AACR American Association for Cancer Research

THOMSON REUTERS

Evaluation of the American Association for Cancer Research (AACR) Methods in Clinical Cancer Research Workshop

November 2014

Contact details

Research Data & Services, **Thomson Reuters**

http://researchanalytics.thomsonreuters.com

1455 Research Boulevard, Rockville MD 20850
USA

77 Hatton Garden, London, EC1N 8JS
UK

Table of contents

1 Executive summary

This report was commissioned by the American Association for Cancer Research (AACR) from Thomson Reuters. It compares the research output, citation impact, clinical trial involvement and research collaboration of participants and unsuccessful applicants to AACR's Methods in Clinical Cancer Research Workshop.

While is it not possible to infer a causative relationship, participation in the Methods in Clinical Cancer Research Workshop is associated with:

- greater publication rates,
- increased citation impact,
- higher levels of clinical trial participation,
- increased collaboration, and
- larger collaborative networks.

The key findings of this report are as follows:

- Searching of the *Web of Science™* found a total of 31,261 publications between 1999 and 2013 which mapped to participants and applicants identified by AACR (456 participants and 480 applicants) (Section 4.1).
- Participants published at a higher rate during the three year period after the workshop, compared to the three year period preceding it (Sections 4.2.1 and 4.2.2).
- The output of workshop participants increased at a faster rate between the periods before and after attendance than the output of unsuccessful applicants in the same cohort (Section 4.2.3).
- The average citation impact of workshop participants increased between the before and after periods from around one-and-a-half times the world average (1.52) to nearly twice the world average (1.95). The citation impact of unsuccessful applicants remained stable between the two periods (1.29).
- Impact Profile® analysis showed the participants' improvements to be driven by an increase in highly-cited publications (from 8.0% of total output to 11.7%) (Section 4.3.1).
- The distribution of changes in citation impact of workshop participants and unsuccessful applicants did not differ significantly (Section 4.3.2).
- We identified a total of 917 trials to which 346 workshop participants and unsuccessful applicants had been linked (205 workshop participants and 141 unsuccessful applicants) (Section 5.1).
- The number of clinical trials conducted by participants following increased significantly between the periods before and after attendance of the workshop (Section 5.2.1).
- The workshop participants collaborated with more co-authors than did the unsuccessful applicants (Section 6.1).
- The workshop participants had larger collaborative networks and these grew more rapidly than those of the unsuccessful applicants (Section 6.2).

- The individuals in the collaborative networks of workshop participants were more closely connected than those in the networks of unsuccessful applicants (Sections 6.3 and 6.4).
- There appears to be no correlation between network clustering with workshop participation (Section 6.5).

2 Introduction

2.1 American Association for Cancer Research (AACR)

The American Association for Cancer Research (AACR) is the world's oldest and largest professional association related to cancer research. It focuses on all aspects of cancer research, including basic, clinical, and translational research into the etiology, prevention, diagnosis, and treatment of cancer. It was founded in 1907 by physicians and scientists, and now has more than 34,000 members in over 90 countries. The AACR publishes eight peer-reviewed journals, hosts an annual meeting with more than 250 invited presentations, funds research through the AACR Foundation for the Prevention and Cure of Cancer, and runs workshops and other education programs.

2.2 Thomson Reuters

Thomson Reuters is the world's leading source of intelligent information for business and professionals. We combine industry expertise with innovative technology to deliver critical information to leading decision makers in the financial, legal, tax and accounting, healthcare, science and media markets, powered by the world's most trusted news organization. Visit our WEBPAGE for more information.

Thomson Reuters Research Analytics is a suite of products, services and tools that provide comprehensive research analysis, evaluation and management. For over half a century we have pioneered the world of citation indexing and analysis, helping to connect scientific and scholarly thought around the world. Today, academic and research institutions, governments, not-for-profits, funding agencies, and all others with a stake in research need reliable, objective methods for managing and measuring performance. Visit our WEBPAGE for more information.

Thomson Reuters Custom Analytics & Engineered Solutions provide reporting and consultancy services within Research Analytics using customized analyses to bring together several indicators of research performance in such a way as to enable customers to rapidly make sense and interpret of a wide-range of data points to facilitate research strategy decision-making. We have extensive experience with databases on research inputs, activity and outputs and have developed innovative analytical approaches for benchmarking, interpreting and visualization of international, national and institutional research impact. Our consultants have up to 15 years' experience in research performance analysis and interpretation.

2.3 Outline of this report

The purpose of this report is to provide AACR with data to support it in understanding the value of its Methods in Clinical Cancer Research Workshop. It comprises a comparison of workshop participants' and unsuccessful applicants' research in the three years before and after the workshop. It analyses data for the cohorts attending (or applying to) workshops held during the years 2002, 2004, 2006, 2008 and 2010. It considers three aspects of research:

- Research publications and citations to those publications
- Clinical trials
- Research collaboration and networks

3 Methodology and data sources

3.1 Bibliometric data and citation analysis

Research evaluation is increasingly making wider use of bibliometric data and analyses. Bibliometrics is the analysis of data derived from publications and their citations. Publication of research outcomes is an integral part of the research process and is a universal activity. Consequently, bibliometric data have a currency across subjects, time and location that is found in few other sources of research-relevant data. The use of bibliometric analysis, allied to informed review by experts, increases the objectivity of and confidence in evaluation.

Research publications accumulate citation counts when they are referred to by more recent publications. Citations to prior work are a normal part of publication, and reflect the value placed on a work by later researchers. Some papers get cited frequently and many remain uncited. Highly cited work is recognized as having a greater impact and Thomson Reuters has shown that high citation rates are correlated with other qualitative evaluations of research performance, such as peer review. This relationship holds across most science and technology areas and, to a limited extent, in social sciences and even in some humanities subjects.

Indicators derived from publication and citation data should always be used with caution. Some fields publish at faster rates than others and citation rates also vary. Citation counts must be carefully normalized to account for such variations by field. Because citation counts naturally grow over time it is essential to account for growth by year. Normalization is usually done by reference to the relevant global average for the field and for the year of publication.

Bibliometric indicators have been found to be more informative for core natural sciences, especially for basic science, than they are for applied and professional areas and for social sciences. In professional areas the range of publication modes used by leading researchers is likely to be diverse as they target a diverse, non-academic audience. In social sciences there is also a diversity of publication modes and citation rates are typically much lower than in natural sciences.

Bibliometrics work best with large data samples. As the data are disaggregated, so the relationship weakens. Average indicator values (e.g. of citation impact) for small numbers of publications can be skewed by single outlier values. At a finer scale, when analyzing the specific outcome for individual departments, the statistical relationship is rarely a sufficient guide by itself. For this reason, bibliometrics are best used in support of, but not as a substitute for, expert decision processes. Well-founded analyses can enable conclusions to be reached more rapidly and with greater certainty, and are therefore an aid to management and to increased confidence among stakeholders, but they cannot substitute for review by well-informed and experienced peers.

