

# Bridging The Inferential Gap: The Electronic Health Record And Clinical Evidence

Emerging tools can help physicians bridge the gap between knowledge they possess and knowledge they do not.

**by Walter F. Stewart, Nirav R. Shah, Mark J. Selna, Ronald A. Paulus,  
and James M. Walker**

**ABSTRACT:** Most clinical decisions involve bridging the inferential gap: Clinicians are required to “fill in” where they lack knowledge or where no knowledge yet exists. In this context we consider how the inferential gap is a product, in part, of how knowledge is created, the limits to gaining access to such knowledge, and the variable ways in which knowledge is translated into decisions. We consider how electronic health records (EHRs) will help narrow this gap by accelerating the creation of evidence relevant to everyday practice needs and facilitating real-time use of knowledge in practice. [*Health Affairs* 26, no. 2 (2007): w181–w191 (published online 26 January 2007; 10.1377/hlthaff.26.2.w181)]

**I**N THIS PAPER WE EXAMINE the inferential gap common to everyday practice: the gap between the paucity of what is proved to be effective for selected groups of patients versus the infinitely complex clinical decisions required for individual patients. Clinicians engage in information gathering and interpretation; they implicitly or explicitly bridge the gap every day to care for their patients. The breadth of the inferential gap varies according to available knowledge, its relevance to clinical decisions, access to the knowledge (that is, what the physician actually knows at the time of a clinical decision), the variable ways in which knowledge is interpreted and translated into a decision, the patient’s needs and preferences, and a host of other factors. Clinicians are required to fill in where their knowledge (or knowledge itself) falls short. These issues are increasingly important for an aging U.S. population where clinical decisions must consider the patient’s entire complement of comorbidities, genetic predispositions, and preferences.<sup>1</sup>

We consider two fundamental means by which electronic health records (EHRs) will narrow this gap. First, EHRs will facilitate the creation of evidence

---

*The authors are with the Geisinger Health System in Danville, Pennsylvania. Walter Stewart (wfstewart@geisinger.edu) is the director of the Center for Health Research, Geisinger Clinic. Nirav Shah is an associate investigator at the center. Mark Selna is associate chief medical officer, Division of Clinical Effectiveness. Ronald Paulus is chief technology and innovation officer; James Walker is chief medical information officer.*

that is directly relevant to everyday clinical decisions.<sup>2</sup> Second, EHRs will greatly increase real-time access to knowledge in the practice setting.<sup>3</sup> We consider five specific scenarios that are relevant to these fundamental shifts in the creation and use of clinical evidence. In the long term, we believe that EHRs will offer a novel approach to the creation of clinical knowledge, in which observing, intervening, and creation of clinical evidence are part of the normal clinical encounter.

We first briefly review the evolution of medical evidence and its limitations.

### **Medical Evidence: How Relevant To Patient Care?**

Historically, the formal creation of clinical knowledge began with case reports; it then evolved to more sophisticated observational studies, then to controlled trials, and, ultimately, to the randomized controlled trial (RCT). To some degree, observational studies and RCTs are complementary. The primary strength of systematic observations (that is, scheduled collection of data) in relatively unselected populations is that the resulting evidence is highly relevant to a broad spectrum of patients (that is, generalizable). However, inference from such studies is prone to challenges with interpretation from bias and confounding.<sup>4</sup> The use of randomization—the key design feature of RCTs—minimizes these biases.

Traditionally, knowledge regarding quality of care was based on clinical experience, case studies, and reasoning from physiological first principles. This knowledge sometimes led to improvements dramatic enough to convince even the most skeptical observers of a treatment's benefits (for example, Jenner's smallpox vaccination): There was minimal variability in outcomes and the cause-effect relationship appeared obvious. The relative dearth of knowledge at the time, the narrowness and simplicity of the questions being asked, the low expectations of patients and other interested parties, and the high benefit-to-risk ratios justified the use of observational methods as a means of discovery. In the early twentieth century, though, when medical knowledge creation itself emerged as a nascent "industry," research methods evolved to include more comprehensive observational studies and, ultimately, the RCT.<sup>5</sup> The increasing complexity of clinical questions and treatment options, diversity of potential clinical outcomes and their alternative scientific explanations, as well as growing concerns with bias and confounding all contributed to a demand for a more rigorous and reliable approach to answering questions, for which the RCT emerged as the gold standard. However, the application of the RCT has evolved in response to the above demands with increasingly narrow-focused interventions applied to increasingly selected populations.<sup>6</sup>

