

# 9. Networks of Statisticians and the Transformation of Medicine

CHRISTOPHER J. PHILLIPS

There is a statistical paradox at the heart of twentieth-century medicine.

In 1900 physicians largely ignored the tools of statistical analysis. Clinicians and laboratory researchers saw themselves as fundamentally opposed to the burgeoning field of academic statistics: they were interested in biomedical causation, statisticians were focused on numerical correlation; they were focused on exceptions and idiosyncrasies, statisticians were focused on norms and averages; they were determinists, statisticians were probabilists. There were essentially no statistical articles in medical journals, no statistical training required for the M.D., no well-known statistical interpretations of laboratory experiments. The American Medical Association lamented that questions about therapeutic efficacy were largely addressed by anecdotal accounts from influential physicians (and drug companies themselves).<sup>1</sup> The burgeoning field of public health (sometimes under the title of “sanitation” or “hygiene”) drew on epidemiological measures of disease, and questions of inoculation and epidemic infection had long been resolved with statistical calculations.<sup>2</sup> But these were seen as limited to large outbreaks where people could be treated as interchangeable; in the clinic, the opposite was true. Patients were unique and the aggregative methods of epidemiology irrelevant.<sup>3</sup>

By 2000 the situation was seemingly reversed. A statistically significant randomized clinical trial was the gold standard of therapeutic efficacy, and such proof was required by the Food and Drug Administration (FDA) prior to licensing drugs.<sup>4</sup> Reformers now promoted “evidence-based” medicine (as if medicine had never before been based on evidence), an initiative which claimed best practices should be determined solely on the basis of statistically

rigorous experiments and meta-analyses of past clinical trials.<sup>5</sup> Pre-diabetes, pre-hypertension, and similar threshold-based diagnoses were now determined on the basis of large studies of correlation and risk factors.<sup>6</sup> The patient experience itself had also been transformed into what Robert Aronowitz termed “risky medicine”: those at risk of disease and those suffering from chronic conditions looked increasingly alike.<sup>7</sup> A range of factors—exercise, diet, environmental exposure—were now linked to an increasing or decreasing probability of disease.<sup>8</sup>

How could the role of statistical practice in clinical medicine have been altered so dramatically? Normally explanations of fundamental change in scientific practice—whether considered as paradigm shifts, revolutions, or otherwise—fall into a few categories.<sup>9</sup> There is the shifting role of schools of thought and training. This doesn’t seem adequate here; the significance of statistics in physicians’ training has not changed dramatically and there are no clearly defined “schools” on the proper role of statistics in medicine. Likewise, the practices within teaching hospitals have remained remarkably stable. Other explanations might rely on the role of charismatic leaders, but again there are no real figureheads, or at least well-known leaders, of any such statistical movement. Some explanations might emphasize powerful new measures that enabled new ways of thinking about the world. There is some of that here—statistical measures largely matured and flourished in the twentieth century—but there is no one measure that was essential or fundamentally transformative. Other explanations rely on high-stakes and visible moments when statistics might prove themselves useful to resolving disputes. Indeed, there is a contender: the use of odds ratios and similar concepts to link smoking to lung cancer in the 1964 Surgeon General’s report on *Smoking and Health*. But there are no clear pre- and post- distinctions centered around 1964; the report itself does not attribute its findings primarily to new statistical measures; and opponents quickly condemned the report as inadequate.

In this chapter I want to suggest another way of explaining the seeming paradox of medical statistics: the increasing use of statistics in clinical medicine was largely invisible because it was accomplished by a *network* of unknown people deep within the federal bureaucracy. Specifically, I will highlight a group of biostatisticians at the National Institutes of Health (NIH) who from the late 1940s pioneered new uses of statistical concepts both by publishing research articles showing possible medical applications and by serving as consultants on projects seeking NIH financial support. Hired by Harold Dorn in 1947–1948 in the “methods” division of the Public Health Service (and soon incorporated into the NIH proper), these biostatisticians showed how formal statistical analysis provided powerful tools for determining efficacy, modeling dose-response curves, and evaluating therapies.<sup>10</sup> As the NIH became the dominant funder of medical research (and science generally) in this period, its model gradually became the dominant mode by which new discoveries in medicine were announced and new practices were established.

Parts of this story are easy to support. The NIH was certainly the dominant funder and gradually became the central organ for American biomedical research in the decades after 1950. Nearly all major medical research went through the institutes and their grant evaluators.<sup>11</sup> Moreover, NIH statisticians were deeply involved not just with the 1964 Surgeon General’s report, but also with the long-running Framingham Heart Study, another crucial site for promoting statistics-based measures of what constitutes health and disease, as well as with the evaluation of drug efficacy and safety through the FDA.

