

Publication patterns in research underpinning impact in REF2014

A report to HEFCE by Digital Science

July 2016

© HEFCE 2016

The Glasshouse Building, 68 Wharfdale Road

LONDON N1 9SR, UK

<http://www.digital-science.com>

consultancy@digital-science.com

Publication patterns in research underpinning impact in REF2014

For 30 years, outputs published by UK researchers have been submitted for cyclical assessment by peer panels, creating a unique longitudinal dataset. This report is the first to describe analyses of all publications submitted for assessment between 1988 and 2013. It covers 921,254 submitted outputs and 36,244 case study references across 25 years, five assessment cycles and both academic impact and economic/societal impact. The report identifies patterns in these data but, as this is the first publication of such analysis, interpretation is intentionally reserved pending wider discussion.

Summary – key outcomes	Reference
Submitted output types shift towards journal articles across time: out of conference proceedings for engineers, out of monographs for social scientists and out of media for arts. This was synchronous with increasing citation awareness.	Figure 1 and Annex 2 – Table A
The time-spread of submitted outputs is skewed to the most recent publication years in early RAE cycles. This skewed pattern was not reported at the time.	Figure 2
The time-skew is modified synchronously for science and engineering in later cycles, but no change in skew is observed for social science or humanities.	Figure 3
The skewed submission patterns, and later changes, are remarkably cohesive across disciplinary UOAs and across HEIs.	Annex 2 – Figures A and B
The time-spread of the earliest of the six references for each impact case study is uniform for science and engineering, although the overall time-spread for references is skewed to recent years as for submitted outputs. Earliest references remain recency-skewed for social science and humanities.	Figure 4
About 42% of case study references with DOIs can be identified as submitted RAE/REF outputs, at about the same rate in every publication year for 1996-2013.	Figure 5
The overlap between case study references and assessment submitted outputs appears to be higher in professional areas and lower in core humanities and sciences.	Table 1 and Figure 6
The overlap between case study references and assessment submitted outputs is higher in early cycles for science and engineering while the overlap with REF2014 is higher in social sciences.	Figure 7
Older, larger HEIs used more case study references but the extent of overlap with outputs submitted for assessment is relatively lower than for other HEIs. More of the overlap is with early RAE cycles whereas post-92 HEIs have a higher overall overlap of which most is with REF2014.	Annex 3 – Figure C and Table C

Contents

Introduction.....	4
Background and methodology	5
Results.....	7
The balance of submitted output types changes over time.....	7
Submitted output publication dates are skewed.....	9
Behaviour in selecting submitted outputs is consistent across UOAs and HEIs.....	11
There is a cultural divide in the selection of submitted outputs.....	11
The earliest case study references underpinning impact are evenly distributed by year for science and engineering.....	12
Overlap between submitted outputs and impact references is even across years	14
Overlap between submitted outputs and impact references varies by UOA.....	15
Overlap between submitted outputs and impact references varies between HEIs	17
Case study analysis provides additional evidence of impact	18
Discussion	19
The relationship between submitted outputs and case study references.....	19
The selection of references used in impact case studies	20
The selection of submitted outputs used in research assessment	21
The homogeneity of response and its implications for peer review	23
ANNEX 1 - Who chooses submitted outputs?.....	25
ANNEX 2 - Supplementary data	26
ANNEX 3 - Overlap between submitted outputs and references: patterns across HEIs.....	28
ANNEX 4 - Hyperlinks in impact case studies	29

Introduction

1. This report describes the patterns, across years and disciplines, of outputs submitted by UK higher education institutions (HEIs) to the 2014 Research Excellence Framework (REF) and to previous Research Assessment Exercises (RAEs)¹. The study originates from the observation that the time-profile of references in impact case studies was more strongly skewed to recent years², particularly after 2008, than had been anticipated from prior studies of the time expected for original research to develop into application and new processes and products.³
2. Researchers publish for many reasons⁴. Less is known about the selection of publications for assessment. This report refers to two overlapping sets of publication data, each presented for distinct assessment purposes. To distinguish these, they are defined as:

Submitted outputs: outputs (primarily printed publications) submitted in assessment cycles as evidence of academic excellence. The RAE/REF requires four publications – or other outputs where appropriate – to be submitted for each researcher. The RA2/REF2 output sections of each assessment form a focus for peer review of academic research achievement.

References: outputs included in REF2014 as impact case study references. REF2014 differed from the prior RAE format in adding four-page case studies of the societal and economic impact⁵ of research. Impact case studies consist of five sections including a description of (Section 2) and six references to (Section 3) the underpinning research.

3. The objective of this report is to publish analyses of submission and referencing patterns to enable a better understanding of the selective use of published outputs by researchers as evidence of their achievement and impact. The questions to address are:
 - a. What is the relationship between submitted outputs (as evidence of academic impact) and case study references (as evidence of societal and economic impact), noting that academic impact has been conceptually associated with exceptional research innovation?⁶
 - b. What do the data say about the selection of outputs for assessment?
 - c. How did selection behaviour respond to assessment?
 - d. How homogeneous was that response, by discipline and institution?

¹ RAE1992, RAE1996, RAE2001, RAE2008

² King's College London and Digital Science (2015). The nature, scale and beneficiaries of research impact: an initial analysis of REF 2014 impact case studies. Research Report 2015/01, Higher Education Funding Council for England.

³ Griliches, Z. (1986). Productivity, R&D, and basic research at the firm level in the 1970s. *American Economic Review*, 76, 141-154. Mansfield, E. (1990). Academic research and industrial innovation. *Research Policy*, 20, 1–12. Wooding, S., Hanney, S., Pollitt, A., Grant, J. and Buxton, M. (2014). Understanding factors associated with the translation of cardiovascular research. *Implementation Science*, 9, 201-209.

⁴ <http://www.rin.ac.uk/system/files/attachments/sarah/Communicating-knowledge-report.pdf>

⁵ For the purposes of the REF, impact is defined as an effect on, change or benefit to the economy, society, culture, public policy or services, health, the environment or quality of life, beyond academia. This was assessed via an Impact template (REF3a) describing the submitted unit's approach and Impact case studies (REF3b) describing specific impacts underpinned by excellent research.

⁶ Breschi, S., Tarasconi, G., Catalini, C., Novella, L., Guatta, P. and Johnson, H. (2006). Highly Cited Patents, Highly Cited Publications, and Research Networks. CESPRI Report to DG Research, European Commission. http://ec.europa.eu/invest-in-research/pdf/download_en/final_report_hcp.pdf

Background and methodology

4. Publication data from successive RAE/REF cycles provide a basis to track changes in decisions made about submitted outputs⁷. The RA2/REF2 outputs can also be compared with the potentially distinct perspective of references supporting REF case studies of societal and economic impact.
5. The RAE started in 1986 and settled into a standardised template from RAE1992⁸. It focused on academic quality. A peer panel reviewed a portfolio of evidence covering a short census period and submitted for a group of researchers, usually a department or school, in subject categories called Units of Assessment (UOAs). That portfolio included data on research income, postgraduate and post-doctoral training, and the research environment. The greatest scrutiny and post-hoc analysis was given to the (typically four) outputs (books, journal articles, conference papers, software, art, videos and other) submitted per researcher.⁹
6. **Output data discussed in this report** were sourced from the RAE archive sites maintained by the Higher Education Funding Council for England (HEFCE, which manages research assessment for all of the UK higher education funding bodies)¹⁰ and supplemented by prior analytical work on RAE data. The dataset is shaped by the varying length of RAE census periods (the acceptable period for publication cover dates) and by technology (e.g. electronic submission).
7. **Early cycles.**¹¹ The RAE1992 census period ran from January 1988 for arts and humanities, January 1989 for other subjects and added a part-year with a closing date of June 30, 1992. The RAE1996 census period ran from January 1990 for arts and humanities, January 1992 for other subjects, and added a part-year to March 31, 1996. Data for both are available only at UOA summary level.
8. **Later cycles.** Census periods for later cycles are mutually exclusive full calendar years. For RAE2001 (January 1996 - December 2000), there is a database of normalised records of submitted outputs¹². RAE2008 introduced a standard electronic format and publication data include Digital Object Identifiers (DOIs)¹³. For RAE2008 (2001-2007) and REF2014 (2008-2013), identifiable, individual records of submitted outputs are available.¹⁴

⁷ See Annex 1: This report is agnostic on where decisions are made about what selections should be submitted and how decision points shift between researchers, departments and institutions.

⁸ Universities Funding Council (UFC) Circular 5/92 setting out procedures and guidelines for RAE1992 was issued in March 1992.

