak{g}^{t^+}$ would not make a change, in terms of limit networks, compared to imposing $(i, j) \in \mathfrak{g}^{0^+}$. Richer dynamics could instead be expected when

1. considering *partially* non-decreasing network - networks in which some previously provided links *can* be destroyed, or

2. introducing time-dependent *negative* constraints.

The growth of the network of citations, which the remaining of the paper focuses on, falls in this second case.

4 The growth of the citations network

The network of citations between scientific papers is an eminent example of an endogenously formed network in which the time component is not just crucial for the endogenous growth mechanism, but also easily observable in the available data. Indeed, papers have a well defined publication date, which imposes a clear temporal hierarchy among them and hence strong restrictions to the set of "actions" - that is, of citations - they can make.

Before diving into the details of the time-related properties of the network of citations among scientific publications, it is worth reviewing why, from a static point of view, an adaptation of the non-cooperative approach *à la* Bala & Goyal (2000), introduced in Section 3, is appropriate for the setting under analysis.

- A citation is a purely one-sided sponsored kind of relation: an author can very well find out *ex-post*, if ever, that she has been cited.

- The fact that being cited can, at least in some cases, represent a gain for a researcher is unanimously recognized, and is part of the motivation for the present study. Less intuitive is the utility obtained from *making* a citation. However, the hypothesis that *there is* some benefit is supported by the obvious fact that the overwhelming majority of scientific articles have a list of bibliographic references (also see the brief survey of the literature on citation behavior in Section 1). This is the only kind of *payoff* the present empirical exercise is interested in - notice that, according to the non-cooperative approach, a paper *cannot* in any way create ingoing links.

- Although there is apparently no cost involved in "sponsoring" a citation, it is evident that the number of bibliographic references contained into a single scientific work is limited: many authors, starting with de Solla Price (1965), have analyzed different aspects of its distribution, evidencing a strong concentration for small values (as will be confirmed in Section 5). While this evidence does not help in disaggregating the implicit costs born by authors in making citations, which may be due partly to editorial/formatting choices and partly to the work involved in processing the literature to be cited, it does provide clear evidence of some implicit costs.

- As best exemplified by the phenomenon of *literature reviews*, it is very natural to assume that the benefit of a citation to a given paper depends in turn also on the citations *included* in that paper. The hypothesis of *perfectly reliable links* - meaning that being connected to another paper through an arbitrarily long path is equivalent to being directly connected - is instead a non-harmful approximation of reality for the present analysis. As will be motivated in Section 4.2, it does not affect its qualitative results, and on the other hand an alternative specification would make the analysis much more complex and require some arbitrary choices.

Having clarified those points, the growth of the citations network is clearly a particular case of a constrained non-cooperative model: citations cannot be removed, so the network model is clearly non-decreasing. At the same time, there are obvious restrictions to the set of available strategies.

It is clear that the choice of adopting a pure value-based model of citation flows implies the exclusion of a wide range of factors that can possibly affect them, and some of which have been already mentioned in Section 1. The aim of the present study is not to convince the reader that such factors are irrelevant: well on the contrary, as will be clear in Section 6, the model is instrumental in showing that they have a significant effect on the citation flows. In particular, the model completely negates the constructive view on the citation behavior, since the strategic behavior of agents (papers) is only directed towards an increase in the quality of the publication: this is clearly an unrealistic assumption. Moreover, many empirical works have focused on the fundamental role of individual researchers and scientific fields as "attractors" of citations: the choice of neglecting such aspects in the model is, again, consciously made in order to *reject* a model according to which bibliometric indicators would be perfect measures of research quality.

For simplicity of exposition, I will assume that the set of nodes of the network of citations is predetermined (i.e. contains all scientific papers which are going to be published in a given time span). The number of instants in time is equal to this number of nodes, so that there is a one-to-one relation between each node i and its instant t_i (and vice-versa, between an instant t and its node n_t). The negative restrictions are then defined as follows:

$$\mathfrak{g}^{t_i^-} = \{(j,k) : j \neq i \text{ or } t_k > t_i\}$$

which means that at each instant in time, only the scientific publication being published can establish links, and it cannot cite works which are yet to be published (Figure 2).[9]

Figure 2: Examples of allowed and forbidden links at time t_i.

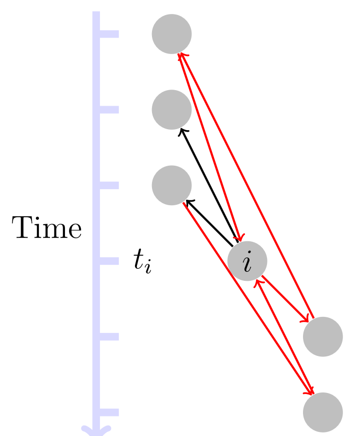

Allowed links in black, forbidden links in red. Links arrows point *from* the citing *to* the cited paper.

In Section 3.2, the fundamental building block of the development of a network with repeated internal constraints was assumed to be the Nash equilibrium of a given step t. Under the specification given for the network of citations, in which at each step only one node acts, such a Nash equilibrium degenerates to the best response of such node. The hypothesis of all

[9] In principle, given the typical publication process, which goes through a period of open discussion in seminars/workshop, an often lengthy referral process, and finally a delay from the definitive acceptance to the publication, it can happen that two papers i and j cite some version of each other. This very special case, which is not admissible under the simplified settings just described, is relatively non influential in the global picture, but would possibly deserve a specific analysis.

nodes existing since time 0 does not influence the strategic choice, which is determined simply as a best response *among allowed links* - because of the direction in which value "flows", *later* links are irrelevant.

The values $v_{i,j}$ and costs $c_{i,j}$ will be assumed to be coming from two probability distributions $\mathcal{V}(v, i, j)$ and $\mathcal{C}(v, i, j)$. The bibliometric literature contemplates different kinds of "*values*": even letting aside measurement issues, while for a public research institution the value of a paper may consist in the efficacy of a new pharmaceutical it describes, for a journal it may lie in its mere *impact* (i.e. copies sold) on the scientific community. Taking this difference at the extreme, if a paper is able to gather much attention, but is based on fabricated evidence, the overall gain for the journal having published it, in terms of additional publicity and possibly subscribers, may still be large, while the overall gains for a public institution financing the research will almost certainly be negative.[10] The present analysis imposes no particular definition of *value*: it could be *originality*, *adherence to scientific standards*, or a combination of those. It must however depend uniquely from the *content* of the paper (including the bibliographic references) rather than i.e. the citations that the paper already received. This is summarized in two assumptions on $v_{i,j}$ and $c_{i,j}$.

> **Assumption 1**: each paper i is characterized by an *intrinsic value* $\bar{v}_i$ such that, for any $t_j > t_i$,
>
> $$\mathbb{E}[v_{j,i}] = f_v(\bar{v}_i, t_j - t_i).$$
>
> where f_v is an age effect.

> **Assumption 2**: the aggregate dynamics of $\mathcal{C}(v, i, j)$ are such that
>
> $$\mathbb{E}[c_{j,i}] = f_c(v_{j,i}, t_j - t_i)$$
>
> where f_c is an age effect.

[10] Although the example is indeed extreme, a recent work by Fang & Casadevall (2011) finds a strong correlation between the "*retraction index*" of journals, and their Impact Factor. Many explanations are possible, but this reminds at the very least that journals may have scarce incentives to ascertain the methodological soundness of studies they accept.

The concept of intrinsic value should not be attributed a normative interpretation: it can be viewed as a mere mean of the idiosyncratic values, or as the target of a funding agency. Also notice that no assumption needs to be made on the shape of f_v and f_c: in particular, they can be non-monotonic (the typical citation flow accruing to a paper *is* indeed non-monotonic, peaking after a few months - see Figure 5).

