

Measuring Influence in Science: Standing on the Shoulders of Which Giants?

Antonin Macé

Measuring Influence in Science: Standing on the Shoulders of Which Giants? *

Antonin MACÉ[†]

November 4, 2017

Abstract

I study the measurement of the influence of scientists based on bibliographic data. I propose a new measure that accounts for indirect influence and allows to compare scientists across different fields of science. By contrast, common measures of influence that “count citations”, such as the h-index, are unable to satisfy either of these two properties. I use the axiomatic method in two opposite ways: to highlight the two limitations of citation-counting schemes and their independence, and to carefully justify the assumptions made in the construction of the proposed measure. (JEL: C43, D85, A11, A12)

1 Introduction

Scientists, the proverb says, are standing on the shoulders of giants. But who are those giants? The metaphor emphasizes that science is a collective endeavor, in which each researcher is *intellectually indebted* to other eminent scholars. This dependency is observable in bibliographic databases, containing the works of various researchers, and the links between these works, as established by the system of references. In this article, I propose to view each scientist as

*I would like to thank Ernesto Dal Bó and Rafael Treibich for useful discussions on this project. All errors are mine.

[†]Aix-Marseille Univ., CNRS, EHESS, Centrale Marseille, AMSE.

a holder of intellectual debt, and to use a bibliographic database to recover the owners of intellectual liabilities. The influence of each scholar is then defined as her total liability, which measures her “gigantism” in the database.

In the proposed measure, each scientist holds a unit of debt, that is divided equally among her articles. The portion of debt of each article is divided in two parts: a fraction $\delta \in (0, 1)$ is owned by the references of the article, and again shared equally, while the remaining fraction $(1 - \delta)$ goes to the (direct and indirect) references of these references. The process is iterated *ad infinitum*. Finally, each scientist owns the liabilities of all of her articles.

I claim that the measure combines two features that are widely recognized as desirable for a measure of influence. First, as the definition above makes clear, the measure accounts for direct as well as *indirect influence* (Palacios-Huerta and Volij, 2004), and this indirect influence is precisely traced back at the level of each article. Second, the measure satisfies a property of *field comparability*: if the database can be divided in two separate fields of science, then the average influence of scientists will be the same in the two fields. This fact quickly derives from the observation that, in such a configuration, the total liability of a field equals its total debt, which is precisely its number of scientists.

The point of view expressed here is at odds with prominent measures of influence that “count citations”, such as the h-index (Hirsch, 2005) or the Euclidean index (Perry and Reny, 2016), to which Section 3 is devoted. I first provide an axiomatic characterization of the class of such measures, that I label *citation-counting schemes*. The result highlights the inability of these schemes to account for indirect influence, and to allow for comparisons of scientists across different fields. I discuss the main remedy proposed in the literature to circumvent this last defect, the idea of field-normalization (Radicchi et al., 2008), and I argue that it leads to serious conceptual difficulties.

The measure of influence described above is formally defined in Section 4, in three consecutive steps. First, a bilateral measure of influence *between papers* is introduced. Second, this first measure is aggregated to define a bilateral measure of influence *between authors*. Third, an influence index is defined for each author. The construction of the measure is justified by

an axiomatic characterization for each aggregation step, followed by a critical discussion of the assumptions made at each step. While the analysis is performed in a simple setting with single-authored papers and without self-citation, I show that these two assumptions can be dropped in a flexible manner.

The present article aims to contribute to the literature on the measurement of influence in science, which started with the introduction of large bibliographic databases in the 60s and 70s. While initial works addressed the measurement of the influence of academic journals (Garfield, 1972; Pinski and Narin, 1976), a more recent literature focused on the measurement of the influence of scientists, following the introduction of the h-index by Hirsch (2005). A sizable part of this literature has been devoted to the introduction and analysis of various citation-counting schemes, proposed as alternatives to the h-index. The axiomatic method has been used to characterize some of these indices, such as the h-index (Woeginger, 2008; Marchant, 2009), the Euclidean index (Perry and Reny, 2016) or the class of step-based indices (Chambers and Miller, 2014), in each case *within the class of citation-counting schemes*. Recently, Bouyssou and Marchant (2016) characterized the fractional citation count¹ within the much broader class of influence measures for authors defined on bibliographic databases. The analysis of Section 3 is developed in a very similar framework, but it provides a characterization of the whole class of citation-counting schemes, rather than that of a particular index.

As common measures of influence create biases across fields, a pervasive theme of the literature on citation metrics is the idea of field-normalization (Ioannidis et al., 2016). For instance, Radicchi et al. (2008) show that, by appropriate normalization of citations in different fields of science, one obtains a distribution of normalized citations per paper that is *universal*. Building on this work, Perry and Reny (2016) introduce the property of *homogeneity* for a citation-counting scheme, which guarantees that re-scaling citations across fields does not alter the relative ranking of scientists within a given field. With regard to this literature, the present article aims at clarifying the notion of comparability across fields. It introduces a property of *field comparability* that is weaker than the concept of field-normalization. Field comparability

¹The fractional citation count is very close to be a citation-counting scheme, in a sense that is made clear in Section 3: it satisfies all the axioms that characterize citation-counting schemes but the Splitting axiom.

requires the average influence to be the same for *disjoint* fields of science, but allows to account for imbalances between related fields. Such imbalances have been empirically documented in the social sciences (Angrist et al., 2017).

Another difficulty of citation metrics has been to account for indirect influence, which is nevertheless decisive because of the cumulative nature of scientific research (Scotchmer, 1991). Palacios-Huerta and Volij (2004) study a measure of influence for journals, the invariant method, for which citations from (endogenously) influential journals weigh relatively more.² This recursive method exploits the fact that journal-citation matrices can be assumed to be irreducible, a property that is likely not to be met by author-citation matrices or paper-citation matrices. Consequently, recursive methods akin to the PageRank algorithm have been proposed for the measurement of the influence of scientists (Radicchi et al., 2009; West et al., 2013) and of papers (Chen et al., 2007). The methods are justified by a process of diffusion of scientific credit in the author-citation network (resp. paper-citation network), that includes a positive probability to be teleported on any given author (resp. paper) at each step. This teleportation component of the diffusion process allows the index to be well-defined, but lacks a proper justification. More importantly, measuring the influence of authors on the basis of an author-citation network creates biases as indirect influence is measured at the aggregate level of authors (Wang et al., 2016). By contrast, the measure proposed here uses the full *paper*-citation network to compute an influence index for *authors*, allowing to trace back indirect influence as precisely as possible, and without creating aggregation biases. In the proposed index, indirect influences of different orders are aggregated through a discounted sum, as suggested by Klosik and Bornholdt (2014) for the measurement of the impact of a single paper. The characterization of the bilateral measure of influence between papers in Section 4.3 provides a justification for this method.

Finally, the article lies in the broader research program of applying the axiomatic method to the measurement and analysis of big data (Patty and Penn, 2015). The axiomatic analysis of the proposed index developed in Section 4 is particularly indebted to the recent work of Bloch et al. (2016). They provide abstract characterization results, that, they show, allow

²An elegant variation on this method is the handicap method proposed by Demange (2014), for which citations from journals citing (endogenously) influential journals weighs relatively more.

to characterize a large class of centrality measures in networks. Suitable adaptations of their results reveal useful for the present analysis.

2 Model

Authors and papers. The set of authors is denoted by \mathcal{A} , and the number of authors is A . Each author $a \in \mathcal{A}$ has written a set of papers \mathcal{P}_a , containing P_a papers. The set of all papers is denoted by $\mathcal{P} = \cup_{a \in \mathcal{A}} \mathcal{P}_a$. The collection of the sets of papers written by each author is denoted by $\mathcal{P}_{\mathcal{A}} = (\mathcal{P}_a)_{a \in \mathcal{A}}$.

Citations and references. Papers in \mathcal{P} are related through a directed network $n : \mathcal{P} \times \mathcal{P} \rightarrow \{0, 1\}$, where $n(p, q) = 1$ means that q cites p . For a given paper $q \in \mathcal{P}$, the set of references is $\mathcal{R}_q = \{p \in \mathcal{P}, n(p, q) = 1\}$, and its number of references is noted R_q . For a given paper $p \in \mathcal{P}$, the set of citations is $\mathcal{C}_p = \{q \in \mathcal{P}, n(p, q) = 1\}$, and its number of citations is noted C_p .