⁹ See: <http://www.rae.ac.uk/pubs/2006/01/>

¹⁰ Higher Education Funding Council for Wales (HEFCW), Scottish Funding Council (SFC) and Department for Education & Learning Northern Ireland (DELNI)

¹¹ No output data are available for the 1986 and 1989 cycles.

¹² Submitted outputs were not in a standard format. Original records were assigned to output type (book, chapter, article, proceedings, other). Books and chapters were processed by library staff at the University of Leeds and an ISBN was assigned to each item. Journal articles were processed by Evidence Ltd in collaboration with Thomson Scientific to clean metadata and assign specific IDs for items indexed in the *Web of Science*.

¹³ <https://www.doi.org/>

¹⁴ In 2008 and 2014, HEIs could request that certain outputs were assessed as 'double weighted', contributing twice to the output sub-profile, and had the option of including a reserve output should the double weighting case not be accepted. Submissions' data do not identify which outputs were treated as double weighted. Therefore a book may have been considered equivalent to two outputs but count only once in analysis.

9. **Data volume.** Each assessment cycle received data for some 50,000 UK-based researchers, across 150 HEIs and a gradually reducing number of between 72 and 36 UOAs. The RA2/REF2 output section of each assessment database contains about 200,000 records of outputs submitted as evidence of research achievement at the time¹⁵. The combined data set of 921,254 outputs selected for assessment is a unique longitudinal perspective on national research activity.
10. **Impact.** REF2014 included four-page case studies of societal and economic impact for a research group, with several groups per UOA submission, each supported by reference publications (or other research outputs). The census period for references (1993-2013) was longer than the census window for submitted outputs and overlapped prior RAE cycles.
11. There are 6,737 non-confidential impact case studies (of 6,975 original documents) available for analysis¹⁶. Output records appear in case studies as 'References to the research' (Section 3: six output references per case study). Reference records, which did not follow a standard format, were extracted and normalised by Digital Science. Thomson Reuters processed items that appeared to be journal articles for matches against *Web of Science*TM. This produced a database of 36,244 identifiable case study journal article references with dates that overlapped assessment cycles from RAE1996.
12. **Evaluating overlap.** Submitted outputs and impact case study references were paired to enable an accurate analysis of their overlap. For journal articles, and some other outputs, this was done by matching publication records in both sources to unique DOIs that can be compared automatically. Matching outputs without DOIs has not been undertaken.
13. The normalised REF case study reference records were compared to CrossRef¹⁷ where possible and matched with DOIs where these were missing. Matching publication records to DOIs was also applied to early RA2 submitted outputs where no DOI had been included in original records. Note that: the DOI system started around 2000, so there is a deficit of DOIs for earlier publications; not all submitted outputs are journal articles with readily discoverable DOIs; and individual output records are not available for RAE1992 and RAE1996.
14. To establish the overlap between the sets of prior submitted outputs and the set of impact references, the augmented RA2/REF2 submission records were searched for the DOIs of those case study references to which a DOI could be matched.
15. Case study references without DOIs were processed manually to assign them to an output typology that followed the RAE/REF system (books, chapters, articles, proceedings, reports and other grey literature, patents, media-based material, and other outputs). Document-type frequency was summarised for comparison with the RA2/REF2 submitted outputs.
16. **Analysis.** Data were initially processed at UOA level (the disciplinary structure of the RAE/REF) and then a set of higher-level categories was used to overcome the problem that UOA count decreased in successive cycles. RAE1992 and RAE1996 data can nominally be reconciled to the REF2014 Main Panels but, where aggregation was required, data were aggregated into domains

¹⁵ A panel chair to RAE1996 asserted that submissions were 'representative' rather than 'selective' but the RAE Guidance refers to excellence and REF2014 Guidance on Submissions is explicit about quality profiles.

¹⁶ <http://impact.ref.ac.uk/CaseStudies/SearchI.aspx>

¹⁷ Crossref (<http://www.crossref.org/>) is a not-for-profit membership organisation for scholarly publishing, which makes content easy to find via metadata tags.

driven by similarity in publication usage¹⁸. The four domains (biomedical and physical sciences; engineering and technological sciences; social and economic sciences; and humanities, and visual and performing arts) differ slightly from the REF Main Panel structure, which combines physical sciences with engineering (REF Panel B), leaving biomedical and clinical sciences as a separate entity (REF Panel A).

17. For HEI analyses, the data were aggregated to the set of HEIs that currently exists. The number and structure of HEIs has changed over the period of analysis, with new foundations, some mergers (e.g. Manchester and UMIST¹⁹) and some splits (e.g. University of London).
18. **Other evidence.** Case study discussion of impact (Section 4) is supported by relevant evidence (Section 5: Sources to corroborate the impact). Some of this is documentary, but not typically in a form that enables ready analysis, and some is testimonial from the impact beneficiaries. There are also links to websites that either reference the beneficiary or provide additional background information supporting the impact description. For this report, all text elements that resemble hyperlinks (for example, with the prefix [http](#) or [www](#)) were extracted, truncated and aggregated by domain. Many examples use URL shortening services (such as [http://bit.ly](#)) and where feasible these were tracked to the originating source.

Results

The balance of submitted output types changes over time

19. Assessment and evaluation have summative and formative effects. There is widespread evidence that the behaviour of individuals and organisations responds to assessment. Indeed, assessment may be intended to stimulate a particular response.²⁰
20. For each RAE cycle, the managing bodies publish a summative report on the research performance of UK institutions and use the indexed outcomes to determine future funding. The reports and funding have a formative influence, encouraging behaviour around more, better research.²¹
21. The present data reveal another response: a changing diversity of output types is observed when the submitted outputs are aggregated at the level of science, engineering, social science and humanities. The format for RAE1992, which established the system of four outputs per person, was published only in March 1992 so the first cycle's data should not be over-interpreted. RAE1992 required only two publications per researcher (plus two other outputs) so there are relatively more books in science and engineering, though the absolute number is similar to RAE1996. Nonetheless, for 1992 and 1996, it is apparent that a preference (in the sense of making a relatively frequent selection) existed for journal articles among scientists; engineers

¹⁸ The analysis underpinning this was originally developed for RAE1996 data and based on clustering UOAs according to similarity in journal frequency. This therefore also subsumes broader output differences.

¹⁹ UMIST- University of Manchester Institute of Science & Technology merged with the Victoria University of Manchester in 2004.

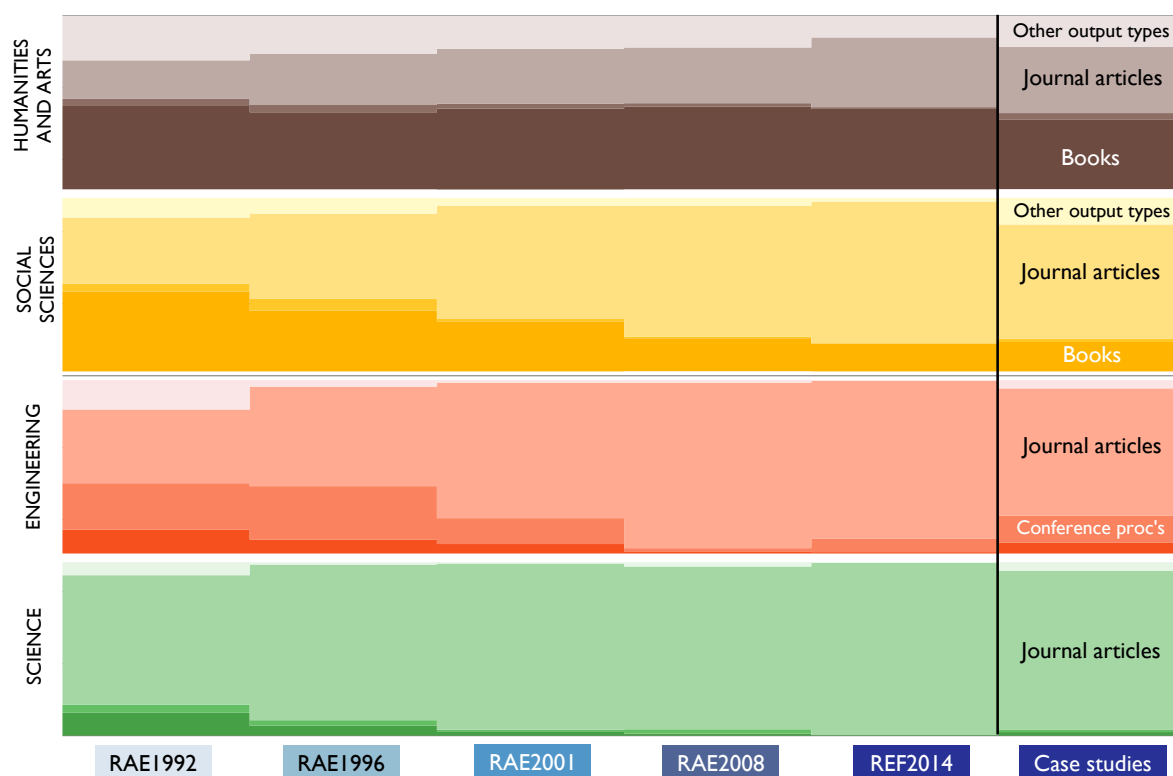
²⁰ Assessment is rarely a neutral process. Response of individuals and organisations under assessment is widely documented (see e.g. sources listed at <http://www.unm.edu/~jka/qual/qualbib.htm>).

