

Ethics and Evidence-Based Medicine

Fallibility and Responsibility in
Clinical Science

Kenneth W. Goodman



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Cambridge University Press 2003

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2003

Printed in the United Kingdom at the University Press, Cambridge

Typeface Minion 10.5/13 pt *System* QuarkXPress™ [SE]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Goodman, Kenneth W., 1954–
Ethics and evidence-based medicine : fallibility and responsibility in clinical science /
Kenneth W. Goodman.
p. cm.

Includes bibliographical references and index.
ISBN 0 521 81933 4 (hbk.) – ISBN 0 521 79653 9 (pbk.)
1. Medical ethics. 2. Evidence-based medicine. I. Title.

R725.5 G66 2002
174'.2–dc21 2002067402

ISBN 0 521 81933 4 hardback
ISBN 0 521 79653 9 paperback

Every effort has been made in preparing this book to provide accurate and up-to-date information that is in accord with accepted standards and practice at the time of publication. Nevertheless, the authors, editors and publisher can make no warranties that the information contained herein is totally free from error, not least because clinical standards are constantly changing through research and regulation. The authors, editors and publisher therefore disclaim all liability for direct or consequential damages resulting from the use of material contained in this book. Readers are strongly advised to pay careful attention to information provided by the manufacturer of any drugs or equipment that they plan to use.

Contents

<i>Preface</i>	ix
<i>Acknowledgments</i>	xii
1 Foundations and history of evidence-based practice	1
2 The research synthesis revolution	23
3 Evidence of evidence, and other conceptual challenges	41
4 Human subjects, the Internet, databases, and data mining	67
5 Evidence at the bedside	91
6 Public health policy, uncertainty, and genetics	113
7 Ethics and evidence	129
<i>References</i>	141
<i>Index</i>	165

Foundations and history of evidence-based practice

It isn't what we don't know that gives us trouble, it's what we know that ain't so.

Will Rogers

This chapter will locate systematic science and evidence-based medicine against the background of biomedical research in the second half of the twentieth century. The growth of this research paralleled and in some ways forced the evolution of current standards for communicating the results of scientific inquiry (i.e., the emergence of peer review and the expansion of the number of research programs, journals, books, etc.). The research raises interesting issues about the role and nature of expertise and medical knowledge, and it has led to a vast tableau of practice guidelines, critical pathways, consensus statements, and assorted other scientifically based imperatives for the care of individual patients. These imperatives are increasingly linked to physician and institution reimbursement. Where the stakes are highest, as in clinical medicine and public health, these forces assume special importance for ethics and public policy.

Before it became a movement, or a cause, evidence-based medicine (EBM) was a kind of cognitive itch: a troublesome doubt that follows from the realizations that humans are fallible, that scientific knowledge increases and that medical decisions sometimes have very high stakes. If you make a mistake, your patient might die. Less ominous, even if they do not die, patients are often paying for physicians' services, either through taxes or by putting cash on the barrel in one way or another. Failure to know what one is doing then becomes a kind of rip-off or scam. These realizations tend to focus the attention of most clinicians.

Patients die or otherwise come to grief all the time, of course. A bad or unhappy outcome can be the result of any of a large ensemble of causes:

- Ignorance, carelessness or inattention (individual)
- Ignorance, carelessness or inattention (collective)
- Futility
- Incompetence

- Conflict of interest
- Deception (by a patient)
- Malign intent

There may be others, but we are concerned in any case with the first two. (Note that the first two are cognate with the increasingly important work on medical *error* or mistakes, in which various forms of inattention and failures of institutional process are identified as responsible for mistakes and consequent harms (see Bosk 1981, the locus classicus; and Kohn, Corrigan, and Donaldson 2000). We shall return to error at the end of the chapter.)

To have evidence is to have some conceptual warrant for a belief or action. This says nothing so far about the *quality* of the evidence and hence the strength of the warrant. To be in a state of ignorance is to have false beliefs or to lack beliefs one way or another about the way the world works. If I do not have the belief that germs cause disease, and if germs do cause disease, then I am ignorant of the fact that germs cause disease. Now, this would be a grave shortcoming in a twenty-first century physician, but not in one practicing in the fourth century BC, say. This is because my twenty-first century ignorance is individual and my fourth century ignorance is collective. Twenty-five centuries ago no-one knew that germs cause disease; now, everyone (at least everyone who is practicing medicine) does, or should.

At any rate, the demand that clinicians know what they are doing, more or less, is an ancient one and it has, from the beginning, been couched as a *moral* imperative. The Hippocratic Oath (likely *not* written by Hippocrates) may be read as a celebration of teachers and education at least as much as an itemization of duties and virtues. When the oath-taker vows to “regard him who has taught me this technē as equal to my parents,” she is celebrating the transmission of knowledge; when she promises not to “cut, and certainly not those suffering from stone, but I will cede [this] to men [who are] practitioners of this activity,” it is a vow not to practice beyond one’s knowledge or capacity (translation by Von Staden 1996).