Other aspects are more difficult to track. The statisticians were not well known outside the field of biostatistics, let alone in medicine. The first generation—Jerome Cornfield, Samuel Greenhouse, Max Halperin, Jacob Lieberman, Nathan Mantel, and Marvin Schneiderman—were self-trained (none initially had doctoral degrees in statistics) and mastered the relevant statistical tools on the job. Though initially based in a single office, after the mid 1950s

they spread out into a variety of new Biometrics Research Branches or Biometric Offices across the NIH.<sup>12</sup> They published prolifically (approximately 650 articles through the 1970s), but remained largely behind the scenes as co-authors, statistical consultants, and advisors, though by the late 1970s had come to assume positions of prominence (head of the American Statistical Association, chair of university departments, etc.).

It is not obvious how to establish an historical argument for the group's influence. No one person or project was responsible for the quantification of clinical medicine. The field and its practices were too diverse and diffuse. We might think of the NIH as causing change, or bureaucratic rule-makers at the FDA as shifting practices, but both claims beg the question of who or what was ultimately responsible, even if it is sensible to focus on the NIH's rules for grant applications or the FDA's regulations for drug approval. Likewise, I'm hesitant to point to the development of odds ratios, Bayesian inference techniques, and the spread of null-hypothesis tests as explanations. Tracking the "successful" concepts on the basis of what turned out to be important risks obscuring what made them attractive in the first place. To twenty-first century observers, it seems obvious that statisticians who developed new measures of efficacy and causality in medicine would be influential. It was not clear in 1946.

I instead want to suggest one way to understand this transformation is to take seriously the way this group functioned as nodes within a network based largely (but not exclusively) at the NIH, and how participants collectively managed to transform standards of practice and spread statistical tools as new ways of defining proof and causality in medicine. I suspect that it is through their research collaborations—often resulting in published papers—that we might look for their influence. Portraying themselves initially as advisers for the design and interpretation of medical experiments and observational studies, they soon showed the worth of their methods. I see them as establishing a network, with people as nodes connected by the projects and papers they

worked on together. Though my use of network tools in this chapter is ultimately more exploratory than conclusive, the reliance on network analysis has the felicitous side-effect that I will study their work using numerical analysis rather than anecdote—precisely the way statisticians thought medical interventions should be assessed.

Thinking of the biostatistics group as a network isn't a replacement for close reading of published materials or deep dives into archival holdings. Rather, thinking in networked terms allows us to take advantage of the ways that researchers and institutions were connected through their projects and papers. This was the era that Derek J. de Solla Price referred to as the dawning of “big science,” and the biostatisticians at the NIH were integral to the rapid expansion of biomedical research, as well as the shift from individual researchers to large teams and collaborations.<sup>13</sup> Both the inclusion of new kinds of experts on projects and the use of ever

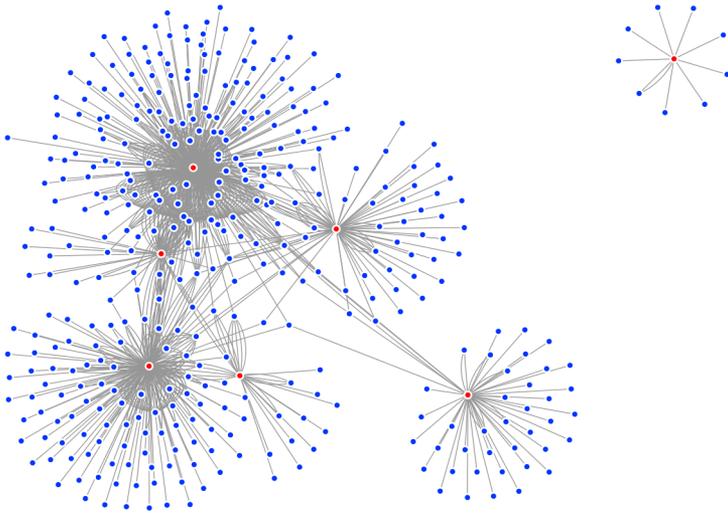


Figure 9.1: Overall publication network, 1930–1980

larger sample sizes in clinical studies in order to establish statistically significant effects often necessitated extensive

collaboration. Mid-century “big science” was not just about giant cyclotrons but also about multicenter studies of therapeutic interventions.

I initially created a network out of every published piece authored or co-authored by one of the first seven members of the NIH’s statistical group. Limiting to publications from 1930–1980 (the key timeframe for the spread of statistical ideas), I found 653 unique articles, abstracts, letters, notes, and reviews. By treating these articles as “edges” and the authors and co-authors as “nodes” I created the network shown in figure 9.1.<sup>14</sup>

The red nodes in figure 9.1 are the seven members of the group, with blue nodes indicating co-authorship. (Clockwise from upper right-hand red node: Dorn, Lieberman, Halperin, Cornfield, Greenhouse, Mantel, and Schneiderman.) Each edge in this image represents a single co-authorship relation, so one article by a member of the statistical group with two co-authors would be represented by two different edges.

