Chapter 5

ETHICAL AND SOCIAL CHALLENGES OF ELECTRONIC HEALTH INFORMATION

Peter S. Winkelstein

Department of Pediatrics and School of Informatics, University at Buffalo, State University of New York, Buffalo, NY 14222

Chapter Overview

The development of modern bioethics has been strongly influenced by technology. Important ethical questions surround the use of electronic health records, clinical decision support systems, internet-based consumer health information, outcome measurement, and data mining. Electronic health records are changing the way health information is managed, but implementation is a difficult task in which social and cultural issues must be addressed. Advice produced by decision support systems must be understood and acted upon in the context of the overall goals and values of health care. Empowering health care consumers through readily-available health information is a valuable use of the internet, but the nature of the internet environment raises the spectre of abuse of vulnerable patients. Outcome studies have inherent value judgments that may be hidden. Data mining may impact confidentiality or lead to discrimination by identifying subgroups. All of these issues, and others, require careful examination as more and more health information is captured electronically.

Keywords

ethics; bioethics; Internet; electronic health records; decision support; consumer health informatics; outcomes; data mining

1. INTRODUCTION

Modern bioethics and health informatics are intimately connected. It is no coincidence that the two fields emerged at the same time. Modern bioethics originated and had evolved in response to increasingly perplexing questions about the goals of health care. These questions arise largely as a result of the advent of modern technology. Prior to the development of the mechanical ventilator and the dialysis machine, medicine was limited in its ability to extend the life span of patients with incurable conditions. With the introduction of these new devices, it suddenly became possible to imagine machine-dependent patients and, more importantly, to imagine patients who would prefer death to such dependency. Hence the beginnings of modern bioethical thought.

Computers and other informatics technologies and techniques have only made matters more complex. Difficult ethical questions surround the use of electronic health records, clinical decision support and prognostic systems, internet-based consumer health information, outcome measurement, and data mining. Some of these questions are new, raised by new technologic capabilities; some are old, but recast in new forms by the use of informatics.

A fundamental question is: what is the proper role of technology in health care? What health care decisions should be entrusted to computers? Are there decisions that computers should not make or roles they should not play? In attempting to answer these questions, it is important to recognize that computers are merely tools. Tools do not, in themselves, change the underlying goals of health care. They are only properly used to advance the goals of the underlying endeavor. Over 20 years ago, Moor (1979) concluded that the one task that should never be assigned to computers is the choice of the goals of medicine itself.

Although it is clear that the use of health informatics tools should be judged in the same way as the use of any medical tool, i.e. by their ability to advance the goals of medicine, that judgment can be very complex because the tools themselves are complex. In addition, the extraordinary usefulness of computer-based technology creates a digital imperative: a strong incentive to adjust all of health care so as to be compatible with computers. It is much less expensive, for example, to move digital information than to move actual patients or laboratory samples. While the use of informatics can have enormous benefits (e.g. telemedicine for geographically isolated patients) the quality of care can also be degraded in the rush to digitization.

There have been attempts to use bioinformatics tools to work on some of the most difficult bioethics problems that face health care, including cost-effectiveness and futility of treatment, especially near the end-of-life. For example, computer-based prognostic systems have been developed for some

critical care settings (Knaus et al., 1991). Their output, however, has not provided answers to questions such as "when should we stop treatment" or "what treatments should be tried" because the answers to those questions are not entirely quantitative and scientific. Families and caregivers often have very personal values and views about what constitutes appropriate health care. For some, extending life by a week through aggressive interventions may be a great blessing; for others, the same treatment may be considered torture. It is not possible to separate personal values from these decisions.

Rather than prescribe particular consequences for particular actions, modern bioethical thought tends to focus on fundamental truths, laws, or motive forces. The language of bioethics often invokes "principles," which are compact statements of fundamental import. The four key bioethical principles developed by Beauchamp and Childress (1994), for example, are autonomy (the right of an individual to determine his or her own health care), beneficence (the duty of health care workers to improve the welfare of their patients), nonmaleficence (the duty of health care workers to avoid doing intentional harm), and justice (both to individuals and society at large). Unfortunately, the term "principles" has become diluted and is often used to describe simple lists of issues.

