

Citation counts reinforce the influence of highly cited papers and nudge us towards undervaluing those with fewer.

*In the context of everyday research assessment citation counts are often taken as a simple indicator of the influence of any particular paper. However, all citations are not the same and can be deployed to achieve different ends. Commenting on a recent study of how researchers across 15 academic fields understand the influence of the work cited in their research **Eamon Duede** shows how citation plays a role both in indicating and shaping the influence of research papers.*

The impacts of science and scholarship are as multifaceted as they are ubiquitous. Research findings shape policy through evidence and principled prediction, refine our lifestyles with new technologies, sources of energy, and entertainment, as well as improve our health and wellbeing through advances in medicine, clinical practice, and education.

Research impact arises, of course, from research. While research varies widely in the breadth and magnitude of its influence on our needs, wants, and governance, it also serves to influence our knowledge and understanding, our future investigations –our curiosity itself. So much so, this is often taken for granted.

In many ways, research papers are repositories of potentially impactful knowledge, each building on, connecting, challenging, or acknowledging one another. And one of the principal means by which these diverse activities are recorded in the ledger of scientific progress is through the practice of citation.

in the majority of cases, the common and reasonable belief that citations account for influence is false.

As a result, it has become common and not entirely unreasonable to believe that the number of times a paper has been cited is a direct measure of how influential it has been. If a work has been cited significantly more times than its peers in the field, then, in all likelihood, that paper has been more influential than the others. So, it isn't any wonder that scientists, administrators, funding bodies, and awards committees use citation counts as one of the main ways of evaluating influence when making decisions.

Yet, in our [recent paper](#), [Misha Teplitskiy](#), [Michael Menietti](#), [Karim R. Lakhani](#), and I find that, in the majority of cases, the common and reasonable belief that citations account for influence is false.

To show this, we conducted a careful, highly personalised survey of nearly 10,000 authors of recent scientific papers. This survey was notable, because rather than inferring, we directly asked authors about specific papers they cited in their recent work. In particular, we were interested in comparing these authors' perceptions of the quality of the papers they have read and referenced with the number of times that those papers had been cited. Additionally, we asked authors to report the degree to which these references influenced their own papers.

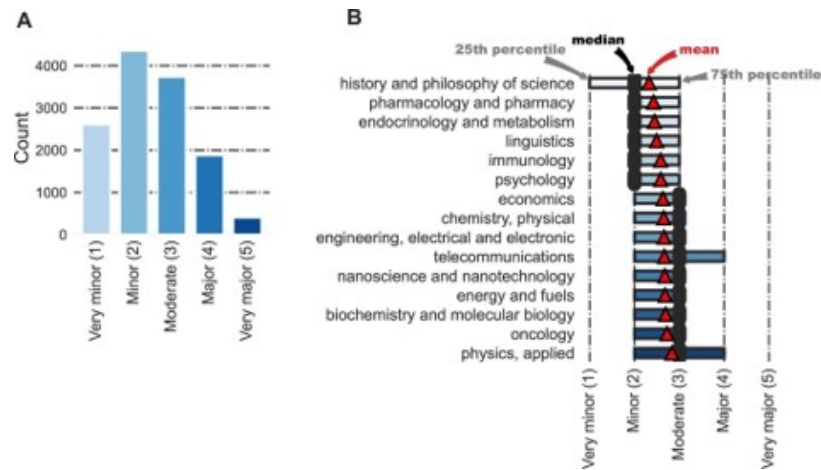


Figure 1: The panel on the left (A) shows the distribution of influence for cited papers. The panel on the right (B) shows that distribution, broken out across 15 scientific and scholastic fields.

Contrary to the idea that citations are a straightforward reflection of academic influence, we found that, across all 15 academic fields represented in our sample of roughly 17,000 paper references, more than half were reported to have had little to no influence on the authors who cited them (see Figure 1).