²¹ Adams, J. and Gurney, K. A. (2010). Funding selectivity, concentration and excellence, pp 1-20. Higher Education Policy Institute, Oxford. <http://www.hepi.ac.uk/2010/03/25/funding-selectivity-concentration-and-excellence-how-good-is-the-uks-research/>

preferred proceedings; the social scientists and scholars in the humanities preferred the monograph; while visual and performing arts used specialised media.

22. Neither the RAE manager nor any UOA panel ever published advice indicating a preference for particular types of output. Nonetheless, evidence of change emerged as early as RAE1996, with engineers' preference shifting from conference proceedings and social scientists shifting from books, the traditional publication modes common in RAE1992, towards journal articles. The shift went further in 2008 and in 2014 when journal articles became the predominant submitted output in all areas except humanities, which remain focussed on books. The humanities/arts shift in frequency from 'other outputs' to 'journal articles' between 2008 and 2014 is driven by changes in arts researchers' submitted outputs.
23. The diversity of output types seen in the REF2014 case study references is more diverse than the REF2014 submitted outputs, in the sense of including a greater proportion of monographs, conference proceedings and non-print outputs. It resembles the balance of output types seen in RAE1996 more closely than it does REF2014. (Figure 1; Annex 2 – Table A)

Figure 1. Output types (darker to lighter colour-bands show: books & chapters, conference proceedings, journal articles, other output types) as a share of total outputs submitted as evidence of quality in five successive UK research assessment cycles and as references in impact case studies. Subject groups are clusters of a variable number of UOAs in different cycles. (Data in Annex 2 – Table A)



24. Among the increasingly frequent journal articles, citation awareness may influence researcher behaviour in selecting outputs for assessment if there is a perception that a journal's citation impact status (Journal Impact Factor or JIF²²) is a substitute 'signal' of article quality²³. There are

²² Pendlebury, D. A. and Adams, J. (2012). Comments on a critique of the Thomson Reuters journal impact factor. *Scientometrics*, 92, 395-401. DOI: 10.1007/s11192-012-0689-6

²³ Among RAE2008 submissions, some high-JIF journals had more submitted records than there were UK-addressed publications for the census period (i.e. every author chose to submit such items). See: Adams, J. and

two questions to test: first, are articles from high-JIF journals submitted disproportionately for assessment; and, second, has this tendency changed in later assessment cycles? To do this, the frequency of journals among UK-authored articles must be compared with the frequency of those journals among submitted outputs. Over the time from 1996 (the first year of the RAE2001 census period) to 2013 (the last year for REF2014), many journals disappeared, research fields evolved, international engagement expanded, new journals appeared, and some leading journals either increased their page count or spawned sister serials. There is thus a challenge in assembling like-for-like literature. An exploratory study, covering RAE2001, RAE2008 and REF2014²⁴, examined a set of journals that were continuously present, that included both high-frequency journals and high-JIF journals (typical JIF varies by field) used by UK researchers, and that collectively included about half of UK journal output in each RAE census. Analysis revealed no statistically significant evidence of behavioural change.

- a. In all cycles, UK researchers ‘preferentially’ submit outputs from high-JIF journals (higher submission/UK publication ratio).
- b. Over time, this preference has not become more marked.

Submitted output publication dates are skewed

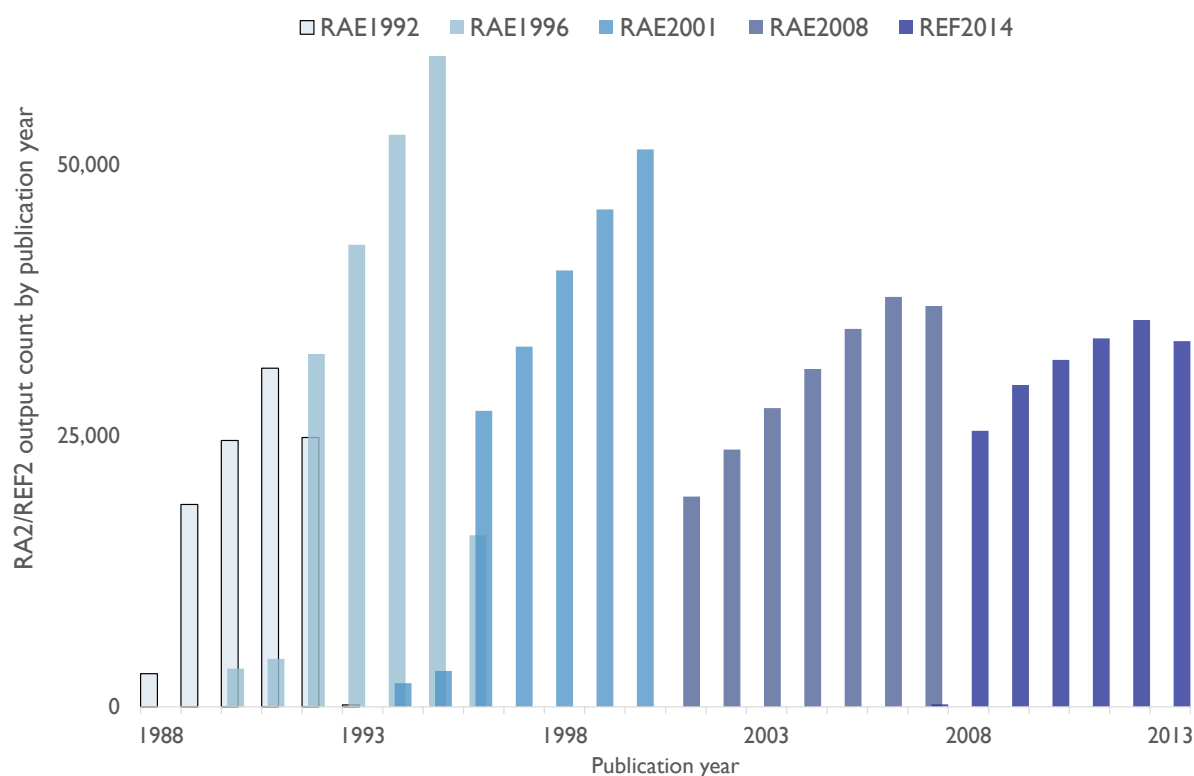
25. The publication dates of RAE submitted outputs are skewed towards the more recent years in each cycle. The time-skew appears to be a general phenomenon, though previously unnoticed.²⁵
26. All other things being equal, the spread of submitted outputs by time within an RAE census period should be uniform. Contrary to this expectation, RAE1992 data reveal a very marked time-skew of submitted outputs towards the most recent publication dates for that cycle (1992 was a ‘half-year’ with a census cut-off at 30 June). This skew persisted in RAE1996 and later cycles but this analysis shows that it was gradually modified. Note that a change in the number of years in each cycle (from a core of four years to five to seven to six) affects the height of the curve and that two – not four – publications plus ‘two other outputs’ were required per researcher in RAE1992.

Gurney, K. A. (2014). Evidence for excellence: has the signal overtaken the substance (analysis of journal articles submitted to RAE2008). Digital Research Report, pp 1-8. Macmillan, London. ISBN 0 992947705

²⁴ Data for RAE1992 and RAE1996 are not available at record level to support this analysis.

²⁵ This is interesting because the publication dates of REF2014 case study references were skewed to recent years (King’s College London and Digital Science (2015). The nature, scale and beneficiaries of research impact: an initial analysis of REF2014 impact case studies. Research Report 2015/01, Higher Education Funding Council for England). This is counter to a general assumption (following prior work based on Griliches and Mansfield) that research linked to impact should draw on an extended period of development.

Figure 2. For each UK research assessment cycle, the count of submitted outputs published in each year of the relevant census period (see text for notes on varying census periods).