3.2 Publications data

For this evaluation, bibliometric data will be sourced from databases underlying the Thomson Reuters *Web of Science*™, which gives access to conference proceedings, patents, websites, and chemical structures, compounds and reactions in addition to journals. It has a unified structure that integrates all data and search terms together and therefore provides a level of comparability not found in other databases. It is widely acknowledged to be the world's leading source of citation and bibliometric data. The *Web of Science* Core Collection is part of

the *Web of Science*, and focuses on research published in journals and conferences in science, medicine, arts, humanities and social sciences.

The *Web of Science* was originally created as an awareness and information retrieval tool but it has acquired an important primary use as a tool for research evaluation, using citation analysis and bibliometrics. Data coverage is both current and retrospective in the sciences, social sciences, arts and humanities, in some cases back to 1900. Within the research community this data source was previously referred to by the acronym 'ISI'.

Unlike other databases, the *Web of Science* and underlying databases are selective, that is: the journals abstracted are selected using rigorous editorial and quality criteria. The authoritative, multidisciplinary content covers over 12,000 of the highest impact journals worldwide, including Open Access journals, and over 150,000 conference proceedings. The abstracted journals encompass the majority of significant, frequently cited scientific reports and, more importantly, an even greater proportion of the scientific research output which is cited. This selective process ensures that the citation counts remain relatively stable in given research fields and do not fluctuate unduly from year to year, which increases the usability of such data for performance evaluation. Within the research community these data are often still referred to by the acronym 'ISI'.

3.3 Clinical trials data

The clinical trials data used in this evaluation were sourced from clinicaltrials.gov, a online registry of clinical trials. It is run by the United States National Library of Medicine (NLM) at the National Institutes of Health, and is the largest clinical trials database, currently holding registrations from over 175,000 studies in over 180 countries. It is a web-based resource which is intended to provide patients, their family members, health care professionals, researchers, and the public with information on publicly and privately supported clinical studies on a wide range of diseases and conditions.

3.4 Bibliometric indicators

Papers/publications: Thomson Reuters abstracts publications including editorials, meeting abstracts and book reviews as well as research journal articles. The terms 'paper' and 'publication' are often used interchangeably to refer to printed and electronic outputs of many types. In our analyses the term 'paper' is used exclusively to refer to substantive journal articles, reviews and some proceedings papers and excludes editorials, meeting abstracts or other types of publication. **Papers** are the subset of publications for which citation data are most informative and which are used in calculations of citation impact.

Citations: The citation count is the number of times that a citation has been recorded for a given publication since it was published. Not all citations are necessarily recorded since not all publications are indexed. However, the material indexed by Thomson Reuters is estimated to attract about 95% of global citations.

Citation impact: 'Citations per paper' is an index of academic or research impact (as compared with economic or social impact). It is calculated by dividing the sum of citations by the total number of papers in any given dataset (so, for a single paper, raw impact is the same as its citation count). Impact can be calculated for papers within a specific research field such as Clinical Neurology, or for a specific institution or group of institutions, or a specific country. Citation count declines in the most recent years of any time-period as papers have had less time to accumulate citations (papers published in 2007 will typically have more citations than papers published in 2010).

Field-normalized citation impact (NCI): Citation rates vary between research fields and with time, consequently, analyses must take both field and year into account. In addition, the type of publication will influence the citation count. For this reason, only citation counts of papers (as defined above) are used in calculations of citation impact. The standard normalization factor is the world average citations per paper for the year and journal category in which the paper was published. This normalization is also referred to as 'rebasing' the citation count. The term 'citation impact' is used in this document as shorthand for Field-normalized citation impact (NCI).

Research field: Standard bibliometric methodology uses journal category as a proxy for research field. Journals are assigned to one or more categories, and every article within that journal is subsequently assigned to that category. Papers from prestigious, 'multidisciplinary' and general medical journals such as Nature, Science, The Lancet, BMJ, The New England Journal of Medicine and the Proceedings of the National Academy of Sciences (PNAS) are assigned to specific categories based on the journal categories of the references cited in the article. The selection procedures for the journals included in the citation databases are documented here http://scientific.thomsonreuters.com/mjl/. For this evaluation, the standard classification of *Web of Knowledge* journal categories have been used to calculate field-normalized citation impact (NCI).

Co-authorship of publications: The metadata associated with every research publication include the addresses of the authors. It is thus possible to develop an analysis of the organizations that co-author publications by extracting and examining these data. Co-authorship is generally accepted as an indicator of collaboration, although there are collaborations that do not result in co-authored publications and co-authored publications which involve limited collaboration. Conceivably other indicators of collaboration such as co-funding and international exchanges could be used but comprehensive and consistent data are not available.

Internationally collaborative publications: Internationally collaborative research publication is increasing rapidly. This is because such collaboration provides access to a wider range of resources, including intellectual resources, and accelerates the rate of discovery as well as increasing the intellectual content and therefore the impact of individual outputs. For this reason, internationally collaborative publications tend to be more highly cited than those that are solely domestic.

4 Analysis of publication output and citation impact

This Section describes a statistical analysis of workshop participants' and unsuccessful applicants' research output and citation impact in the periods before and after attendance.

4.1 Identification of publications produced by workshop participants or applicants

AACR provided Thomson Reuters with a list of 1000 workshop participants (480) and unsuccessful applicants (520) for five cohort years (2002, 2004, 2006, 2008 and 2010) with institutional affiliations and e-mail addresses.

Table 4.1 Number of workshop participant and unsuccessful applicant names provided by AACR by cohort year

Cohort year	Participants	Applicants
2002	87	113
2004	100	100
2006	96	114
2008	98	92
2010	99	101
Total	**480**	**520**

Algorithmic searching of Thomson Reuters databases underlying the *Web of Science* found a total of 31,261 publications over the period 1999 to 2013 which mapped to 936 of the individuals identified by AACR (456 participants and 480 applicants). Of the publications identified, 17,471 were within the three years before or after the relevant workshop year.

4.2 Analysis of workshop participation and publication output

4.2.1 Were workshop participants more productive before or after the workshop, as measured by number of publications?

For each of the 456 participants who attended the workshop, an "increase score" was calculated as the total number of publications he/she authored after the workshop less the total number of publications he/she authored before the workshop. For example, if a researcher authored 5 publications before the workshop and 7 publications after the workshop, his/her increase score would be 7 - 5 = 2. Publications that were written more than 3 years before or after the workshop date were excluded from the analysis. All publication types were counted, such as articles, meeting abstracts, book reviews, letters, etc. Table 4.2.1 provides descriptive statistics for the increase scores in the row "participant publications."

Outliers were removed from the set of increase scores by setting a lower threshold = (1^{st} quartile) - (length of interquartile range) = -9, and an upper threshold = (3^{rd} quartile) + (length of interquartile range) = 21. Figure 4.2.1 shows the distribution of increase scores after these limits were enforced. The Shapiro-Wilk test confirmed that the increase scores were not from a normal distribution.