As the dominant method for creating clinical knowledge, the RCT is not typically used to address questions directly relevant to the practice setting; for cost, practical, and logistical reasons, the traditional means of using the RCT in this context is not sensible. Instead, RCTs have been and will continue to be used to test for treatment benefits in highly selected populations with a low comorbid disease burden. More specifically, RCTs are used primarily to define whether a

*“In the past three decades, questions that have to be answered in making clinical decisions have become more intricate.”*

.....

drug or intervention is beneficial in an artificially optimized clinical setting, not whether it makes sense for most patients or is suitable for one patient subgroup but not another. Strict selection criteria exclude high-risk patients. Furthermore, commonly used run-in phases prior to randomization select for the most adherent patients. The RCT can be easily manipulated to influence what is observed and discoverable.<sup>7</sup> Even seemingly trivial changes in inclusion criteria, intervention characteristics, or follow-up duration (for example, the recent publicity over the cutoff date for the Vioxx trial's cardiac events) can lead to different conclusions.<sup>8</sup> For example, RCTs on the treatment of mild hypertension and coronary disease have alternatively shown both benefit and harm, as a result of the differential patient exclusion criteria among different trials.<sup>9</sup>

Other important changes have also occurred that contribute to growing challenges with the inferential gap. In the past three decades, patient populations have become older and more heterogeneous, and the questions that have to be answered in making clinical decisions have become more intricate. Patients with complex medical needs tend not to be eligible for RCTs. Consequently, for the growing population of patients with multiple comorbidities, medication intolerances, limited cognition, and diverse insurance coverage, the knowledge needed to support objective clinical decision making is largely nonexistent.<sup>10</sup>

To some degree, systematic observational studies—especially in clinical settings—address some of the limits to RCT-based evidence.<sup>11</sup> Outcomes in “real-world” clinical practice reflect the clinician's knowledge, skills, preferences, and interaction with the patient as well as patient factors (such as self-management, adherence, and willingness and ability to pay) and system features (such as care-coordination assets and ease of access to care). Systematic observations of these and other factors in unselected populations address a key need: generalizability. The primary challenge, however, is dealing with the increased likelihood of confounding, specifically confounding by indication (that is, when a medical condition both triggers the use of a specific treatment and is associated with a risk of the outcome under study).<sup>12</sup> Confounding by indication is inherent to many clinical decisions. For example, disease severity can influence the choice of treatments, making it difficult to separate the indication for a treatment from the risks/benefits of the treatment itself. This problem is commonplace when a new drug comes to market. For example, because of preestablished practice patterns, a new class of medication might implicitly be used only for patients with more-severe disease (for example, uncontrolled hypertension). Analysis of the benefits of the new drug class versus older classes will be confounded because, in part, the treatment decision will link disease severity to when a drug came to market.<sup>13</sup>

## EHRs And Clinical Evidence

The creation of evidence using traditional research designs is extremely time-consuming, costly, and, as previously noted, limited. We consider five different scenarios (Exhibit 1) for how EHRs will likely influence the traditional research paradigm. The scenarios are presented in sequence from ones that primarily offer logistical advantages to those directly relevant to clinical practice needs. Together, these scenarios describe unique features of research enabled by the advent of EHR-based clinical practice. Namely, the process of engaging in practice-based research will motivate improvements in data quality and the specificity of questions that can be answered, which will in turn influence the ability to monitor and improve patient outcomes and vice versa.