4.1 Hypotheses

The null hypothesis which we want to test with the present econometric exercise is the following.

> **Hypothesis H0**: the growth of the citation network can be modeled as a non-decreasing model based on the non-cooperative approach *à la* Bala & Goyal (2000), with unidirectional flows, and with a distribution of values satisfying assumptions 1 and 2.

It is worth stressing that this assumption *does not resort* to stating that the probability of any link from j to i, with $t_j > t_i$, is only a function of $\bar{v}_i$: it will also depend on the links already present in the network, and namely on the paths *starting* from i and going backward in time (from now on, the *reference network*). In other terms, the benefit of a citation also depends on the bibliographic references of the cited paper. Also, recall that the values $v_{i,h}$ possibly include, by assumptions 1 and 2, an age effect.

The rest of the present study is devoted to testing on real data the Hypothesis H0, against the possibility that other *environmental characteristics* of papers may be influencing the choice to cite them. The main such characteristic I will focus on is the scientific activity of the author of i: does it have an indirect influence on the attention, authority, and flow of citations that i receives? Ideally, to answer such a question one would compare the real network of citations with the same, hypothetical, network where a given publication $\hat{a}$ was removed, and all other value and cost coefficients are left unchanged. A feasible alternative to this thought experiment exists: we can compare the citations inflow of *different*, but comparable, papers. Consider two scientific publications i, j of two different authors, with similar observable characteristics, including year of publication and citations flow in the first years after publication, but such that i's author (and *not* j's author) publishes another work after x years, as shown in Figure 3. According to

Hypothesis H0, we should expect the number of citations for i *after* those x years to increase less (Figure 4a) than those for j (Figure 4c), because of a substitutability effect: the reference network of $\hat{a}$ can include all or part of the reference network of i.[11]

Figure 3: Ex-ante comparable papers (same year of publication, same initial flow of citations)

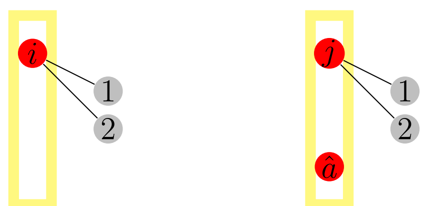

The yellow rectangle visually regroups works by a same author.

Figure 4: Effect of an additional paper $\hat{a}$ by the author of i

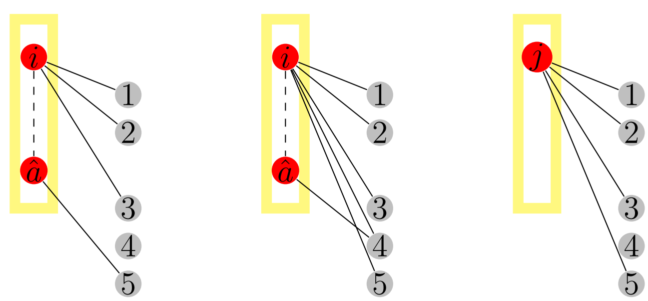

(a) (b) (c)

A new publication $\hat{a}$ by the same author of a paper i should cause citations to i to decrease (a) rather than increase (b) compared to an otherwise similar paper j (c), i.e. because now a link to $\hat{a}$ provides an indirect connection to i.

If instead the number of citations to i tends to increase *more* (Figure 4b) after the publication of $\hat{a}$, we have a clear sign that the choice of linking to a

[11]This is even more true if we make the (realistic but unnecessary) assumption that papers of a same author tend to have overlapping reference networks, or, in the extreme case, cite each other.

node is affected by factors which are not considered by the value/cost model. The above reasoning can be translated into the following regression:

$$cit_{ij}^{T} = \beta_0 + \beta_1 pub_j^{T-1} + \beta_2 cit_{ij}^{\rightarrow T-1} + u_{ij} \tag{1}$$

where cit_{ij}^{T} is a dummy variable denoting whether the ith paper of author j is cited by the paper published at time T, $cit_{ij}^{\rightarrow T-1}$ is the indegree of paper i before time T (that is, the cumulated sum of cit_{ij}^{t} for $t = t_i, \ldots, T-1$), and pub_j^{T-1} is a dummy variable taking value 1 if and only if author j published another paper at time $T-1$. While in the model presented so far there is a one-to-one match between the periods of time and the published papers, so that cit_{ij}^{T} is indeed a dummy variable, the empirical analysis will have to be adapted to the temporal resolution of available data, and cit_{ij}^{T} will hence be a discrete variable *counting* the number of citations to the ith paper of author j at time T.

If the aforementioned "success breeds success" phenomenon only had an interpretation in terms of correlation, we would observe $\beta_2 > 0$, but $\beta_1 \leq 0$: the existing citations are a predictor for the future success of a paper, in terms of citations, while the publication of additional papers has a negative effect, if any. Hence Hypothesis H0 can be translated into the following:

Hypothesis H1: $\mathcal{H}_0 : \beta_1 \leq 0$ vs $\mathcal{H}_1 : \beta_1 > 0$.

Moreover, if recent publications of a same author do eclipse, at least in part, previous ones, we should expect the effect to be even stronger in those cases in which the new paper receive much visibility or prestige. To test for this additional assumption, the following expanded formulation of Equation (1) can be considered:

$$cit_{ij}^{T} = \beta_0 + \beta_1 pub_j^{T-1} + \beta_2 c_{ij}^{\rightarrow T} + \beta_3 infl_j^{T-1} + u_{ij} \tag{2}$$

where $infl_j^{T-1}$ is a proxy for the influence of the journals on which author j published her papers (if any) at time $T-1$, such as a function of their Impact Factor at the time of publication. The prediction would then be the following:

Hypothesis H2: $\mathcal{H}_0 : \beta_3 \leq 0$ vs $\mathcal{H}_1 : \beta_3 > 0$.

4.2 On perfectly reliable links

The hypothesis of *perfectly reliable links*, intrinsic in the basic model of Galeotti *et al.* (2006), and hence in the approach by Haller (2012) which the present study builds upon, is at odds with both the intuition and the evidence coming from the citations network (for instance because, under such hypothesis, networks formed by rational nodes would never contain cycles). This assumption is however not by chance widespread in the literature on social networks: any alternative requires complicated and arbitrary formalizations of the *dispersion* of value along the paths, such as imposing a maximum path length through which value can flow, or a progressive decay. The aim of this paragraph is not to convince the reader that perfectly reliable links are a realistic approximation of the flow of information "over" bibliometric networks, but rather to show that the hypotheses formulated above remain valid under imperfect value flows. The fundamental assumption of the present approach is that citations toward some node i only depend on the private value it has for the subsequent nodes, and on its *reference network*. Now, the already mentioned substitution effect could loose relevance due to dispersion, but not reverse it sign: pushing the imperfection to its limit, if all paths of length higher than 1 conveyed *no benefit* to nodes, the publication of $\hat{a}$ would have no influence whatsoever on the citation flows to i, and β_0 would eventually become 0.