The Wilcoxon signed-rank test (a non-parametric alternative to Student's t-test) was used to test the null hypothesis (H_o) that participants did not change their publication rate after the workshop. Outliers do not have undue influence in the Wilcoxon signed-rank test but do increase its power. Hence, all 456 increase scores were included in the analysis.

The Wilcoxon signed-rank test rejected the null hypothesis ($p = 2.2 \times 10^{-16}$). There was a statistically significant, positive change in the publication rate of workshop participants following the Methods in Clinical Cancer Research Workshop. Participants published at a higher rate during the three year period after the workshop, compared to the three year period preceding it. A feasible explanation for this result is a natural maturation process, whereby researchers in general publish more often as they mature in their profession. This possibility is controlled for in section 4.2.3 by comparing the increases made by workshop participants with those made by workshop applicants.

Figure 4.2.1 Distribution of increase scores for participant publications (number of publications authored after workshop – number of publications authored before workshop)

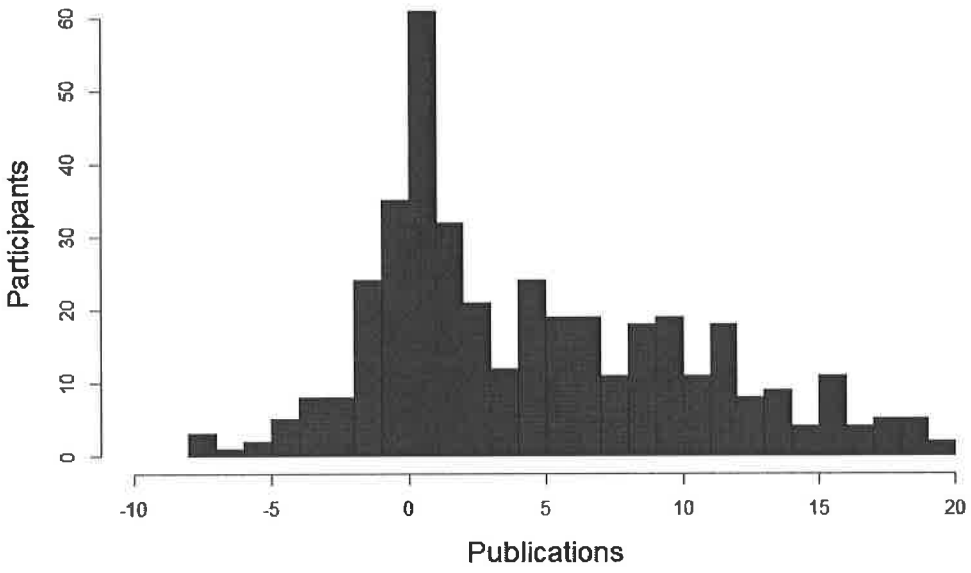

Table 4.2.1 Distribution of participant publication output increase scores before outliers have been removed

	Distribution of increase scores				
Research unit	Minimum Score	1st quartile (q1)	Median (q2)	3rd quartile (q3)	Maximum score
Participant publications	-29	1	5	11	67

4.2.2 Were workshop participants more productive before or after the workshop, as measured by number of published articles/reviews?

For each of the 456 participants who attended the workshop, the total number of articles/reviews he/she authored before the workshop was subtracted from the total number of articles/reviews he/she authored after the workshop. Articles/reviews that were not published within a three year range of the workshop date were excluded from the analysis. Table 4.2.2 shows descriptive statistics for the increase scores.

Outliers were removed from the set of increase scores by setting a lower threshold = (1^{st} quartile) - (length of interquartile range) = -5, and an upper threshold = (3^{rd} quartile) + (length of interquartile range) = 10. Figure 4.2.2 shows the distribution of increase scores after these limits were enforced. The Shapiro-Wilk test confirmed that the increase scores were not from a normal distribution.

The Wilcoxon signed-rank test was applied to test the null hypothesis that the number of participant articles/reviews published before the workshop was equal to the number published after the workshop. Based on data from all 456 participants, (H_o) was rejected ($p = 2.0 \times 10^{-16}$). There was a statistically significant, positive change in the number of articles/reviews authored by participants in the years following the Methods in Clinical Cancer Research Workshop. Participants wrote articles/reviews at a higher rate during the three year period after the workshop than they did during the three year period preceding it.

Figure 4.2.2 Distribution of increase scores for participant articles/reviews (number of articles/reviews authored after workshop – number of articles/reviews authored before workshop)

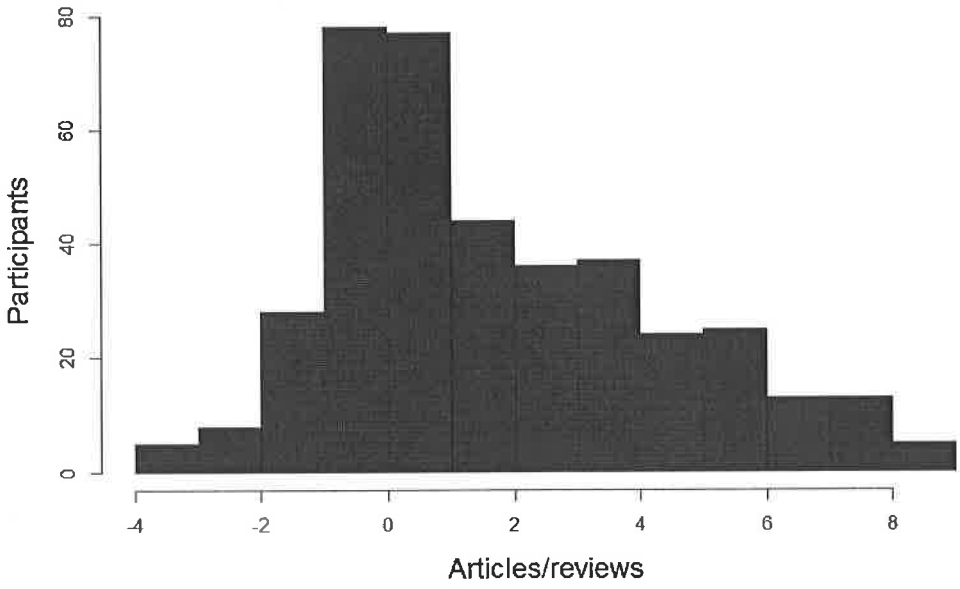

Table 4.2.2 Distribution of participant article/review output increase scores before outliers have been removed

	Distribution of increase scores				
Research unit	Minimum Score	1st quartile (q1)	Median (q2)	3rd quartile (q3)	Maximum score
Participant articles/reviews	-91	0	2	5	35

4.2.3 Which group increased their publication rate the most, workshop participants or workshop applicants?

For each of the 480 applicants who did not participate in the workshop, an increase score for publications was created as in Section 4.2.1. Table 4.2.3 shows descriptive statistics for the increase scores. Figure 4.2.3 shows the applicant and participant score distributions.