■ **Facilitating practice-based RCTs.** There are regulatory, logistical, and cost

**EXHIBIT 1**  
**Profiles Of Scenarios Involving The Electronic Health Record (EHR)**

Scenario	Dominant application	Example	Strengths/advantages	Disadvantages
Facilitating RCTs in practice	Traditional practice-based RCTs	Large multicenter clinical trials	Acceleration of RCT timeline, understanding of selection bias	Favor EHR adopters
Retrospective analysis of EHR data	Retrospective analysis of typical clinical interventions that have been in use for several years; etiology, diagnosis, and prognosis studies	Evaluate effectiveness of different treatments and combinations of treatment for the management of hypertension	Highly relevant to clinical practice: offers the only means for comparative treatment analysis for the same outcome	Measurement error, bias, confounding, confounding by indication, and completeness of data
Testing of self-sustaining health services models	Use of patient-completed questionnaires, systematic use of behavioral health interventions, patient self-management training, systematic adherence to care per guidelines	For example, like the 1960s HIP study of mammography: introducing a new diagnostic by varying eligibility <sup>a</sup>	Testing of new models of care must fit the workflow of clinical practice; if the model improves outcomes and is successful, then the “translation to practice” challenge is simultaneously addressed and solved	Requires a strong alignment of the research, IT, and practice groups
Real-time use of decision support and all patient data in practice	In everyday practice to directly integrate CPGs and other codified knowledge with practice decisions and to make use of data on other patients as needed to tailor treatment for a particular patient	What is the best medication to add to an already complex regimen for a patient with multiple comorbidities and insurance restrictions on treatment options?	Provides clinical guidance in complex situations; directly provides the physician with knowledge as needed; makes effective use of longitudinal data on other patients	Limits to knowledge and the above methodological limits to use of retrospective data
Randomizing uncertainty in real time	Clinical decision making in situations of equipoise	What is the real-world NNT for various statins?	Simple to do and replicate in numerous settings	Generalizability might be overestimated if relying on one population

**SOURCE:** Derived from the authors’ own work.

**NOTES:** RCT is randomized controlled trial. IT is information technology. CPG is clinical practice guideline. NNT is number needed to treat.

<sup>a</sup>See S. Shapiro et al., “Current Results of the Breast Cancer Screening Randomized Trial: The Health Insurance Plan (HIP) of Greater New York Study,” in *Screening for Breast Cancer*, ed. N.E. Day and A.B. Miller (Toronto: Hans Huber, 1988), 3–15.

challenges in conducting RCTs that together have greatly increased trials' average time to completion. For a number of reasons, expansion of the number of ambulatory practices with EHRs will mitigate various causes of delay. First, pretrial analysis of various inclusion/exclusion criteria against an EHR database will improve the protocol used to optimize safety, increase the number of eligible patients, and speed enrollment. Second, the growth of EHR-based practices will influence the "reach" of clinical trials and help move from the typical "high recruiting" clinics to a broader population. Third, access to EHR data will provide a more specific understanding of one dimension of the inferential gap: differences between patients who choose to participate in trials versus peers who would have qualified but opted out versus those who were excluded. Lastly, the efficiency of data capture and data quality monitoring is likely to improve as EHRs and electronic data capture (EDC) tools become increasingly integrated, providing for seamless transmission of data from the EHR to a digital case report form, billing record, or real-time adverse-event alert.

Several years ago we at Geisinger Health System embarked on a systems approach to clinical trials, providing centralized support services (for example, protocol review, contracting, institutional review board [IRB] application and management, and trial setup and monitoring) to both primary and specialty care. In the past year we have begun to use EDC tools to seamlessly extract data from the EHR to the digital case report form. Although we believe that EDC tools offer tactical advantages in the management of individual trials, they also address a fundamental challenge at a system level in accurately tracking progress in multiple trials (that is, number of patients, specific visits completed, specific data captured). Finally, implementation and use of EDC tools at a system level will likely facilitate data collection and study management for traditional epidemiologic and health services research. Clinical departments at Geisinger may develop their own resources; however, the centralized resources we are developing are offered as a competitive option. Increasingly, the system-level resources and capabilities exceed what is possible within any given research unit or clinical department.

■ **Retrospective analysis.** Given the relative ease with which access to longitudinal EHR data can be gained, retrospective data analysis will be a dominant focus of interest in addressing questions of treatment benefit and harm, in mining for new treatment indications, and in answering the complex but common questions that arise in practice. We consider a few examples of retrospective data analyses with a specific focus on blood pressure—a common and important clinical measure.