The Oath of Maimonides, after the twelfth century physician, rabbi, and philosopher Moses Maimonides, entreats, “Grant me the strength, time and opportunity always to correct what I have acquired, always to extend its domain; for knowledge is immense and the spirit of man can extend indefinitely to enrich itself daily with new requirements” (translation by Friedenwald 1917). Surely this should be read as a plea not to fall too far behind in monitoring the shifting landscape of medical evidence, a medie-

val anticipation of the importance – the *moral* importance – of continuing medical education.

It could not be otherwise. The intersection of knowledge and health points to a moral imperative because idiosyncratic ignorance causes or allows people to be harmed. An individual clinician's ignorance becomes blameworthy in part because his or her patients would have fared better elsewhere, in the hands of another whose greater knowledge (or lesser ignorance) would have saved the day. Now, this raises interesting questions about how far behind one might lapse without blame, and we will return to them later. The notion is important to us now because we want to distinguish idiosyncratic ignorance from community or collective ignorance.

What physicians don't know¹

At its core, evidence-based practice rests on a supposition which, while probably true, itself has unclear evidentiary support.

The demand that clinicians make the most of evidence – or even high-quality scientific evidence – in patient care is a demand that catches our attention only if it can be shown that they were not doing so already. One might try a little experiment with a friend or family member whose work is wholly outside health care . . . mention that evidence-based medicine has become a profession-wide movement to try to get doctors to practice in accord with, well, the evidence. The hypothesis is that you will be met by confusion if not outright incomprehension: “What were they basing their decisions on before all this?” Well, what indeed? In fact, of course, the healing professions have always, albeit in one degree or another, been based on evidence. It is just that there has been precious little of it, and it hasn't been any good. In some (but not all) cases, it has been just enough to distinguish physicians from shamans.

Thomas Beddoes and Pierre Louis

The likely true but undersupported supposition at the core of evidence-based medicine is that most health care is (or, until comparatively recently, was) not evidence based. This means either of two things: (1) There is no (adequate) evidence available to support clinical inferences, or there is, somewhere, but clinicians have no access to it; (2) the truth may be out

there, but it doesn't matter if we cannot lay our hands on it. The latter is among the targets of post-Enlightenment medical epistemology and a celebration of human experience and open communication. We can probably trace the intellectual birth of evidence-based practice to Thomas Beddoes (1760–1808), the English physician known in part for his “Pneumatic Institution” for the study of the medical use of gases² and, perhaps more importantly, for his criticism of turn-of-the-century medical practice.

Beddoes argued that eighteenth century medicine had become hide-bound, stagnant, and secretive (Porter 1992). Here is how the medical historian Roy Porter summarizes the points Beddoes makes in his 1808 “Letter to the Right Honourable Sir Joseph Banks . . . on the Causes and Removal of the Prevailing Discontents, Imperfections, and Abuses, in Medicine”:

Beddoes proposed two solutions. First, systematic collection and indexing of medical facts. “Why should not reports be transmitted at fixed periods from all the hospitals and medical charities in the kingdom to a central board?” Other “charitable establishments for the relief of the indigent sick” must also supply information, as should physicians at large. Data should be processed by a paid clerical staff, and made freely available. Seminars should be held. The stimulus to comparison and criticisms would sift good practice from bad. “What would be the effect”, Beddoes mused, of “register offices, not exactly for receiving votive tablets, like certain ancient temples, but in which attestations, both of the good and of the evil, that appears to be done by practitioners of medicine, should be deposited?” Without effective information storage, retrieval and dissemination, medicine would never take its place amongst the progressive sciences. “To lose a single fact may be to lose many lives. Yet ten thousand, perhaps, are lost for one that is preserved; and all for want of a system among our theatres of disease, combined with the establishment of a national bank of medical wealth, where each individual practitioner may deposit his grains of knowledge, and draw out, in return, the stock, accumulated by all his brethren.” . . . Second, to complement his medical bank, Beddoes urged his fellows to publish more . . . (Porter 1992: 10, notes omitted)

Data sharing . . . collecting and archiving . . . analysis and reporting . . . publishing . . . It seems that the good Dr. Beddoes was calling for a comprehensive system of medical information management. Moreover, he was calling for such a system because he believed, with good warrant, that the medical science of his day was shortchanging – was harming – patients, and that it could be better. Information becomes evidence when it applies to, bears on, or constitutes a reason for (dis)believing the truth of a proposition. If we are talking about propositions related to life, death, pain, disability, and so

forth, then it is just a few short steps until we identify a *duty* to collect and share information that bears on those propositions.

Thomas Beddoes is suggesting a moral link between information management and medical practice. He is proposing outcomes research and fantasizing about systematic reviews; demanding databases and hoping for data mining; insisting on broader dissemination, and doing so two centuries ago, or before the World Wide Web would, at least in principle, put every publication on every desktop.