27. The last full year of the census period provided the greatest number of publications submitted for assessment in RAEs 1992, 1996 and 2001. RAE1992 allowed publisher-accepted outputs in the year of assessment and RAE1996 allowed submissions to March 1996, but 1991 and 1995 are the peak publication years: 1995 has twice the output volume of 1992, the first year of the RAE1996 cycle, and, in RAE2001, 2000 has almost twice the volume of 1996. In RAE2008, however, the penultimate year (2006) is the peak publication point for submitted outputs. In REF2014, the time profile is more evenly distributed and 2011 and 2012 publication volumes both exceeded 2013. (Figure 2)

28. RAE administrators did not report the pattern in Figure 2 at the time (paragraph 22); it was not recognised by university research managers; and it does not appear to be an organised submission practice²⁶. Nor is there evidence that a shortage of material constrained the selection of outputs for submission. First, HEIs reported that submitted outputs for RAE1992 were drawn from a total pool of 787,138 potentially eligible outputs²⁷; second, the relatively selective Thomson Reuters *Web of Science* database records 90,000 UK-authored journal articles indexed per year so the sum of these across each census period would exceed requirements for that RAE cycle. Third, non-indexed journals, conference proceedings, books and other forms of output would add to the indexed article total. The selection of submitted outputs appears therefore to be deliberate choice, initially skewed towards most recent outputs in RAE1992 and yet subject to modification over successive cycles.

²⁶ Several senior research and institutional managers were consulted about the results reported here in summer 2015. They confirmed that the recency skew in submission publication dates was unknown to them and had not been reported in institutional data or discussed in subject-based conferences on the RAE.

²⁷ RAE1992 - RAE Data RAI - Active Research Staff Return (HEFCE)

Behaviour in selecting submitted outputs is consistent across UOAs and HEIs


29. The net values for the whole database in each cycle (Figure 2) capture the aggregate of many independent HEI and UOA submission choices. The expressed preferences may be homogeneous and generic or the profile may be the smoothed outcome of combining many heterogeneous variants from disciplines with varying cultures and from HEIs with different missions.
30. To explore whether the observed profile is homogeneous across UOAs and HEIs, two comparative analyses were created by calculating the percentage of total outputs by publication year for each UOA in each cycle and for each of the 150+ submitting HEIs. The median value in each year and the upper and lower quartiles bounding that median indicate the spread across units (UOAs and HEIs).
31. This analysis reveals that the 1989-2013 profiles for UOAs (Annex 2 – Figure A) and HEIs (Annex 2 – Figure B) are similar both to one another and to the overall picture (Figure 2). There is remarkably low variance, with quartiles around the median that are relatively tightly bounded compared to the year-to-year change in medians, and with a narrowing inter-quartile spread in later cycles.
32. Recall that these submitted outputs are the combined selections of four outputs for each of 50,000 individual researchers across institutions and subject areas. As noted (paragraph 28), the skew was not reported and institutional managers recalled no institutional awareness. Nonetheless, many independent choices have produced a well-defined time-skew, consistent across disaggregated UOA and HEI units, that moderates in a common fashion across subjects, institutions and cycles.

There is a cultural divide in the selection of submitted outputs

33. Initial cultural distinctions in the selection of outputs submitted for assessment are displaced by a convergence on journal articles (Figure 1). However, the analyses underpinning Figure 2 provide statistical information that throws light on cultural divergence in other aspects of behaviour.
34. For outputs submitted by year of publication, the average across the dataset is lower than median values for individual UOAs and HEIs (Annex 2 – Figures A and B). This is due to a volume factor, where the relatively large volumes of submitted outputs for a small number of UOAs influence the overall average. The disparity implies some UOA-related difference that then affects HEIs according to the subject balance of their portfolio. There is also greater variance towards the end of later cycles (the inter-quartile range of medians increases), implying an emerging difference.
35. To explore the source of these disparities, the data were re-aggregated into the four broad subject groups (medical and natural sciences; engineering and technology; social sciences and business; and humanities and arts, as in Figure 1). At this group level (see Figure 3):
 - a. In RAE1992, profiles are similar but humanities and social sciences are more ‘recent’ to a similar extent, with the six-month 1992 output count approaching the full year of 1991.
 - b. In RAE1996, the profiles appear similar across subject groups (the humanities line is lower only because of data spread across two additional census years).

- c. In RAE2001, a slight terminal inflection appears for science and engineering.
- d. In RAE2008, a divergence between ‘arts’ and ‘sciences’ becomes clear.
- e. In REF2014, science and engineering peak across two penultimate years and dip in the final census year while social sciences and humanities/arts retain the skew of earlier cycles.

Figure 3. The percentage of outputs submitted in each UK research assessment cycle that were published in each relevant census year, disaggregated by major subject groups (as Figure 1). [Curve heights are affected by variable year count in each cycle (see Figure 2); terminal drop in RAE1992, RAE1996 due to publication cut-off at 30 June 1992 and 31 March 1996; early years allowed for art/humanities in RAE1992-2001].



36. The ‘science’ profile, submitting less recent outputs in the early cycles and later shifting towards a more even spread across years, does not occur in social sciences or in humanities. The separation between ‘science/technology’ and ‘social science/humanities’ is unexpected, as is the marked similarity within each pair. And, again, it should be recalled that both skew and subsequent change are the outcome of many independent submission choices.


The earliest case study references underpinning impact are evenly distributed by year for science and engineering

37. An overall time-skew has been reported for the case study references²⁸ and it is now clear that this skew is similar to that demonstrated for submitted outputs (Figure 2). There is more subtle differentiation, however, in the dates of the earliest references in science and engineering.

²⁸ King’s College London and Digital Science (2015). The nature, scale and beneficiaries of research impact: an initial analysis of REF 2014 impact case studies. Research Report 2015/01, Higher Education Funding Council for England.

38. Impact case study references could draw on a census window to 1993. Those selected might be spread over an extended period, or they might emphasize earlier research from which later developments arose, or instead they might focus on the recent achievements of current staff²⁹. In fact, the general profile of references (combining all output types) was time-skewed. Of 36,244 unique references in the case studies, there were 20,779 with publication dates in the same census period as REF2014. Another 11,000 were published in the RAE2008 census period. Fewer than 7,000 (20%) were published before 2003, in the first half of the available window.
39. However, impact ‘origin’ is better signalled by the earliest reference publication date for each case study, which can be identified separately. The complete time-spread for all case study references does not describe the specific origin of intellectual property underpinning any particular case study. REF2014 required case studies of impact occurring within the 2008-2013 period, but impact emerges as research develops and matures over time.

Figure 4. The spread by publication year of the earliest of six supporting references in each case study grouped by the four main panels of REF2014; case studies were expected to report impact that occurred during the REF2014 census period from 2008, which explains the sharp drop in earliest references with a later publication date.



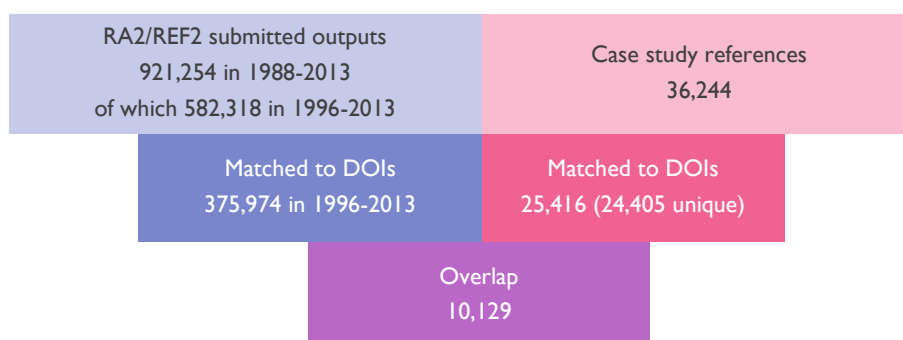
40. Disaggregation of the earliest reference dates in an impact case study by the four REF Main Panels (Figure 4) reveals a cultural divide akin to that for submitted outputs’ data. The earliest dates for science and engineering are spread relatively uniformly across 1993-2008 whereas those for social sciences and humanities are skewed to relatively recent years (as in standard assessment, Figure 3).

²⁹ The REF2014 eligibility criteria dictated that impact case study references referred to research conducted at the submitting institution only. This contrasts with submitted outputs, which were *ad personam* so outputs travelled with authors. Consequently, there is a possibility that case study references are more likely to be the recent output of the research staff currently in post who developed the case study document.

Overlap between submitted outputs and impact references is even across years

This section refers to documents, usually journal articles, matched algorithmically to DOIs to enable analysis of the overlap between the sets of DOIs in submitted outputs and references underpinning impact case studies. A more intensive scrutiny would increase the match rate and thus the overlap so this is a minimum estimate and may be more effective for disciplines that have a particular reliance on journals as a reporting mode. The REF2014 eligibility criteria dictated that submitted outputs were ad personam, so outputs travelled with authors. Impact case study references referred to research conducted at the submitting institution only. Consequently, there will be some research conducted by researchers who moved HEI and were thus eligible for submission to the REF but not eligible for inclusion in an impact case study for the new HEI. Such data disparities may affect the likelihood of an automated match.