Bibliographic databases. With the above notations, a database is a collection $d = (\mathcal{A}, \mathcal{P}_{\mathcal{A}}, n)$. For the clarity of the exposition, I focus in this article on the set of databases \mathbb{D} such that each paper is single-authored, there is no self-citation, and each author cites at least one other author. Formally:

$$\mathbb{D} = \left\{ d = (\mathcal{A}, \mathcal{P}_{\mathcal{A}}, n) \quad , \quad \left. \begin{array}{ll} \forall a, b \in \mathcal{A}, a \neq b, & \mathcal{P}_a \cap \mathcal{P}_b = \emptyset \\ \forall a \in \mathcal{A}, \forall p, q \in \mathcal{P}_a, & n(p, q) = 0 \\ \forall a \in \mathcal{A}, \exists q \in \mathcal{P}_a, \exists p \in \mathcal{P} \setminus \mathcal{P}_a, & n(p, q) = 1 \end{array} \right\} \right\}.$$

Measures of influence. A *measure of influence for authors* I (or *influence index* I) assigns to any database $d = (\mathcal{A}, \mathcal{P}_{\mathcal{A}}, n) \in \mathbb{D}$ a vector of scores $I[d] \in \mathbb{R}^{\mathcal{A}}$. The number $I[d](a)$ is the influence of author a in the database d , as measured by the index I .

Neutral measures. A measure I is *neutral* if the allocation of influence to authors is independent of their names and of the names of the papers. Formally, neutrality requires that for any bijection of authors $\pi : \mathcal{A} \rightarrow \mathcal{A}'$, for any bijection of papers $\sigma : \mathcal{P} \rightarrow \mathcal{P}'$, if $d = (\mathcal{A}, \mathcal{P}_{\mathcal{A}}, n)$, if $d' = (\mathcal{A}', \mathcal{P}'_{\mathcal{A}'}, n')$ with $\mathcal{P}'_{\mathcal{A}'} = (\sigma(\mathcal{P}_{\pi^{-1}(a)}))_{a \in \mathcal{A}'}$ and $n'(\sigma(p), \sigma(q)) = n(p, q)$ for all $p, q \in \mathcal{P}$, then

$I[d'](\pi(a)) = I[d](a)$ for all $a \in \mathcal{A}$. Neutrality is a very weak property that will be satisfied by all measures considered in the article.

Example 1: the Euclidean index. This index, introduced by Perry and Reny (2016), is defined by:³

$$I[d](a) = \left[\sum_{p \in \mathcal{P}_a} \left(\sum_{q \in \mathcal{P}} n(p, q) \right)^2 \right]^{1/2} = \left[\sum_{p \in \mathcal{P}_a} (C_p)^2 \right]^{1/2}.$$

This index does not account for indirect influence, and it is not comparable across fields. To see the latter point, consider a database divided in two independent fields, with the same number of authors. Assume that each author of the second field can be mapped to an author of the first field with twice as many citations for each paper. The average index in the first field will be twice the one of the second field.

Example 2: a comparable measure. Consider the following measure:

$$I[d](a) = \sum_{p \in \mathcal{P}_a} \sum_{b \in \mathcal{A}} \frac{1}{P_b} \sum_{q \in \mathcal{P}_b} \frac{1}{R_q} n(p, q).$$

As the Euclidean index, this measure is a simple index that only accounts for direct influence, as only direct citations to a 's papers are counted to compute a 's score. However, contrary to the previous example, I claim that this index is comparable across fields. Indeed, we have $\sum_{a \in \mathcal{A}} I[d](a) = A$.⁴ Therefore, if the database is divided in two disjoint fields, the average influence will be the same in the two fields. In Section 4, this index is extended to account for both direct and indirect influence.

Example 3: a comprehensive measure. Consider the following measure:

$$I[d](a) = \sum_{p \in \mathcal{P}_a} \sum_{q \in \mathcal{P}} \sum_{r \in \mathcal{P}} n(p, q) n(q, r).$$

As the Euclidean index, this measure is a simple index that is not comparable across fields.

³To be precise, Perry and Reny (2016) characterize the Euclidean index as an ordinal measure. The formula exhibited here coincides with their recommended cardinalization of the ordinal measure.

⁴Note that $\sum_{a \in \mathcal{A}} I[d](a) = \sum_{p \in \mathcal{P}} \sum_{b \in \mathcal{A}} \frac{1}{P_b} \sum_{q \in \mathcal{P}_b} \frac{1}{R_q} n(p, q) = \sum_{b \in \mathcal{A}} \frac{1}{P_b} \sum_{q \in \mathcal{P}_b} \left(\sum_{p \in \mathcal{P}} \frac{1}{R_q} n(p, q) \right) = A$.

However, the index accounts for some indirect influence, as $I[d](a)$ corresponds to the total number of citations obtained by papers citing a 's papers.⁵

In the sequel, I show in Section 3 that the inability of the Euclidean index to allow for comparisons across fields, and to account for indirect influence extends to all indices that “count citations”. Moreover, these two defects of citation-counting schemes are independent, and each one can be dropped independently of the other, as hinted in Examples 2 and 3. Then, I define a new measure in Section 4 that drops these two deficiencies at the same time, and I carefully justify the construction of the measure.

3 Citations-counting schemes

Following the introduction of the h-index (Hirsch, 2005) to measure the productivity of scientists, a large number of other indices that “count citations” have been proposed. The popularity of these indices probably lies on the parsimony of information used: an author is identified with the collection of the citations counts of each of her papers. This collection can be viewed as a multiset on \mathbb{N} , this is, a function $m_a[d] : \mathbb{N} \rightarrow \mathbb{N}$, defined by $\forall k \in \mathbb{N}, m_a[d](k) = \#\{p \in \mathcal{P}_a, C_p = k\}$. The number $m_a[d](k)$ represents the number of papers written by a and receiving exactly k citations in the database d . The set of all multisets on \mathbb{N} is denoted by \mathcal{M} .

Citation-counting schemes. A measure of influence I is a *citation-counting scheme* if there exists a function $F : \mathcal{M} \rightarrow \mathbb{R}$ such that: $\forall d \in \mathbb{D}, I[d](a) = F(m_a[d])$.

Citation-counting schemes include, but are not limited to, step-based indices (Chambers and Miller, 2014), such as the h-index, and measure-based indices (Palacios-Huerta and Volij, 2014), such as the Euclidean index.

⁵Note that with this crude measure, if a paper cites two of a 's papers, it is counted twice.

3.1 Axioms

This section introduces five axioms on measures of influence. The first axiom applies when we have two disjoint databases $d, d' \in \mathbb{D}$, such that $\mathcal{A} \cap \mathcal{A}' = \emptyset$ and $\mathcal{P} \cap \mathcal{P}' = \emptyset$. In that case, we note $d'' = d \oplus d'$ the union of the two disjoint databases, defined by $d'' = (\mathcal{A}'', \mathcal{P}''_{\mathcal{A}''}, n'')$, where $\mathcal{A}'' = \mathcal{A} \cup \mathcal{A}'$, $\mathcal{P}''_{\mathcal{A}''} = ((\mathcal{P}_a)_{a \in \mathcal{A}}, (\mathcal{P}'_a)_{a \in \mathcal{A}'})$ and $n'' = n \mathbb{1}_{\mathcal{P} \times \mathcal{P}} + n' \mathbb{1}_{\mathcal{P}' \times \mathcal{P}'}$.⁶ The first axiom requires the score of each author to be unaffected by the addition of a disjoint database.

Separability. For any two disjoint databases $d, d' \in \mathbb{D}$, we have:

$$\forall a \in \mathcal{A}, \quad I[d \oplus d'](a) = I[d](a).$$

The following property expresses the fact that the references of an author should not affect her score. Note that this property, while primarily desirable from a normative point of view, also prevents an author from increasing her score by manipulating her references.⁷

Reference Independence. For any author $a \in \mathcal{A}$, we have:

$$\forall q \in \mathcal{P}_a, \forall p \in \mathcal{P} \setminus \mathcal{P}_a, \quad n(p, q) = 0 \quad \Rightarrow \quad I[d' = (\mathcal{A}, \mathcal{P}_a, n + \mathbb{1}_{\{(p, q)\}})](a) = I[d](a).$$

The next axiom states that splitting an uncited paper into two papers with disjoint reference lists should not affect the score of any author. This property limits the possibility of normalizing the source of citations.