Somewhat after Beddoes, in 1834, Pierre Charles Alexandre Louis (1787–1872), published his *Essay on Clinical Instruction*. The foundation of what was for a while called the “Numerical Method,” it, along with his other works, also constituted a cornerstone in the history of clinical evaluation. Louis apparently performed the first chart reviews (“outcomes research” if you like) and thereby produced evidence to undermine beliefs about blood-letting, not least the notion that bleeding cured cholera³ (Porter 1996; Weatherall 1996):

As to different methods of treatment, it is possible for us to assure ourselves of the superiority of one or other . . . by enquiring if the greater number of individuals have been cured by one means than another. Here it is necessary to count. And it is, in great part at least, because hitherto this method has not at all, or rarely been employed, that the science of therapeutics is so uncertain. (Louis 1834: 26–28)

It is simple arithmetic, but it is systematic enough to be credited as an intellectual parent by Osler (1985 [1897]) and as a key antecedent of evidence-based practice in a major text and exposition (Sackett et al. 2000). Indeed, it might be possible to blame Osler for contributing to, if not originating, the belief that the acquisition of evidence is a simple matter:

Louis introduced what is known as the Numerical Method, a plan which we use every day, though the phrase is not now very often on our lips. The guiding motto of his life was “*Ars medica tota in observationibus*”, in carefully observing facts, carefully collating them, carefully analysing them. To get an accurate knowledge of any disease it is necessary to study a large series of cases and to go into all the particulars – the conditions under which it is met, the subjects specially liable, the various symptoms, the pathological changes, the effects of drugs. This method, so simple, so self-evident, we owe largely to Louis. (Osler 1985 [1897]: 193)

We see at any rate the several forces at work as medical science moved from innocence to awareness of the varied and gorgeously useful data to be teased from clinical experience.

From Beddoes and Louis to Cochrane

What followed, albeit as a result of a number of forces disconnected from Beddoes and his dicta and Louis and his data, was an industry that published the cases and observations of clinicians, often in journals linked to professional societies (Bynum and Wilson 1992). Still, it took until the middle of the twentieth century before medical science was to evolve the tool we call the randomized clinical trial and which we tend to regard as the gold standard for generating the information which we then turn into evidence (British Medical Journal 1998).⁴ But what had failed to evolve was a system for making the information-evidence alchemy reliable and broadly available. The failure was frank and unavoidable: Clinicians needed help in muddling through the vast and often contradictory mess of information that might or might not drift across the transom.

What has emerged over the past quarter-century is a series of pronouncements about the percentage of health care that is based on (high-quality or even gold standard) evidence. This percentage is always very low – it ranges from 10% to 25% of medical decisions. The numbers leave us slack-jawed. If clinicians' decisions are based on (high-quality) evidence only 10% or 25% or even 50% of the time, then what on earth is guiding the rest of the decisions in which pain, suffering, disability, and life hang in the balance?

The numbers matter, because if they are wrong we have less to worry about, and if they are right we had better get moving.

The origin of the numbers is obscure. At a delightful and illuminating UK-based website, “What proportion of healthcare is evidence based?”, the question elicits pointers to a broad variety of sources that try to establish the percentage with some degree of, well, evidence. The origin of the minimal-evidence claim is worth savoring:

“The 10–25% of medical decisions are evidence-based” comes from a series of conjectures, many of them humorous, starting back in the 70's. For example, in an exchange between giants of epidemiology, Kerr White . . . and Archie Cochrane . . . in Wellington, NZ, Kerr had just suggested that “only about 15–20% of physicians' interventions were supported by objective evidence that they did more good than harm” when Archie interrupted him with: “Kerr, you're a damned liar! [Y]ou know it isn't more than 10%.” (Booth et al. 1999)

Other sources of the claim include the US Office of Technology Assessment (10–20%; Office of Technology Assessment 1978) and the Institute of

Medicine (2–25% varying by strength of evidence and consensus; Field and Lohr 1992). But we might as well trace the contemporary unease, the epistemological gap, the cognitive itch to Archie Cochrane’s observation and complaint: “It is surely a great criticism of our profession that we have not organised a critical summary, by specialty or subspecialty, adapted periodically, of all relevant randomized controlled trials” (Cochrane 1979; cf. Cochrane 1972). There is a sense in which Cochrane, born in Galashiels, Scotland, in 1909, a century after Beddoes’ death, has done as much as any other individual to reshape health education and practice. Since his death in 1988, his core idea has mutated from insight to tribute to movement to professional imperative.

Such a simple idea, so straightforwardly expressed, so intuitively compelling: It was a great criticism indeed that there were, two millennia after the physicians at Kos, few collective accounts or reports that what clinicians did had any demonstrable scientific traction. It is surely one of the great oversights in the history of human inquiry and applied epistemology. Indeed, even Beddoes’ proposed system for rudimentary data collection and outcomes research had not been instituted in some quarters, two centuries later.

To be sure, the fact that a claim or practice is not supported by (adequate) evidence does not mean that it is false. But it does mean that practitioners have inadequate grounds for believing it to be true or effective. Put differently: The remaining 90% or 75% or whatever percentage of health care that is said not to be evidence based is not necessarily false – only that clinicians lack (adequate) justification to assert it to be true and perhaps to practice as if they knew it to be.

Efforts to warrant such assertions have, in less than 25 years, shaped everything from medical reimbursement to resource allocation to managed care to public health to individual decisions by individual clinicians – faced and challenged by vast amounts of health research and by partisan claims about how to make sense of it all.