Some interpretations are immediately apparent. Dorn is entirely isolated, whereas Lieberman shares only a few edges with the main cluster. Indeed, Dorn was head of the group, but was trained as a sociologist and never published extensively in biostatistics (though he did have an ongoing role managing surveys of the prevalence of cancer across the country). Lieberman also had relatively few connections because he did not co-author any articles with other members of the initial group. Among the remaining five statisticians, Mantel and Cornfield have by far the most publications (over 250 and 150 unique publications, respectively) and the largest number of connected edges. Greenhouse, interestingly, is far more connected to Mantel and Cornfield as a co-author (and in the visualization appears directly between them), than to either Schneiderman or Halperin.

Different visualizations of the network can help refine different aspects of the group’s influence. First, by dividing the data into two temporal groups (1945–1960 and 1961–1975), it is clear that there is

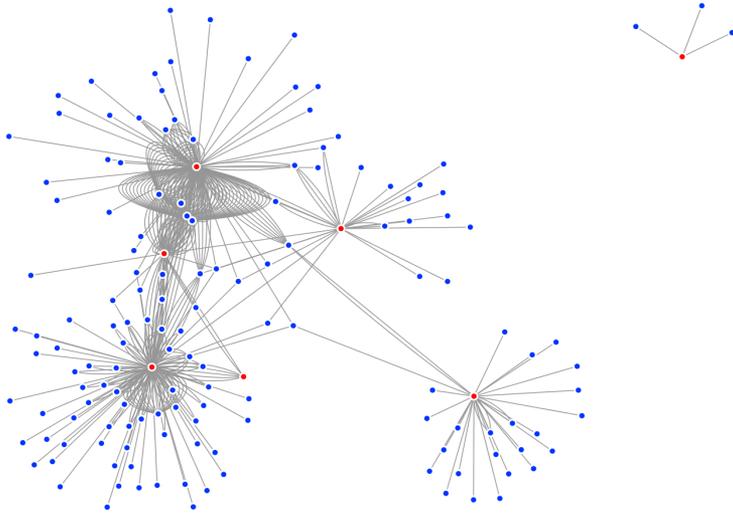


Figure 9.2: Publications, 1945–1960

little difference in publication practice (with the exception that Dorn’s untimely death in 1963 removes him). Figures 9.2 and 9.3 portray the networks created respectively by this temporal division.

Rather than dividing by time, it is also possible to look at the entire timespan, labeling edges by the discipline of the publication’s journal. This gives a quick estimate of the various fields in which the group was publishing.

The group was publishing widely, with the greatest number of publications in the fields of cancer research (edges colored light green, ~130 publications), medicine (blue, ~150), and statistics (orange for biometrics journals, ~75; pink for general statistics journals, ~125). There were also publications in general biology and chemistry (white, ~55), social science (purple, ~35), and epidemiology and public health (red, ~60).<sup>15</sup> Essentially every member of the group was publishing in both statistics and medical journals, serving as intellectual links between the disciplines. Each author had different disciplinary emphases, but it was not the case

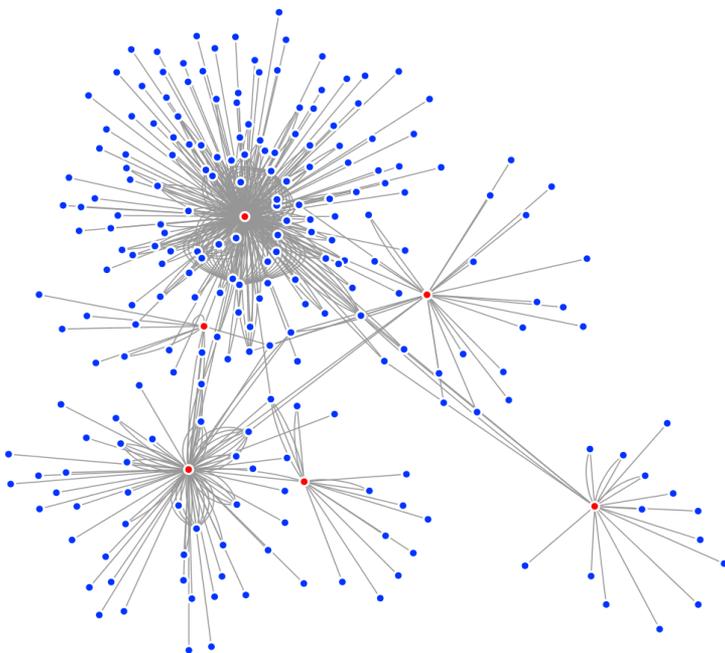


Figure 9.3: Publications, 1961–1975

that authors started publishing in statistics journals and then transitioned to medical journals. The entire group published widely across disciplines over time. Moreover, the relative lack of publication in epidemiological journals (the traditional locus of numerical analysis within medicine) suggests an explicit attempt to popularize statistical methods in medicine, and particularly in cancer research. Even as biostatistics and epidemiology were finding more established institutional homes in medical and public health schools in these years, early practitioners were establishing the field's prominence by publishing *elsewhere*.