Splitting. Let $b \in \mathcal{A}$ and $q \in \mathcal{P}_b$ such that $C_q = 0$. If $\mathcal{P}'_b = \mathcal{P}_b \cup \{q'\}$, with $q' \notin \mathcal{P}$, and n' is such that $\mathcal{R}'_q \cup \mathcal{R}'_{q'} = \mathcal{R}_q$, $\mathcal{R}'_q \cap \mathcal{R}'_{q'} = \emptyset$ and $\forall r \neq q, \mathcal{R}'_r = \mathcal{R}_r$, then:

$$\forall a \in \mathcal{A}, \quad I[d' = (\mathcal{A}, \mathcal{P}'_a, n')](a) = I[d](a).$$

The following axiom requires the score of an author to be independent of the identity its citations. As we shall discuss in Section 3.2, this property prevents the measure to account for

⁶Note that with this definition, $d'' \in \mathbb{D}$.

⁷In the context of the measurement of influence for journals, Kóczy and Strobel (2009) observes that the invariant method is vulnerable to precisely this type of manipulation.

indirect influence at the level of each paper.

Citation Anonymity. For any permutation $\sigma : \mathcal{P} \rightarrow \mathcal{P}$ such that $\forall a \in \mathcal{A}, \sigma(\mathcal{P}_a) = \mathcal{P}_a$, if $\forall p, q \in \mathcal{P}, n'(p, q) = n(p, \sigma(q))$, then:

$$\forall a \in \mathcal{A}, \quad I[d' = (\mathcal{A}, \mathcal{P}_A, n')](a) = I[d](a).$$

The last axiom of this section requires the score of an author a to be independent of the assignment of papers not written by a to other authors. As we shall discuss in Section 3.2, this property prevents the measure to be comparable across fields.

Author Anonymity. For any two databases $d = (\mathcal{A}, \mathcal{P}_A, n)$ and $d' = (\mathcal{A}, \mathcal{P}'_A, n)$ in \mathbb{D} such that $\mathcal{P} = \mathcal{P}'$, for any a such that $\mathcal{P}_a = \mathcal{P}'_a$, we have:

$$I[d' = (\mathcal{A}, \mathcal{P}'_A, n)](a) = I[d](a).$$

3.2 Result and discussion

The main result of this section is the following characterization of citation-counting schemes.

Theorem 1. *A neutral measure of influence I satisfies Separability, Reference Independence, Splitting, Citation Anonymity and Author Anonymity if and only if I is a citation-counting scheme. Moreover, the 5 axioms are independent.*

Theorem 1 clarifies the assumptions that are made when one uses a citation-counting scheme, such as the h-index or the Euclidean index, to measure the influence of authors in a database. On the one hand, such an index satisfies two properties that seem particularly appealing. First, Separability means that the score of an author does not depend on the pattern of citations in a field that is independent from the author's field. Reference Independence implies that the references of an author are not useful to assess her own influence in the database. On the other hand, it appears that citation-counting schemes combine three other independent properties that are problematic for measuring the influence of scientists.

The first problematic assumption, and perhaps the more benign one, is Splitting. Because

splitting any uncited paper does not change the index of any author, this property suggests that papers with more references will be relatively more important in the assignment of credit to authors. For instance, a paper with 50 references would be equivalent to 10 papers with 5 references each. This runs into contradiction with the idea of source-normalization, according to which each source of citations must account for the same level of influence. As an example, Splitting is violated by the fractional citation count (Bouyssou and Marchant, 2016).⁸

The second problematic axiom is Citation Anonymity. When it holds, indirect influence between papers cannot be taken into account for the index. Consider for instance a paper p , cited by a paper q , which is itself cited by a paper r . By Citation Anonymity, one can permute the citations of p without altering the index. It follows that the statement “ p is indirectly cited by r ” is not a meaningful statement for the computation of the index. It seems that this limitation is a serious one, as it implies that being cited by an uncited paper is equivalent to being cited by a paper with 100 citations.

Note that, in all rigor, Citation Anonymity still permits to account for some indirect influence, at the level of authors. Indeed, the permutation of papers used to define the axiom respects the partition of papers into authors, and therefore preserves the global flow of citations between authors. This allows to use network-based methods akin to the PageRank algorithm to measure the influence of authors recursively (Radicchi et al., 2009; West et al., 2013). However, these procedures will inevitably lead to biases in the assessment of indirect influence. Consider for instance an author A, with two papers, one that is uncited and one that is highly cited. If author B gets a citation from A, she will gain a lot of influence in a recursive index (A is a highly cited author), even if she is cited by the uncited paper. Assessing indirect influence at the author level therefore seems too crude if one has access to the citation network between papers.

The last problematic property is Author Anonymity. It implies that the assignment of authors other than a to papers not written by a will not affect the influence of a . As an illustration, a paper with 100 citations will be judged equally influential whether these citations

⁸Fractional citation count: $I[d](a) = \sum_{p \in \mathcal{P}_a} \sum_{q \in \mathcal{P}} \frac{1}{R_q} n(p, q)$.

all come from the same author or from 100 different authors. This seems disputable from a normative perspective, and also raises the issue of the possibility of manipulating the measure. But more importantly, this single axiom prevents the measure of influence to be comparable across fields, as discussed in the next section.

3.3 Comparisons of authors across fields

The idea that an influence index should be comparable across different fields of science is pervasive in the literature on bibliometrics. Nevertheless, common citation-counting schemes are known to produce significantly different results across fields, because of their varying sizes and of their different traditions in terms of publications and citations. As a result, there seems to be a consensus that citations should be normalized by field, so as to allow for meaningful comparisons across fields (Ioannidis et al., 2016). For instance, Perry and Reny (2016) propose to divide the number of citations of a paper by the average number of citations per paper in the paper’s field, before computing the Euclidean index.⁹ In this section, I argue that field-normalization leads to serious conceptual difficulties. I thus propose a weaker notion of *field comparability*. I show that this property is incompatible with Author Anonymity, and is therefore violated by any citation-counting scheme.

There are at least three important issues with the notion of field-normalization. First, field-normalization *creates biases* between fields. Indeed, even if most citations are issued within fields, there are important flows of citations across fields, and these flows may not be balanced. For instance, Angrist et al. (2017) report substantial and unbalanced flows of citations in recent decades across fields in the social sciences.¹⁰ In such circumstances, field-normalization underweights the influence of influential fields, and overweights the influence of fields “under influence”.

Second, field-normalization is sensitive to the level of aggregation retained to perform the

⁹It has also been suggested to normalize a given index by field (Kaur et al., 2013). Perry and Reny (2016) provide a discussion of the relative merits of the two approaches.

¹⁰In Figure 2 of Angrist et al. (2017), one observes that, in 2010, less than 3% of citations from economics journals are given to political science journals, while more than 10% (almost 15%) of citations from political science journals are given to economics journals. Note that, in this figure, citations from a given journal are weighted by the importance of the journal within the journal’s field.

normalization. Consider two authors who obtain the same list of citations counts. Suppose that the first is an economic theorist, while the second is a graph theorist. As papers are more cited in economics than in mathematics, it seems that the second author should be considered more influential, if one admits the notion of field-normalization. At the same time, if one observes that papers are more cited in graph theory than in economic theory, one must also conclude that the first author is more influential than the second one. It follows that field-normalization is an ambiguous notion, that crucially depends on the level at which the normalization is performed (Zitt et al., 2005).

The third difficulty concerns the categorical definition of fields. Depending on the database under consideration, the partition of papers into fields can be obtained either exogeneously by an external observer, or endogenously from the database, thanks to suitable network methods (Copic et al., 2009).¹¹ In any case, the problem arises that many authors and papers lie at the intersection of different fields. This difficulty seems to limit the applicability of field-normalization at the level of the whole database.¹²

To avoid the pitfalls associated with field-normalization, I propose a weaker property of field comparability. The axiom requires the average influence to be equal for two different fields of science, *only in the extreme case in which the two fields are completely disjoint*. With this property, imbalances between interrelated fields are allowed, and the definition of fields is unambiguous.