5 Data

Citation databases have existed at least since the 1961 *Science Citation Index* produced by Eugene Garfield. Today, the most prominent ones for bibliometric studies are ISI Web of Knowledge (which is the evolution of Garfield's database) and Elsevier Scopus.[12] However, for the present exercise the possibility to analyze the *whole* network - or some large and densely connected subset of it - is crucial. Hence, the analysis was ran on a specialized database, provided by the American Physics Society to researchers in an aggregate form. The APS database (available at `http://publish.aps.org/datasets`) contains metadata for all articles published over the years on any of the 11

[12]Google Scholar is getting much momentum, but compares substantially worse than the alternatives for what concerns the filter on what is considered "scientific publishing": this is due to the automatized process which is behind it, which has been shown to be easily manipulable (as exemplified by the extreme case of Labbé, 2010).

highly popular journals of the society, as well as a record of the citations among them, making it possible to reconstruct a relevant bibliographic network: overall, it covers 464.817 articles and 4.710.547 of citations among them. The metadata, provided in XML format, allows to extract, for each article, the names of the authors and the date and journal on which it was published. Citations are instead reported in a single file containing two *doi* identifiers per line: the first corresponds to the paper which contains the citation, the second to the cited paper. Given its quality and availability, the APS database has been the subject of many studies on the social aspects of scientific research (see for instance Radicchi *et al.* , 2009, Deville *et al.* , 2014; the website at `http://www.physauthorsrank.org` provides a tool entirely based on such dataset). Although the network it represents is truncated, in the sense that it misses citations to and from papers published on other journals, the journals of the APS are not simply among the most important (although clearly a small minority) in the realm of physics: they also represent a sort of *ecosystem* (each paper having an average of 10 citations *to other papers of the sample*) inside of which it is possible, as will be clear in Section 6, to detect social effects. Since the data covers a specific research discipline, it can also be expected that the noise due to authors homonyms (the APS database does not help in distinguishing researchers with the same name and surname) is lower than when running a search on generalist bibliographic databases. Still, it is worth observing that the effect of homonyms, if any, will be to *dilute* the "author effect", and hence the results will be conservative.

In light of the availability of a large amount of high quality data, and of the possibility on the other hand that the characteristics of interest may be changing with time, it is possible, and convenient, to decompose the analysis in intervals of time. While the definition of *a period* will simply be "a month", the analysis itself will focus on data from a single year, in order to have an optimal tradeoff between temporal localization and samples numerosity. For all papers published in 1990, Equation (1) will be estimated on the 120 months following the publication - that is, approximately until 2000. Sensitivity tests will then be ran for the following years, until 1999 - since the data ends in 2010, it would not be possible to analyze the ten years *following* the publication for works published in the new millennium.

Table 1 reports some descriptive statistics for the sample of interest. The number reported as "Total" of publications (841207) is much larger than the number of articles published in 1990: this reflects the fact that (here, as in the rest of the analysis) multiple authored papers are counted once for each author . As can be observed, and expected, the distribution of citations per month is very uneven, as well as very skewed, with an average of 0.8 but a median of 0 and a maximum of 11: even more so the number of subsequent publications by coauthors. This can be better observed in figures 6 and 7, which could be compatible with not just citations (as suggested by Redner, 1998), but also publications being distributed according to a *power law* distribution.[13] A precise statistical characterization of such empirical distributions is out of the scope of the present work: suffice to say that the growing literature on inference with fat tailed distribution recommends caution before ruling out alternative distributions such as the log-normal or the stretched exponential (both Laherrere & Sornette, 1998 and Clauset *et al.* , 2009 analyze, among many others, data sets on citations, respectively at the author and paper level, and do not find strong support for power law distributions).

The large amount of zero observations could raise some suspect of zero inflated data. This can be ruled out for two reasons. From the theoretical point of view, there is no obvious rationale suggesting that some categories of papers should be *excluded* from the realm of "citable" papers:[14] well on the contrary, there is widespread evidence that even when considering a single journal issue (Campbell, 2005) a huge heterogeneity in the number of citations received by individual articles can be found. From the empirical point of view, Figure 6 shows that the curve describing the frequency of citations per month is remarkably smooth at least in the range from 0 (included) to 6, exhibiting no discontinuity between 0 and 1.

In order to test Hypothesis H2, a proxy for the impact, or prestige, of a given publication is needed. Although it is possible, for recent years, to use the well known Impact Factor as a rough measure of the gain in visibility from publishing on a given journal, the characteristics of the database make it preferable to develop the analysis using only data coming from it, in order to avoid confounding errors due to the "truncatedness" of the network.

In particular, the choice of relying only on APS data neutralizes the risk of the analysis being distorted by important citation flows coming from

[13]The clear non-monotonicity which can be spotted between 100 and 1000 publications per month is presumably a sign of the very peculiar publication habits of large teams of experimental physicists, which will be considered again in Section 6.3.

[14]Compare with the observation of Baccini *et al.* (2014) that, at the *researcher* level, zeros could signal *non-active* individuals rather than a mere failure to attract citations.

Table 1: Descriptive statistics

	Observations	**Total**	**Min**	**Median**	**Mean**	**Max**
Citations	9121	89559	0	4	9.82	326
Citations (non-self)	9121	73255	0	3	8.03	308
Publications	9121	841207	0	21	92.23	30751
Cit./month	1094520	89559	0	0	0.08	11
Cit./month (non-self)	1094520	73255	0	0	0.07	11
Pub./month	1094520	841207	0	0	0.77	1173

Note: Data on publications from year 1990, observed over the 10 years after publication. The unit of observation is the individual work, "*Publications*" are subsequent works by same author(s). A citation is "non-self" if the set of authors of the citing paper and of the cited paper are disjoint.

Figure 5: Average flow of citations during the first 10 years after publication.

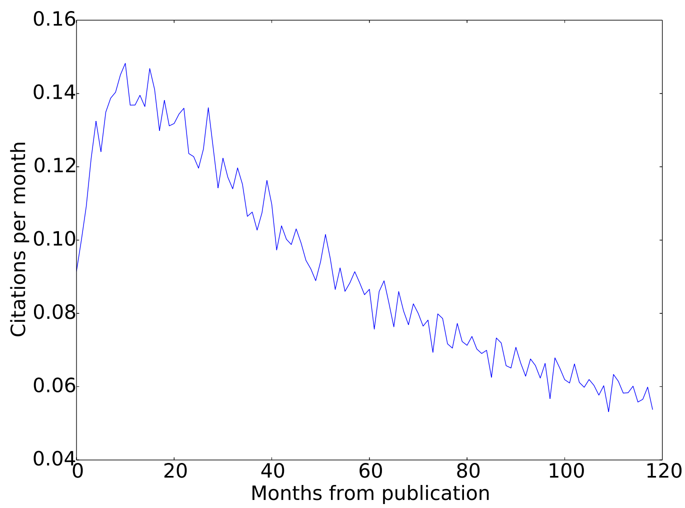

Figure 6: Frequency distribution of citations per month per paper.

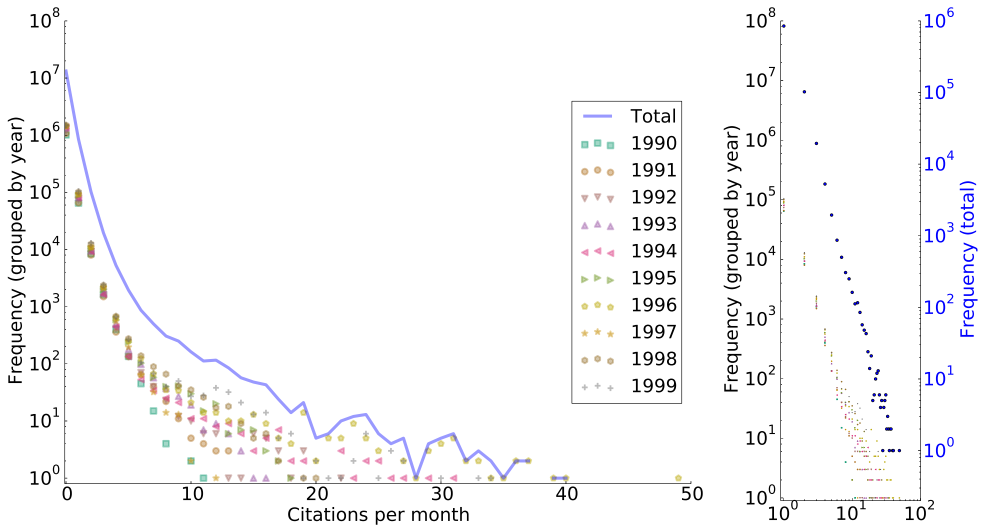

Note: frequencies of citations per month in the time span 1990-1999, plotted as a semilog plot including the origin of the x axis (left), and as a log-log plot (right).

other disciplines (i.e. engineering, mathematics), which could have a different response to the publication of a new work in the realm of physics (i.e. because different theoretical works of a same author can have drastically different practical applications). From a more pragmatic point of view, the ISI Web of Knowledge only makes Impact Factor measurements available for years from 2002 onwards.