41. The specific overlap between the time-based profile of the REF2014 case study references and the submitted outputs of their own and earlier cycles can be compared for outputs that have matched DOIs. DOIs became common around 2000 but it was feasible to look for some matches back to 1996 from the RAE2001 data. DOIs were included more comprehensively in the RAE2008 and REF2014 data. Despite early sparsity, and non-journal output types, it was possible to link 25,416 reference records to DOIs. Of these, around 1,000 are duplicates between UOAs or HEIs. This means that there are DOIs for 24,405 (about two-thirds) of the unique case study references.

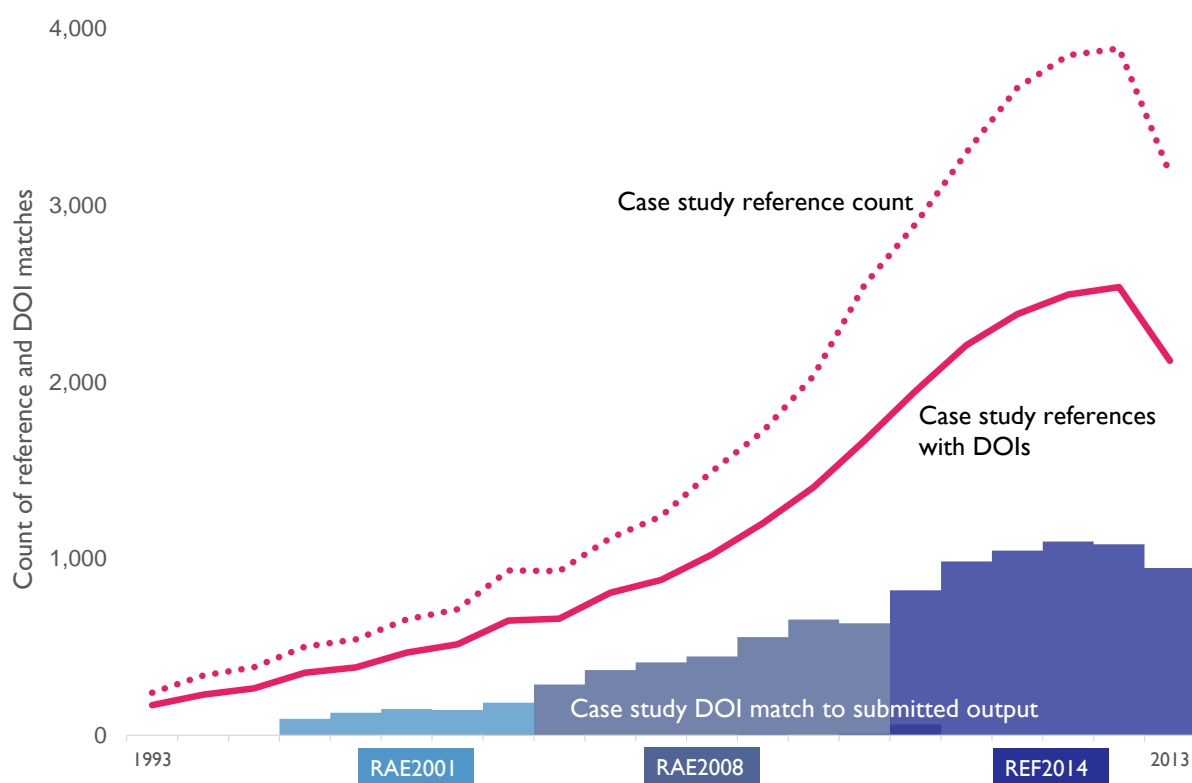


42. Analysis of the records matched to DOIs shows that when impact case studies drew on outputs previously submitted as evidence of academic excellence they did so at a remarkably consistent rate across 18 years (Figure 5):

- a. If all case study references with DOIs were matched to submitted outputs, then the maximum overlap would be 6.76% of the submitted outputs with DOIs. The actual overlap is 2.7%.
- b. Overall, 42% of the case study references with matched DOIs were also submitted outputs in one of the RAE/REF cycles.
- c. By assessment cycle, the overlap of total case study references is:
 - i. 3% overlap with outputs submitted to RAE2001
 - ii. 14% overlap with RAE2008
 - iii. 25% overlap with REF2014.

- d. By year, for documents with DOIs, the average annual overlap is about 42% of the total unique count of DOI-matched documents for 1996-2013, with a small annual variance.
 - i. For 2001-2013 (the RAE2008 and REF2014 census periods where DOI assignments are comprehensive) the median overlap is 44% and the inter-quartile range is just 43.5 to 45.7.
 - ii. For 1996-2000, a lower overlap rate extends the interquartile range down to 34.3% overlap but there are only sparse DOI matches for the outputs submitted to RAE2001.

Figure 5. Analysis of publication dates for references included in impact case studies submitted to REF2014. Bars show: the total number of references by publication year; the number of references to which DOIs could be matched; and the number of those references with DOIs that could be paired with outputs submitted in each assessment cycle.



Overlap between submitted outputs and impact references varies by UOA


- 43. The overlap between submitted outputs and references varies by UOA between less than 10% and over 60% of the case study references, with an average of 42%. The greatest degree of overlap appears to be in subjects allied to professional practice whilst overlap in fundamental research areas in the humanities and core natural sciences is smaller. (Table 1, Figure 6).³⁰

³⁰ Analysis of overlap by UOA for case study references uses their relevant REF2014 UOAs. Calculation of overlap by UOA for submissions is not meaningful because the UOA list varied between cycles.

Table 1. For each REF2014 Main Panel: the UOAs with the highest and lowest percentage overlap between REF impact case study references and RAE/REF submitted outputs, for outputs with DOIs.


Main Panel	UOA with greatest overlap	%	%	UOA with least overlap
A	Public health, health services and primary care	51.7	36.9	Biological Sciences
B	Civil and construction engineering	56.5	31.5	Chemistry
C	Education	53.9	30.9	Law
D	Communication, culture, media, library and informatics	40.0	8.5	Classics

Figure 6. Overlap by UOA (in standard sequence by Main Panel) for outputs that could be matched via DOIs. Reference line: average percentage (42%) of case study references also submitted as outputs to RA2/REF2.



44. When the data are disaggregated by assessment cycle, case studies from UOAs in science and engineering had the greatest overlap with RAE2008 whereas for UOAs in the social sciences the greatest overlap was with REF2014 submitted outputs. The highest relative overlap between impact references and outputs submitted to RAE2001, the oldest for which record-level data are available, is again in the science and engineering UOAs. (Figure 7)

Figure 7. Overlap by assessment cycle for each UOA, where a pairing could be made via matched DOIs, between case study references and outputs submitted to each of the research assessment cycles.



45. The pattern in Figure 7 may reflect the likelihood of an output submitted prior to 2001 being a journal article with a DOI. Electrical engineering had 7% overlap with RAE2001 and both Biological sciences and Physics were over 6% compared to the 3% average. No social science or humanities UOA exceeded 2% and most had less than 1% overlap with the earliest cycle.

Overlap between submitted outputs and impact references varies between HEIs

46. At HEI level, the average 42% overlap between case study references with matched DOIs and submitted outputs varied between 30% and 50% for about two-thirds of institutions including most of the pre-1992 universities. There is a spread of higher overlap margins for a few institutions, of which some are specialist, at well over 50%.

47. There is a relationship between an HEI's volume of case study references with matched DOIs and the overlap by assessment cycle between references and submitted outputs. (Annex 3 – Figure C). This trend will be influenced by institutional subject mix, research mission and the availability of candidate outputs. A dearth of choice would necessarily lead to greater overlap, but so would a research mission focussed on development and application (greater reference/submission intersection seen in professionally linked UOAs: Table 1). A research mission focussed on fundamental research, by contrast, may tend to reduce overlap.

48. Data disaggregated at sub-sector group level reveal higher overall overlap and a greater concentration of case study references among REF2014 submitted outputs for post-1992 HEIs compared with 'big civic' HEIs. (Annex 3 –Table B)

- a. Older, research-intensive HEIs have less than 40% total overlap, so the references on which they draw for the case studies are relatively distinct from their submitted outputs.

- b. Case study references of older HEIs overlap slightly more with RAE2001 than for the sector as a whole.
- c. Post-1992 HEIs have a total overlap close to 48% of which two-thirds is with recent REF2014 submitted outputs.
- d. There are outliers: several post-92 HEIs had greater than 6% overlap with RAE2001 submitted outputs, as did technology-focussed HEIs such as Imperial College and Strathclyde.