Field Comparability. Let $d \in \mathbb{D}$ be divided in two disjoint fields: there is a partition $\mathcal{A} = \mathcal{A}^1 \cup \mathcal{A}^2$ such that for all $(p, q) \in (\cup_{a \in \mathcal{A}^1} \mathcal{P}_a) \times (\cup_{a \in \mathcal{A}^2} \mathcal{P}_a)$, $n(p, q) = n(q, p) = 0$. Then:

$$\frac{1}{A^1} \sum_{a \in \mathcal{A}^1} I[d](a) = \frac{1}{A^2} \sum_{a \in \mathcal{A}^2} I[d](a).$$

In the sequel, I prove the incompatibility between this property and Author Anonymity. For that, I introduce the following weak axiom, guaranteeing that the measure is non-degenerate.

¹¹A common strategy is to obtain a categorization of journals into fields with either method, and then to deduce the field of a paper from the journal where it has been published.

¹²For instance, Perry and Reny (2016) restrict their empirical analysis to economists working in a single field.

Null Author.¹³ Let $d \in \mathbb{D}$ and $a \in \mathcal{A}$. We have $I[d](a) > 0 \Leftrightarrow \exists p \in \mathcal{P}_a, C_p > 0$.

We obtain the following impossibility result.

Proposition 1. *There is no measure of influence satisfying Null Author, Author Anonymity and Field Comparability.*

As a corollary of Proposition 1 and Theorem 1, we obtain that citation-counting schemes are not field comparable. It remains possible to perform field-normalization before or after computing such an index, but we have argued that such normalization is inherently arbitrary and potentially biased. On the contrary, I propose in the next section to move beyond the idea of counting citations: I introduce a novel influence index that is comparable across fields.

4 An alternative measure of influence

In this section, I define a new measure of influence, I_δ , that is computed in three consecutive steps. As we look for a measure that is field comparable, we know from Proposition 1 that the measure should not only assess the influence of an author on papers, but also on authors. The index I_δ is thus constructed from a bilateral measure of influence *between authors*, AI_δ . Moreover, we want to account for indirect influence, but without creating the aggregation biases arising when indirect influence is assessed at the level of authors. Therefore, the measure AI_δ is based on a bilateral measure of influence *between papers*, PI_δ . This last measure is built from the network n to combine direct and indirect influence.

For the clarity of the exposition, I define and study the measure of influence I_δ for a database $d \in \mathbb{D}$ such that each paper has at least one reference: $\forall q \in \mathcal{P}, R_q \geq 1$. I explain in Section 4.4 that this assumption is not substantive and can be dropped in a flexible manner.

4.1 Definition of the measure of influence I_δ

Bilateral influence between papers. For two papers $p, q \in \mathcal{P}$, the direct influence exerted by paper p on paper q is defined by $PI_1(p, q) = \frac{1}{R_q}n(p, q)$. This measure is normalized: the

¹³The terminology follows the *Null Player* property from the theory of cooperative games (Shapley, 1953).

total influence exerted on any given paper is equal to 1.

The indirect influence of order $k \geq 1$, exerted by p on q , is defined recursively by $PI_{k+1}(p, q) = \sum_{r \in \mathcal{P}} PI_k(p, r)PI_1(r, q)$. For instance, the influence of order 2 exerted by p on q will be positive if there exists a paper r , cited by q and citing p . The measure is again normalized.¹⁴

Finally, for a given discount factor $\delta \in [0, 1)$, the δ -discounted influence of p on q is given by $PI_\delta(p, q) = (1 - \delta) \sum_{k=1}^{+\infty} \delta^{k-1} PI_k(p, q)$. In this definition, the factor δ reflects the relative importance of indirect versus direct influence.

Bilateral influence between authors. For two authors $a, b \in \mathcal{A}$ and a discount factor $\delta \in [0, 1)$, the δ -discounted influence exerted by author a on author b is defined by

$$AI_\delta(a, b) = \sum_{p \in \mathcal{P}_a} \frac{1}{P_b} \sum_{q \in \mathcal{P}_b} PI_\delta(p, q).$$

The measure is normalized: the total influence exerted on any given author is equal to 1.

Influence. For an author $a \in \mathcal{A}$ and a discount factor $\delta \in [0, 1)$, the δ -discounted influence of a is defined by $I_\delta(a) = \sum_{b \in \mathcal{A}} AI_\delta(a, b)$.¹⁵

How does this new measure relate to the analysis performed in the last section? It is easy to see that the measure is neutral and satisfies the appealing axioms of Separability and Reference Independence. However, the three other axioms are violated. Splitting is violated, as the measure PI_δ is normalized: adding references to an uncited paper does not increase its weight in the allocation of influence. Citation Anonymity is also violated, as the construction of PI_δ accounts for indirect influence at the level of papers. Finally, as AI_δ is normalized, and by definition of I_δ , the property of Author Anonymity is not satisfied.

As we have seen in the previous section, the violation of Author Anonymity opens the door to comparability of authors across fields. Indeed, the measure I_δ satisfies Field Comparability,

¹⁴Note one subtlety in the definition of PI_k : all paths of length k in the network n are not equal. Indeed, the strength of a path depends on the outdegree of the nodes composing the path: the indirect influence of p on q is more important if all papers along the path have a unique reference than if they have multiple references.

¹⁵As the database d is kept fixed in this section, it is omitted in the notations, and $I(a)$ stands for $I[d](a)$.

since we have:

$$\sum_{a \in \mathcal{A}} I_\delta(a) = \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{A}} \frac{1}{P_b} \sum_{p \in \mathcal{P}_a} \sum_{q \in \mathcal{P}_b} P I_\delta(p, q) = \sum_{b \in \mathcal{A}} \frac{1}{P_b} \sum_{q \in \mathcal{P}_b} \left(\sum_{p \in \mathcal{P}} P I_\delta(p, q) \right) = A.$$

In the sequel, I properly justify each step in the construction of the measure I_δ . For this, I introduce six axioms that apply to arbitrary aggregation functions, following the approach proposed by Bloch et al. (2016) to characterize a large variety of centrality measures in networks.

4.2 Axioms on aggregation functions

In the definition of the measure of influence I_δ , each step performs an aggregation, by which a collection of statistics is merged into a single statistic. It is therefore convenient to define axioms directly at the level of aggregation functions. Starting from a sequence of statistics $s = (s_l)_{l \in \mathcal{L}} \in [0, 1]^\mathcal{L}$, where $\mathcal{L} \subseteq \mathbb{N}$ is an index set, an *aggregation function* $f : [0, 1]^\mathcal{L} \rightarrow \mathbb{R}$ associates to s an aggregate statistic $S = f(s)$. We define six properties on aggregation functions.

Monotonicity. If for all $l \in \mathcal{L}$, $s'_l \geq s_l$, then $f(s') \geq f(s)$.

Homogeneity. For all $\lambda \geq 0$, $f(\lambda \cdot s) = \lambda f(s)$.

Independence. For any partition $\mathcal{L} = \mathcal{J} \cup \mathcal{K}$, $f(s_\mathcal{J}, s_\mathcal{K}) - f(s_\mathcal{J}, s'_\mathcal{K}) = f(s'_\mathcal{J}, s_\mathcal{K}) - f(s'_\mathcal{J}, s'_\mathcal{K})$.

In the sequel, all aggregation functions that we will consider satisfy these three properties. To simplify the notations, we say that a function f is **regular** whenever it is monotonic, homogeneous and independent.

Symmetry. For any permutation $\pi : \mathcal{L} \rightarrow \mathcal{L}$, we have $f(s_\pi) = f(s)$.

The next two properties only apply to the case $\mathcal{L} = \mathbb{N}^*$.

Recursivity. For all $s, s' \in [0, 1]^{\mathbb{N}^*}$, $f(s'_1, s'_2, \dots) f(0, s_1, s_2, \dots) = f(s_1, s_2, \dots) f(0, s'_1, s'_2, \dots)$.