For this reason, an "*Internal Impact Factor*" is reconstructed; for each of the 11 journals appearing in the database, and for each year considered in the analysis, it is calculated analogously to the well known Impact Factor. Given the set A_j^y defined as containing the articles published on journal J in the two years before y, we have hence

$$IIF_J^y = \frac{\text{number of citations in year } y \text{ to articles in } A_j^y}{\text{number of articles in } A_j^y}.$$

Finally, the influence of a publication will be proxied with the IIF for the journal on which it is published, in the year it is published. This does not introduce endogeneity issues, since the citations *to* a given publication do not enter in the calculations for the IIF at the date of publication.

Figure 7: Frequency distribution of publications of an author of a given paper, in a given subsequent month.

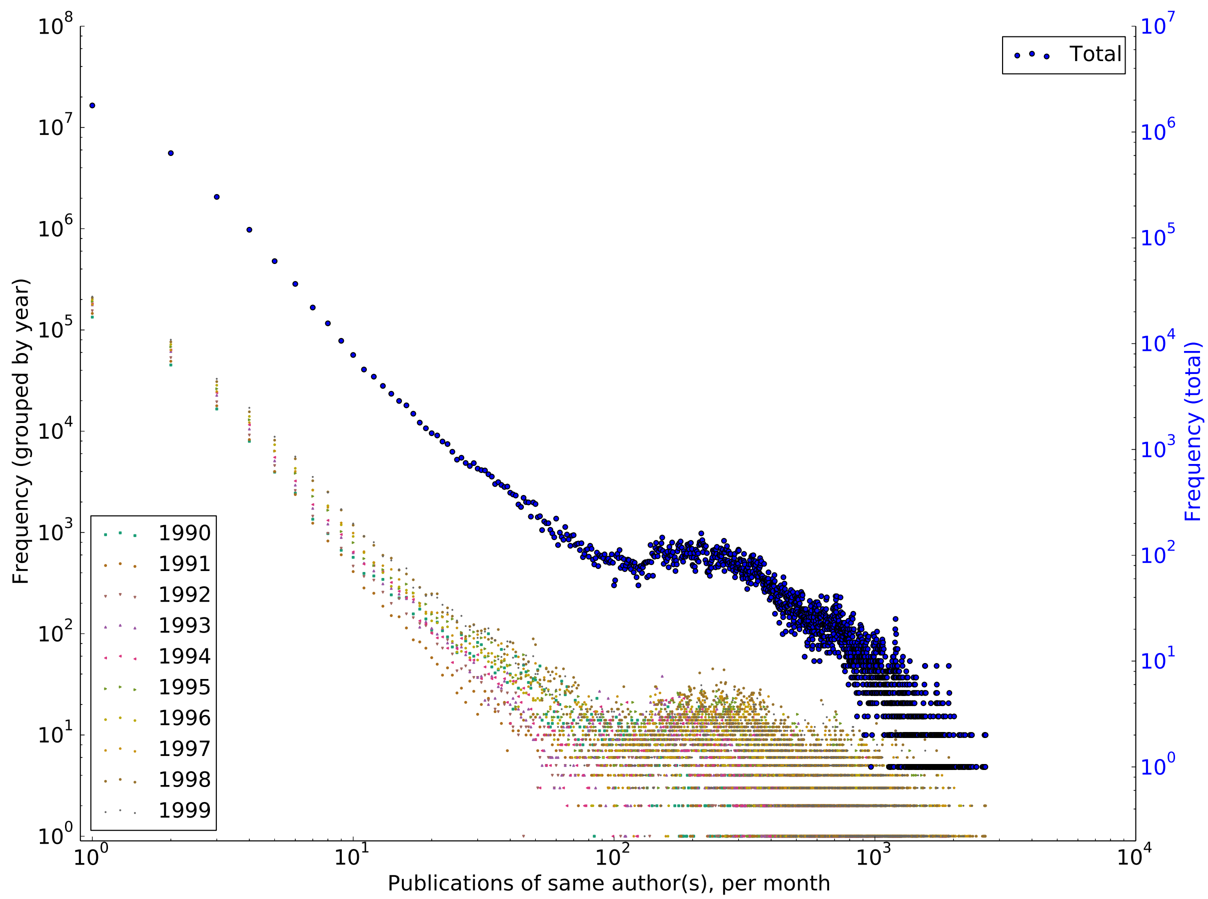

Table 2: Main results

	Dependent variable: cit				
	all	avg(cit)<1	max(cit)<2	avg(pub)<1	max(pub)<2
Cit/month (pre)	1.328***	1.346***	1.489***	1.330***	1.291***
	(0.005)	(0.005)	(0.007)	(0.005)	(0.010)
Pub. in past 6 m.	0.080***	0.078***	0.051***	0.083***	0.103***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
R^2	0.078	0.073	0.069	0.076	0.063
Observations	1085399	1081710	778141	1000195	272153

Note: Results of the two-way OLS estimation. "all" refers to all papers published in 1990; the second and third columns report estimates calculated only using those papers with average (respectively: maximum) number of citations per month lower than 1 (respectively: 2); the fourth and fifth columns are constructed analogously but filtering instead based on the number of publications per month of a same author.

6 Results

Table 2 shows the result of a two-way implementation of Equation 1 ran on different subsets of the data, for papers published in the year 1990. For each paper p considered, and each month T in the 10 years after the paper's publication, the number of citations *per month* received by p *at time* T is regressed on the number of citations per month received until then, and on a dummy variable which takes value 1 if and only if one of the authors of the papers published a paper in the six months *before* T. Time effects are defined based on the number of months elapsed since the publication of a paper (rather than on the effective month of the year), in order to account for the possibility of intrinsic "age effects" affecting the life cycle of a paper. On top of that, the two-way approach involves paper-specific fixed effects, which are strongly recommended given the high level of heterogeneity among papers.

As expected, the *past* citations flow seems to be the most important predictor, since it proxies the impact of the publication. But what is most interesting for the current analysis is the coefficient for "Pub in the last 6 months": its interpretation is that, on average, if a new paper gets published, each previous paper of a same author gains between 0.05 and 0.1 citations per month in the following six months. While the statistical significance is evident, the economic relevance of its estimated value must be judged in light of Table 1, showing that the average number of citations per month is 0.08. In other words, on a random article published in 1990, this effect would consist at least temporarily in an *increase* between 63% and 125%. This allows us to state the following.

> **Result 1:** The null hypothesis of H1, $\beta_1 \leq 0$, is rejected: when new articles are published, citations flows to previous ones by the same authors *increase*.