Because this network was constructed by taking the publications of members of the group, it naturally places them at the center of the graph;

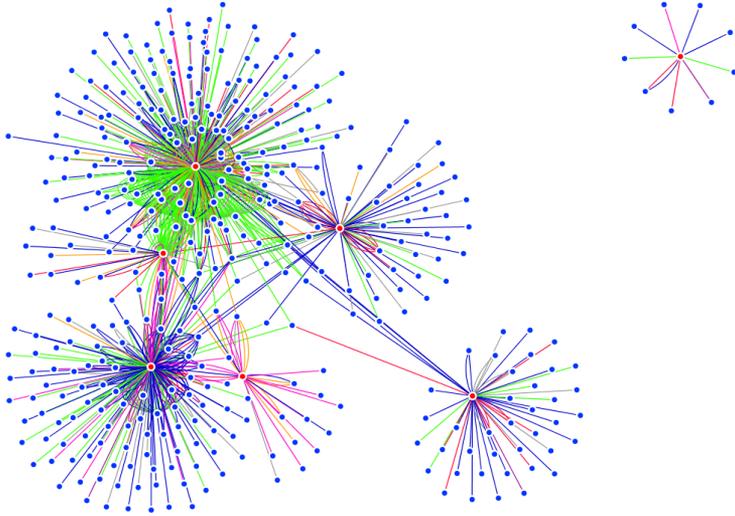


Figure 9.4: Publications (edges) labeled by discipline

a research collaboration that didn't involve one of them is simply missing. To get a wider sense of their influence, we need to situate their work within that of the biometrics and biomedical community. This is not easy, however, as the number of medical articles in this period quickly overwhelm most statistical software packages or network visualization tools. There are nearly six million articles in the PubMed collection between 1930 and 1980, and even when limited to topics involving cancer (using the Medical Subject Headings [MeSH] “neoplasm”), there are still a half-million articles. Given that many of these were co-authored, creating a network of co-publication would quickly make an unwieldy mess.

As a preliminary approach I took what I understood as one key case for the group's influence, namely epidemiological studies of cancer between 1950 and 1965. (A similar claim could be made for influence upon studies of heart disease with slightly later dates, but this search is at least consistent with the group's original location in the National Cancer Institute.) By limiting the articles to those

in English labeled with the MeSH terms “neoplasm” and “epidemiologic methods” between 1950 and 1965, I produced a network with 7585 nodes (authors) and 9116 edges (articles).<sup>16</sup>

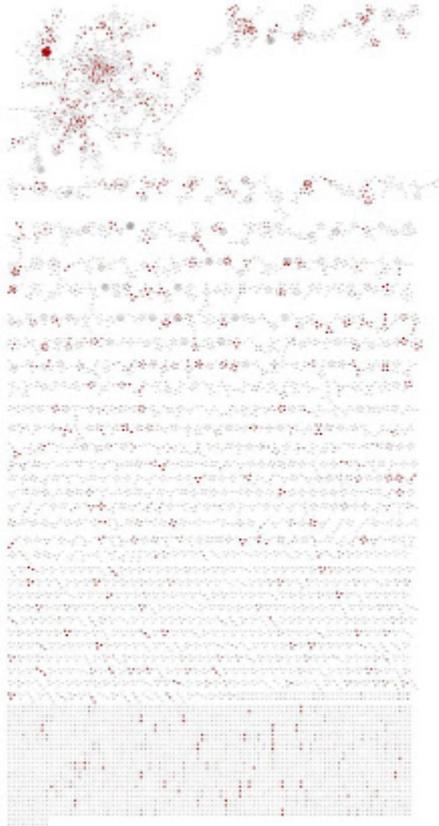


Figure 9.5: *Cancer and epidemiological methods articles, 1950-1965*

There is one large and well-connected network of articles in the upper left hand of the image, and then decreasingly small networks until at the bottom we see many articles with two co-authors who never published with anyone else. If the group I’m looking at had

influence, surely they'd be in the main network in the upper left *and* would be, statistically speaking, important or central members of that network.

