The number behind the number: suggesting a truer measure of academic impact

10/5/2016

The limitations of simple 'citation count' figures are well-known. **Chris Carroll** *argues that the impact of an academic research paper might be better measured by counting the number of times it is cited within citing publications rather than by simply measuring if it has been cited or not. Three or more citations of the key paper arguably represent a rather different level of impact than a single citation. By looking for this easily generated number, every researcher can quickly gain a greater insight into the impact (or not) of their published work.*

The academic research and policy agenda increasingly seeks to measure and use 'impact' as a means of determining the value of different items of published research. However, there is much debate about how best to define and quantify impact, and any assessment must take into account influence beyond the limited bounds of academia, in areas such as public policy. However, within academia, it is generally accepted that the number of times a paper is cited, the so-called 'citation count' or 'citation score', offers the most easily measured guide to its impact. The underlying assumption is that the cited work has influenced the citing work in some way, but this metric is also viewed as practically and conceptually limited. The criticism is loud and frequent: such citation scores do not tell us how a piece of research has actually been used in practice, only that it is known and cited.

Image credit: **beyond measure** **by frankieleon. This work is licensed under a** **CC BY 2.0** **license.**

Is there a better solution?

So, the obvious answer is to see if there is an easy way to measure how a paper has been used by its citing

publications, rather than simply recording whether it is cited (the standard 'citation score'). I did this by using one of my own frequently-cited papers as a case study and applying a basic metric: impact of the case study paper was considered to be 'high' if it was cited three or more times within the citing publication; 'moderate' if cited twice; and 'low' if cited only once.

The case study paper had 393 unique citations by November 2015: 59% (230/393) of publications cited the paper only once, suggesting its impact on these publications was low; 17% (65/393) cited it twice, suggesting moderate impact; but 25% (98/393) cited it three or more times, suggesting that this paper was having a genuine influence on those studies. The citation frequency within this 'high impact' group of publications ranged from three to as many as 14 in peer-reviewed studies and 17 in academic dissertations (with their longer word count). Primary research studies published in peer-reviewed journals were the principal publication type across all levels of impact.

I also noted where these single or multiple citations appeared within these publications. Single citations tended to appear only in the introduction or background sections of papers. These instances appeared to be simple citations 'in passing', the necessary 'nod' to existing literature at the start of a publication. However, when there were three or more citations, they tended to appear across two or more sections of a paper, especially in the methods and discussions, which suggests some real influence on the justification or design of a study, or the interpretation of its results. These findings on the location of the citations confirmed that the number of times a paper was cited within a publication really was a good indicator of that paper's impact.

Pros and cons

A metric based on within-publication citation frequency is unambiguous and thus capable of accuracy. The assessment of where the citations appear in the citing publications also provided contextual information on the nature of the citation. In this case, it confirmed the viability of within-publication citation frequency as an impact metric. The approach is transparent and data can be verified and updated. The proposed metric, the 'citation profile', is not a measure of research quality, nor does it seek to measure other forms of impact or to address issues such as negative findings producing relatively lower citation scores. It merely seeks to contribute to the debate about how citation metrics might be used to give a more contextually robust picture of a paper's academic impact.

Of course, even this 'citation profile' is not without its problems, but it is arguably less of a "damned lie" or "statistic" than some other metrics: only a more in-depth analysis of each citing publication can more accurately gauge a paper's influence on another piece of research or a policy document. Yes, there are also questions concerning this metric's generalisability: as with other bibliometrics, publications in different disciplines are also likely to have different citation profiles. However, these issues can be easily addressed by future research, given the simplicity of the metric.

If academics or funders want to improve their understanding of how research is used, then this easy-to-generate metric can add depth and value to the basic, much-maligned, and 'damned' 'citation score'.

Note: This article gives the views of the author(s), and not the position of the LSE Impact Blog, nor of the London School of Economics. Please review our comments policy if you have any concerns on posting a comment below.

About the author

Chris Carroll is a Reader in Systematic Review and Evidence Synthesis at the University of Sheffield. His role principally involves synthesising published data to help inform policymaking, particularly in the fields of medicine and health, as well as the development of methods to conduct such work.