Case study analysis provides additional evidence of impact

49. Section 5 of the impact case studies ('Sources to corroborate the impact') contains many hyperlinks. These point to over 12,000 unique domain references, of which about 3,200 are referenced more than once (not necessarily in different case studies as these are truncated domain names and may point to different source pages within the same domain). These data reflect a diversity of beneficiary organisations, notably in the public sector but including a large number of companies. They are discussed further in Annex 4 –Table C. The many other forms of evidence cited in case study Section 5 have not been analysed.

Discussion

50. This is the first analysis of evolving patterns in the use of published outputs as evidence of research achievement where comprehensive data were available for research published over more than two decades, covering all disciplines and institutions in a national system, with selective submitted outputs to five assessment cycles and for two (complementary but distinct) assessment objectives.
51. The questions set at the outset were:
 - a. What is the relationship between submitted outputs (as evidence of academic impact) and case study references (as evidence of societal and economic impact), noting that academic impact has been conceptually associated with exceptional research innovation?
 - b. What do the data say about the selection of outputs for assessment?
 - c. How did selection behaviour respond to assessment?
 - d. How homogeneous was that response, by discipline and institution?

The relationship between submitted outputs and case study references

52. There is considerable overlap between references to research underpinning the REF2014 impact case studies and outputs submitted in prior and current assessment cycles. The overlap between submitted outputs (the RA2/REF2 section of the RAE/REF submissions) and case study references is evidence that work of fundamental quality regularly leads to a wider range of benefits for wealth creation and the quality of life.
53. The overlap accounts for 42% of case study references matched to a DOI, and it is reasonable to infer that a similar overlap would apply to other outputs if they were manually curated. The overlap is very consistent across years, close to 43% in every year between 2001 and 2013. This remarkably uniform time-spread of overlap occurs despite the evolving research priorities of different agencies and the fact that researchers in early cycles had no inkling of impact assessment. (Figure 5)
54. The case study references are from more diverse output types than recent submitted outputs (Annex 2 –Table A; the reduction in diversity among submitted outputs is discussed below). This comparative diversity is a signal that the selection of outputs as evidence is deliberate and is modified according to objective. In particular, whereas journal articles dominate submitted outputs, case study references include a diversity of other output types such as grey literature (e.g. reports) and patents, both of which are likely to be allied to application and to socio-economic and technological impact.
55. The overlap between submitted outputs and references varies by UOA. It is lower in some basic research areas and is higher in areas related to professions (Table 1). It is also lower in the humanities (Figure 6). Divergence may indicate different concepts of linking research to impact, but further work is required to develop an understanding of the ways in which each discipline constructs this relationship.
56. There is a relative abundance of grey literature among references in social sciences, which points to the possibility that the degree of overlap between submitted outputs and case study

references may be associated with impact types. The hypothesis that overlap could be – for example – higher for technology and health impacts and lower for cultural and policy impacts is entirely testable.

The selection of references used in impact case studies

57. There is an interaction between submission profiles and publication dates, for both case study references and submitted outputs. The submitted outputs deployed in early RAE cycles show a marked recency-skew in selections (Figure 2) echoed later in the overall profile of choices for impact references (Figure 5). However, more detailed analysis shows that this broad picture contains important and informative components about the selection of evidence.
58. The skew for impact references in science/engineering disappears when the references are filtered so just the earliest ‘source’ publications are identified and profiled, though it remains for social science/humanities (Figure 4). We may therefore infer that the earliest references better reflect the historical origins of the impactful intellectual property (IP), while the bulk of references – which also overlap later assessment cycles – are affected by the general recency influence, and are likely to refer much more to explication and development of that source IP.
59. There is a 20-year period allowed for REF impact case study references, which is longer than the census period for RA2/REF2 submitted outputs. This longer time-frame was based on the classic work of Griliches and Mansfield³¹ on innovation cycles and the time indicated for research and development (R&D) to move from innovative knowledge to new processes and products. There have been rather few authoritative studies on this³² because the many-to-many relationship between research discoveries and commercially successful outcomes makes it extremely difficult to track, and apportion value across, the interlocked pathways.
60. The data on REF case study references only partly support this classic expectation, that impactful research typically has a distant origin. (Figure 4)
 - a. The lag between innovation and impact is evidently variable and there are indications that it is context-related.
 - b. The time span of original ‘source’ IP in the earliest references supporting case studies covers the full assessment period and does so at a steady rate for science and engineering. (Figure 7)
 - c. Almost exactly the same proportion – around 42% – of REF2014 case study references for any publication year had previously been submitted (quite independently) as evidence of excellence, back to 1993.
 - d. The timeframe for evidence seems to interact with impact type and discipline: drawn-out development might be the case for longitudinal social research as much as for science, whereas rapid transition to impact may occur for social sciences, visual/performing arts and professional health research closely engaged in studying, analysing and reflecting the dynamics of their domains.

³¹ See references in Introduction, footnote 3

³² But see: Wooding, S., Hanney, S., Pollitt, A., Grant, J. and Buxton, M. (2014). Understanding factors associated with the translation of cardiovascular research. *Implementation Science*, 9, 201-209.

61. This may influence thinking about what can and should be seen as impact exemplars. Further analysis of case studies with older and younger origins might be informative. There is evidently more work to be done to deconstruct and develop the expectations of both researchers and beneficiaries about the dynamics of the beneficial impact of academic research. No prior analysis of the profile of research impact across many disciplines has previously been available. These analyses will be useful to broader policy development but they are likely also to influence academic thinking about achievable impact for future assessment.

The selection of submitted outputs used in research assessment

62. Two characteristics to the selection of submitted outputs have progressively changed across RAE cycles. These are most readily interpreted with reference to disciplinary differences that help to throw light on the motivation for the changes.
- a. A marked time-skew towards the most recent publication years, which moderated in later cycles. (Figure 2)
 - b. A distinctive diversity of output types across subject domains, which generally converged on a common preference for journal articles. (Annex 2 –Table A)
63. First: the **skew across publication years** within RAE1992 and RAE1996. This peaked at 30% of the total items selected from just the last full years (1991 and 1995). Recency, a tendency to choose recently experienced items, is a widely recognised phenomenon, but there is no obvious reason for academics to select their most recent publications. It may be the case, however, if each research step is built on earlier steps, then later work is more ‘advanced’ (absorbing the content of earlier outputs) as well as having novelty appeal.
64. The time-skew in output choice is narrowly bounded when disaggregated by discipline and institution, indicating a strong homogeneity across academic culture. Such homogeneity has been widely supposed but it is rarely evidenced³³. The analyses here show that, without discussion or coordination, cognate selection patterns were made across 150 institutions and several dozen UOA subject categories. The patterns progressively changed in each RAE cycle, while remaining entirely coherent and with a narrowing spread across UOAs and HEIs in later cycles. (Annex 2 – Figures A and B)
65. Despite this general ‘dynamic homogeneity’, a cultural divide then appears and grows between science/technology and social science/humanities. Within these two cultural pools, the correlation in time-spread of output selection between subjects remains very tight. (Figure 3)
66. Second: the **diversity of output types** submitted to RAE1992. Recall the initial relative frequency of output types: scientists favoured journal articles; engineers favoured conference proceedings; and social scientists and humanities researchers submitted more monographs. These early differences indicate positive selection and were presumably driven by culturally prioritised publication modes. In other words, while such results are not unexpected, they create a reference point which shows that RAE1992 submitted outputs were active and discipline-specific choices.

³³ Bourdieu, P. and Passeron, J-C. (1990). *Reproduction in Education, Society and Culture*. Sage Publications. ISBN 0803983204

67. During later assessment cycles, the selection of submitted outputs shifted, and did so across all areas, towards journal articles until they became the most common output type in almost every UOA. (Figure 1, Annex 2 –Table A)
68. The RAE datasets may reflect a change in what was chosen for submission, a change in how researchers went about publishing their work, or a signal of a wider shift in research culture to a science model of ‘impactful’ publication – possibly influenced by a global view of communication. Discussions with engineers suggested that the change from proceedings to articles was partly cosmetic, in that many outputs were published in both formats³⁴. Computer scientists continue to be less journal-oriented and less change in submitted outputs occurs in that UOA. Social scientists were, by the mid-1990s, already engaged in a shift towards larger projects involving teams, centres and longitudinal studies; articles were a more flexible output mode than monographs. Journals also increased visibility outside the UK, notably in North America, and were part of a general cultural shift across many countries towards international engagement and a rising tide of international research collaboration³⁵. Research assessment may have beneficially accelerated such change.
69. However, it seems likely that influence also came from the growing awareness of citation analysis and its focus on journal articles (because other outputs are not included in commercial citation databases), which links changes in choice of output type and choice of publication date.
- a. The first major bibliometric analyses of the UK’s comparative international research performance were published after RAE1996.³⁶
 - b. The analyses drew attention to citation data as an indicator of relative academic impact (frequently over-interpreted as achievement or excellence). This raised general awareness of the availability and use of such indicators.³⁷
 - c. A metric value was attached to academic journal articles that was not available for other output types, and in practice was only available for the journals selected for inclusion in commercial databases.
 - d. ‘Quantitative prioritisation’ of a particular type of output almost certainly accounts for the change in submission choices (not the same as actual publication patterns) made by engineers.
 - e. Cultural pressure may also have been felt by social scientists, whose indexed journal space was dominated by the USA. It was also apparent that citation rates varied by discipline and were generally so much lower outside mainstream science that analytical data were less relevant in, for example, social science.³⁸

³⁴ Discussion with IEEE Secretary, Professor Roger Pollard.