The Mann-Whitney U test was applied to test the null hypothesis that the distribution of increase scores for applicant publications did not differ from the distribution of increase scores for participant publications. The Mann-Whitney U test is not sensitive to outliers and thus all 480 applicant/456 participant scores were entered into the analysis. The resulting U statistic required the null hypothesis to be rejected ($p = 2.5 \times 10^{-4}$). The two distributions of increase scores were significantly different, either in terms of their shape or median.[1]

Table 4.2.3 Distribution of participant and applicant publication output increase scores before outliers have been removed

	Distribution of increase scores				
Research unit	Minimum Score	1st quartile (q1)	Median (q2)	3rd quartile (q3)	Maximum score
Participant publications	-29	1	5	11	67
Applicant publications	-5	1	2	6	14

[1] Hart, A (2001) Mann-Whitney test is not just a test of medians: differences in spread can be important, *British Medical Journal* 323:391-393

Figure 4.2.3 Distribution of increase scores for applicant/participant publications (number of publications authored after workshop – number of publications authored before workshop)

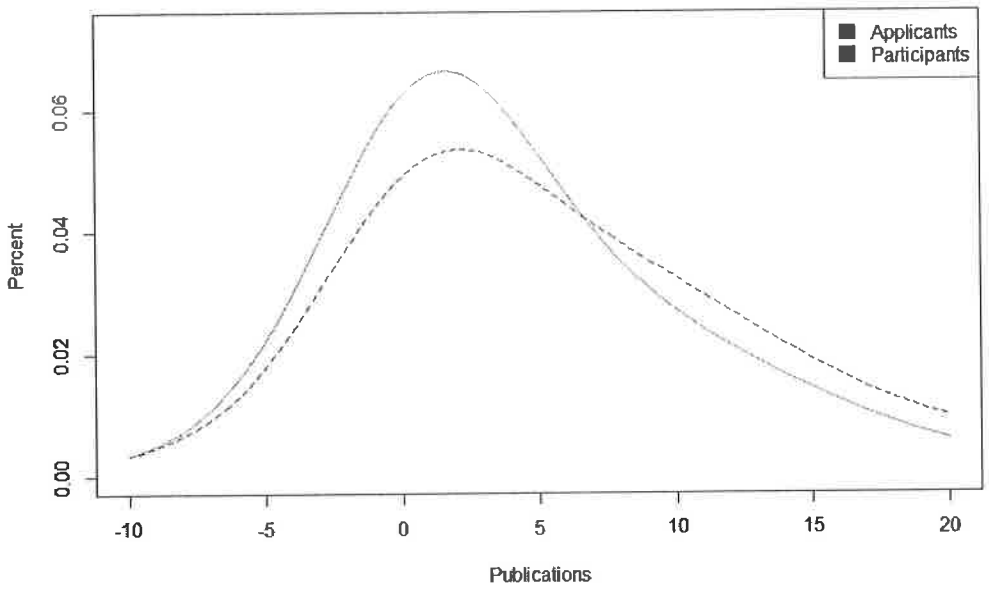

4.3 *Analysis of workshop participation and citation impact*

4.3.1 Did workshop participant have a higher citation impact that the applicants?

The average normalized citation impact of workshop participants increased between the before and after periods from around one-and-a-half time world average (1.52) to nearly twice the world average (1.95). The average normalized citation impact of unsuccessful applications, however, did not change (remaining at 1.29).

While the average citation impact describes the overall performance of a body of research it does not describe the underlying distribution of citation within that research. Thomson Reuters has developed Impact Profiles® to provide a visualization of that distribution which provides much more information about the basis and structure of research performance than the average. Impact Profiles® enable an examination and analysis of the balance of published outputs relative to world average and relative to comparators.

An Impact Profile® shows the proportion of papers that are uncited and the proportion in each of eight categories of relative citation rates, normalized to world average (which becomes 1.0 in this graph). Normalized citation rates above 1.0 indicate papers cited more often than world average for the field in which that journal is categorized and in their year of publication.

Attention should be paid to:

- The proportion of uncited papers on the left of the chart
- The proportion of cited papers either side of world average (1.0)
- The location of the most common (modal) group near the centre
- The proportion of papers in the most highly-cited categories to the right, (≥4 x world, ≥8 x world).

The higher citation impact of workshop participants after attendance is primarily due to a substantial increase in highly-cited papers (those cited at least four-times the word average) compared to the period before attendance – from 8.0% of total output to 11.7% (Figure 4.3.1). While the profile for the applicants shifted toward the right between the before and after periods, the percentage of highly-cited papers was unaffected, hence the stable average citation impact mentioned above.

Figure 4.3.1 Impact Profile® comparing the citation impact distribution of workshop participants and unsuccessful applicants before and after attendance/application for all cohort years

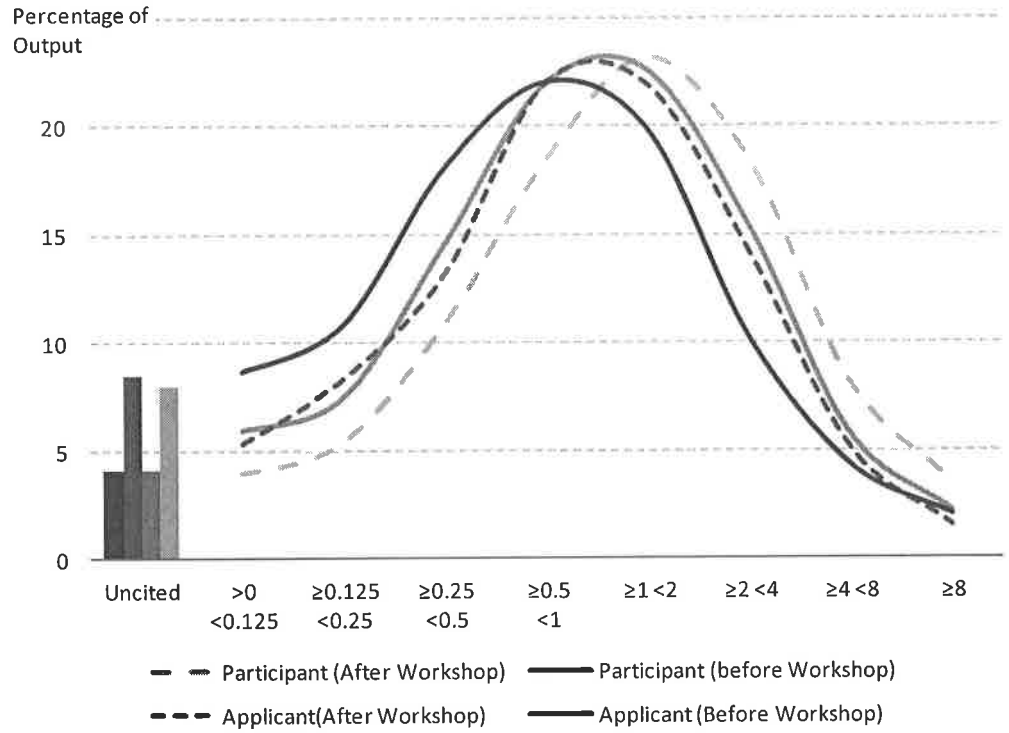

4.3.2 Did workshop participants show a greater increase in their citation rate than workshop applicants?

Another way to consider the citation data is to ask if the distribution of change in individual participants' performance between the before and after periods is different from the distribution of changes in the unsuccessful applicants' performance?