However, many papers do exert meaningful influence on those who read and cite them. There are two common and contradictory hypotheses about which articles are actually influential. According to one view, most citations to papers do, to a certain, context relative extent, denote influence on their readers. However, as articles become more highly cited, authors stop reading them closely and tend to cite them for the rhetorical and argumentative benefits they bring to the authors' own papers. On the other hand, the competing hypothesis claims that, as articles become more highly cited, their citation status signals to potential readers that those papers are of higher quality. This, in turn, leads authors to invest more time and cognitive effort into reading them. As a result of that investment, readers are more likely to be influenced.

We find that the most highly cited papers are, indeed, the most influential. Figure 2 shows that, as the number of times a given paper has been cited increases, the probability increases that a given paper will have significantly influenced a reader who references it. However, the data show that the likelihood of significant influence really spikes only for the most highly cited papers (e.g., papers with citation counts in the high 100s and 1000s). But are these highly cited papers viewed as being of higher quality?

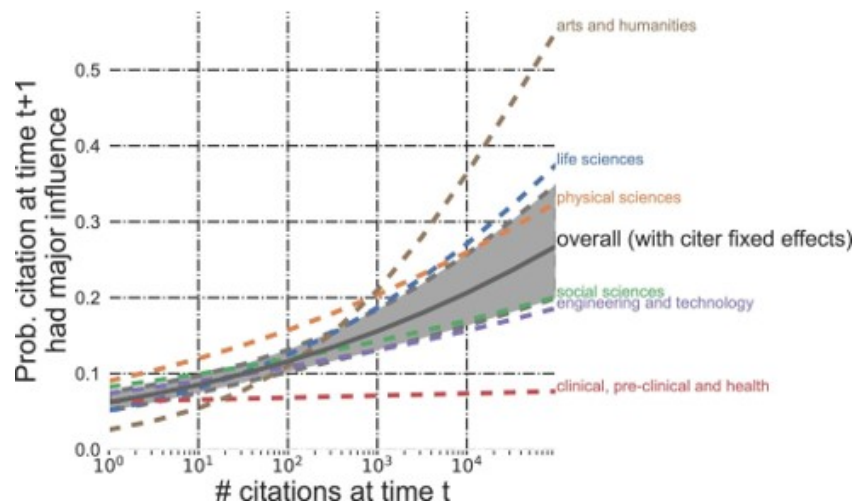


Figure 2: The probability that a given paper exerts a high degree of influence on an author who references it as a function of the number of times it has been cited.

To answer this question, we conducted an experiment in which we revealed (treatment) or hid (control) the citation status of the reference about which we asked respondents. Surprisingly, showing or hiding the citation status of the most highly cited papers did little to change respondents' perceptions of their quality. However, showing respondents paper citation status significantly reduces their perceptions of the quality of all but the most highly cited papers (see Figure 3A).