Recursivity implies that $f(s')/f(s) = f(0, s'_1, s'_2, \dots)/f(0, s_1, s_2, \dots)$, when the two ratios are well-defined. As Bloch et al. (2016) put it, “recursivity requires that the calculation being done based on later stages of the [sequence of statistics] look ‘similar’ (in a ratio sense) to those done earlier in the [sequence of statistics]”.

For the next property, we introduce the set $\mathcal{Z} \subset [0, 1]^{\mathbb{N}^*}$ of sequences that converge to zero.

Formally $\mathcal{Z} = \{z \in [0, 1]^{\mathbb{N}^*} \mid \lim_{k \rightarrow +\infty} z_k = 0\}$.

Long Run. For all $s \in [0, 1]^{\mathbb{N}^*}$, for all $\varepsilon > 0$, there exists a sequence $z \in \mathcal{Z}$ such that $f(s \cdot z) \geq f(s) - \varepsilon$.

Long Run emphasizes that what happens at the very end of a sequence of statistics should not matter too much for the computation of the aggregate statistic.

4.3 Characterization of each aggregation step

This section aims at formally justifying the proposed measure of influence. For each of the three steps, a characterization result is provided, describing conditions under which the aggregation is realized.¹⁶

Bilateral influence between papers. A bilateral measure of influence between papers is a function $PI : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$. It is normalized if $\forall q \in \mathcal{P}, \sum_{p \in \mathcal{P}} PI(p, q) = 1$. The following result characterizes the δ -discounted influence PI_δ .

Proposition 2. *Let PI be a normalized bilateral measure of influence between papers, defined as $PI(p, q) = f((PI_k(p, q))_{k \geq 1})$. Then, f is regular, recursive and satisfies long run if and only if there exists $\delta \in [0, 1)$ such that $PI = PI_\delta$.*

Proposition 2 highlights the assumptions made to define the δ -discounted influence PI_δ . First, the influence between papers is supposed to depend on the normalized influences of various orders. Accounting for all possible orders of indirect influence is a prerequisite here, as we want to avoid the too narrow focus on direct influence, that characterized citation-counting schemes. Moreover, taking as arguments normalized measures of influence is important, as it allows to weight the various orders of influence in a transparent manner.

Second, the important axiom at stake in Proposition 2 is recursivity. It implies that the relative importance of the indirect influence of order $k + 1$ relative to that of order k is independent of k .

¹⁶Ideally, one would want to characterize the class of influence indices that are computed in these three steps. As this class is more complex than the class of citation-counting schemes, I don't attempt to solve this exercise here, but I rather choose to carefully justify each aggregation step.

Bilateral influence between authors. A bilateral measure of influence between authors is a function $AI : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}$. It is normalized if $\forall b \in \mathcal{A}, \sum_{a \in \mathcal{A}} AI(a, b) = 1$.

Proposition 3. *Let PI be a normalized bilateral measure of influence between papers and let AI be a normalized bilateral measure of influence between authors, defined as $AI(a, b) = g((PI(p, q))_{p, q \in \mathcal{P}_a \times \mathcal{P}_b})$. Then, g is regular and symmetric if and only if*

$$AI(a, b) = \sum_{p \in \mathcal{P}_a} \frac{1}{P_b} \sum_{q \in \mathcal{P}_b} PI(p, q).$$

Proposition 3 justifies the construction of the measure of influence between authors, based on that between papers. First, the property of symmetry is imposed, implying that all papers written by b are judged equally important as receivers of influence. Second, the aggregation function is supposed to be regular, and in particular independent. It requires that, when the influence of a paper p , written by a , on a paper q , written by b , increases, the rise in the influence of author a on author b must be independent of the level of influence of p' on q' , for any couple $(p', q') \neq (p, q)$ such that p' is written by a and q' is written by b .

Influence. A measure of influence I is normalized if $\frac{1}{A} \sum_{a \in \mathcal{A}} I(a) = 1$.

Proposition 4. *Let AI be a normalized bilateral measure of influence between authors and let I be a normalized measure of influence, defined as $I(a) = h((AI(a, b))_{b \in \mathcal{A}})$. The function h is regular and symmetric if and only if $I(a) = \sum_{b \in \mathcal{A}} AI(a, b)$.*

Proposition 4 describes several assumptions made when constructing the measure I_δ from the bilateral measure AI_δ . The functional form given by h prescribes that the influence of a only depends on the bilateral influence of a on other authors, and not on the bilateral influence between other authors. This is an important restriction, which prohibits the use of auto-referent measures, such as the invariant method (Palacios-Huerta and Volij, 2004) or PageRank-like methods (Radicchi et al., 2009; West et al., 2013),¹⁷ for which exerting influence on an

¹⁷Note that the invariant method would not be readily available to measure the influence of authors here, as the author-citation matrix would likely not be irreducible, contrary to what can be assumed for journal-citation matrices.

(endogenously) influential author weighs more than on an (endogenously) uninfluential author. In the current model, this assumption is reasonable, as the bilateral measures of influence already account for indirect influence. In fact, for the proposed measure, a citation from an influential author is important only insofar as the citing article exerts influence on other authors, and not because the influential author is deemed more important. As a consequence, the proposed measure tracks indirect influence more precisely than those based on aggregates.

A related assumption is given by the axiom of symmetry. It implies that all authors are judged equally important as receivers of influence. This is a substantive normative assumption, reflecting an egalitarian perspective. We discuss how to adapt this assumption to real databases in the next section.

Finally, the function h is assumed to be independent. This means that increasing the influence of a on a given author b has an effect on the score of a that is independent of the levels of influence of a on authors other than b . This seems to be a natural assumption to obtain a simple measure, although one could prefer to reward a limited influence on many authors more than a high influence on a small group of authors.¹⁸

4.4 Application of the measure to real databases

A first assumption made at the beginning of Section 4 was that each paper in the database has at least one reference. If this is not the case, the measure PI_δ will not be normalized. The bilateral measure of influence between authors can then be adapted as follows:

$$AI(a, b) = \sum_{p \in \mathcal{P}_a} \frac{1}{P'_b} \sum_{q \in \mathcal{P}_b} PI(p, q), \quad \text{with} \quad P'_b = \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{P}_b} PI(p, q).$$

With this modification, the measure is well-defined for any database in \mathbb{D} , and the property of field comparability is still satisfied.¹⁹ I now turn to the application of the measure to databases

¹⁸Under this requirement, an alternative influence index could be defined as $I(a) = \sum_{b \in \mathcal{A}} \frac{(AI(a, b))^\alpha}{\sum_{c \in \mathcal{A}} (AI(c, b))^\alpha}$, with $0 \leq \alpha \leq 1$. This alternative index will keep the property of field comparability.

¹⁹As each author cites at least one other author for $d \in \mathbb{D}$, we have that $P'_b > 0$ for all $b \in \mathcal{A}$. Moreover, field comparability is preserved because AI remains normalized.

that do not belong to \mathbb{D} .

In modern research, a growing number of papers are written by a group of authors, rather than by a single author.²⁰ It is therefore important to adapt the measure to databases with such features, and this can be done flexibly. Suppose that for each paper $p \in \mathcal{P}$, we have a distribution of weights $\omega_p = (\omega_p^a)_{a \in \mathcal{A}}$, reflecting the various contributions of authors in \mathcal{A} to p , and such that $\sum_{a \in \mathcal{A}} \omega_p^a = 1$. For instance, for each co-author a of p , ω_p^a could be defined as the inverse of the number of co-authors, while $\omega_p^b = 0$ for all other authors $b \in \mathcal{A}$ (Radicchi et al., 2009). Then, the bilateral measure of influence between authors can be adapted as follows:

$$AI(a, b) = \sum_{p \in \mathcal{P}} \frac{1}{P_b''} \sum_{q \in \mathcal{P}} \omega_p^a \omega_q^b PI(p, q), \quad \text{with} \quad P_b'' = \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{P}} \omega_q^b PI(p, q).$$

With such a natural modification, the property of field comparability is preserved. Note that in the proposed modification, the shares of co-authors of a paper are exogenous. Alternatively, the influence index could be combined with endogenous sharing rules, for which the relative share of a co-author on a paper depends on its overall influence in the database (Szwagrzak and Treibich, 2017). Note also that if the shares of a given paper did not sum up to one, for instance because co-authorship was not accounted for, then the property of field comparability would be lost, as a field with a higher number of authors per paper would also have a higher average influence index (Waltman and van Eck, 2015).