Policymakers have an interest in the relative value of different treatment regimens for managing blood pressure.<sup>14</sup> Access to this type of comparative information will address a fundamental gap relevant to the cost and effectiveness of various drug formulary designs. Although meta-analysis applied to traditional RCT data might reveal comparative information on treatment benefits across RCTs, conclusions are almost always uncertain, given numerous methodological challenges (for example, heterogeneity in measurement protocols, populations, and

treatment regimens) and previously noted limitations.<sup>15</sup>

For individual patients, the study of treatment response in practice and interactions between treatment status and patient traits will advance clinical practice guidelines toward the numerous and complex issues common to everyday practice. Treatment decisions for a particular patient will likely be influenced by a host of system (for example, formulary options), socioeconomic, treatment (for example, adverse events and response), clinical, genetic, and other factors. Making treatment decisions with evidence relevant to these and other nuances of needs and individual attributes both closes the inferential gap and supports a more patient-centered approach to care.

*High blood pressure.* A large population example we have begun to consider concerns the management of high blood pressure in the elderly. Since 1991, large-scale studies suggest that antihypertensive drugs that greatly lower diastolic blood pressure in men older than age seventy-five and possibly in older women are associated with higher mortality rates, despite their apparent benefit in lowering systolic pressure.<sup>16</sup> The relative impact on systolic and diastolic pressure likely differs by drug class, some of which have not been evaluated; comorbidities; and other risk factors such as smoking status. Today we still know relatively little about who is actually at risk, aside from the broad category definitions such as males older than age seventy-five. We will be using comprehensive longitudinal EHR data (that is, sequential blood pressures, medications prescribed, nonfatal and fatal events, diagnoses, smoking status, and so on) to investigate the relative safety and benefits of antihypertensive medications relative to probable risk factors. The resulting knowledge will be used to develop decision-support logic for blood pressure management in older patients.

*Data limitations.* Limitations to the quality and completeness of EHR data will be the “Achilles’ heel” that constrains evidence that can be extracted from retrospective analysis. In research settings, rigorous standards are used to measure and record data. In practice, though, this is not always the case. Ultimately, we believe that data quality and completeness challenges will have to be resolved through standards of practice that satisfy different stakeholders. For example, improving the quality of blood pressure data would require that the numerous clinical “habits” that result in bias and error are minimized.<sup>17</sup> Research protocols designed for this purpose are not logistically or financially suitable for practice settings. Rather, a more sensible standard of practice could mean that blood pressure is measured only by automated cuffs and obtained sequentially at defined intervals, and that results are directly transmitted (perhaps wirelessly) to the EHR. Improvements to data quality will accelerate as common interests are identified among key stakeholders (that is, clinical effectiveness monitoring, clinical operations, and research) and, in particular, where new methods improve data quality without imposing a burden on the practice and, more likely, where such methods improve efficiency.

Notions of data completeness—a routine obsession of research protocols—poses challenges for retrospective analysis. The schedule for when and what data are collected on patients in practice must, by necessity, be linked to what is sensible for appropriate care. In general, patients with more health problems will have more visits and more data. EHRs can facilitate and improve the likelihood that a patient is seen when appropriate (for example, via automated visit reminder letters). But the notion of completeness itself raises numerous questions. Are there minimal data needs (for example, height, repeated measures of weight, blood pressure, or lipids) for any retrospective analysis? Does the optimal schedule for collecting data differ for each clinical measure? What design, analytic, and inference challenges are created when the amount of data is related to a patient's health status? These and other questions will pose challenges, as previously noted, to aligning the interests of stakeholders regarding the data to be routinely collected during encounters. In particular, the data needs of researchers, practitioners, and those who manage quality of care will create tensions that can best be balanced through protocols that both improve the completeness and standardization of data captured and reduce the cost of obtaining such data. Such technical solutions already exist in a number of areas and are used with increasing frequency at Geisinger. For example, we have begun to develop workflow and data capture models for patient-completed questionnaires. Computerized order entry of prescriptions, tests, and procedures that require selection from a menu of predefined options can be set to require that one or more diagnostic codes be selected, indicating the intention behind the order. When used properly, predefined order sets, structured notes with defaults, and consultation templates standardize the content and organization of data input and can even enable structured, codified data capture. Even these rudimentary EHR protocols have the potential to contribute meaningful evidence that will complement knowledge gained from randomized trials.