6.1 Interpreting the effect

The work of researchers is characterized by large fixed costs in terms of exploring new fields, and their literature: this often causes strong forms of specialization and an imperfect knowledge of the existing work done by other scholars. It is hence only natural that self-citations - that is, citations from a paper A to previous work of one or more authors of paper A - are a frequent phenomenon in any scientific discipline, even leaving aside any form of strategic behavior. Their quantitative importance was already evidenced in Table 1, where the "non-self" variables are defined net of self-citations: the difference is around 18%. For comparison, Aksnes (2003) reports a rate 26% for a sample of physics articles published between 1981 and 1996.[15] Self-citations can be expected to explain some of the value of β_1: if an author publishes in a given month, this is, all else equal, a sign of high productivity, and more productive authors will likely have more occasions to cite previous papers of them. While a cumulated advantage process driven uniquely by self-citations would still be at odds with a purely value-based interpretation of citations and bibliometric indicators (and specifically, go against the null of Hypothesis H1), it would largely alter the policy implications, since it is relatively easy to consider indicators which discard them (although the benefits of doing so are debated). This justifies a specific analysis of "non-self" citations, which can be of minor interest for what concerns the accuracy of bibliometric measures, but helps gaining a better understanding of the mechanisms at work. The first column Table 3 features the results of estimating Equation (1) discarding self citations: it shows that indeed the previously found estimate of β_1 is largely driven by self-citations, but that even discarding them,

[15] Given the numerous differences in the sample selection, it is unclear to what extent the difference can be attributed to different citing behaviors.

Table 3: Results of different specifications of Equation (1).

	Dependent variable:				
	cit (n.s.)	cit (n.s.)	cit	cit (n.s.)	cit
	all	max(pub)<2	all	all	ISI dataset
Cit/period (pre)			1.328^{***}		0.623^{***}
			(0.005)		(0.023)
Cit/per. (pre, non-self)	1.358^{***}	1.309^{***}		1.358^{***}	
	(0.005)	(0.011)		(0.005)	
Pub. in past 6 periods	0.003^{***}	0.004^{*}	0.090^{***}	-0.000	0.209^{***}
	(0.001)	(0.002)	(0.002)	(0.001)	(0.032)
Max. IIF in past 6 m.			-0.008^{***}	0.002^{**}	
			(0.001)	(0.001)	
R^2	0.064	0.053	0.078	0.064	0.257
Observations	1085399	272153	1085399	1085399	2596

Note: Dependent variable "cit": all citations; "cit (n.s.)": only non-self citations. First four columns: APS data, the temporal unit of observation (a *period*) is the month. Last column: ISI data, the temporal unit of observation is the year.

the effect is highly significant (the coefficient of 0.003 still corresponds to a percentage increase of 3.75% in the number of citations per month).

Most importantly, the modeling of the network of citations in Section 4 was based on the assumption that the idiosyncratic values reflect some *objective* value, and that citation flows reflect such idiosyncratic values. The natural interpretation of $\beta_1 > 0$ is that citation flows also reflect other factors - namely, that new papers getting published alter the citation flows of previous ones. This is not, however, the only possible interpretation: it could be that papers by highly skilled researchers, who write on average high quality papers, tend to initially get relatively low amounts of citations. Still, since their works are indeed of high quality, they sooner or later get noticed, and cited; and since the authors are skilled, they manage to get other papers published. In other words, β_1 could be detecting a downward bias of citations in the first periods after the publication of a scientific work, rather than an upward bias after the publication of a new work by the same author(s).

Although such effect is plausible, we can verify that it is not the only one driving the results, by implementing a *regression discontinuity design*. Figure 8 features the frequency distribution of the *distance between a publication and*

Figure 8: Distribution of citations in time

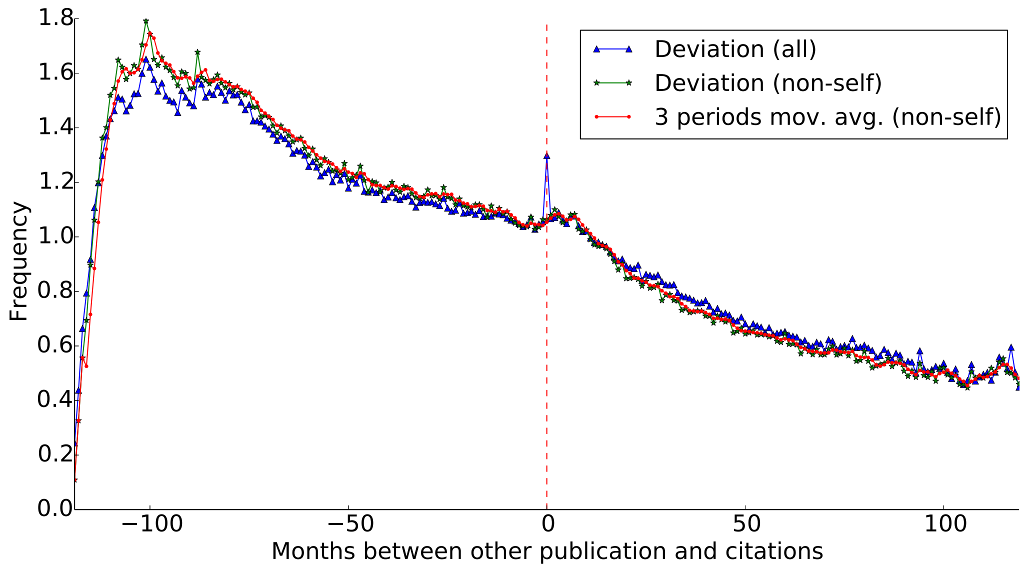

Note: Frequency of citations, as a function of the distance from the publication of a new work by a same author (time 0). The lines represent the normalized difference from the prediction of a null model with citations distributed homogeneously in time.

citations to previous publications of the same author, as deviations from a null model with citations distributed homogeneously in time. After time 0 - that is, the month of the publication - there is an evident and unanticipated increase in the flow of citations. When considering all citations, the effect is large and immediate (the spike corresponding to 0 is mostly due to the large number of cases in which a new paper cites a previous one by a same author), but the increase is evident also when discarding self-citations. The figure can also be interpreted as a difference-in-differences model, comparing the citation increases to "ordinary" papers with those to papers the authors of which just published a new work.

6.2 Journals prestige

Since the null $\mathcal{H}_0$ of Hypothesis H1 was unambiguously rejected, the enriched model 2 potentially introduces two countervailing effects. On one hand, assuming the "getting noticed" effect is at work, a paper with a high impact will presumably have a stronger influence in getting the author to be cited. On the other hand, if a new paper gets more attention than the previous, it may in part *eclipse* it, effectively reducing its flow of citations.

When the measure of journals influence IIF (see Section 5) is introduced as explanatory variable, the results do in fact depend on the exact specification, as summarized in the third and fourth columns of Table 3. If all citations are taken into account, having published on a better journal in the last 6 months tends to *decrease* the flow of citations to previous papers, as postulated in Hypothesis H2. But the results change radically if self-citations are discarded: in this case, the estimate of β_2 is positive, while the estimate of β_1 is no more significant.

> **Result 2:** the null hypothesis of H2 is valid (and $\beta_2 < 0$) only if self-citations are taken into account; otherwise, the prestige of a journal is correlated with an *increase* of citations to previous works.

6.3 Sensitivity tests

The high heterogeneity characterizing the data was already mentioned in Section 5. In particular, the small set of papers with extremely high citations and/or publications of the same author could be to due very specific kinds of research organizations (in the field of high energy physics, for instance, publications with hundreds of authors are not uncommon), and even to homonyms. Columns 2 to 5 of Table 2 report the result for the main model ran only on sets of "low volume" papers, defined in terms of citations or publications of a same author. The results are in line with the full estimation, in some cases actually increasing the estimate of β_1. The lowest value, 0.05, is obtained when restricting to a maximum of 1 citation per month - a criterion which can be easily falsified even by "unexceptional" publications. In general, those sensitivity tests show that the results are not driven by exceptional papers or scientific practices.