³⁵ Adams, J. (2013). The fourth age of research. *Nature*, 497, 557-560.

³⁶ May, R. M. (1997). The Scientific Wealth of Nations. *Science*, 275, 793-796. Adams, J. (1998). Benchmarking International Research. *Nature*, 396, 615-618.

³⁷ Adams, J., Bailey, T., Jackson, L., Scott, P., Small, H. and Pendlebury, D. (1998). Annex and summary report. Pp 5-19 in *Research funding: introduction of a policy factor*. HEFCE Consultation paper 98/54, HEFCE, Bristol. Adams, J., Cook, N., Law, G., Marshall, S., Mount, D., Smith, D., Wilkinson, D., Bayers, N., Pendlebury, D., Small, H. and Stephenson, J. (2000). *The role of selectivity and the characteristics of excellence - report of a consultancy study for the Higher Education Funding Council for England*. HEPU, University of Leeds. 72 pp. ISBN 0 85316 211 5. Adams, J. (2002). *Research Assessment*. *Science*, 296, 805.

³⁸ Adams, J. and Jackson, H. L. (2004). *Bibliometric profiles for RAE outputs in the Social Sciences*. Report to the Economic and Social Research Council, September 2004 (41 pp).

70. The moderation and flattening of the time-skew is affected by citation awareness because citations draw attention to more and less cited articles within journals as well as highlighting high averages (analysis showed no statistically significant change in selection in favour of articles from journals with a high impact factor: see paragraph 24). Citations accumulate over time and recently published items have no or few citations whereas those published earlier are demonstrably valued, or ignored. Attention – and thus choice – shifts towards individual articles that can be given a measure of ‘objective’ relative impact within their field.
71. When the attention is on socio-economic impact rather than academic impact, in case study references, the greater diversity of output type compared to submitted outputs at REF2014 (Annex 2 –Table A) captures a ‘reversion’ to the balance of output types seen in RAE1996. There is also a substantial overlap between case study references and submitted outputs throughout the earlier as well as recent years (Figures 5 and 7). This raises the possibility that the historical output type portfolio was in practice relatively successful in capturing a balance of impacts, and that is why these publications are recaptured now in impact case studies.

The homogeneity of response and its implications for peer review

72. The homogeneity of selection made across HEIs and UOAs (Annex 2 –Table I; Figures A and B) suggests a powerful consensus within the sector, and perhaps more widely across academia internationally, about what selection of outputs would best represent academic impact. There is also an evident consensus in the flexible response that appears as information sources change, with citation data becoming available and accepted.
73. Sector homogeneity may be evidence of the good sense of peer review, in that a community with such broad agreement can reasonably be subject to a peer evaluation that diverse parties would accept as equitable. This acceptance is a key part of entraining the supportive response that is essential for assessment of a sector where output and effectiveness are largely self-motivated and self-managed.
74. The data provide evidence of coherence across many academic units but distinct choices in ‘science’ and in ‘arts’. There is the apparent divide in response to assessment in terms both of academic impact (‘science’ UOAs shift towards less skewed time choices) and economic/societal impact (‘arts’ UOAs choose more recent evidence of origins in impact case studies). Thus, evidence of coherence and yet also of divergence can support the validity of peer-led processes while raising new questions about the processes that are applied in judging what is, and what appears to be, evidence of excellence.

The consultants

Digital Science is an operating division of the Holtzbrinck Publishing Group. Its consultants work closely with Nature Publishing Group, a sister company. Digital Science supports a growing range of analyses drawing on a diversity of data in the research process. Portfolio companies include Altmetric.com (on media attention to research publications), ÜberResearch (on funding), ReadCube (on publication content), figshare (on data, tables and figures within publications), BioRAFT (research management, compliance and productivity), LabGuru (on laboratory management) and Symplectic (on research management information).

The consultancy capability is led by Jonathan Adams, formerly Director of Research Evaluation at Thomson Reuters and founder of UK-based consultancy *Evidence* Ltd. Digital Science consultants focus on projects around research evaluation, whilst maintaining close links to academic research and institutions. Recent projects include a study for Nesta on the use of alternative metrics in relation to identification of innovative biomedical research; a review of data around conference proceedings for the Australian Research Council; work for HEFCE to analyse and host the REF2014 impact case studies; analyses on the interpretation of interdisciplinary research for the UK Research Councils; and on-going work on Science 2.0 for the European Commission.

A series of other reports discuss research data and analytical possibilities in a practical, applied context, drawing on the data and expertise within Digital Science, see <http://www.digital-science.com/blog/posts/digital-science-launches-digital-research-reports>.

Glossary

DELNI	Department of Education & Learning Northern Ireland
DOI	Digital Object Identifier
HEFCE	Higher Education Funding Council for England
HEFCW	Higher Education Funding Council for Wales
HEI	Higher Education Institution
IP	Intellectual Property
ISBN	International Standard Book Number
IEEE	Formerly, the Institute for Electrical and Electronic Engineering
JIF	Journal Impact Factor
RAE	Research Assessment Exercise
RA2 and REF2	RAE/REF Section 2: outputs submitted for assessment
REF	Research Excellence Framework
R&D	Research and Development
SFC	Scottish Funding Council
UFC	Universities Funding Council
UOA	Unit of Assessment
URL	Uniform Resource Locator

ANNEX 1 - Who chooses submitted outputs?

In this report, terminology has been used to avoid implications that the choice of outputs to be submitted for research assessment is made at any consistent point.

Decisions about the selection of outputs may be made by researchers, departmental managers, institutional managers or some combination of these. They may rely on peer views of excellence or draw on a variety of quantitative indicators. The balance of influence has likely changed over time and may operate in different ways by institution and perhaps by discipline.

No comprehensive survey of recent submission choice processes exists, and it would be difficult now to establish exactly what processes were used in RAE1992 or RAE1996. A rapid telephone survey of research managers produced a sufficient spread of responses to suggest that any generalisation would be unwise.


The terms 'choice' and 'selection' are therefore applied indicatively.

ANNEX 2 - Supplementary data


Annex 2 –Table A. The percentage of output types submitted by UK-based researchers in five successive research assessment cycles, including REF2014 impact case studies. There were two submitted publications per researcher in 1992 and four submitted outputs thereafter. Subject groups are clusters of a variable number of UOAs in different cycles. The output typology for the REF case studies uses a manual assignment.