Within the bounds of these limitations, retrospective analyses of EHR data offer enormous potential value and will inevitably advance methods relevant to causal inference. It is likely that a body of science (for example, validation studies of established previous findings and knowledge of the potential influence of different confounders) will emerge to focus on just this issue as it has in other areas and will give way to standards of practice relevant to the interpretation of EHR-based evidence. Specific methods will ultimately need to be developed linking the types of questions to the analytic methods most suitable to answer them.<sup>18</sup>

■ **Translation of health services models to practice.** The products of research on a new health services protocol rarely get widely translated into practice. Inherent constraints to traditional practice settings limit both the usability and the complexity of protocols that can be tested and the sustainability, exportability, and scalability of proven solutions. The EHR-based practice offers a paradigm shift for how research moves beyond these traditional constraints. We specifically consider the use of patient-completed questionnaires to exemplify differences in the method

and impact of research.

Questionnaires have been developed to facilitate diagnosis, improve patient-physician communication, standardize patient-reported outcomes, and possibly save time.<sup>19</sup> Many excellent questionnaires are used in research; few are actually used in practice. Simple protocols (for example, patients completing a self-scoring questionnaire) do not do enough to influence outcomes in a multistep care process. Complex, idealized protocols might influence outcomes but are costly to deploy and not logistically feasible outside of a research framework. The failure to make effective use of questionnaires speaks to practical limits in a paper-based environment and how this tempers what is sensible to imagine.

Several years ago we began to examine how questionnaires could be used, specifically focusing on our Pediatric Neurodevelopmental Clinic, for the diagnosis and management of disorders such as autism and language delay. In this project, parents completed a seven-page questionnaire before the encounter, with a specific focus on rudimentary questions—which did not require a physician—on eight domains of behavior. Questionnaire data were digitized and imported to EHR templates. The challenge was in designing a process that would allow the physician to rapidly interpret parent-reported data during a clinical encounter and build on these data with standardized structured and semistructured probes. Finally, access to structured patient and physician data was used to generate highly tailored after-visit summaries, including educational material. It is impossible to develop a workable model like this in a paper-based world.

Although the model we created was narrowly focused, the lessons learned are generalizable to clinical settings where behavioral or symptom-based conditions are commonplace. Questionnaires are more valuable in an EHR environment than they can ever be in a paper-based environment. Some values (such as access to structured data from patient) are common to both environments; however, most are unique to an EHR environment (efficiency, seamless use of data and links to structured probes, tracking outcomes over time, and tailored patient education). The EHR provides the means to make effective use of questionnaire data in creating a practical workflow model (that is, meaningful display of data and higher-level probing), solving a fundamental barrier in paper-based clinics. The questionnaire is part of and motivates thinking about a sustainable systems-based approach to new health services models—ingredients that are essential to exportability and scalability of new solutions.

■ **Decision support: CPGs in real time and more.** To facilitate access to clinical knowledge, efforts have been under way over the past decade to codify what we know through the development of clinical practice guidelines (CPGs). This activity is essential to but insufficient for translating knowledge into practice. Sophisticated clinical decision-support technology in combination with EHRs will be required to make effective use of CPGs.

CPGs, based on explicit methods for summarizing established evidence with

*“Patient data themselves will become a critical practice asset, motivating the need to generate high-quality, complete data.”*

expert clinical consensus, have expanded at an accelerating rate.<sup>20</sup> Today there are 1,970 active CPGs listed in the U.S. National Guideline Clearinghouse.<sup>21</sup> This codification of knowledge is essential to increasing its usefulness. However, CPGs represent only one step in a complex process to translate knowledge into practice.

For a number of reasons, CPGs have had only a modest impact on care. First, relatively primitive means such as publication and education are used to promulgate them. Second, clinicians do not have the time to begin to learn even a fraction of them. Third, CPGs rarely attempt to describe how to operationalize the recommended tasks (that is, CPGs represent knowledge but not accountable, manageable work flows). Even if complete, relevant, perfect, and codified knowledge existed today, there is no effective means of accessing such knowledge in real time even with existing EHRs. What we are lacking is a decision-support capability that can assimilate detailed relevant information about the patient, evaluate such data in real time against existing knowledge, and then yield recommendations that the physician can act upon.