All results presented in Table 2 come from data about papers published in 1990 only. In order to enhance the external validity of such results, the analysis was repeated for each year between 1990 and 1999. Figure 9 shows that the estimation for the relevant coefficient of the base model is always significantly different from 0 - and fairly stable across time. The available

Figure 9: Evolution of β_1 over the years.

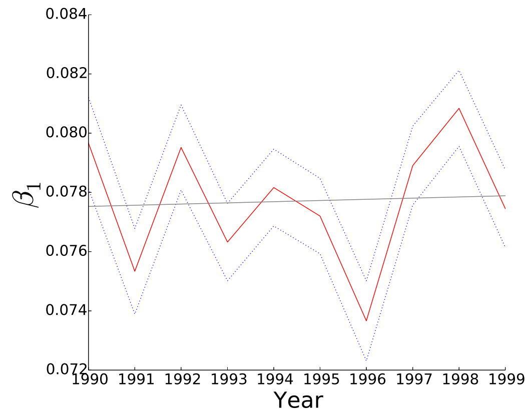

The dashed lines represent the 95% confidence intervals, the solid straight line is the best linear fit.

data, which ends in 2010, does not allow to run the same analysis for years after 1999.

Figure 10 investigates the impact of new publications on citation flows as a function of months elapsed since the original publication. Notice that the model is ran on cross sections, so it is a simple OLS rather than the two-way model presented in Table 6: this implies that the magnitude of the coefficient is not directly comparable, and that the result could be deeply affected by other confounding factors. All that said, the fact that the coefficient is clearly decreasing in time suggests that the "getting noticed" (or self-citing) effect is particularly strong in the few months after a paper has been published.

Finally, in order to ascertain the sensitivity of the results to the discipline under analysis, the same analysis was applied to a smaller sample of papers in the field of Economics, extracted from ISI Web of Science. The criterion adopted in order to have a meaningful sample for the present analysis was the following: all papers published in 1995 on the Scandinavian Journal of Economics (chosen because of its middle range status in terms of Impact Factor) were considered (those are the units of observation), and for each of them the flow of citations was retrieved, as well as the flow of publications of the same authors. The resulting network, in the scope of the present analysis, is almost free from truncation problems[16]. Because of the much lower numerosity of

[16]The ISI Web of Science database is obviously truncated, but in a milder way than the

Figure 10: Variation of β_1 over the life of a publication.

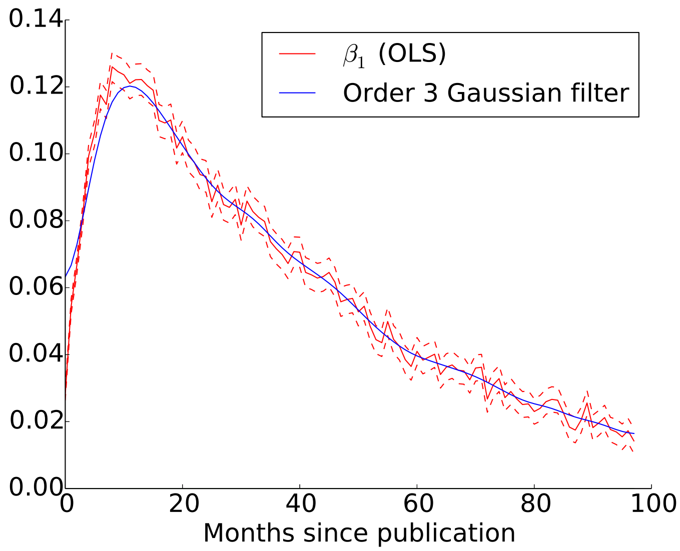

Result of cross-sectional OLS versions of Equation (1), as a function of the number of months elapsed since publication (in dashed lines, the 95% confidence intervals).

the data, however, the whole year was adopted as fundamental unit of time, rather than the month, and the estimation was made with cross-sectional fixed-effects only, without time effects. The result is reported in the last column of Table 3: the estimated effect (0.21) is actually much stronger than what measured so far. Although its entity is not directly comparable, due to the different design, (in particular, the APS journals are presumably more homogeneous, in terms of impact, than the Economics journals appearing in the ISI data set), it provides a strong signal that the cumulative advantage process being observed is a phenomenon not restricted to the field of physics.

7 Conclusions

Determining the extent to which bibliometric indicators, and in particular citation flows, reflect some intrinsic quality of research is an important task. This paper approached the problem via a novel framework for the study of network formation under constraints. Constraints can be *positive* (sets of links appearing in any possible network) and *negative* (sets of forbidden links), and can change in time, depending on the choices of agents. Some theoretical results by Haller (2012) concerning the existence of Nash networks were generalized to the presence of negative constraints. The model was then specialized in order to study the formation of the network of citations under the assumption that citations only "point towards quality", and the resulting predictions were tested against bibliometric data. The hypothesis by which citation flows truthfully reflect some kind of intrinsic value of cited papers was rejected, showing that instead the influence of *environmental factors* on them is relevant. In particular, it was shown that *after* the publication of a scientific paper, there is an immediate raise in the flow of citations to previous publications by the same author(s), even when discarding self-citations. The status of the journal in which the new work is published, measured with a specialized version of the Impact Factor, has different consequences depending on whether self-citations are considered: if they are discarded, it implies an increase in the flows of citations to previous papers. Although the intuition behind such results is far from being new, the novelty of the approach lies in its ability to distinguish a *causal* effect of environmental factors which goes beyond anecdotal evidence, or mere correlations: thus it is an important instrument for science policy.

APS database.

The methodology adopted captures a possibly very limited portion of the environmental factors which shape citation flows (for instance, the effect on the popularity of *subsequent* papers is presumably much larger, but impossible to measure due to endogeneity issues), but it can be applied to any given network of citations (for instance, on networks regrouping works in the same disciplines, or in the same period) and give an indication of how the importance of such factors varies with the context.

Finally, the theoretical model can be specialized to study many different kinds of social networks, and provide empirical researchers with tools to go beyond what the mere *static* analysis of networks allows to identify.

Appendix A Proofs

In order to prove Proposition 1, let us first introduce the following definition.

Definition 1. *$g \in \mathcal{G}$ is a* Nash extension *for* $\Pi(\mathfrak{g}^+, \mathfrak{g}^-, \cdot)$ *if it is disjoint from* $\mathfrak{g}^+$ *and* $g \oplus \mathfrak{g}^+$ *is a Nash network.*

That is, a Nash extension is the set of non-exogenously given links contained in a given Nash network. We can then state the following simple result.