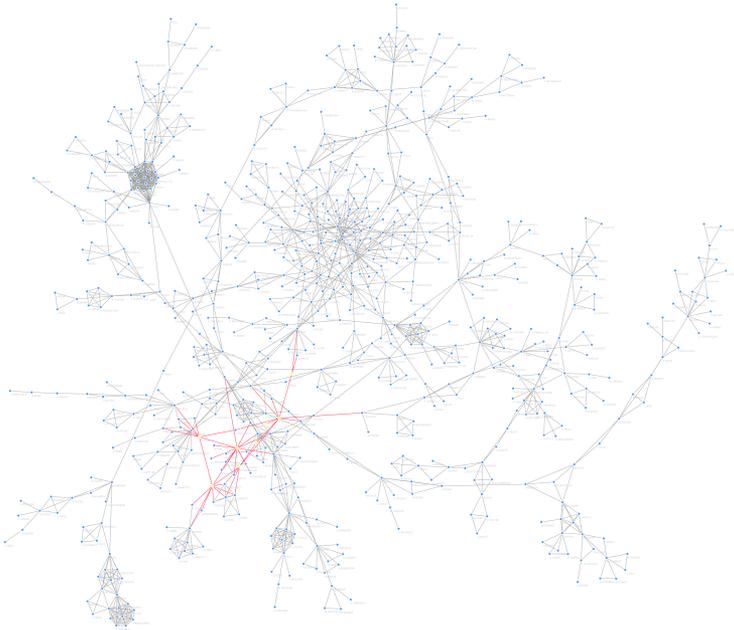


Figure 9.6: Sub-network of articles on cancer and epidemiological methods

Figure 9.6 shows the main sub-network (including 771 authors) from the upper-left corner of figure 9.5, with NIH statisticians listed as yellow nodes. Indeed, by taking statistics of only this sub-network, we can see how important the NIH group was to the publication of articles. If we take the “closeness centrality” or “shortest path length,” then out of these nearly 800 authors, Greenhouse has the fourth highest value, Mantel the twelfth, Schneiderman the thirteenth, and Dorn the twenty-eighth. While the “closeness” metric looks at shortest paths within the whole network, “betweenness” looks also at subgroups within the network, and for this latter measure, Mantel’s value ranks 21<sup>st</sup>, Greenhouse 26<sup>th</sup>, Dorn 272<sup>nd</sup>, and Schneiderman 265<sup>th</sup>. (As noted earlier, one problem

of this smaller network is the elimination of other members of the group despite their contribution to the topic of cancer epidemiology; nevertheless at least this gives a first approximation assuming that the other statisticians would have only increased the group's influence.) If we include two members who joined the statistical group slightly later, Sidney J. Cutler and Fred Ederer, the influence is even more impressive. Of the nearly 800 authors, Cutler had the highest score for “closeness” and the second highest for “betweenness” while Ederer had the third-highest score overall for both. Even with the obvious simplifications such an analysis entails, this is rather clear-cut evidence for the influence of the NIH group within the larger publication network concerning cancer and epidemiological methods.<sup>17</sup>

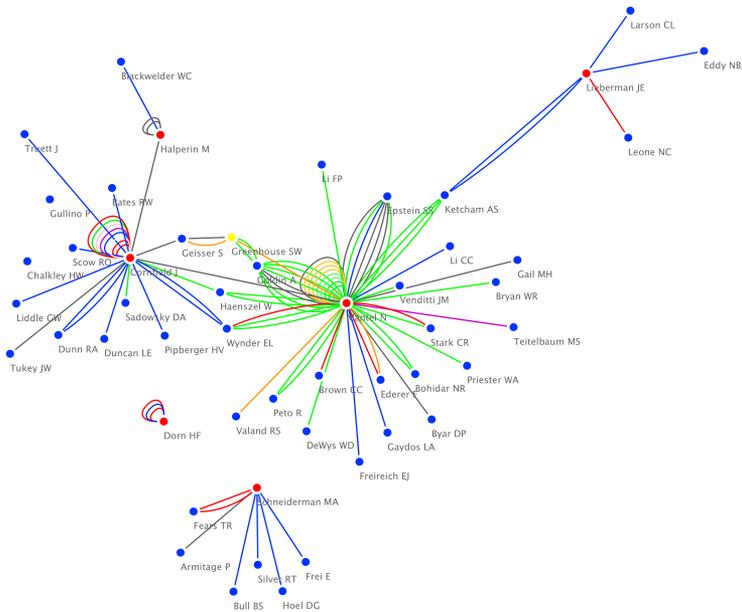


Figure 9.7: Edges and nodes that correspond to publications with over 50 citations

Another measure of influence would be to simply examine whether and how the initial group's publications were cited. Returning to

only those articles that had one of the original members as an author, we can also visualize only articles with substantial numbers of citations.

