Colloquium GAP 10
Köln 2017

⋆

Epistemic diversity and editor decisions

⋆

Remco Heesen and Jan-Willem Romeijn
University of Cambridge | University of Groningen

Contents

1 Diversity and bias in science

Next to social diversity, epistemic diversity in the sciences and the humanities is a good thing that merits institutional support.

> We sometimes want to maintain cognitive diversity even in instances where it would be reasonable for all to agree that one of two theories was inferior to its rival, and we may be grateful to the stubborn minority who continue to advocate problematic ideas (Kitcher 1990, p. 7).

> The history of science has been and should be a history of competing research programmes (or, if you wish, 'paradigms'), but it has not been and must not become a succession of periods of normal science: the sooner competition starts, the better for progress (Lakatos 1978, p. 69).

Diversity and bias in science

A systematic bias of peer review towards a mono-culture is detrimental to scientific progress.

- Disciplines that harbour a greater variety of methods and theories will offer more balanced viewpoints and are better equipped to respond to challenges.
- This epistemic diversity stands alongside a social diversity in science that is worthy of pursuit in its own right.
- The organization of the process of science must facilitate diversity in research programmes. This holds in particular for the organization of the peer-review system.

Diversity and bias in science

A systematic bias towards more prominent research programmes is inherent in the selection decisions that editors make.

- Editors may suffer from confirmation bias and anchoring when assessing the quality of papers, and choose papers conservatively with an eye on journal reputation.
- But even if editors manage to purge their decision procedures of unconscious biases, they will be left with biases of a strictly statistical nature!

The current paper is aimed at clarifying and ultimately counteracting this “statistical Matthew effect”.

Diversity and bias in science

Our analysis of statistical biases in editorial decision making proceeds in two steps.

- A consequence of a recent argument from Heesen (2017) is that editors may favour submissions from a more prominent programme through a mere information asymmetry.
- Even if we avoid this asymmetry, a generalization of Borsboom et al. (2008) shows that a more established programme will typically suffer fewer false positives and false negatives in the selection of papers.

Hence editors will inadvertently contribute to an already existing tendency towards a mono-culture in science.

2 Information asymmetry

We first consider the decision situation from the editor's viewpoint, inspired on Heesen (2017).

- We suppose that the overall distribution of quality in the output from two rivalling research programmes is the same.
- The editor will know more researchers from her own programme, and hence she will have information on the quality of their papers more often.
- This information is at least sometimes relevant to the decision to accept or reject the paper.
- Editors can use the author information to improve on their judgment. This is justified if the goal is simply to maximize the quality of accepted papers.

Information asymmetry

More precisely, we make the following modeling assumptions, expressed in the editor's subjective probability assignment π.

- The distribution of paper quality, $\pi(Q)$, is independent of the program from which the paper originates: $\pi(Q|P) = \pi(Q)$ and $\pi(Q|P, K = k) = \pi(Q|K = k)$ for $P = \{0, 1\}$ and $K = 0, 1, \ldots N$.
- The expectation of the quality for known authors, $K > 0$, is the same as for unknown authors, $K = 0$: $\mathbb{E}_{k>0}[\pi(Q|P, K = k)] = \pi(Q|P, K = 0)$ for $P = \{0, 1\}$.
- The reviews R are also independent of program and author status, so $\pi(R|Q, K, P) = \pi(R|Q, P) = \pi(R|Q)$ for $P = 0, 1$ and any K.
- But if the identity of the author is known, this may impact on the posterior distribution over quality, $\pi(Q|K = k) \neq \pi(Q)$.

Information asymmetry

The decision of the editor is then based on the posterior subjective probability.

- If the identity of the author is known, it is more probable that the paper belongs to the dominant progam, $\pi(P = 1|K > 0) > \pi(P = 0|K > 0)$.
- The posterior is constructed by conditioning π on the known values for K, R, and P. The latter of these only impacts the posterior via the variable K.
- The utility of accepting a paper is its quality $Q = q$, while rejecting it always leads to a fixed q^*. The editor accepts a paper from a programme, $D_P = A$, if $\mathbb{E}[Q|P, R, K] > q^*$.

Information asymmetry

We can then arrive at the following results, based on Good (1967).