We recently completed a pilot project to determine how such a process could be created with a specific focus on cardiovascular risk management in primary care. The process itself led to specific protocols for ordering measures to fully determine a patient’s cardiovascular risk, including questionnaire data on behavioral risk factors. A decision-support rules “engine,” external to the EHR, was used to extract patient data in real time, evaluate the data in relation to rules, and generate and return a recommended order. This project demonstrated that it is technically possible to create a real-time decision-support workflow that translates CPGs into practice. It also revealed a conceptual and practical challenge: Translation of all CPGs into practice will result in too much care. CPGs often do not recognize the confounding aspects and logistic complexities of comorbid conditions. For example, the blind application of all relevant CPGs to a typical hypothetical elderly patient would lead to twelve prescribed medications, costs of \$400 per month, and numerous potential side effects.<sup>22</sup> Codified knowledge does not exist that can guide decision priorities among diverse sets of CPGs.

■ **Determining optimal management protocol: clinically ranked data.** Finally, we expect advances in the use of retrospective data analysis to influence clinical decision support. The real-time use of longitudinal EHR data to guide clinical decision making will be a logical extension of retrospective data analysis but will represent a conceptual leap—well beyond traditional notions of evidence and decision support. In this framework, patient data themselves will become a critical practice asset, motivating the need to generate high-quality, complete data. For example, consider the needs of a patient with hypertension, diabetes, and atrial fibril-

lation. Different classes of antihypertensive medication might be recommended if each disease were clinically considered in isolation: thiazide for hypertension, angiotensin-receptor blocker (ARB) for diabetes, and a beta-blocker for atrial fibrillation. Traditional efficacy evidence cannot provide explicit guidance on a decision that seeks to simultaneously optimize the relative benefits of each medication, along with ease of the regimen, formulary coverage, interactions with other medications, importance of side effects, and patient frailty.<sup>23</sup> A logical extension of retrospective analysis of EHR data will be the real-time, rapid processing of longitudinal, population-based EHR data to determine the optimal management protocol given a patient's overall profile and individual preferences. In the example above, the EHR might simplify decision making by presenting data in clinically relevant rankings, taking into account known factors from earlier experience with the patient and others similar patients. For example, a patient's strong preference for generic medications and lower copayments (even at the expense of more frequent dosing) might lead the EHR's ranking of medications to list twice-daily generic medications over once-daily brand-name medications.

■ **Randomizing decision uncertainty: randomized database studies.** The inferential gap in medicine will continue to be an everyday occurrence where the needs and questions asked are always changing and evolving. An everyday solution will be required to meet the perpetually growing demand for new knowledge in medicine. One solution might be to use the power of the EHR to randomize clinical decisions in the face of uncertainty and to evaluate the outcomes accordingly.<sup>24</sup>

In the patient scenario described above, there will inevitably and frequently be true clinical equipoise between two or more potential decisions. In such a situation, a real-time protocol might be embedded within the system to randomly prioritize one decision pathway or another. Over time, the knowledge created by such mini-RCTs could shed light on many questions that could not be addressed by traditional randomized trials or observational studies.

## Concluding Comments

In this paper we have characterized a notion labeled the “inferential gap” and considered the future role of the EHR in closing this gap. In part, the gap is the product of knowledge being created at a faster rate than we can use it, and, importantly, clinical questions growing at a faster rate than can be answered through traditional research methods. We recognize that numerous challenges exist in the widespread application and effective use of EHRs. We expect that these challenges will be minimized as the technology, the data, and their application in practice evolve. From our own recent experience, it is clear that an EHR-based practice environment engenders an unavoidable shift in thinking about clinical evidence and how to create and use it and, importantly, a loss of distinction between clinical practice, quality management, and the creation of knowledge.

.....  
 The authors appreciate the input of Marc Gourevitch, who commented on earlier drafts of this paper.