Normalized citiation impact (NCI) was calculated for each of the participant/applicant articles that were published within three years of the workshop date. Two average NCIs were calculated for each researcher; one based on the articles published before the workshop date, and one based on the articles published after the workshop date. An increase score was then created by subtracting the average before workshop NCI from the average after workshop NCI. Any researcher who had a NCI of 0 (i.e., did not publish) either before or after the workshop was excluded from the analysis. All average NCI increase scores are summarized in Table 4.3.2 and Figure 4.3.2, for both participants and applicants.

The Mann—Whitney U test was applied to test the null hypothesis that the distribution of increase scores for applicants did not differ from the distribution of increase scores for participants. The null hypothesis could not be rejected at the 95% confidence level (p= 0.51). The two distributions of increase scores did not differ in a statistically significant way, both in terms of their shape and their medians. A limitation of this test was the restricted sample size. Of the 480 applicants, 183 published in the three year period leading up the conference and in the three year period following the conference. 202 of the 456 participants published in the

three year period leading up to the conference and in the three year period following the conference. This meant that 297 applicants and 254 participants were excluded from the citation rate analysis. Including these researchers would have confounded publication quality (i.e., average NCI) with publication quantity. In other words, including the unpublished researchers in the analysis would have equated low publication quantity with low publication quality.

Table 4.3.2 Distribution of participant and applicant publication NCI increase scores before outliers have been removed

	Distribution of increase scores				
Research unit	Minimum Score	1st quartile (q1)	Median (q2)	3rd quartile (q3)	Maximum score
Participant average NCI	-1.28	-0.15	0.15	0.77	4.74
Applicant average NCI	-0.95	-0.14	0.15	0.44	1.23

Figure 4.3.2 Distribution of increase scores for applicant/participant average NCIs (average NCI after workshop – average NCI before workshop)

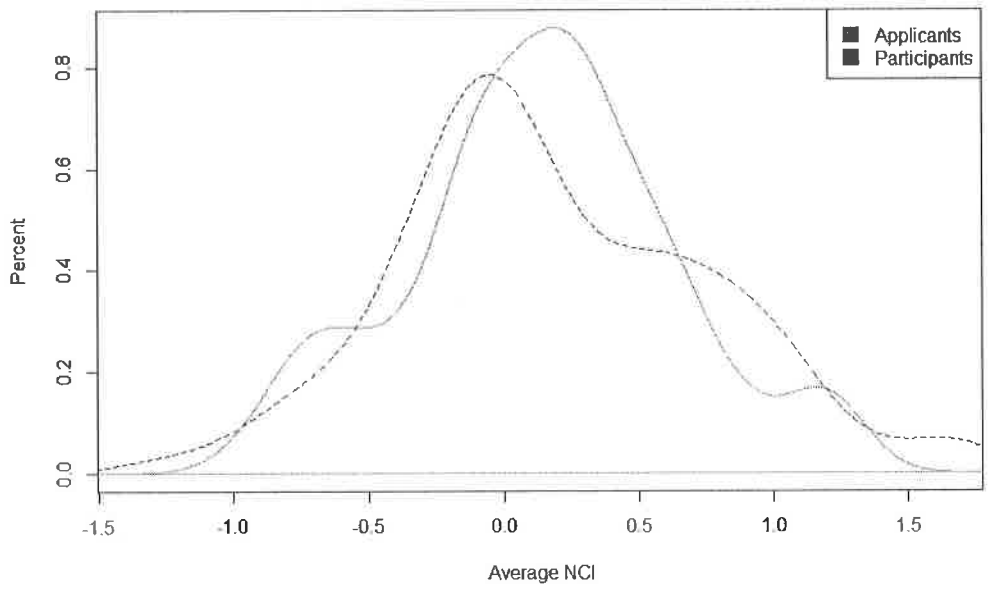

5 Analysis of clinical trials involvement

This Section describes a statistical analysis of workshop participants' and unsuccessful applicants' involvement in clinical trials in the periods before and after attendance.

5.1 Identification of clinical trials involving a workshop participant or applicant

Of the 1,000 individuals identified by AACR, searching clinicaltrials.gov identified a total of 917 trials in which 346 of them had been involved (205 workshop participants and 141 unsuccessful applicants).

5.2 Analysis of workshop participation and involvement in a clinical trial

5.2.1 Were workshop participants more productive before or after the workshop, as measured by number of clinical trials?

Based on the list of clinical trials provided by AACR, 205 of the 456 workshop participants were identified as having been involved in a clincal trial. The total number of clinical trials that each of these 205 participants conducted before the workshop was subtracted from the total number of clinical trials he/she conducted after the workshop. Clinical trials that were not within three years of the workshop date were excluded from the analysis. Table 5.2.1 summarizes the increase scores in the row "Clinical trials." Figure 5.2.1 presents the distribution graphically.

Outliers were removed from the set of increase scores by setting a lower threshold = (1^{st} quartile) - (length of interquartile range) = -2, and an upper threshold = (3^{rd} quartile) + (length of interquartile range) = 4. The Shapiro-Wilk test confirmed that the increase scores were not from a normal distribution.

The Wilcoxon signed-rank test was applied to test whether the number of clinical trials conducted by participants before the workshop was equal to the number of clinical trials conducted after the workshop. The null hypothesis was again rejected ($p= 5.4 \times 10^{-14}$). There was a statistically significant, positive change in the number of clinical trials conducted by participants following the Methods in Clinical Cancer Research Workshop.