Observe that Cochrane’s complaint here is not that most individual clinicians’ beliefs are not supported by evidence. It is that, for all they know, they are! It is therefore a much broader challenge: If the beliefs of clinicians can be linked to evidence, then those links need to be accessible – in an “organized, critical summary.” In this sense, Archie Cochrane was not asking for more research or more evidence; he was merely observing that the evidence we already have is removed or disconnected from the people who ought to be using it to take care of sick people.

In terms of the two kinds of ignorance we included in our earlier list of causes of bad outcomes, Cochrane is therefore talking about the first, or individual ignorance (call it “I-I”; collective ignorance will be “C-I”). Now, I-I must, if it is to make any sense, mean something like this:

A clinician is *individually ignorant* of evidence that would affect her practice if she were aware of it, if:

- such evidence has already been acquired by someone, somewhere; and
- the evidence is not secret, and has not been hidden or unpublished/unshared.

We want these two conditions to apply in order to insulate us from cases (as suggested in Chapter 2) in which, say, a discovery had been made but unreported – ignorance of such a discovery would be ignorance that (almost) everyone has and so would be no different in salient respects from collective ignorance:

A clinician is *collectively ignorant* of evidence that would affect her practice if she were aware of it, if:

- such evidence has not been acquired by anyone yet, where “acquired” may include evidence requiring either primary research or knowledge discovery as is attributed to database research, meta-analysis, etc.

In other words, the set or class of people who are collectively ignorant comprises members none of whom has evidence that would affect practice decisions if the evidence were known. Ignorance, like misery, loves company.

There are several ways in which our simple distinction does not cover all eventualities. For instance, in trying to distinguish between having and not having some evidence, we have adopted a vague and imprecise picture of evidence. It makes no mention of the quality of the evidence, for instance. When we say “evidence that would affect her practice if she were aware of it,” we have said nothing about whether our clinician *should* alter her practice in light of it, whether it would be *rational* to do so, etc. These are important questions and we will return to them in Chapter 2. For now, it is adequate to point out that the historical thrust of evidence-based practice owes much to men who were asking for something very basic and obvious.

That something is basic or obvious, however, does not always make it easy to believe.

Health science and the growth of knowledge – the role of “experts”

Even as scientific progress entails scientific change, the reverse is not the case. Some changes are trivial, off the mark, or even destructive or regressive. Evaluating candidates for, and changes in, scientific corpora requires a broad cluster of attributes, most especially knowledge of the science in which an inquiry is conducted and in which the progress is alleged. The conflicting forces of specialized knowledge and interdisciplinary knowledge create a vast cluster of problems for those who would assess scientific change and progress. Such an assessment is precisely what evidence-based medicine demands.

On one obvious reading, the evolution of specialization in medicine has been driven by the great and rapid accumulation of information about the human organism. The effects of this accumulation force us to abandon the hope that individuals might become complete masters of particular disciplines:

In every subject of scientific study the progress of investigation and the accumulation of knowledge must reach a point where it becomes a serious task to master all its facts, or to be acquainted with all that has been written about it. When a great number of zealous observers are bending their energies in a common pursuit, it happens after a time that not the oldest and most eminent among them can possibly attain to a perfect acquaintance with all that is known about it. (Noyes 1865: 59)⁵

In the century-and-a-third since Dr. Noyes reflected thus on the heavy weight of information that increasingly attaches to the good fortune of knowledge, the situation has become somewhat more complicated. Not only has medicine progressed, it has found itself closely allied with disciplines not then imagined. All the while and as ever, going back to Plato and Aristotle, the questions of how to assess medical and other scientific claims, and of who is most fit to communicate the claims, have tended to turn to experts.⁶ Now, this could be a problem because it is probably too much to require that ordinary clinicians become experts. If evidence-based practice required this, it would be doomed to failure. It is a more democratic enterprise, requiring that all clinicians take responsibility for their own epistemic warrants. In other contexts, we would call this “education.”

Fortunately, there are a number of ways to defang calls to expertise. One is to make the case that “expertise” is akin to “narrowness” or “rigidity.”

Another is to undermine the very idea that there is such a thing as expertise, or that it is needed for the purpose at hand.

For instance, regarding the former, the philosopher Paul Feyerabend has condemned experts “who quite naturally confound knowledge with mental rigor mortis” (Feyerabend 1975: 182; cf. Feyerabend 1978). Feyerabend is also reading “authority” for “expert” and this bit of economy links social position with (narrow) epistemological status: better to plague both houses.

Can we successfully set aside the very concept of “expert?” Or, better, does evidence-based medicine even need any expertise?⁷ Since a large part of the evidence-based engine requires the synthesis and communication of information by ordinary clinicians, we should look a little more closely at this process of evaluation.