## NOTES

1. T. Bodenheimer, E.H. Wagner, and K. Grumbach, "Improving Primary Care for Patients with Chronic Illness," *Journal of the American Medical Association* 288, no. 14 (2002): 1775–1779.
2. R.B. Haynes, R.S. Hayward, and J. Lomas, "Bridges between Health Care Research Evidence and Clinical Practice," *Journal of the American Medical Informatics Association* 2, no. 6 (1995): 342–350.
3. R.B. Haynes, "Using Informatics Principles and Tools to Harness Research Evidence for Patient Care: Evidence-Based Informatics," *Medinfo* 9, Part 1 Supp. (1998): 33–36.
4. K.J. Rothman and S. Greenland, *Modern Epidemiology* (Philadelphia: Lippincott Williams and Wilkins, 1998).
5. A. Flexner, "Medical Education in the United States and Canada: From the Carnegie Foundation for the Advancement of Teaching, Bulletin Number Four, 1910," *Bulletin of the World Health Organization* 80, no. 7 (2002): 594–602; and J.P. Vandenbroucke, "A Short Note on the History of the Randomized Controlled Trial," *Journal of Chronic Disease* 40, no. 10 (1987): 985–987.
6. N. Black, "Why We Need Observational Studies to Evaluate the Effectiveness of Health Care," *British Medical Journal* 312, no. 7040 (1996): 1215–1218.
7. K.M. Mattocks and R.I. Horwitz, "Placebos, Active Control Groups, and the Unpredictability Paradox," *Biological Psychiatry* 47, no. 8 (2000): 693–698.
8. J.M. Drazen, "Cox-2 Inhibitors—A Lesson in Unexpected Problems," *New England Journal of Medicine* 352, no. 11 (2005): 1131–1132.
9. S. MacMahon et al., "Blood Pressure, Stroke, and Coronary Heart Disease, Part I: Prolonged Differences in Blood Pressure: Prospective Observational Studies Corrected for the Regression Dilution Bias," *Lancet* 335, no. 8692 (1990): 765–774.
10. S.H. Greenblatt, "Limits of Knowledge and Knowledge of Limits: An Essay on Clinical Judgment," *Journal of Medical Philosophy* 5, no. 1 (1980): 22–29.
11. R.I. Horwitz et al., "Developing Improved Observational Methods for Evaluating Therapeutic Effectiveness," *American Journal of Medicine* 89, no. 5 (1990): 630–638.
12. Greenblatt, "Limits of Knowledge and Knowledge of Limits."
13. M. Salas, A. Hofman, and B.H. Stricker, "Confounding by Indication: An Example of Variation in the Use of Epidemiologic Terminology," *American Journal of Epidemiology* 149, no. 11 (1999): 981–983.
14. P.J. Neumann, "Emerging Lessons from the Drug Effectiveness Review Project," *Health Affairs* 25 (2006): w262–w271 (published online 6 June 2006; 10.1377/hlthaff.25.w262).
15. *Ibid.*
16. F. Forette et al., "Does the Benefit of Antihypertensive Treatment Outweigh the Risk in Very Elderly Hypertensive Patients?" *Journal of Hypertension* 18, no. 3 Supp. (2000): S9–S12.
17. U.O. Andersen, J.H. Henriksen, and G. Jensen, "Sources of Measurement Variation in Blood Pressure in Large-Scale Epidemiological Surveys with Follow-up," *Blood Pressure* 11, no. 6 (2002): 357–365.
18. J.D. Seeger, P.L. Williams, and A.M. Walker, "An Application of Propensity Score Matching using Claims Data," *Pharmacoepidemiology and Drug Safety* 14, no. 7 (2005): 465–476; and N.R. Shah, "What Is the Best Evidence for Making Clinical Decisions?" *Journal of the American Medical Association* 284, no. 24 (2000): 3127–3128.
19. S.M. Downs et al., "Human and System Errors, Using Adaptive Turnaround Documents to Capture Data in a Busy Practice," *American Medical Informatics Association Annual Symposium Proceedings* (2005): 211–215.
20. D.A. Grimes et al., "The Good Clinical Practice Guideline: A Bronze Standard for Clinical Research," *Lancet* 366, no. 9480 (2005): 172–174.
21. See the National Guideline Clearinghouse home page, <http://www.guideline.gov>.
22. C.M. Boyd et al., "Clinical Practice Guidelines and Quality of Care for Older Patients with Multiple Comorbid Diseases: Implications for Pay for Performance," *Journal of the American Medical Association* 294, no. 6 (2005): 716–724.
23. M.V. Maciosek et al., "Priorities among Effective Clinical Preventive Services Methods," *American Journal of Preventive Medicine* 31, no. 1 (2006): 90–96.
24. J.A. Sacristan et al., "Randomized Database Studies: A New Method to Assess Drugs' Effectiveness?" *Journal of Clinical Epidemiology* 51, no. 9 (1998): 713–715.