Table 5.2.1 Distribution of participant clinical trials increase scores before outliers have been removed

	Distribution of increase scores				
Research unit	Minimum Score	1^{st} quartile (q1)	Median (q2)	3^{rd} quartile (q3)	Maximum score
Participant clinical trials	-6	0	1	2	16

Figure 5.2.1 Distribution of increase scores for participant clinical trials (number of clinical trials conducted after workshop – number of clinical trials conducted before workshop)

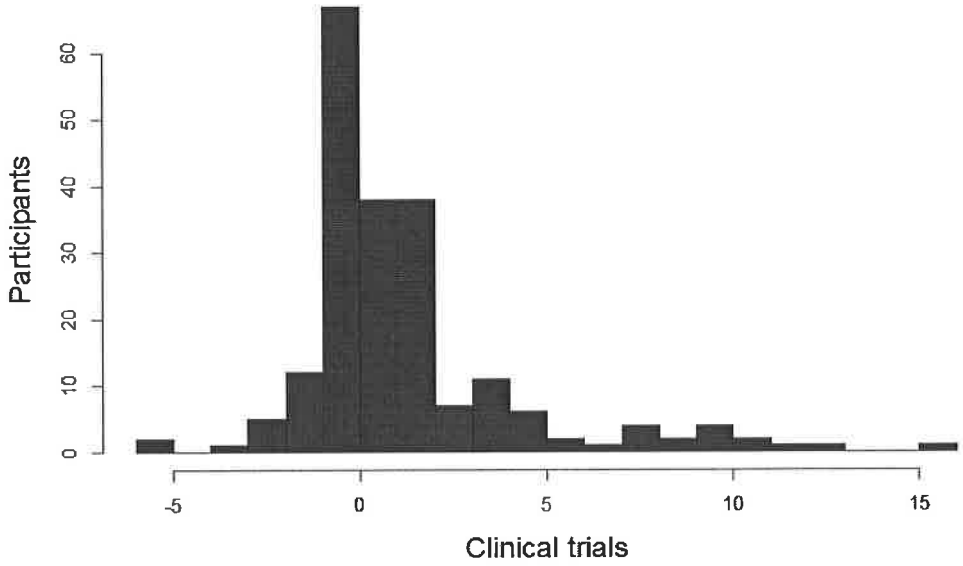

6 Analysis of research collaboration

To enable the effect of the AACR workshop participation on participants' research collaboration to be understood, we have analyzed their patterns of co-authorship and collaborative networks for the periods before and after attendance and compared this to the unsuccessful workshop applicants.

This Section provides an analysis of several indicators of collaboration and network parameters:

- The average number of authors per publication
- The average number of collaborators per participant/applicant
- The largest distance between two connected individuals within the collaborative network of participants and applicants (network diameter)
- The average distance between two connected individuals within the collaborative network of participants and applicants (average distance)
- The mean probability that two coauthors will also be coauthors of one another (clustering coefficient)

Diagrammatic representations of the collaborative networks of participants and applicants both before and after attendance/application are provided in Annex 1.

6.1 *Analysis of the average number of authors per publication*

The average number of authors per publication is an indication of how collaborative an author's research is.

The workshop participants collaborated with more co-authors than did the unsuccessful applicants (Table 6.1).

The workshop participants had, on average, a greater number of authors per publication than the unsuccessful applicants in all cohort groups both before and after attendance/application. In all cases there is an increase in the number of authors per publication between the periods before and after attendance/application, however, there is no discernible pattern in the magnitude of these changes.

Table 6.1 Average number of authors per publication for workshop participants and applicants, by cohort year

Cohort year	Participants/ applicants	Before	After	Change
2002	participants	7.50	8.50	+1.0
	applicants	6.20	6.80	+0.6
2004	participants	5.90	8.10	+2.2
	applicants	3.20	5.00	+1.8
2006	participants	7.40	8.10	+0.7
	applicants	6.20	6.90	+0.7
2008	participants	7.30	8.30	+1.0
	applicants	6.70	7.90	+1.2
2010	participants	7.90	9.10	+1.2
	applicants	7.10	8.60	+1.5

6.2 *Analysis of the average number of collaborators per participant/applicant*

The average number of collaborators and individual has worked with over a period of time is an indication of the size of their collaborative networks.

The workshop participants had larger collaborative networks and these grew more rapidly than those of the unsuccessful applicants (Table 6.2).

The workshop participants had a greater number of collaborators than the unsuccessful applicants both before and after attendance/application, except in the period before the 2004 workshop. In this period the applicants had a slightly larger number of co-authors (30.2) than the participants (28.7). In all cohort years the participants' average number of collaborators increased to a greater extent than did the unsuccessful applicants'.

Table 6.2 Average number of collaborators for workshop participants and applicants, by cohort year

Cohort year	Participants/ applicants	Before	After	Change
2002	participants	24.0	37.0	+13.0
	applicants	14.7	20.2	+5.5
2004	participants	28.7	80.9	+52.2
	applicants	30.2	34.0	+3.8
2006	participants	49.0	80.3	+31.3
	applicants	18.0	25.3	+7.3
2008	participants	34.9	69.8	+34.9
	applicants	18.7	32.4	+13.7
2010	participants	31.5	97.1	+65.6
	applicants	17.4	35.6	+18.2

6.3 *Analysis of the largest distance between two connected individuals within the collaborative network of participants and applicants*

The largest distance between two connected individuals within a network (known as the network diameter) is an indicator of the linear size of a network.

The individuals in the collaborative networks of workshop participants are more closely connected that those of unsuccessful applicants (Table 6.3).

Except in the period before the attendance/application of the 2002 cohort, the workshop participants had a smaller network diameter than the unsuccessful applicants. And except for the unsuccessful applicants in the 2002 cohort, the networks of both participants and applicants were more closely connected in the years after attendance/application.

Table 6.3 Network diameter of workshop participants and applicants, by cohort year

Cohort year	Participants/ applicants	Before	After	Change
2002	participants	16	14	-2
	applicants	6	18	12
2004	participants	12	10	-2
	applicants	15	13	-2
2006	participants	11	6	-5
	applicants	15	13	-2
2008	participants	15	9	-6
	applicants	20	12	-8
2010	participants	16	10	-6
	applicants	23	13	-10

6.4 *Analysis of the average distance between two connected individuals within the collaborative network of participants and applicants*

The average distance between two connected individuals within a network shows the number of steps it takes to get from one member of the network to another.

In most cases the individuals in the collaborative networks of workshop participants are more closely connected that those of unsuccessful applicants (Table 6.4).

In the 2002 and 2004 cohorts the collaborative networks of workshop participants before attendance/application were less closely connected that those of unsuccessful applicants. However, in all other cases, both before and after attendance/application, the collaborative networks of workshop participants were more closely connected that those of unsuccessful applicants. The collaborative networks of workshop participants have become more closely connected following attendance, for applicants there is no discernible pattern.

Table 6.4 Average distance between two connected individuals within the collaborative networks of workshop participants and applicants, by cohort year

Cohort year	Participants/ applicants	Before	After	Change
2002	participants	6.0	5.9	-0.1
	applicants	2.6	7.2	4.6
2004	participants	5.0	4.4	-0.6
	applicants	4.7	5.1	0.4
2006	participants	5.2	2.0	-3.2
	applicants	5.9	5.2	-0.7
2008	participants	6.0	4.5	-1.5
	applicants	6.8	5.4	-1.4
2010	participants	5.5	4.0	-1.5
	applicants	8.3	5.6	-2.7

6.5 *Analysis of clustering within the collaborative network of participants and applicants*

The clustering coefficient of an individual in a network is the ratio of existing links connecting an individual's neighbors to each other to the maximum possible number of such links. It is sometimes described as the extent to which the friends of my friends are also my friends. A high clustering coefficient indicates that the individuals in the network are closely-linked to one another.

There appears to be no correlation of network clustering with workshop participation (Table 6.5).