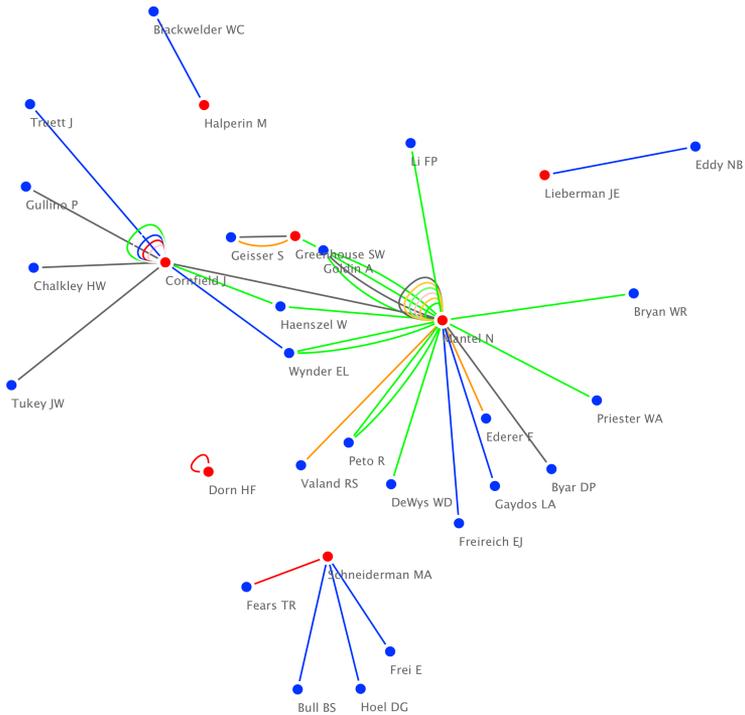


Figure 9.8: Edges and nodes that correspond to publications with over 100 citations

Some of these publications were certainly widely influential. There are about 100 articles with more than 50 citations, and about 50 of those articles have more than 100 citations. About 10 articles have more than 500 citations, according to the Web of Science citation index. On one level, this is to be expected; the articles are of interest precisely because they were influential. But it does also reveal the nature of their influence, and perhaps explain the network's relative invisibility. There was no one article or author

among this group that took the lead in establishing the field; rather, as the visualization suggests, their efforts were distributed. This is unlike, for example, a traditional laboratory model in the sciences in which publication authorship reflects institutional hierarchy. Furthermore, the majority of highly cited publications were in cancer and heart disease research, suggesting that it was the study of those fields in which the relevance of statistical analysis became most widely visible. The highly cited papers also range from the 1950s through the 1970s, suggesting that there was not one moment of influence, but rather a sustained program of interest to colleagues.

It is also possible, using Clarivate Analytics’s Web of Science citation indexing service, to track all the articles which cited those initial publications. Cornfield’s work, for example, has been cited in 4889 papers, with the peak of citation occurring around 1980. Cornfield’s most cited article (over 750 times since its publication) is on the analysis of patients enrolled in the Framingham Heart Study, a paper which in turn became a central model of the methodological basis of the “risk factor” approach.<sup>18</sup> Similar analyses can be made for the other members of the group: This data, however simplified,

Author	Cited by <sup>19</sup>	Peak of citations	Topic of most cited paper
Cornfield, Jerry	4889 papers	Late 1970s	Multivariate risk analysis of observation study (762 citations, published 1967)
Dorn, Harold	573 papers	mid-1960s	Cancer mortality (227 citations, publ. 1959)
Greenhouse, Samuel	4348 papers	Continuing to grow	Methods for analyzing profile data, such as tests given to individuals (3065, publ. 1959)
Halperin, Max	2812 papers	Around 1980	Estimating risks of diseases (218, publ. 1971)
Lieberman, Jacob	645 papers	Late 1970s	Testing of synthetic analgesics (212, publ. 1950)
Mantel, Nathan	35,724 papers	In late 1980s, then again around 2014	Statistical analysis of data from retrospective studies of disease (11,584, publ. 1959)
Schneiderman, Marvin	1417 papers	Around 1980	Methods of counting platelets (431, publ. 1965)

does suggest some clear aspects of the influence of these original seven members of the group. Their most cited work was originally published between 1950 and 1971, with the peak of citations of the group around 1980. This would be consistent with a general timeline of work in the 1950s and 1960s establishing the basic research that would coalesce in the 1970s into the established role of statistical methods in clinical work. Also, though it is somewhat arbitrary to focus only on the most cited paper by each author (because in some cases that particular paper was not much more cited than others), it is indicative that their most cited work was in interpreting observational data, particularly data around cancer and heart disease. This was indeed how this group was seen. They were known to have invented new measures for making causal claims about complex diseases of unknown origin. Future research might explore whether tools that focus on the content of their papers—epistemic network analysis, for example—might reveal the ways they shifted the conversation on a more granular level.<sup>20</sup>

There are some obvious problems with the network approach. Citation analysis is susceptible to criticism given the possibility of unreliable metadata, as well as the presumption that citation is a direct measure of influence. In addition, it ignores connections and collaborations that did not result in co-authorship. Other influential biostatisticians (including Donald Mainland and A. B. Hill) were in dialogue with this group (we know this because there is correspondence in their papers, as well as many citations in their published papers), but they were not co-authors and so are absent in the network. Moreover, by “flattening” collaborations into nodes and edges, nuances are erased, not least of which is the fact that there are many reasons for including (or excluding) another scholar as a co-author. Co-authored articles may reflect genuine collaboration or may simply reflect a primary author giving credit to others who made minor contributions to the project. Such distinctions are ignored when all co-authors are treated symmetrically.