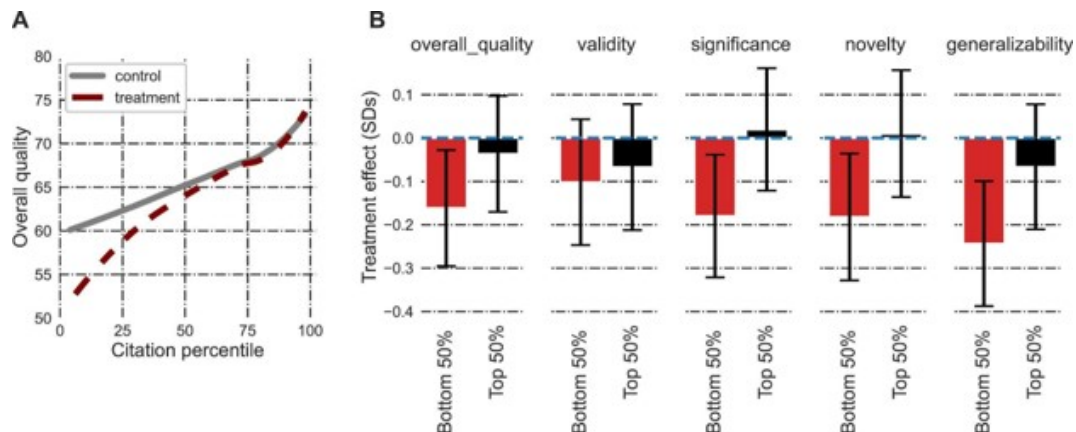


Figure 3: Panel A (left) shows the effect of displaying (red line) or not showing (grey line) the citation status of papers on reader's evaluation of those papers' quality. Panel B (right) shows the effect of showing or not showing citation status on perceptions of the quality, validity, significance, novelty, and generalizability of papers.

Finally, to answer whether perceptions of quality lead to more or less cognitive investment, we asked readers questions about how and when they found papers that they ended up citing in their work and how well they know the content of those papers. We found that papers with lower perceived quality are discovered later in the research cycle for a given project, are read more superficially, and, in turn, confer significantly less influence. The reverse is true for the most highly cited papers. These are found earlier, read more closely, and have a higher probability of influencing their readers (see Figure 4).

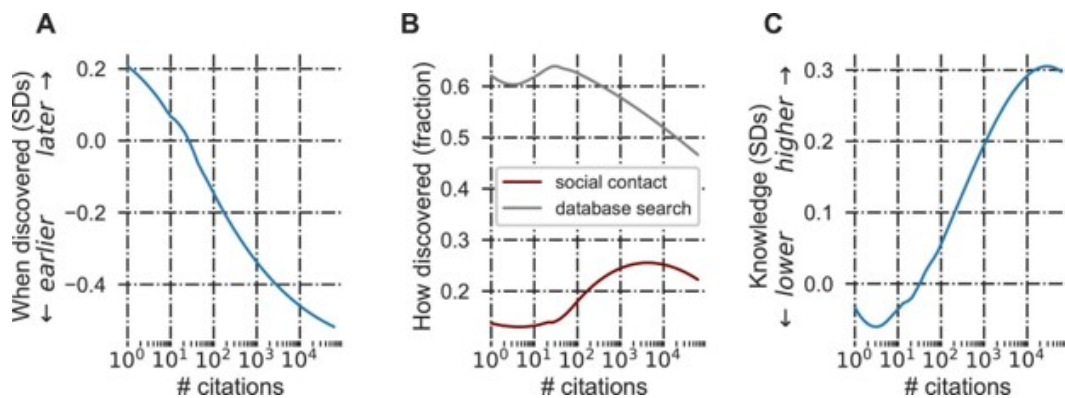


Figure 4: Panel A (left most) shows that the more highly cited a paper is, the earlier it is discovered in the course of researching a project. Panel C (right most) shows that the more highly cited a given paper, the more readers invest into knowing and understanding its content. Panel B (centre) shows that the more highly cited a paper, the less likely it is to have been discovered by using a database search.

These findings have enormous implications for how we conduct and assess the impact of research. We found that most of the papers in our sample were discovered by their readers using common database searches. Such searches commonly display citation counts alongside results. Yet, our study shows that this affects how and what we choose to read. As a result, a significant amount of potentially influential work may go unread for reasons having nothing to do with its quality. As we note in the paper, even though citations are a poor proxy for the quality of a paper, the “losers” of citation status signals outnumber the “winners” roughly 9-to-1.

*This post draws on the author's co-authored article, [How status of research papers affects the way they are read and cited](#), published in *Research Policy*.*

The content generated on this blog is for information purposes only. This Article gives the views and opinions of the authors and does not reflect the views and opinions of the Impact of Social Science blog (the blog), nor of the London School of Economics and Political Science. Please review our [comments policy](#) if you have any concerns on posting a comment below.

Image Credit: Adapted from [Y. Peyankov](#) via Unsplash.