The clustering coefficient of the collaborative networks of workshop participants does not appear to differ markedly from that of unsuccessful applicants. And there does not appear to be any pattern in the changes of the clustering coefficient between the before and after time periods.

Table 6.5 Clustering coefficient of the collaborative network of workshop participants and applicants, by cohort year

Cohort year	Participants/ applicants	Before	After	Change
2002	participants	0.4	0.3	-0.17
	applicants	0.9	0.9	-0.02
2004	participants	0.2	0.2	0.03
	applicants	0.2	0.2	0.00
2006	participants	0.5	0.5	0.01
	applicants	0.4	0.4	0.01
2008	participants	0.5	0.5	0.00
	applicants	0.4	0.4	0.01
2010	participants	0.8	0.5	-0.33
	applicants	0.5	0.5	-0.05

Annex 1 Visualisation of the collaborative networks of AACR workshop participant and applicants

This Annex provides visualizations of the collaborative networks of the workshop participants and unsuccessful applicants. The network is constructed using nodes and edges, where nodes represent authors, and edges publication(s) between two scientist:

- Each network has 3 nodes types: participants (P), applicants (A), and external co-authors (E). When two participants of the same workshop year collaborate a blue edge is assigned. In a similar fashion, when two unsuccessful applicants of the same year collaborate, a blue edge is assigned. But when a participant and an applicant collaborate, a pink edge is given to differentiate the crossing of the groups. Finally, we characterize collaborations between two external co-authors with green edges.
- Thickness of the edges is proportional to the number of publications co-authored by a pair of researchers.
- Size of the nodes is also scaled proportionally to the number of connected edges
- Nodes are not ordered among groups, i.e. orange nodes may or may not be positioned next to each other. The nodes are arranged based on the number of edges they share with each other, this is done as an effort to reduce the volume of edges criss-crossing the circle.

Note: There were too many external co-author to external co-author edges for the year of 2004. We placed a threshold, and only considered the pairs of external co-authors who have collaborated at least 3 papers together.

These diagrams were produced using the Cytoscape software package (www.cytoscape.org).

Figure A1.1 Collaborative networks of workshop participants and unsuccessful applicants in the 2002 cohort

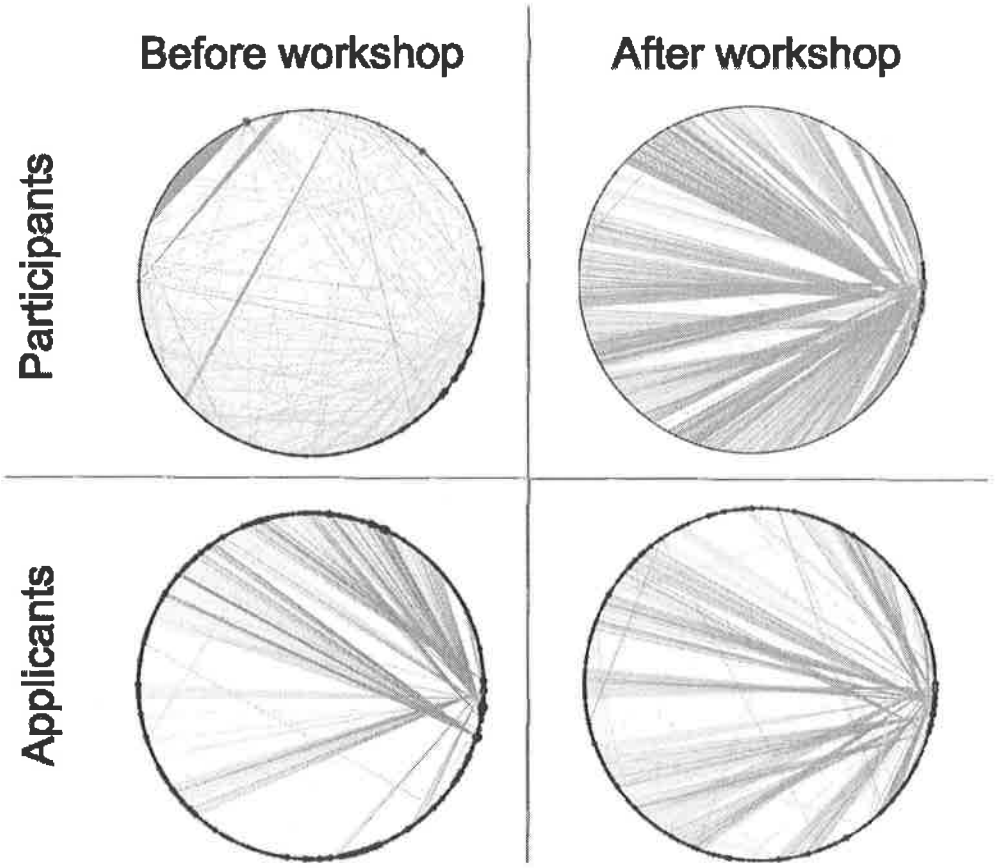

Figure A1.2 Collaborative networks of workshop participants and unsuccessful applicants in the 2004 cohort