The development of modern bioethics has been strongly influenced by technology. But technology itself does not determine the ethics of medicine. Technological advances need to be seen in, and judged by, the light of health care goals. The digital imperative must be resisted unless a clear benefit of computerization can be demonstrated. Technology cannot give us answers to questions that require personal and social value judgments.

2. OVERVIEW OF THE FIELD

2.1 Electronic Health Records

Electronic health records (EHRs), also known as electronic medical records or computerized patient records, are found in an increasing number of physician offices and hospitals. One estimate puts current penetration at 10%, increasing to 25% in 3 years (private communication, 2004). EHRs represent more than a simple computerization of the traditional paper chart. They provide the ability to manage health information using modern information techniques that are impossible to apply to paper record keeping. The use of these techniques has the potential to dramatically change how both individuals and society view health care. Dramatic changes in health care have always been accompanied by equally important ethical challenges, and adoption of EHRs is no exception.

Before enumerating the ethical issues raised by EHRs, it is reasonable to ask why we need EHRs at all. Fundamentally, the answer is that we cannot expect to provide quality medical care without optimal information management. Proper medical treatment depends on timely access to accurate patient medical histories, laboratory results, and many other pieces of data. Problems such as missing or misplaced charts, paper-based laboratory reporting, and illegible handwriting are common roadblocks to care. We are in an era where health care providers can generate enormous amounts of information about a patient but have only antiquated and inadequate methods of managing that information. This discrepancy will only get worse as health care skills get better, especially with the advent of readily-available genomic data. In addition, health care providers cannot begin to empower patients with access to their own medical information if the providers can't manage that information themselves. EHRs provide the tools that can be used to begin to solve these problems.

In addition to improving the quality of care, adoption of EHRs holds out other promises, including improvements in the efficiency with which health care is provided, increased patient satisfaction, opportunities for research, quality improvement and, especially, reduction of errors. Although the EHR has been touted as the key tool for reducing medical errors, the evidence of success is still scant. While recognizing that EHRs are a key component of error reduction in health care the Institute of Medicine (2000) cautions "ALL technology introduces new errors, even when its sole purpose is to prevent errors."

Another promise for which there is as yet no empirical support is return on investment (ROI). EHRs are expensive and complex to install. Savings from reduced file room staffing and square footage, along with savings from reduced dictation expenses, may not fully cover the cost of the EHR. As a result, one can imagine a "have" and "have not" condition, where some practices (perhaps specialty practices in affluent areas) can afford EHRs, but other practices (perhaps inner-city primary care) cannot. That disparity could jeopardize the quality of care for the patients of the "have not" practices. It could also interfere with public health issues because data from the "have not" patients would be much less readily available and therefore underrepresented in public health databases.

The implementation of an EHR system is no easy task (Ash 2003). The conversion from a paper chart to an EHR system puts a great deal of stress on the complex social systems that exist within health care institutions. It requires reconceptualizing the medical record and medical communication, including organizational-level changes in workflow. Resistance to even minor changes is a normal response, especially in complex environments,

and a change of the magnitude represented by the EHR engenders resistance of the same order.

Although EHRs are an information technology (IT) product, the decision to implement an EHR and the selection of an appropriate vendor are not solely within the IT realm. The end-users of these systems must be included in the decision process. Clinicians, especially physicians, are the *de facto* arbiters of EHR acceptance in any health care institution. EHR implementation therefore requires a strong, committed physician champion with the time to devote to the project. It is also critical to manage expectations. Clinicians often wish to believe that an EHR will immediately and completely eliminate all perceived barriers to access to clinical data, and they become frustrated when they find that it does not.

Clinicians are also most concerned about clinical data entry, which is the component least improved by the EHR. Most EHR benefits initially accrue to back-office staff at the perceived cost of clinician's time. In addition, clinicians may perceive the EHR as a barrier to provider-patient communication and family-centered care. Clinicians are also focused on the content of an EHR, such as clinical documentation templates, alerting capabilities, and patient lists. Unfortunately, EHRs are far from turn-key at this time. They come with very limited content and require a great deal of customization in order to function in a particular clinical environment, customization that must largely be done by the end-user. This is a significant barrier to EHR adoption. EHR vendors are beginning to understand that and are devoting more resources to content development.