The model also neglected self-citations, which are pervasive in real bibliographic databases. This issue can be dealt with in two easy steps. First, we erase from the database all authors that cite only themselves. Second, one can adapt the bilateral measure of influence between authors, in such a way that (i) each author holds the same amount of intellectual debt (ii) this debt is allocated only among other authors. Formally,

$$\forall b \neq a, \quad AI(a, b) = \sum_{p \in \mathcal{P}} \frac{1}{P_b'''} \sum_{q \in \mathcal{P}} \omega_p^a \omega_q^b PI(p, q), \quad \text{with} \quad P_b''' = \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{P}} (1 - \omega_p^b) \omega_q^b PI(p, q),$$

²⁰This trend has been documented for science in general by Wuchty et al. (2007) and for economics in particular by Hamermesh (2013).

and $AI(a, a) = 0$. The property of field comparability is still satisfied.

Another important assumption made in the definition of the measure was the symmetry property imposed in the third aggregation step, as emphasized in Proposition 4. In real databases, this would lead to consider a very young scholar as equally indebted as a scholar already having a long academic life. This seems too strong and one could for instance weight each author according to the number of years for which she has published at least one paper. For a given weighting of authors $\alpha = (\alpha_a)_{a \in \mathcal{A}}$, the measure of influence can be modified as:

$$I(a) = \sum_{b \in \mathcal{A}} \alpha_b AI(a, b).$$

This modification preserves field comparability under the assumption that the distribution of weights α is the same for different fields of science.²¹ On the contrary, if one weighted the intellectual debt of each author by her number of publications, the distribution of weights would be highly asymmetrical across fields, and the property of field comparability would be lost.

Finally, there are two more instances of symmetry that could be relaxed. First, as stated in Proposition 3, it has been assumed that all papers of an author have the same importance in the allocation of scientific credit. A natural modification would be to weight papers according to their number of pages. Second, it has been assumed in the construction of PI_1 , and then of PI_δ , that all the references of a paper count equally. This needs not be the case: for instance, if one has access to the text of each paper, references could be weighted according to their number of occurrences in the citing paper.

5 Conclusion

This article has presented a new measure of the influence of scientists in bibliographic databases.

The measure allows for comparisons of scientists across different fields of science and accounts

²¹This may be true if α_a is the number of “time-windows” for which author a has at least one paper, for a well-chosen duration length of the time-window. Note that the classification of authors in fields would then be useful to choose the appropriate time-window.

for indirect influence, traced back at the level of papers. I should mention a number of important caveats of the approach.

Since the objective has been to compute a single number for each scientist, it is clear that the index only provides very limited information on the work of each scientist. Because of the sparse information used to compute the index, it cannot be a measure of the quality of a scientific work. Moreover, note that the influence index is not meant to be an index of productivity, contrary to the h-index, as a scholar may be highly influential with a single piece of work.

While the discussion focused on the measurement of the influence of scientists, the method proposed in this article is flexible and could be adapted to measure the influence of other entities, such as teams of authors, journals, departments, universities, countries, or fields of research. Another range of applications would be to subset the holders of scientific debt and/or the holders of scientific credit, in order to analyze the flow of influence between different fields of science and/or periods of time.

Finally, one potential avenue for further research would be to explore the susceptibility of the influence index to manipulation by a scientist or by a group of scientists. By construction, as the total debt of a given group of authors is fixed, its potential impact on the index is limited. This runs opposite to what happens for citation-counting schemes, for which the leverage of such a group seems unbounded.

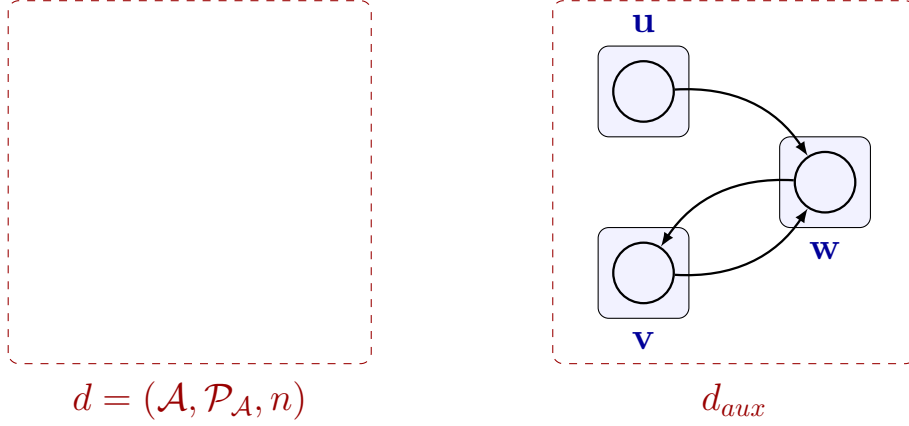
6 Appendix

6.1 Proof of Theorem 1


We show a first lemma, establishing that Author Anonymity extends to the case where the number of authors change when we reallocate authors to papers.

Lemma 1. *If a neutral measure of influence satisfies Separability and Author Anonymity, then it satisfies Extended Author Anonymity: for any databases $d = (\mathcal{A}, \mathcal{P}_{\mathcal{A}}, n)$ and $d' = (\mathcal{A}', \mathcal{P}'_{\mathcal{A}'}, n)$ in \mathbb{D} such that $a \in \mathcal{A} \cap \mathcal{A}'$ and $\mathcal{P} = \mathcal{P}'$, we have $I[d'](a) = I[d](a)$.*

Proof. For $\#\mathcal{A} = \#\mathcal{A}'$, the result is a direct consequence of neutrality and Author Anonymity. To show the claim, it is sufficient to prove that it holds for $\#\mathcal{A}' = 1 + \#\mathcal{A}$. By neutrality, we may even assume that $\mathcal{A} \subset \mathcal{A}'$, so that we have $\mathcal{A}' = \mathcal{A} \cup \{u\}$, with $u \notin \mathcal{A}$. The strategy of the proof consists in adding an auxiliary database d_{aux} to d , and in applying Author Anonymity to $d \oplus d_{aux}$.



The auxiliary database d_{aux} contains three authors u, v, w , and each author is represented by a blue (rounded) square. Each author has a single paper, represented by a circle. Each citation is represented by an arrow, for instance u 's paper cites w 's paper. By Author Anonymity, we can reallocate authors to papers, while keeping the set of authors, the set of papers and the network fixed. We obtain the following database.



In the database d'_{aux} , there are only two authors v and w , and v has written two papers. As

we have $d_{aux}, d'_{aux} \in \mathbb{D}$, we can write, using Separability and Author Anonymity:

$$I[d](a) = I[d \oplus d_{aux}](a) = I[d' \oplus d'_{aux}](a) = I[d'](a).$$

□

We now prove Theorem 1.


Proof. It is clear that any citation-counting scheme is neutral and satisfies the five axioms. Conversely, let I be a neutral measure of influence satisfying the five axioms. Let $d = (\mathcal{A}, \mathcal{P}_{\mathcal{A}}, n)$ be a database in \mathbb{D} , and let $a \in \mathcal{A}$. We will show that the influence of a in d , $I[d](a)$, equals the influence of a in any database $d' = (\mathcal{A}', \mathcal{P}'_{\mathcal{A}'}, n')$ such that:

- $\mathcal{A}' = \{a, b\}$
- $\mathcal{P}'_a = \mathcal{P}_a$ and $P'_b = \sum_{p \in \mathcal{P}_a} C_p$ (b has written a number of papers equal to the total number of citations of a in d)
- $\forall p \in \mathcal{P}_a, C'_p = C_p$ (each of a 's papers keeps the same number of citations), while $\forall q \in \mathcal{P}'_b, R'_q = 1$ (each of b 's papers has a unique reference)
- $\exists (p, q) \in \mathcal{P}'_b \times \mathcal{P}'_a, n'(p, q) = 1$ (there is at least a citation from a to b).