Evaluating progress in medicine

Evidence of progress in medicine and nursing is in some respects less controversial than evidence of progress in physics, genetics, astronomy, psychology, and other sciences. Where quarks, genes, black holes, super-egos, and other entities have raised, and continue to raise, difficult problems for those who postulate their existence, structure, and function, the entities of modern medicine lend themselves somewhat less readily to philosophical scavenging. This is emphatically not to argue that medicine offers no or even only few special difficulties for the analysis of new evidence – indeed, as we will see later, scientific uncertainty poses the greatest ethical challenge to evidence-based practice – it is merely to suggest that these difficulties are in some respects harder to come by or more tractable than issues in other sciences. The observation has this to recommend it: To the extent that medicine reduces to chemistry and physics, its deepest problems will not be uniquely medical at all but rather chemical, physical, and so forth.

Still, there are fundamental difficulties in the task under analysis, namely assessing and communicating facts and allegations of progress. No matter how we join the old philosophical debate over progress in science, the question of whether there has been any of it in a given domain will be answered only or best by those who have some set of skills and/or some amount of knowledge.

Consider provisionally that an assessment of progress in a science will require knowledge of the (at least short-term) history of the science and of

the events alleged to be progressive. For now, call a person with that knowledge an “expert.”

Now suppose there are two kinds of progress – technological and theoretical. Technological progress might be found in the invention of a surgical tool or a drug, even perhaps in the identification of a disease. Theoretical progress will be the acquisition of explanatory knowledge about the structure and function of organs, diseases, systems, and so on. Thus the ancient discovery *that* certain substances had antiseptic properties was an instance of the former; work by Pasteur and Lister increased our understanding of *why* some substances are antiseptic, and so was progressive at the deeper, theoretical level.

We seem therefore to require two different kinds of expert. One teases phenomena from the world, the other learns its secrets. Of course, both of these oversimplified attributes might be found in a single person, as for example in William Harvey, who embodied experimental insights as well as explanatory and theoretical genius. Likewise, if we are lucky, the phenomenon comes to us kit-and-caboodle *with* its explanation and we get both kinds of progress at once. The world is rarely so accommodating, however. (Many errors in the history of science – geocentrism, phlogiston, and the theory of bodily humors come to mind – result from attempts to link too quickly initial phenomena, which might not be genuine, to theories and posits, which enjoy inadequate support. They might be unavoidable.)

The evaluation of medical progress is then a two-fold affair: It is the evaluation of technological progress and theoretical progress. Because every claim or report of a scientific advance ought not to be taken at face value (any individual claim might be wrong), we require some way to judge the truth content and importance of scientific claims. Viewed as a meta-scientific job description, this requirement underlies the work done by journal referees and editors, journalists (popular and specialized), grant administrators, department heads, government officials, and others who must gauge the activity and claims of scientists. Do we need experts for these tasks? In the other direction: Are those who perform these tasks experts?

There are few concepts that are so frequently used and appealed to, and yet so poorly understood, as expertise. This is too bad, because it assumes something we have yet to learn, namely, what it is to be an expert, that is, what constitutes expertise. There are different views about this and they point to different qualities. Psychologists have for some time studied the cognitive bases of expertise and many of these bases have been applied in

crucial ways by researchers in artificial intelligence. Computational expert systems seek largely to model human problem-solving abilities. What is important is the overwhelmingly instrumental or procedural nature of expertise that emerges: “Experts solve complex problems considerably faster and more accurately than novices do” (Larkin et al. 1980: 1335).

Similarly, several common features characterize the progression to expertise:

- Learn and apply basic rules.
- Generalize the rules.
- Develop concepts and additional rules.
- Evaluate and assimilate many situations and cases.
- Apply experience and synthesize. (Hankins 1987: 303)⁸

Thus the novice who would be expert learns the procedural *skills* that have already been mastered by experts. This emphasis on procedure does not overlook content; it just gives it a smaller role.

Sociologists shift the focus from a person’s skill to his or her standing in a scientific community. This standing or reputation is the result of a vast, multi-articulated system of publications, citations in the literature, awards, appointments and so forth (Garfield 1979; cf. Chubin 1976).

Eugene Garfield cites several studies that purport to demonstrate the expert status of Nobel Prize winners by virtue of the fact that they were more frequently cited in their respective literatures than non-prize winners. The affirmation of expertise as such is only tacit in the citation approach, but the implication is clear: Solid work by good or expert scientists is recognized. In any event, Garfield is aware of potential shortcomings of the approach and says of citation counts that “They very definitely are an interpretive tool that calls for thoughtful and subtle judgments on the part of those who employ them” (p. 249).

Of course, this has the effect of starting us off on a circle as vast as the citation network itself – for who picks the Nobel Prize winners if not scientists who we hoped were affirmed *independently* as experts? Garfield mentions one study that assessed “the accuracy of citation counts as a measure of quality in the field of psychology by asking a panel of experts to list the people who they felt had made the most significant contribution to their specialties” (p. 64). Of course, the “experts” confirmed that the most-often cited work was of the highest quality. The appeal to the panel of experts is as simple and straightforward as it is circular and regressive.

Independent of psychology and sociology, we have the intuitively powerful picture of an expert as one who just knows a lot: “In principle, a scientific specialist knows *everything* about a particular subject” (Ziman 1986: 97; original emphasis; note the apparent use of “specialist” as synonymous with “expert”). A problem with this view, call it “strong expertise,” is that even “in principle” it is probably false. Truths can be generated willy-nilly in any field we like and there is no sense in supposing anyone might ever know them all. To be sure, many of these truths will be trivial, but the strong-expert advocate must give us a principled way to separate truths worth knowing from others that are not. This is not a trivial feat and, yet, a core idea of evidence-based practice – that ordinary clinicians may (must!) come to be competent vetters of the science in their domains – requires something like this at the outset.