There is, however, good evidence that co-authorship *was* precisely how the statisticians thought about their work. They initially functioned as a single group in a large office on the NIH campus, and when a call for statistical advice came into the office, whoever answered the phone would take on the consultation.<sup>21</sup> Though at the time they were not very concerned about turning every project into a published article, the group quickly realized that the statistical tools and techniques deployed in their consultations could be published to allow others to know how to approach this kind of problem. In this sense the diagram captures an essential feature of these statisticians' practice—that they served as physical and intellectual links from the NIH out into other researchers' labs (and into other institutes of the NIH). The edges here are not just articles, but true connections between statisticians and the wider biomedical, scientific, and statistical worlds. By setting themselves as an “on call” service, the group's publications serve as a written legacy of the projects to which they contributed.

There are many ways to expand this preliminary work. Some of the most important early clinical trials were conducted abroad, particularly in Great Britain, and it might also be worth trying to analyze more precisely how nodes within this network might be connected in other ways to co-authorship networks based in other nations. Perhaps a particular member of the NIH group served as a conduit to statistical researchers abroad, or perhaps there were many connections across multiple people. It would also be useful to label not just publications by discipline but also nodes by institutional affiliation. This would require a great deal of time, because institutional affiliations shift over a half-century (and some research projects might span multiple affiliations, etc.), but this might also help reveal the pathway of influence out from this initial group. Alternatively, nodes might be institutions rather than authors, and alternative network constructions would certainly provide different views of the phenomena. Moreover, I might include statisticians who joined the NIH after these first seven, or see if new hires changed the direction of the publishing effort.

There is also much to be done to clean up the data. I have checked Cornfield's publication list against a bibliography compiled late in his life, for example, but have not tried to do this yet for any of the other primary nodes.

In the end this analysis is preliminary, both in the sense that the corpus of medical documents is too big a network to examine easily and in the sense that it is still not obvious how, precisely, to add network analysis to traditional archival work. Nevertheless, given the way in which statistical ideas spread at mid-century, changing the entire way medicine is conducted without a clear person or reason driving the transformation, publication networks are useful tools for thinking about how research practices change. We have long known about the key role scientific journals played in the dissemination of research, and that played by funding agencies like the NIH in medicine, but there is surprisingly little historical analysis of how, precisely, novel methods and techniques spread. This chapter, at a minimum, suggests ways that a small group of statisticians hidden away at the NIH could still have an outsized and visible presence in the literature, introducing novel methods for analysis which connect medicine, statistics, and the physical and social sciences.

## Endnotes

1. Harry M. Marks, *The Progress of Experiment: Science and Therapeutic Reform in the United States, 1900–1990* (Cambridge: Cambridge University Press, 1997), 15–70.

2. William G. Rothstein, *Public Health and the Risk Factor: A History of an Uneven Medical Revolution* (Rochester: University of Rochester Press, 2003); Gérard Jorland, Annick Opinel, and George Weisz, eds., *Body Counts: Medical Quantification in History and Sociological Perspectives* (Montreal: McGill–Queen’s University Press, 2005).

3. Of course, there were exceptional clinicians who used numbers—e.g., the so-called “Paris School” of the mid-nineteenth century focused on comparing treatments through the measurements of responses. But these were exceptions that proved the rule (and which emphasized the application of epidemiological tools precisely in large hospitals and charity wards wherein people could be treated as interchangeable). See J. Rosser Matthews, *Quantification and the Quest for Medical Certainty* (Princeton: Princeton University Press, 1995), 1–61.

4. Laura Bothwell, “The Emergence of the Randomized Controlled Trial: Origins to 1980” (PhD diss., Columbia University, 2014); Daniel Carpenter, *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA* (Princeton: Princeton University Press, 2010), 269–280.

5. Jeanne Daly, *Evidence-Based Medicine and the Search for a Science of Clinical Care* (Berkeley: University of California Press, 2005).

6. Jeremy A. Greene, *Prescribing by Number: Drugs and the Definition of Disease* (Baltimore: Johns Hopkins University Press, 2007); David Shumway Jones and Gerald M. Oppenheimer, “If the Framingham Heart Study Did Not Invent the Risk Factor, Who Did?” *Perspectives in Biology and Medicine* 6, no. 2 (Spring 2017): 131–150.

7. Robert Aronowitz, *Risky Medicine: Our Quest to Cure Fear and Uncertainty* (Chicago: University of Chicago Press, 2015).

8. The penetration of statistics into everyday clinical medicine has not meant that physicians are deeply trained in statistics, or that physicians and patients understand the statistical measures that underlie their interactions. See the work of Gerd Gigerenzer, including his *Calculated Risks: How to Know When Numbers Deceive You* (New York: Simon and Schuster, 2002).