	Science		Engineering		Social science		Arts, humanities	
	Outputs	%	Outputs	%	Outputs	%	Outputs	%
RAE1992								
Books and chapters	5,718	13.2	2,148	13.7	11,463	46.0	9,553	48.1
Conference proceedings	2,021	4.7	4,207	26.9	1,109	4.5	808	4.1
Journal articles	32,271	74.6	6,650	42.5	9,520	38.2	4,411	22.2
Other	3,258	7.5	2,658	17.0	2,806	11.3	5,106	25.7
RAE1996								
Books and chapters	5,013	5.8	2,405	8.1	16,185	35.1	22,635	44.4
Conference proceedings	2,657	3.1	9,117	30.8	3,202	6.9	2,133	4.2
Journal articles	77,037	89.8	16,951	57.3	22,575	49.0	15,135	29.7
Other	1,104	1.3	1,122	3.8	4,154	9.0	11,128	21.8
RAE2001								
Books and chapters	1,953	2.5	1,438	5.4	12,972	28.6	25,217	46.5
Conference proceedings	751	0.9	3,944	14.9	857	1.9	1,619	3.0
Journal articles	76,182	95.8	20,657	78.1	29,449	65.0	17,074	31.5
Other	618	0.8	408	1.5	2,008	4.4	10,345	19.1
RAE2008								
Books and chapters	1,048	1.2	216	1.2	12,632	19.0	21,579	47.6
Conference proceedings	2,164	2.5	326	1.8	614	0.9	897	2.0
Journal articles	80,203	93.8	17,451	95.4	50,163	75.5	14,543	32.1
Other	2,125	2.5	301	1.6	3,018	4.5	8,287	18.3
REF2014								
Books and chapters	228	0.3	197	0.8	8,307	15.9	18,168	46.3
Conference proceedings	81	0.1	2,056	7.9	233	0.4	380	1.0
Journal articles	73,039	99.1	23,521	90.9	42,545	81.5	15,749	40.2
Other	331	0.4	108	0.4	1,105	2.1	4,914	12.5
REF case studies								
Books and chapters	274	2.1	282	6.3	1,819	16.9	3,409	40.0
Conference proceedings	150	1.2	686	15.4	195	1.8	334	3.9
Journal articles	11,752	91.7	3,263	73.4	7,102	65.9	3,251	38.1
Other	631	4.9	213	4.8	1,649	15.3	1,523	17.9

Annex 2 –Figure A. For data analysed by UOA, the median and upper (Q3) and lower (Q1) quartile percentage of outputs submitted in each UK research assessment cycle that were published in each year of the relevant census period. There were 72 UOAs in RAE1992, 69 UOAs in RAE1996, 67 UOAs in RAE2001 and RAE2008, and 36 UOAs in REF2014. Despite this diminishing number, the inter-quartile spread is relatively small. It is relatively greater in the final years, particularly in the last two cycles.




Annex 2 –Figure B. For data analysed by HEI, the median and upper (Q3) and lower (Q1) quartile percentage of outputs submitted in each UK research assessment cycle that were published in each year of the relevant census period. There were about 150 HEIs in each cycle varying from large, research-intensive universities founded over 100 years ago to smaller specialist colleges and recent foundations. Despite this, the inter-quartile spread is relatively small throughout.



ANNEX 3 - Overlap between submitted outputs and references: patterns across HEIs

Annex 3 –Figure C. For every HEI that submitted to REF2014: variation across in overlap between case study references and (a) outputs submitted for assessment in the RA2/REF2 for RAE2001, RAE2008 or REF2014 and (b) REF2014 submitted outputs only (linear trend line is illustrative, not analytical).



The gap between the lines (in Figure C) reflects the share of DOI-matched references that were documents submitted in earlier assessment cycles, i.e. a greater gap means relatively more case study references were documents submitted for assessment prior to REF2014.

The older universities had lower overall overlap but drew relatively more on submitted outputs made in earlier RAE cycles and less on REF2014 submitted outputs. (Table B)

Annex 3 –Table B. The overlap for each RAE cycle between RA2/REF2 submitted outputs and impact case study references: system average plus the averages for three HEI sub-sectors, grouped by age of foundation and research intensity.

<u>Overlap (% of references with DOIs)</u>	RAE2001	RAE2008	REF2014	Total %
UK average	3	14	25	42.5
Large civic institutions (pre-1960)	3.20	14.8	21.5	39.5
Institutions established in the Robbins expansion	3.11	14.4	25.9	43.4
Institutions that became universities after 1992	2.31	13.2	32.2	47.7

ANNEX 4 - Hyperlinks in impact case studies

Section 5 (Sources to corroborate the impact) of each REF impact case study lists sources, usually external to the submitting HEI ³⁹, that corroborate specific claims made in the case study about the impact of the research. These sources were provided by the HEI and include published reports, reviews, accounts from individual beneficiaries and hyperlinks to websites.

The hyperlinks point to over 12,000 unique domain references, of which about 3,200 are referenced more than once. Table C (below) lists the 40 URLs that occurred at least 50 times and shows how they spread across REF Main Panels. Counts shown for one of the shortening services refer to links that it was not possible to resolve.

Public sector organisations and media sources are strongly represented and there are associations between domains and panels that make intuitive sense. For example, 'publications.parliament' is the second most frequent URL (there is also www.parliament.uk - row 18) and occurs mainly in Panel C (social sciences), presumably in relation to policy, whereas 'NICE' (and guidance.nice) are links to clinical excellence that are clustered almost entirely in Panel A. The most frequent link is the BBC for which almost half the references are in Panel D case studies, but the BBC is found in significant numbers in all panel areas.

Additional manual work could be done to aggregate variant URLs that represent the same parent organisation. For example, www.bbc.co.uk at the top of the table could be consolidated with news.bbc.co.uk (row 14), and the variant 'nice' and 'parliament' domain names were noted in text above. These references are independent in the case studies, so these are not duplicates of the same information but different routes to specific but related information sources. This occurs quite frequently with public sector sites, presumably because there are a variety of pages or areas on those sites that hold specific information.

URLs referring to private sector organisations are individually less frequent but there are many such links as only a few impact case studies refer to the same company. There are 23 links to a Rolls Royce domain and 12 to a GSK domain among the 200 most frequent links, but variants have not been consolidated so the full corporate totals will be higher. It is likely that significant supporting evidence of technology impact would emerge.

Table C is not an exhaustive indication of all the diverse evidence of research impact. It refers only to instances where a claim was corroborated with hyperlinks. A review of other documents listed in case study Section 5 is very likely to reveal further information of particular relevance to beneficiaries and thus to potential economic and social analyses of research impact.

³⁹ Links to university sites are quite common and may be referencing sources of additional explication or illustration of the impact discussion in the case study. The University of Edinburgh (row 4), the LSE (row 6) and Imperial College London (row 10) are particularly frequent.

Annex 4 – Table C. Frequent URLs found in impact case study (Section 5: Sources to corroborate impact). Data are disaggregated by Main Panel. In some cases, two or more URLs refer to the same parent organisation but these are all original references to specific items, not duplicates. Where the URL refers to bit.ly or another shortening service this has been tracked to source; the residual counts for these sites (e.g. row 17) refer to links that could not be resolved. This table lists only hyperlinks; other forms of evidence were also cited.

Row	Domain	Total	Panel A	Panel B	Panel C	Panel D
1	http://www.bbc.co.uk/	749	131	107	153	358
2	http://www.publications.parliament.uk/	494	95	45	311	43
3	https://www.gov.uk/	494	170	60	244	20
4	https://www.wiki.ed.ac.uk/	441	183	76	80	102
5	http://www.youtube.com/	392	60	66	55	211
6	https://apps.lse.ac.uk/	338	0	18	248	72
7	http://webarchive.nationalarchives.gov.uk/	291	125	23	132	11
8	http://www.nice.org.uk/	279	251	10	18	0
9	http://www.guardian.co.uk/	274	37	36	91	110
10	https://www.imperial.ac.uk/	259	53	206	0	0
11	http://www.theguardian.com/	243	24	20	66	133
12	http://ec.europa.eu/	190	41	27	101	21
13	http://www.scotland.gov.uk/	179	40	17	107	15
14	http://news.bbc.co.uk/	176	50	26	47	53
15	http://www.who.int/	166	139	8	12	7
16	http://www.telegraph.co.uk/	160	38	19	33	70
17	http://bit.ly/	147	16	20	72	39
18	http://www.parliament.uk/	114	16	14	65	19
19	http://guidance.nice.org.uk/	103	100	3	0	0
20	http://www.bis.gov.uk/	97	11	22	58	6
21	http://clinicaltrials.gov/	93	92	1	0	0
22	http://tinyurl.com/	93	36	23	14	20
23	http://www.independent.co.uk/	92	8	7	28	49
24	http://www.dailymail.co.uk/	90	43	17	13	17
25	http://onlinelibrary.wiley.com/	84	71	3	7	3
26	http://www.ncbi.nlm.nih.gov/	72	64	4	3	1
27	http://www.scottish.parliament.uk/	68	10	5	39	14
28	http://www.fda.gov/	68	56	11	1	0
29	http://whqlibdoc.who.int/	64	59	0	4	1
30	http://wales.gov.uk/	63	20	3	30	10
31	https://www.youtube.com/	60	8	19	12	21
32	http://www.nytimes.com/	58	8	10	14	26
33	http://www.official-documents.gov.uk/	57	11	7	33	6
34	http://www.nhs.uk/	56	47	3	4	2

Row	Domain	Total	Panel A	Panel B	Panel C	Panel D
35	http://www.dh.gov.uk/	56	36	4	13	3
36	http://www.esrc.ac.uk/	54	4	4	35	11
37	http://www.nature.com/	53	31	14	6	2
38	http://www.ema.europa.eu/	53	52	1	0	0
39	http://archive.defra.gov.uk/	52	19	5	28	0
40	http://www.economist.com/	50	0	11	21	18