If we modify this account we might get “weak expertise.” Here, the expert knows everything, or much, of *importance* in his or her field. We still do not have a way to distinguish the wheat-truths from the chaff-truths, but the difference is at least acknowledged.

Sometimes weak expertise is relativized to the “rest of us” (cf. Johnson 1983) so that the expert might just know more than others, even if he or she does not know a lot. But this is certainly expertise on the cheap.

Experts might be supposed to have been the first to find out or discover some of what is publicly known: “Their authority is what they are supposed to know and to have measured” (Calder 1965: 1). Expertise is thus placed at the intersection of knowledge and knowledge-gathering. This picture embodies the one sketched earlier, the one that distinguished between theoretical and technical knowledge, with the crucial point here being that the two kinds of knowledge are conflated. This is tantalizing, but surely it is too strong: clinician-researchers are what we want as peer reviewers – we cannot insist that all clinicians become investigators.

The issue we are considering relates, in fact, to one of the most important and perennial distinctions in the history of human inquiry: the distinction between knowledge and skill. It begins with Plato (*epistēmē* and *technē*), is developed by the philosopher Gilbert Ryle (knowing *that* and knowing *how*), is crucial to linguist Noam Chomsky (*competence* vs. *performance*), and has vexed countless philosophers who have tried to come to terms with the difference between moral philosophy and applied ethics.⁹ Caution is sagely urged: Applied to medical evidence and expertise, the distinction can propel us into vast and unproductive disputes about whether the practicing

physician has more skill but less knowledge than the research biologist, whether the advanced-practice nurse has less knowledge but potentially more skill than the generalist physician, whether the specialist in private practice has less skill and less knowledge than the MD/PhD program graduate, and so on. Better to set this aside and look at how information is communicated, for this is where evidence-based practice must first hang its hat.

Communicating progress – peers and specialists

Evidence changes, accretes. Evidence-based practice requires that this flux and growth be shared, transmitted, understood. Even if we had uncontroversial means for harvesting our scientific crops, the delivery system leaves much to be desired. It is often assumed that scientific publication should be understood as a model of rigor and accuracy – for what does the communication of knowledge entail at first if not accurate dissemination of information within and among scientific communities? But look at the kinds of things that can go wrong:

Fifty randomly selected references from a single monthly issue of *The American Journal of Surgery*; *Surgery, Gynecology and Obstetrics*; and *Surgery* were evaluated for citation and quotation errors. Thirteen major and 41 minor citation errors were found in the three journals. Thirty-seven major quotation errors were identified. The data support the hypothesis that authors do not check their references or may not even read them. This hypothesis may be expanded to maintain that reviewers do not check references. (Evans, Nadjari, and Burchell 1990: 1353)

Other studies¹⁰ (from the first of several international congresses on peer review) found that most manuscripts in a set that had been refereed, revised, and accepted for publication in a leading medical journal were found, when re-submitted to other referees, to require further revisions – but for problems that were different than those found during the first review (Garfunkel et al. 1990), and that “articles citing invalid, retracted work are abundant and ubiquitous in the scientific literature” (Pfeifer and Snodgrass 1990: 1422).

Careful, now. The point is not that any *particular* authors, editors, or journals have erred or fallen short of the mark or some gold standard . . . it is that *any* investigator, reviewer, or publisher is susceptible to a broad array of failures between the science and its uptake by clinicians. This is a problem – of course! – but not in the ways that are first supposed. It is not that any

individual patient will be harmed by such failures (though that is certainly not impossible); it is for our purposes most interesting that such failures are systematic and color the fabric of scientific literature and discourse. Evidence-based medicine, evidence-based nursing, evidence-based public health, evidence-based psychology – all the sciences that have come to rely on the research-literature nexus must be quite clear about the fallibility and uncertainty of their enterprises.

Has anyone been injured by a physician who relied on reports thus flawed? Has a public health policy moved forward or been impeded because of erroneous publication? Has a government or insurance company been unintentionally deceived about treatment efficacy as a result of a scientific mistransmission? It can be quite difficult to answer these questions. Rather than try to answer them, we will do better to come to terms with the idea that, for all our insistence on more and better evidence, uncertainty and not closure is a more constant feature of our inquiries. This point is not to provide comfort to the skeptic or relativist. It is, rather, to try to set out clearly that the growth of knowledge is rarely orderly and stepwise; that fallibility and uncertainty in one degree or another shape beliefs about the way the world works; and that facile pictures about evidence provide more succor for our conceits than warrant for our beliefs.

Specialization and the growth of knowledge

What happened to the old experts? Were there specialists or experts in bodily humors or phlogiston, in vital spirits or flat earths? If I know all there is to know presently about something that does not exist, then what do I know? The answer seems to be that I know nothing; but when false beliefs are cast aside and true ones are adopted, we have nothing less than the growth of knowledge. This does not give us expertise of any sort, but it is well to remember that the search for knowledge does not entail the obligation to become a “know-it-all.”