Note that the last line is required to have $d' \in \mathbb{D}$. Moreover, the identities of the papers from a citing b , as well as these of the papers cited by a , should not matter for the score of a , because I satisfies Reference Independence. Therefore, all such databases give the same score to a (note that, by neutrality of I , the names of papers in \mathcal{P}'_b do not matter).

Let $b \in \mathcal{A}$ be an author citing a in the database d : there exists a paper $p_b \in \mathcal{P}_b$ with $\mathcal{R}_{p_b} \cap \mathcal{P}_a \neq \emptyset$. We represent graphically a portion of the database d containing p_b and all the references from p_b to papers written by a .²² We add an auxiliary database with two authors, u and v , as we know by Separability that this does not affect the score of a .


²²In the graphical representation, b has written a single paper, although it will be clear that the construction equally applies if b has written more than one paper.




Then, by Author Anonymity, we can modify the database to transfer the paper with no citations and no references from u to b .




By Citation Anonymity, we can permute the identities of the citations originating from b . We modify the network so that the citations from b now originate from the uncited paper that has been taken from u .



Note that the paper from b citing papers from a has no citation. We can thus apply Splitting, to split this paper into three: one that cites a 's upper paper, one that cites a 's lower paper, and one that cites other papers (not written by a).



Then, by applying Author Anonymity, we can isolate the papers from b citing a , by having v absorbing the two other papers written by b .



Finally, let's observe that we can iterate this construction for any paper citing a paper written by a . Moreover, we can then merge all the authors citing at least one of a 's papers to a single author (say b), by application of Extended Author Anonymity (the property holds by Lemma 1). Following this method, we obtain a database whose papers can be partitioned in three parts: the papers written by a ($\mathcal{P}^1 = \mathcal{P}_a$), the papers citing a (\mathcal{P}^2), all being written by the same author b , and all the other papers (\mathcal{P}^3). The main observation here is that there is no citation from \mathcal{P}^3 to either \mathcal{P}^1 or \mathcal{P}^2 . There is no citation from \mathcal{P}^2 to \mathcal{P}^3 , and all the citations

from \mathcal{P}^1 to \mathcal{P}^3 can be removed, by Reference Independence, if there exists (or if we add) a citation from \mathcal{P}^1 to \mathcal{P}^2 . With this last provision, the sub-database containing the set of papers $\mathcal{P}^1 \cup \mathcal{P}^2$ remains in the domain \mathbb{D} . Moreover, we obtain by construction that the sub-database containing the set of papers \mathcal{P}^3 belongs to \mathbb{D} .

We obtain a partition of the database into two disjoint fields, and by application of Separability, we can only keep the sub-database containing the set of papers $\mathcal{P}^1 \cup \mathcal{P}^2$. Finally, we obtain the desired database. To conclude, by neutrality of I , we obtained that if two databases are such that $m_a[d] = m_a[d']$, we must have $I[d](a) = I[d'](a)$. Thus, I is a citation-counting scheme.

Independence of the axioms (each example satisfies all but the corresponding axiom):

- Separability. The score of an author a equals the ratio between the number of citations received by a and the total number of references of authors other than a .
- Reference Independence. The score of each author equals her number of references.
- Splitting. The score of each author is her fractional citation count.
- Citation Anonymity. The score of an author a is the total number of citations that papers citing a paper from a receive, and that are not issued by a . Formally:


$$I[d](a) = \sum_{q \in \cup_{p \in \mathcal{P}_a} \mathcal{C}_p} \#(\mathcal{C}_q \cap (\mathcal{P} \setminus \mathcal{P}_a)).$$

- Author Anonymity. The score of a is the number of authors citing a .

□

6.2 Proof of Proposition 1

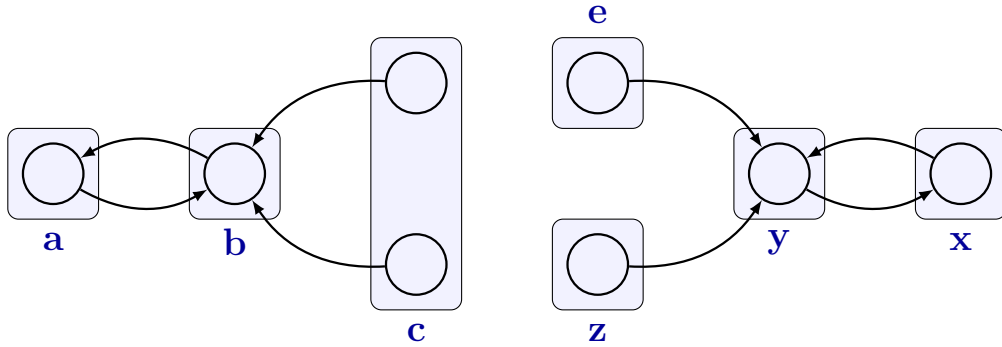
Proof. Let I be a measure of influence satisfying Null Author, Author Anonymity and Field Comparability, and consider the following database $d \in \mathbb{D}$.



By Null Author, we must have $I[d](c) = I[d](e) = I[d](z) = 0$. By Field Comparability, we have:

$$\frac{I[d](a) + I[d](b)}{4} = \frac{I[d](x) + I[d](y)}{3},$$

and by Null Author, these two numbers must be strictly positive. Now, by Author Anonymity, the scores of a , b , x and y must be the same in the database d' represented below, obtained by having c taking a paper from e and e taking a paper from z .



By Field Comparability, we must also have:

$$\frac{I[d](a) + I[d](b)}{3} = \frac{I[d](x) + I[d](y)}{4}.$$

We obtain a contradiction with the previous equality. □

6.3 Proof of Proposition 2, Proposition 3, Proposition 4

The three propositions are corollaries of the following two lemmas.²³

²³The two lemmas and their proofs follow the proofs of Theorems 1 and 2 in Bloch et al. (2016), with two differences. First, we assume that the measures are homogeneous in the two lemmas, and thus obtain linear

Lemma 2. *An aggregation function $f : [0, 1]^L \rightarrow \mathbb{R}$, with L finite, is symmetric, monotone, homogeneous and independent if and only if there exists $\lambda \geq 0$ such that $f(s) = \lambda \sum_{l \in L} s_l$.*

Proof. It is straightforward that the sum satisfies the properties enunciated in the statement. Conversely, let $f : [0, 1]^L \rightarrow \mathbb{R}$ be symmetric, monotone, homogeneous and independent. For $s \in [0, 1]^L$, we note $s_i^j = (s_l \mathbf{1}_{\{i \leq l \leq j\}})_{l \in \{1, \dots, L\}}$. Since $f(0) = 0$ (by homogeneity), we can write:

$$f(s) = \sum_{l=1}^L [f(s_1^l) - f(s_1^{l-1})].$$

By independence, we have $f(s_1^l) - f(s_1^{l-1}) = f^l(s_l)$, where $\forall l \in \{1, \dots, L\}$, $f^l : [0, 1] \rightarrow \mathbb{R}$ is non-decreasing and homogeneous. Therefore, $\forall l \in \{1, \dots, L\}$, $\exists \lambda_l \geq 0$ such that $\forall x \in [0, 1]$, $f^l(x) = \lambda_l x$. As f is symmetric, we must have $\forall l \in \{1, \dots, L\}$, $\lambda_l = \lambda \geq 0$. \square

Lemma 3. *An aggregation function $f : [0, 1]^{\mathbb{N}^*} \rightarrow \mathbb{R}$ is monotone, homogeneous, independent, recursive and satisfies the long run property if and only if there exists $\delta \in [0, 1)$ and $\lambda \geq 0$ such that $f(s) = \lambda \sum_{k=1}^{+\infty} \delta^{k-1} s_k$.*

Proof. It is straightforward that the discounted sum satisfies the properties enunciated in the statement. Conversely, let $f : [0, 1]^{\mathbb{N}^*} \rightarrow \mathbb{R}$ be monotone, homogeneous, independent, recursive and satisfy the long run property. As in the proof of Lemma 2, we write:

$$f(s) = f(s) - f(s_1^k) + \sum_{l=1}^k [f(s_1^l) - f(s_1^{l-1})].$$

Let us show that $[f(s) - f(s_1^k)]$ converges to 0. Let $\varepsilon > 0$ and, by the long run property, let $z \in \mathcal{Z}$ be such that $f(s \cdot z) \geq f(s) - \frac{\varepsilon}{2}$. As $z \in \mathcal{Z}$, let k_0 be such that for all $k \geq k_0$, $z_k \leq \frac{\varepsilon}{2f(\mathbf{1})}$, where $\mathbf{1}$ denotes the sequence whose coefficients are all equal to 1.²⁴ Then, we can write, using

functions. Second, the proof of Lemma 3 justifies the statement “[$f(s) - f(s_1^k)$] tends to 0”, which is obtained here by a domain condition on f , and the long run property.