9. This is not the place to review the vast philosophical and historical literature on how and why science changes, but some of the classic conceptual arguments about influence include that of research schools: Gerald L. Geison and Frederic L. Holmes, eds., “Research Schools: Historical Appraisal,” *Osiris* 8 (1993): 1–248; pedagogy: David Kaiser, ed., *Pedagogy and the Practice of Science* (Cambridge: Massachusetts Institute of Technology Press, 2005); thought collectives: Ludwik Fleck, *Genesis and Development of a Scientific Fact*, trans. Fred Bradley and Thaddeus J. Trenn (Chicago: Chicago University Press, 1979 [1935]); and paradigm shifts and revolutions: Ian Hacking, ed., *Scientific Revolutions* (Oxford: Oxford University Press, 1981) and Thomas S. Kuhn, *The Structure of Scientific Revolutions* (Chicago: University of Chicago Press, 1962).

10. Though there is little published information on the group, there is a brief discussion of their origins in Jonas H. Ellenberg, Mitchell H. Gail, and Nancy L. Geller, “Conversations with NIH Statisticians: Interviews with the Pioneers of Biostatistics at the United States National Institutes of Health,” *Statistical Science* 12, no. 2 (May 1997): 77–81, and of their influence in Sejal Patel, “The Benevolent Tyranny of Biostatistics: Public Administration and the Promotion of Biostatistics at the National Institutes of Health, 1946–1970,” *Bulletin of the History of Medicine* 87 (2013): 622–647.

11. For a retrospective on the influence of the post-war NIH, see a special issue of *Science* upon the institution’s centenary: *Science* 237, no. 4817 (August 21, 1987), especially two retrospective articles in which NIH directors surveyed the institution’s growth leading up to the 1980s: James A. Shannon, “The National Institutes of Health: Some Critical Years, 1955–1957,” *Science* 237, no. 4817 (August 21, 1987): 865–868 and James B. Wyngaarden, “The National Institutes of Health in its Centennial Year,” *Science* 237, no. 4817 (August 21, 1987): 869–874.

12. Ellenberg, Gail, and Geller, “Conversations with NIH Statisticians.”

13. Derek J. de Solla Price, *Little Science Big Science* (New York: Columbia University Press, 1963); more broadly, Peter Galison and Bruce Hevly, eds., *Big Science: The Growth of Large-Scale Research* (Stanford: Stanford University Press, 1992). Most of the literature focuses on “bigness” in physics.

14. I primarily used Clarivate Analytics’ Web of Science collection, including Medline, Social Sciences Citation Index, and Biosis Citation Index. The network visualization was produced with Cytoscape, using tables created with OpenRefine.

15. Publication numbers are given as approximate because choices about “unique” publications are difficult—an abstract for a project also later published as a research article counts twice, sometimes multiple letters to the editor are grouped and only count once, etc. All such problems seem to be evenly distributed across disciplines, suggesting that while the absolute numbers are approximate, the *relative* distribution is largely stable.

16. Because of the size of the network, I didn’t download and edit the data myself but rather used the “Social Network” plugin for Cytoscape, which auto-populates authorship networks directly from PubMed data. This may introduce problems with metadata (including the duplication of names or misspellings, for example) but I ignored them given the size of the data set and its role in my work primarily for situating the main co-authorship network. There are also, of course, substantial complications using MeSH terms given that (1) they were often introduced at different times by the National Library of Medicine and applied retroactively to articles (epidemiologic methods and neoplasm were both introduced in the mid-1960s) and (2) they may not capture what contemporary actors considered under those terms. This has unfortunate narrowing effects: Jerome Cornfield, despite authoring an influential article on using statistical methods (odds ratios in particular) to examine cancer rates in the 1950s, was filtered out using these search terms. Either one would have to go through by hand, or simply expand the network finding other more inclusive MeSH terms.

17. One suspects, in fact, that the large, interconnected network under discussion is in fact just the NIH itself. That’s rather to be expected, as the first and largest single funder of cancer research. However, there was no guarantee that the statistical group would be publishing in that particular network, let alone have such central places within it, suggesting that their influence was still impressive.

18. Patel, “The Benevolent Tyranny of Biostatistics,” 629–630.

19. All numbers are approximate given that new citations are still appearing.

20. David Williamson Shaffer, Wesley Collier, and A.R. Ruis, “A Tutorial on Epistemic Network Analysis: Analyzing the Structure of Connections in Cognitive, Social, and Interaction Data,” *Journal of Learning Analytics* 3, no. 3 (2016): 9–45; see also Ruis, this volume.

21. Samuel W. Greenhouse, “Some Reflections on the Beginnings and Development of Statistics in ‘Your Father’s NIH,’” *Statistical Science* 12, no. 2 (1997): 82–87, on p. 84.