There was a time when medical knowledge was increased by making observations in gross anatomy; indeed, dissecting cadavers was once progressive. Now, however, medicine is increasingly and necessarily allied with chemistry, genetics, computing, physics, engineering, and so on. Moreover, this is the case even as many or most practicing physicians have forgotten much of what they once knew about chemistry, calculus, or Kant, and decry

proposals that they re-take board certification examinations periodically or be tested to determine familiarity with the current literature in their specialty. The point is that, as medicine has become increasingly interdisciplinary, the ability to gauge progress has in concert become increasingly difficult.

The forces at work are therefore countervailing and powerful. Specialization in some sense narrows the scope of medicine, while interdisciplinary requirements broaden it. Evidence-based dicta require that we get a handle on both.

The growth of medical knowledge entails that individuals will come to learn more than previously. And some people do know more than others, perhaps a great deal more. The above arguments against expertise may be understood in support of a picture that seeks to reject expertise in the limit. There are degrees of knowledge between utter ignorance and “omniscience-in-a-domain”; so, individuals will know x such that $0 < x < 1$, where ‘0’ is ignorance and ‘1’ expertise, and x is not too close to either. (This seems to capture ordinary intuitions about the difference between competence and expertise or, perhaps, the Aristotelian distinction between *acquaintance* and *knowledge*.) Further, knowledge is open to public scrutiny by several means. Minimally, clinicians or biomedical researchers can be, and are, judged informally by the very instrumental and procedural criteria we earlier disdained as tools for judging expertise. This is not to retreat; it is merely to recognize that performance can be taken as evidence of knowledge.

For taking care of sick people, this might be evidence enough.

Defining “evidence-based practice”

Before we proceed, we must state with at least some clarity what is meant by “evidence-based medicine” (or public health or nursing or psychology or physical therapy . . . or, better, *practice*). This requirement is imposed because, with the growth of the movement, there are a number of careless or loose formulations floating about. When the stakes are high, conceptual clarity itself can exert moral pressure.

We want such a definition to make a commitment regarding what will be accepted as *evidence*. Well, of course! – but the devil in the details here will do mischief later if we are too broad or too narrow in the kinds of warrant we accept as the evidence for evidence-based practice. The definition should

also say something about *clinical experience*, which concept will vex us in subsequent chapters when it is proffered as a foil or kind of counterargument to practice guidelines and the other instantiations of an evidence-based philosophy. Similarly, a definition that is somehow *patient-centered* will help to clarify that the reason to go to all this trouble is not epistemic purity, political correctness or medical fashion; it is because there are good reasons to believe that health care is improved if it is evidence based. Here is what we are getting at, more or less:

Evidence-based medicine . . . is the integration of best research evidence with clinical expertise and patient values.¹¹ (Sackett et al. 2000: 1).

To be sure, this will require a little unpacking, but it will not go as well as we would like:

By *best research evidence* we mean clinically relevant research, often from the basic sciences of medicine, but especially from patient-centered clinical research into the accuracy and precision of diagnostic tests (including the clinical examination), the power of prognostic markers, and the efficacy and safety of therapeutic, rehabilitative, and preventive regimens. New evidence from clinical research both invalidates previously accepted diagnostic tests and treatments and replaces them with new ones that are more powerful, more accurate, more efficacious, and safer. (Sackett et al. 2000: 1)

This is a good start, but imperfect. We might have “clinically relevant research” but it might be of low quality; this would occasion no disappointment except that we are parsing a definition relying on “best research evidence.” Put differently, the concept of “best (research) evidence” should eventually tell us something about what makes it the best. That the research covers a lot of ground does not mean it is any good, let alone the best. The debate over evidence often turns not on how to separate the wheat from the chaff, but how to separate the good wheat from lesser grains. To continue:

By *clinical expertise* we mean the ability to use our clinical skills and past experience to rapidly identify each patient’s unique health state and diagnosis, their individual risks and benefits of potential interventions, and their personal values and expectations. (Sackett et al. 2000: 1)

This is the direction we want, but it seems to be conflating “clinical expertise” with “clinical skill.” The question of expertise, as we saw in a somewhat different context, does not lend itself to facile appropriation; do not appeal to experts unless you are certain you know what you are getting into. At any rate, surely we do not want to insist of clinicians that they be experts before

they can adopt an evidence-based practice. On the contrary, we should want the adoption of an evidence-based practice to be a requirement of all clinicians, and we want this to begin with students. (Of course, it might be that Sackett et al. would reply that by “expertise” they mean “excellence” or “aptitude” or something else that might be demanded of all clinicians and trainees.) Moving on:

By *patient values* we mean the unique preferences, concerns and expectations each patient brings to a clinical encounter and which must be integrated into clinical decisions if they are to serve the patient. (Sackett et al. 2000: 1)

This sets the bar very high, especially if the failure to integrate all these unique preferences, concerns, etc., would entail the failure of evidence-based practice. The advantage of setting the bar high, though, is that the players’ failure becomes attributable in larger part to the players themselves and not to the designers of the game. As above, though, a key motivation for making evidence-based medicine patient centered is that it underscores those for whose sake we are toiling so hard.