Lemma 1. *If $\vec{g}$ is a Nash extension for* $\Pi(\mathfrak{g}^+, \mathfrak{g}^-, \cdot)$*, all of its links are bridges for* $\vec{g} \oplus \mathfrak{g}^+$*.*

Proof. Let $i, j \in N$ be such that (i, j) is in $\vec{g}$ and it is not a bridge for $\vec{g} \oplus \mathfrak{g}^+$. By definition of Nash extension, $(i, j) \notin \mathfrak{g}^+$. Then, by removing this link, the connected components (of $\vec{g}_{-i,j} \oplus \mathfrak{g}^+$) remain the same as those of $\vec{g} \oplus \mathfrak{g}^+$. Hence, the benefit of node i, which is only determined by the extent of its component, is unchanged:

$$B_i(\vec{g} \oplus \mathfrak{g}^+) = B_i(\vec{g}_{-i,j} \oplus \mathfrak{g}^+),$$

while the total costs for the node have decreased:

$$C_i(\vec{g} \oplus \mathfrak{g}^+) > C_i(\vec{g} \oplus \mathfrak{g}^+) - c_{i,j} = C_i(\vec{g}_{-i,j} \oplus \mathfrak{g}^+).$$

So finally the payoff for i has increased: $\mathfrak{g}^+ \oplus \vec{g}_{-i,j} \succ_i \mathfrak{g}^+ \oplus \vec{g}$. This implies that $\mathfrak{g}^+ \oplus \vec{g}$ is not a Nash equilibrium, and hence $\vec{g}$ is not a Nash extension. □

Proof of Proposition 1. Consider any $\mathfrak{g}^+$. I will show that, given any g^{m-} composed by m links, if there exists a Nash network g^* for $\Pi(\mathfrak{g}^+, g^{m-1^-}, \cdot)$ with $g^{m-1^-} \subset g^{m-}$ of numerosity $m-1$, then one also exists for $\Pi(\mathfrak{g}^+, g^{m-}, \cdot)$. g^{m-} is $g^{m-1^-} \oplus (i, j)$, for some $(i, j) \notin g^{m-1^-} \oplus \mathfrak{g}^+$. If $(i, j) \notin g^*$, then g^* itself is a Nash network for g^{m-} (i's strategies set having being restricted, and all of the others nodes' ones staying unchanged, the equilibrium is still such), and hence this step is trivial. So let us assume that $(i, j) \in g^*$. The link (i, j) is contained in the Nash extension $g^* \ominus \mathfrak{g}^+$ (since it is by assumption not in $\mathfrak{g}^+$), so Lemma 1 guarantees that it is a bridge in g^*. Two cases are possible:

Figure 11: Link (h,k) replaces link (i,j).

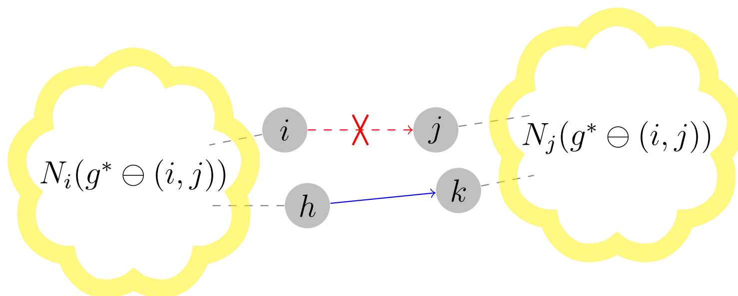

A) there exist h,k in the union of $N_i(g^* \ominus (i,j))$ and $N_j(g^* \ominus (i,j))$ such that $h \notin N_k(g^* \ominus (i,j))$ (see Figure 11), $(h,k) \notin g^{m-}$, and

$$c_{h,k} < \sum_{k' \in N_k(g^* \ominus (i,j))} v_{h,k'}. \tag{3}$$

B) There are no such h,k.

Let us consider case A, and denote as g^{*A} the network $g^* \ominus (i,j) \oplus (h,k)$. For any node $l \in N \setminus \{i,h\}$, the strategies set is unchanged from g^* to g^{*A}, as well as the payoffs. For i, all *available* strategies now deliver a payoff increased by $c_{i,j}$ (since the cost of connecting the two components is now borne by h), so their preference ordering does not change. For what concerns h, since costs are owner-homogeneous, she does not strictly prefer to replace the link (h,k) with a link to any other $k' \in N_k(g^* \ominus (i,j))$. Apart from that, all her available strategies now deliver a payoff decreased by $c_{h,k}$, except the ones where $g_{h,k} = 0$, which however are strictly dominated because of (3).[17] So g^{*A} is a Nash equilibrium.

Consider case B instead, and denote as g^{*B} the network $g^* \ominus (i,j)$. For any node $l \neq i$, the strategies set is unchanged from g^* to g^{*B}, while the payoffs are decreased for all nodes in $N_i(g^* \ominus (i,j))$ and $N_j(g^* \ominus (i,j))$, but the preference ordering over available strategies does not change (except for those connecting the two components, which are however dominated). For what concerns i, in virtue of (3) and of the assumption of owner-homogeneity of costs, we know that (i,j') is in g^{m-} for each $j' \in N_j(g^* \ominus (i,j))$. The

[17] This line of reasoning holds with minimal changes in the case in which $i = h$.

preference ordering of available strategies then clearly corresponds to the preference ordering of the same strategies with the addition of link (i, j), and in particular $g^{*B} = g_i^* \ominus (i, j)$ is optimal. So g^{*B} is a Nash equilibrium.

Whatever the case, the existence of a Nash network for $\Pi(\mathfrak{g}^+, g^{m-}, \cdot)$ was proved, by assuming that one exists for $\Pi(\mathfrak{g}^+, g^{m-1^-}, \cdot)$. The case $\Pi(\mathfrak{g}^+, g^{0^-}, \cdot)$, that is $\Pi(\mathfrak{g}^+, e, \cdot)$, is Proposition 1 by Haller (2012). The result is hence proved for any possible $\mathfrak{g}^-$ by induction.

□

Proof of Proposition 2. Let $\mathfrak{g}^+$ be Pareto optimal. The case $\mathfrak{g}^- = e$ is proved by Haller (2012). When considering $\mathfrak{g}^- \neq e$, the actions set of some nodes is restricted, but the links in $\mathfrak{g}^+$ are left untouched (recall that $\mathfrak{g}^+$ and $\mathfrak{g}^-$ are disjoint). Hence, the empty network is still a strict Nash network, because the preference ordering on available strategies does not change.

Suppose next that some $g^* \neq e$ is a Nash network. The proof develops as in the original result: given some player i with $g_i^* \neq 0$, it must be that g_i^* is a best response against g_{-i}^*. But then g^* is strictly preferred to $\mathfrak{g}^+$ by at least i, while it is at least equally preferred by all other agents. This contradicts the Pareto optimality of $\mathfrak{g}^+$. □

Appendix B Estimation through count data models

The results presented in Section 6 where obtained through two-way OLS estimations. This choice was taken for several reasons: it makes results immediately comparable across different specifications, allows them to be intuitively interpreted as "aggregate flow of spurious citations", and makes them consistent with the RDD approach (Figure 8). On the other hand, this was not motivated by the belief that variables of interest are distributed as normal variables: well on the contrary, this can be excluded by looking at Table 1, or by simply observing that they are positive, integer valued and concentrated in small values.

In order to obtain estimates which can be more easily interpreted at the individual level, Table 4 and Table 5 provide the equivalent of, respectively, Table 2 and Table 3 when estimated through a maximum likelihood estimation under the assumption of *negative binomial* distribution of citations per paper per month, rather than through OLS minimization. It can be observed

Table 4: Main results - negative binomial estimation

	Dependent variable: cit			
	all	avg(cit)<1	avg(pub)<1	max(pub)<2
(Intercept)	1.923***	3.868*	2.073***	1.644***
	(0.271)	(2.330)	(0.341)	(0.565)
Cit/month (pre)	3.327***	4.616***	3.563***	4.535***
	(0.030)	(0.035)	(0.035)	(0.109)
Pub. in past 6 m.	0.738***	0.753***	0.775***	1.089***
	(0.008)	(0.008)	(0.008)	(0.028)
Observations	1085399	1081710	1000195	272153p

Note: Analogue of Table 2, estimated through maximum likelihood assuming a negative binomial distribution, rather than through OLS. The case "max(*cit*) $<$ 2" reduces the dependent variable to a Boolean variable, and was hence excluded due to the insufficient number of degrees of freedom.

that the signs of significant variables are always coherent between the two tables. Still, the changes of significance are worth considering. The coefficient for past publications in the second column of Table 5 is not significant, while it was weakly so in Table 3. The absolute value of the coefficient (0.069) is *larger* than the corresponding coefficient obtained when looking at the whole sample (0.046), which is consistent with the lack of significance being due to data limitations. On the other hand, the significance gets swapped between the "past publications" coefficient to the "Max. IIF" one in column 4. This is consistent with the two variables being strongly correlated.