²⁴We suppose $f(\mathbf{1}) > 0$, otherwise the claim is trivial.

monotonicity, independence, and then homogeneity, that for all $k \geq k_0$,

$$\begin{aligned}
f(s) - f(s_1^k) &\leq f(s \cdot z) + \frac{\varepsilon}{2} - f(s_1^k) \\
&\leq f(s \cdot z) + \frac{\varepsilon}{2} - f((s \cdot z)_1^k) \\
&\leq f((s \cdot z)_{k+1}^{+\infty}) + \frac{\varepsilon}{2} \\
&\leq \frac{\varepsilon}{2f(\mathbf{1})} f(\mathbf{1}_{k+1}^{+\infty}) + \frac{\varepsilon}{2} \\
&\leq \varepsilon.
\end{aligned}$$

Finally, we obtain that $[f(s) - f(s_1^k)]$ converges to 0. We can write, as in the proof of Lemma 2:

$$f(s) = \sum_{l=1}^{+\infty} [f(s_1^l) - f(s_1^{l-1})] = \sum_{l=1}^{+\infty} f^l(s^l) = \sum_{l=1}^{+\infty} \lambda_l s_l.$$

As f is recursive, taking $s = (\mathbf{1}_{\{l=k\}})_{l \in \mathbb{N}^*}$ and $s' = (\mathbf{1}_{\{l=k+1\}})_{l \in \mathbb{N}^*}$, we obtain that for all $k \geq 1$, $\lambda_{k+1}^2 = \lambda_k \lambda_{k+2}$. Finally, we have three possible cases. If $\lambda_1 = 0$, we obtain that $f(s) = 0$. If $\lambda_1 > 0$ and $\lambda_2 = 0$, we obtain that $f(s) = \lambda_1 s_1$. If $\lambda_1 > 0$ and $\lambda_2 > 0$, we note $\delta = \frac{\lambda_2}{\lambda_1}$, and we obtain that $\forall k \geq 1$, $\lambda_k = \lambda_1 \delta^{k-1}$, and thus $f(s) = \lambda_1 \sum_{k=1}^{+\infty} \delta^{k-1} s_k$. As $f(\mathbf{1}) < +\infty$, it must be that $\delta < 1$. \square

References

- ANGRIST, J., P. AZOULAY, G. ELLISON, R. HILL, AND S. F. LU (2017): “Inside Job or Deep Impact? Using Extramural Citations to Assess Economic Scholarship,” Tech. rep., National Bureau of Economic Research.
- BLOCH, F., M. O. JACKSON, AND P. TEBALDI (2016): “Centrality measures in networks,” *Mimeo*.
- BOUYSSOU, D. AND T. MARCHANT (2016): “Ranking authors using fractional counting of citations: An axiomatic approach,” *Journal of Informetrics*, 10, 183–199.
- CHAMBERS, C. P. AND A. D. MILLER (2014): “Scholarly influence,” *Journal of Economic Theory*, 151, 571–583.
- CHEN, P., H. XIE, S. MASLOV, AND S. REDNER (2007): “Finding scientific gems with Google’s PageRank algorithm,” *Journal of Informetrics*, 1, 8–15.

- COPIC, J., M. O. JACKSON, AND A. KIRMAN (2009): “Identifying community structures from network data via maximum likelihood methods,” *The BE Journal of Theoretical Economics*, 9.
- DEMANGE, G. (2014): “A ranking method based on handicaps,” *Theoretical Economics*, 9, 915–942.
- GARFIELD, E. (1972): “Citation Analysis as a Tool in Journal Evaluation,” *Science*, 178, 471–479.
- HAMERMESH, D. S. (2013): “Six decades of top economics publishing: Who and how?” *Journal of Economic Literature*, 51, 162–172.
- HIRSCH, J. E. (2005): “An index to quantify an individual’s scientific research output,” *Proceedings of the National Academy of Sciences*, 16569–16572.
- IOANNIDIS, J. P., K. BOYACK, AND P. F. WOUTERS (2016): “Citation Metrics: A primer on how (not) to normalize,” *PLoS biology*, 14, e1002542.
- KAUR, J., F. RADICCHI, AND F. MENCZER (2013): “Universality of scholarly impact metrics,” *Journal of Informetrics*, 7, 924–932.
- KŁOSIK, D. F. AND S. BORNHOLDT (2014): “The citation wake of publications detects Nobel laureates’ papers,” *PloS one*, 9, e113184.
- KÓCZY, L. AND M. STROBEL (2009): “The invariant method can be manipulated,” *Scientometrics*, 81, 291–293.
- MARCHANT, T. (2009): “An axiomatic characterization of the ranking based on the h-index and some other bibliometric rankings of authors,” *Scientometrics*, 80, 325–342.
- PALACIOS-HUERTA, I. AND O. VOLIJ (2004): “The measurement of intellectual influence,” *Econometrica*, 72, 963–977.
- (2014): “Axiomatic measures of intellectual influence,” *International Journal of Industrial Organization*, 34, 85–90.
- PATTY, J. W. AND E. M. PENN (2015): “Analyzing big data: social choice and measurement,” *PS: Political Science & Politics*, 48, 95–101.
- PERRY, M. AND P. J. RENY (2016): “How to count citations if you must,” *The American Economic Review*, 106, 2722–2741.
- PINSKI, G. AND F. NARIN (1976): “Citation influence for journal aggregates of scientific publications: Theory, with application to the literature of physics,” *Information processing & management*, 12, 297–312.
- RADICCHI, F., S. FORTUNATO, AND C. CASTELLANO (2008): “Universality of citation distributions: Toward an objective measure of scientific impact,” *Proceedings of the National Academy of Sciences*, 105, 17268–17272.

- RADICCHI, F., S. FORTUNATO, B. MARKINES, AND A. VESPIGNANI (2009): “Diffusion of scientific credits and the ranking of scientists,” *Physical Review E*, 80, 056103.
- SCOTCHMER, S. (1991): “Standing on the shoulders of giants: cumulative research and the patent law,” *The journal of economic perspectives*, 5, 29–41.
- SHAPLEY, L. S. (1953): “A value for n-person games,” *Contributions to the Theory of Games*, 2, 307–317.
- SZWAGRZAK, K. AND R. TREIBICH (2017): “Co-authorship and the measurement of individual productivity,” Tech. rep.
- WALTMAN, L. AND N. J. VAN ECK (2015): “Field-normalized citation impact indicators and the choice of an appropriate counting method,” *Journal of Informetrics*, 9, 872–894.
- WANG, H., H.-W. SHEN, AND X.-Q. CHENG (2016): “Scientific credit diffusion: Researcher level or paper level?” *Scientometrics*, 109, 827–837.
- WEST, J. D., M. C. JENSEN, R. J. DANDREA, G. J. GORDON, AND C. T. BERGSTROM (2013): “Author-level Eigenfactor metrics: Evaluating the influence of authors, institutions, and countries within the social science research network community,” *Journal of the Association for Information Science and Technology*, 64, 787–801.
- WOEGINGER, G. J. (2008): “An axiomatic characterization of the Hirsch-index,” *Mathematical Social Sciences*, 56, 224–232.
- WUCHTY, S., B. F. JONES, AND B. UZZI (2007): “The increasing dominance of teams in production of knowledge,” *Science*, 316, 1036–1039.
- ZITT, M., S. RAMANANA-RAHARY, AND E. BASSECOULARD (2005): “Relativity of citation performance and excellence measures: From cross-field to cross-scale effects of field-normalisation,” *Scientometrics*, 63, 373–401.