This definition has a number of virtues, not least of which is that it strives to be comprehensive (although in so doing it encounters a number of difficulties, as pointed out). It is noteworthy that, perhaps because of this comprehensiveness, it is the definition cited in a key US Institute of Medicine report on quality improvement (Committee on Quality of Health Care in America, Institute of Medicine 2001). Like other IOM reports, this one includes a number of recommendations to government authorities. One such, to the US cabinet secretary for Health and Human Services, reads like the voice of Archie Cochrane from beyond the grave. Listen:

The Secretary of the Department of Health and Human Services should be given the responsibility and necessary resources to establish and maintain a comprehensive program aimed at making scientific evidence more useful and accessible to clinicians and patients. In developing this program, the Secretary should work with federal agencies and in collaboration with professional and health care associations, the academic and research communities, and the National Quality Forum and other organizations involved in quality measurement and accountability. (Sackett et al. 2000: 1)

Actually, Dr. Cochrane would likely have put it somewhat differently, but the point is clear. From progress to communication to evidence to expertise to the recommendation that government itself take responsibility for improving health care, the idea that clinicians might not have available the

conceptual tools of the trade becomes a source of moral offense: Ignorance is blameworthy.

Error and evidence

We are, of course, talking about “individual ignorance” of the sort Cochrane was tacitly bemoaning, and which we were discussing earlier. A person has “individual ignorance,” recall, when she or he does not know of, or is unaware of, evidence that others *do* know about or are aware of. The distinction between individual and collective ignorance – a simple affair – is mainly that if no-one has evidence to support a particular belief, then it is not blameworthy if you are among them. But if most people have evidence for a belief, and you are not among them, then you are, or should be, in trouble. Failing to stay abreast of progress in one’s field becomes a moral failing. Indeed, as we earlier noticed, the duty to stay up to date was identified in antiquity.

The reason for insisting on everything from adequate training to continuing education to government policies on the accessibility of evidence is that we have varying degrees of warrant to believe that without them, people will be hurt, or, rather, will be hurt more frequently than with them.¹² This is especially true of training; it would be absurd if otherwise. It is almost always bad news when people with no training attempt to practice medicine or nursing. Continuing education is also essential for clinical practice. The clinician who allows educational or pedagogic time to stand still as of the moment of completion of a residency program, say, has arranged things so that after a few years his patients are no longer seeing a physician, but visiting a museum. In both these cases – training and continuing education – the reason we insist on obedience, compliance, and fealty is that people are injured, harmed, or suffer other untoward and generally undesirable outcomes. Ignorance leads to error. Error can lead to medical harm. Preventable harms can be morally blameworthy.

But where does evidence-based medicine fit on this continuum? What about the evidence- and outcomes-based practice guidelines? Do we know that use of, or adherence to, guidelines reduces particular errors – or is that we have reason to believe it will improve outcomes, a general measure? This distinction, between error reduction and outcome optimization, captures

one of the great tensions in the debate over evidence-based practice. It is, in many respects, a debate as old as any of the health professions, as old as Hippocrates, Paracelsus, or Florence Nightingale; it is a debate that shapes the relationship between epidemiology and public health on the one hand and Mr. Jones's appointment on Tuesday at 2; it is a debate between individual and group, patient and community, specific rights and collective duties. The debate over evidence-based practice is a reprise of one of the oldest conflicts in the history of civilization and rational inquiry: between individual persons on one hand and groups of them on the other. In philosophical ethics, a parallel debate sets a duty-and-rights-based morality against a utilitarianism that defines good itself as good-for-the-community.

The Institute of Medicine report on quality improvement was actually the second on aspects of quality. The first, *To Err is Human: Building a Safer Health System*, addressed aspects of patient safety (Kohn, Corrigan, and Donaldson 2000). The report suggested that adverse events occur in 2.9%–3.7% of US hospitalizations, with 44 000–98 000 preventable deaths annually, findings that have been contested and disputed (Brennan 2000, Sox and Woloshin 2000). What is striking about the first report is that it says almost nothing about the role of evidence-based practice, about outcomes, about guidelines.

A comprehensive research program that looked at the effect of evidence-based practice on error reduction would be a worthy and interesting project – if we didn't have our hands full already. Having identified the vast amounts of research and the intractable number of publications as overwhelming the poor twenty-first century clinician, we need to set our sights on the various attempts to make good on Archie Cochrane's dictum. This leads to many and interesting attempts to reduce clinicians' empirical fallibility with more and better research. Aimed at rescuing clinicians from menses of information, we have produced a new kind of evidence, and it has the effect of increasing the cognitive and intellectual demands on professionals who take care of sick people. It is by no means clear yet that in the process it has reduced their fallibility.

NOTES

- 1 This heading appropriates and alters the title of an important early criticism of outcomes research, "What physicians know" (Tanenbaum 1993).
- 2 Beddoes hired the chemist Humphrey Davy to conduct some of this research. Sir