> Despite the programmes being on a par in terms of quality, the likely acquaintance of the editor with the more prominent one gives it an advantage in terms of either acceptance rates or average quality of the published work:

$$\mathbb{E}[Q|D_{P=1}=A] > \mathbb{E}[Q|D_{P=0}=A] \quad \text{or} \quad P(D_{P=1}=A) > P(D_{P=0}=A).$$

The editor is caught in a dilemma, having to choose between two ways of advantaging the dominant programme. Moreover, if quality and reviews are normally distributed, both inequalities hold, as per Heesen (2017).

Information asymmetry

The model adopts a number of idealizations and assumptions that need critical scrutiny.

- It is assumed that there is some distribution over a paper quality Q. But this is typically a highly multi-dimensional characteristic.
- The driver of the result is information asymmetry. If we want to protect epistemic diversity at the cost of some scientific quality, we must impose double blind reviewing.

3 Latent quality differences

We require that no prior information about the author of a paper may be taken into account in evaluating it. Unfortunately, we may still run into statistical Matthew effects.

- Plausibly, a more dominant and more established research programme will on average produce work of higher quality. Novel research programmes may suffer from startup problems.
- We do not need to suppose that the editor takes this latent difference between the programmes into account, merely that such differences occur.
- Editors simply base their judgments on the review reports of the papers, which are themselves only dependent on the actual quality of the papers.

Latent quality differences

More precisely, we make the following modeling assumptions, deriving from Borsboom et al. (2008).

- The quality Q follows the same log-concave and hence unimodal distribution in either research programme, with means $\mu_0 < \mu_1$.
- For now we assume equal variances for the two programmes. We leave aside further results that drop this assumption.
- The editor will accept a paper if its review reports R pass a threshold t. We make no assumptions on the distribution of R except that the conditional probability of exceeding the threshold $Pr(R > t|Q = q)$ is increasing in q.

Latent quality differences

We can then arrive at the foregoing results, now in conjunction.

> Under fairly weak assumptions about latent differences between programmes, less well established programmes will be disadvantaged in terms of acceptance rates *and* average quality of published work:

$$\mathbb{E}[Q|D_{P=1} = A] > \mathbb{E}[Q|D_{P=0} = A] \quad \text{and} \quad P(D_{P=1} = A) > P(D_{P=0} = A).$$

We can also derive results for unequal variances but they depend on the location of the cut-off point t.

Latent quality differences

The above result follows from stronger claims on comparisons between probabilities of paper acceptance and paper quality in the two programmes.

- If we label papers for which $Q > q^*$ suitable and for which $R > t$ accepted, we can talk about two types of errors: $Pr(\neg A|P, S)$ and $Pr(\neg S|P, A)$.
- The problem is that these errors are both larger for the developing program. Hence the quality of the test performed by the editor is lower for that program.
- Owing to the symmetry in the distributions, there are two other types of errors. But we argue that they are less harmful to the reputation of the program.

Latent quality differences

Especially noteworthy is a link to recent public debate over fairness in AI.

- This result may not seem very disconcerting because we have assumed latent differences between the programs, and these will lead to different publication rates.
- The problem is not that the developing program is at a disadvantage, but that it is disproportionately so.
- A good illustration of this point is in the recent public debate over automated risk scores for recidivism, as used by the US judicial system. Kleinberg et al. (2017) make essentially the same point as Borsboom et al. (2008).
- While the focus is often on the ethical consequences of this point, our current concern is epistemic diversity.

4 What can be done?

Our findings underline the challenges involved in safeguarding a diversity of research programmes.

- If a journal promotes a dominant programme to the detriment of others, this cannot be ascribed to cognitive biases alone.
- Nothing suggests that we may cease to critically assess our proneness to unconscious bias. But. . .

What can be done?

The insights might also lead to reorienting our efforts to counteract bias.

- Special attention must be paid to papers whose reviews place it close to the cut-off for acceptance and rejection.
- This means that we will comparatively more often ask additional reviews for papers from underrepresented research programmes.
- Hence we suggest actively encouraging diversity at the expense of treating all papers equally.

Thank you

The slides for this talk will be available at http://www.philos.rug.nl/˜romeyn. For comments and questions, email j.w.romeijn@rug.nl.