Table 5: Alternative specifications - negative binomial estimation

	Dependent variable:				
	cit (n.s.)	cit (n.s.)	cit	cit (n.s.)	cit
	all	max(pub)<2	all	all	ISI dataset
(Intercept)	1.498***	1.321***	1.916***	1.498***	−2.540***
	(0.164)	(0.417)	(0.269)	(0.164)	(0.186)
Cit/period (pre)			3.330***		1.697***
			(0.030)		(0.092)
Cit/per. (pre, non-self)	3.419***	4.770***		3.419***	
	(0.037)	(0.125)		(0.037)	
Pub. in past 6 periods	0.046***	0.069	0.916***	0.045**	1.019***
	(0.010)	(0.043)	(0.014)	(0.019)	(0.129)
Max. IIF in past 6 m.			−0.127***	0.000	
			(0.009)	(0.011)	
Observations	1085399	272153	1085399	1085399	2596

Note: Analogue of Table 3, estimated through maximum likelihood assuming a negative binomial distribution, rather than through OLS.

References

Aksnes, Dag W. 2003. A macro study of self-citation. *Scientometrics*, **56**(2), 235–246.

Azoulay, Pierre, Stuart, Toby, & Wang, Yanbo. 2014. Matthew: Effect or Fable? *Management Science*, **60**(1), 92–109.

Baccini, A., & Barabesi, L. 2010. Interlocking editorship. A network analysis of the links between economic journals. *Scientometrics*, **82**(2), 365–389.

Baccini, A., Barabesi, L., Cioni, M., & Pisani, C. 2014. Crossing the hurdle: the determinants of individual scientific performance. *Scientometrics*, 1–28.

Baccini, Alberto. 2010. *Valutare la ricerca scientifica*. Bologna: Il Mulino.

Bala, Venkatesh, & Goyal, Sanjeev. 2000. A Noncooperative Model of Network Formation. *Econometrica*, **68**(5), 1181–1229.

Baldi, S. 1998. Normative versus social constructivist processes in the allocation of citations: A network-analytic model. *American Sociological Review*, **63**(6), 829–846.

Cainelli, Giulio, Maggioni, Mario A, Uberti, T Erika, & De Felice, Annunziata. 2010. *The strength of strong ties: Co-authorship and productivity among Italian economists*. Working Paper. Dipartimento di Scienze Economiche "Marco Fanno".

Campbell, P. 2005. Not-so-deep impact. *Nature*, **435**(77045), 1003–1004.

Clauset, Aaron, Shalizi, Cosma Rohilla, & Newman, Mark EJ. 2009. Power-law distributions in empirical data. *SIAM Review*, **51**(4), 661–703.

de Solla Price, D.J. 1965. Networks of Scientific Papers. *Science*, **149**, 510–515.

de Solla Price, D.J. 1976. A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science*, **27**(5), 292–306.

De Stefano, Domenico, Fuccella, Vittorio, Vitale, Maria Prosperina, & Zaccarin, Susanna. 2013. The use of different data sources in the analysis of co-authorship networks and scientific performance. *Social Networks*, **35**(3), 370 – 381.

Deville, Pierre, Wang, Dashun, Sinatra, Roberta, Song, Chaoming, Blondel, Vincent D, & Barabási, Albert-László. 2014. Career on the Move: Geography, Stratification, and Scientific Impact. *Scientific reports*, **4**.

Dutta, Bhaskar, & Mutuswami, Suresh. 1997. Stable Networks. *Journal of Economic Theory*, **76**(2), 322 – 344.

Fang, Ferric C, & Casadevall, Arturo. 2011. Retracted science and the retraction index. *Infection and immunity*, **79**(10), 3855–3859.

Galeotti, A. 2006. One-way flow networks: the role of heterogeneity. *Economic Theory*, **29**(1), 163–179.

Galeotti, Andrea, Goyal, Sanjeev, & Kamphorst, Jurjen. 2006. Network formation with heterogeneous players. *Games and Economic Behavior*, **54**(2), 353 – 372.

Garfield, Eugene. 1955. Citation indexes for science. A new dimension in documentation through association of ideas. *Science*, **122**(3159), 108–111.

Garfield, Eugene, *et al.* . 1965. Can citation indexing be automated? *Pages 189–192 of: Statistical association methods for mechanized documentation, symposium proceedings.*

Goeree, Jacob K., Riedl, Arno, & Ule, Aljaž. 2009. In search of stars: Network formation among heterogeneous agents. *Games and Economic Behavior*, **67**(2), 445 – 466.

Goyal, S., Van Der Leij, M.J., & Moraga-González, J.L. 2006. Economics: An emerging small world. *Journal of Political Economy*, **114**(2), 403–412.

Gross, P. L. K., & Gross, EM. 1927. College libraries and chemical education. *Science*, **66**(1713), 385–389.

Haller, H. 2012. Network extension. *Mathematical Social Sciences*, **64**(2), 166–172.

Haller, H., Kamphorst, J., & Sarangi, S. 2007. (Non-) existence and Scope of Nash Networks. *Economic Theory*, **31**(3), 597–604.

Herings, P. Jean-Jacques, Mauleon, Ana, & Vannetelbosch, Vincent. 2009. Farsightedly stable networks. *Games and Economic Behavior*, **67**(2), 526 – 541.

Kirchsteiger, Georg, Mantovani, Marco, Mauleon, Ana, & Vannetelbosch, Vincent. 2011 (Feb.). *Myopic or Farsighted? An Experiment on Network Formation.* CEPR Discussion Papers 8263.

Labbé, Cyril. 2010. Ike Antkare one of the great stars in the scientific firmament. *International Society for Scientometrics and Informetrics Newsletter*, **6**(2), 48–52.

Laherrere, Jean, & Sornette, Didier. 1998. Stretched exponential distributions in nature and economy: "fat tails" with characteristic scales. *The European Physical Journal B-Condensed Matter and Complex Systems*, **2**(4), 525–539.

Merton, Robert K. 1968. The Matthew effect in science. *Science*, **159**(3810), 56–63.

Moed, Henk F, & Garfield, Eugene. 2004. In basic science the percentage of "authoritative" references decreases as bibliographies become shorter. *Scientometrics*, **60**(3), 295–303.

Pratelli, Luca, Baccini, Alberto, Barabesi, Lucio, & Marcheselli, Marzia. 2012. Statistical analysis of the Hirsch Index. *Scandinavian Journal of Statistics*, **39**(4), 681–694.

Radicchi, Filippo, Fortunato, Santo, Markines, Benjamin, & Vespignani, Alessandro. 2009. Diffusion of scientific credits and the ranking of scientists. *Physical Review E*, **80**(Nov), 056103.

Redner, Sidney. 1998. How popular is your paper? An empirical study of the citation distribution. *The European Physical Journal B-Condensed Matter and Complex Systems*, **4**(2), 131–134.

Toivonen, Riitta, Kovanen, Lauri, Kivelä, Mikko, Onnela, Jukka-Pekka, Saramäki, Jari, & Kaski, Kimmo. 2009. A comparative study of social network models: Network evolution models and nodal attribute models. *Social Networks*, **31**(4), 240 – 254.

West, J.D., Bergstrom, T.C., & Bergstrom, C.T. 2010. The Eigenfactor Metrics$^{\mathrm{TM}}$: A network approach to assessing scholarly journals. *College & Research Libraries*, **71**(3), 236–244.