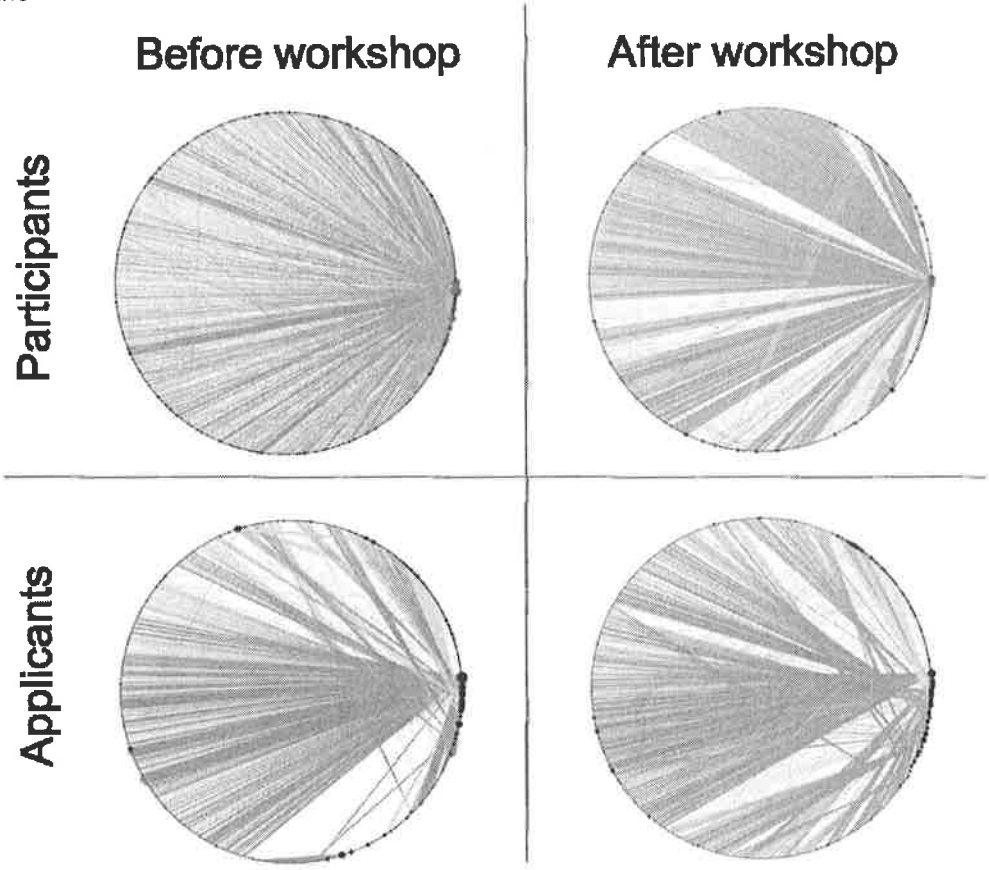

Figure A1.3 Collaborative networks of workshop participants and unsuccessful applicants in the 2006 cohort

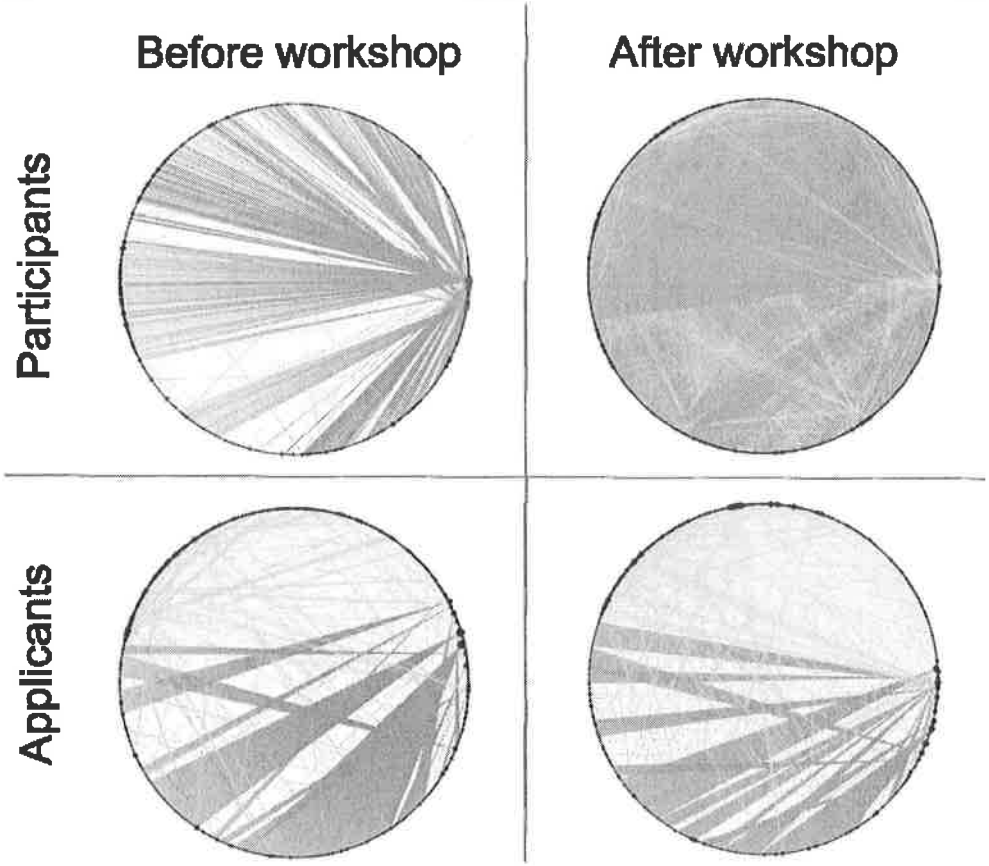

Figure A1.4 Collaborative networks of workshop participants and unsuccessful applicants in the 2008 cohort

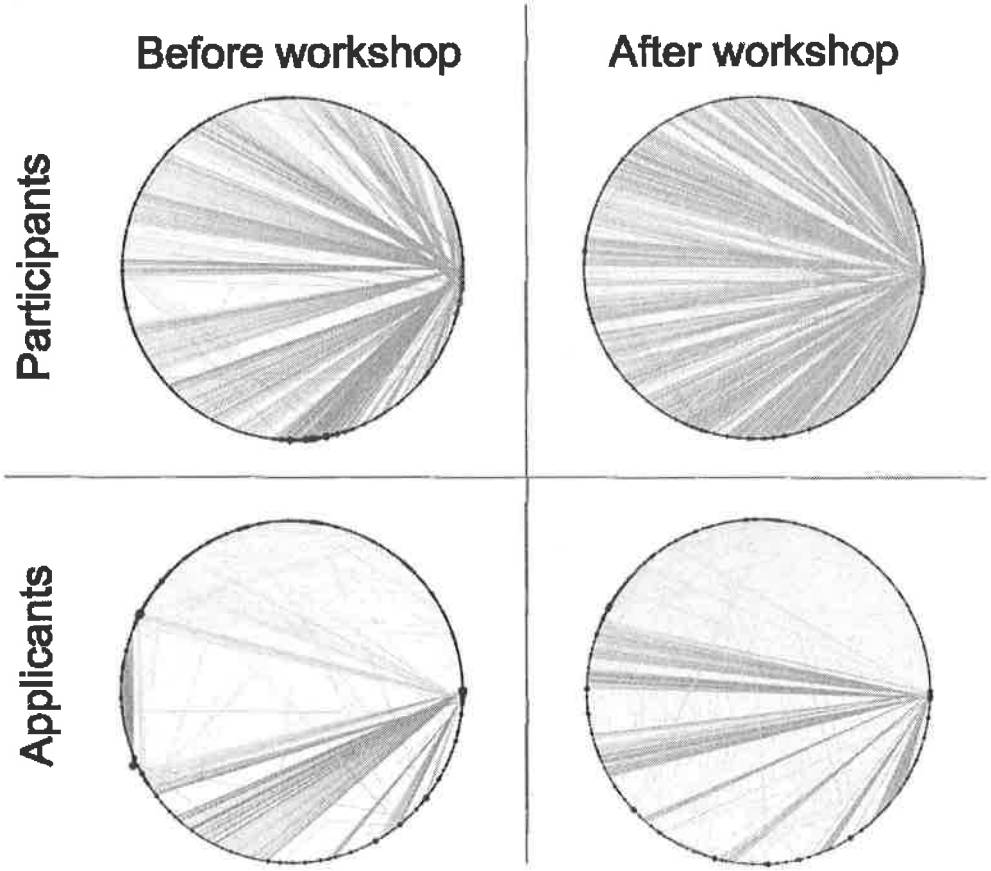

Figure A1.5 Collaborative networks of workshop participants and unsuccessful applicants in the 2010 cohort