Another barrier to adoption is that current EHRs are largely stand-alone systems. They typically interface with billing systems and, in a hospital setting, may have connections to laboratory and other systems. But EHRs do not typically communicate with each other. There are currently no well-accepted standards for EHR interoperability. In order to apply the tools of modern information management to health data, especially for population-based studies, there must be a way to aggregate data from many EHRs. The National Health Information Infrastructure project is beginning to address some of these issues (Yasnoff 2004).

The need for EHR interoperability, along with the expense of EHR systems, is likely to drive fundamental changes in how medical records are stored. Centralized third-party medical record keeping, in the form of data "banks," may supplant the current model of record keeping by individual practices. Centralized record keeping would enable health care workers, and patients themselves, to access medical records where and when needed. It would also, of course, require strong security measures.

No discussion of EHRs can ignore the concerns of privacy, confidentiality, and security. Privacy is the ability of a patient to control the

information about him or herself. Confidentiality is the commitment of another person or organization to the patient to control information about the patient. Security measures are safeguards against inadvertent or malicious breaches of confidentiality. Security measures also include protections against loss of information. It is generally accepted that privacy of medical data is an important right of the individual. Privacy may be viewed either as a utilitarian concept (i.e. patients will not honestly and completely discuss their medical problems without assurances of confidentiality) or as a right in and of itself. Privacy is also essential to the exercise of autonomy in medical decision making, just as a secret ballot is fundamental to the exercise of democracy.

Maintaining privacy and confidentiality through appropriate security is one of the key challenges of EHRs. It has long been recognized for related uses of electronic media, such as email (Kane 1998). Aside from technical issues, there are a number of factors that contribute to the challenge. Determining the proper security measures for medical records must be done in the context of the goals for the records. For instance, an important goal of EHRs is to improve access to medical records, for both providers and patients (Delbanco 2004). A perfectly secure EHR would be one to which no access was allowed, so a balance between security and access must be struck. However, the answer to the question of what is the "correct" balance is not a technical or scientific one but rather a social and political one. The answer depends on the values of the participants. These values vary widely. Some people see the benefits of access and are perfectly comfortable with their medical data recorded on computers while others are concerned about of breaches of confidentiality and resist such record-keeping.

EHRs are not the only systems where this balance must be struck. Electronic toll badges, for example, allow for the convenient payment of tolls without actually stopping at the tollbooth via electronic identification. Many people use such devices without concern. Others refuse to use them, fearing the use of the data to track their movements. One feature of the electronic toll badge is that a driver can opt-out of its use at any time simply by leaving it at home and paying the toll in cash. The ability to opt-out of technology on an as-desired basis is important to the acceptance of the technology. At the moment, it is difficult for patients to opt-out of having at least medical billing information entered into a computer. As EHRs become more ubiquitous, opting-out may become impossible.

The issues of privacy, confidentiality, and security have attracted the attention of government regulators. In 2002, the Health Insurance Portability and Accountability Act (HIPAA) privacy regulations went into effect (c.f. http://www.hhs.gov/ocr/hipaa/). The HIPAA regulations are designed to restrict the inappropriate flow of medical information without disrupting

medical care. In particular, the regulations target health data belonging to specific, identifiable patients (PHI). They regulate data flow by dividing medical information use into categories. PHI that is used for payment and other health care operations is subject to the "minimum necessary" restriction, which simply means that only the minimum amount of information necessary (as determined by a reasonable person) to accomplish the task should be used. PHI that is used for medical treatment of the patient is not subject to the "minimum necessary" restriction. Use of PHI for any other purpose requires explicit authorization from the patient.

The structure of the HIPAA regulations puts the burden of determining the "minimum necessary" amount of information and of detecting inappropriate disclosures on the health care provider. This structure is consistent with the current model of record keeping, namely that medical records are largely held by health care providers. With a shift to centralized record-keeping, it becomes possible to give patients more control over and responsibility for the confidentiality of their records. Patients could receive periodic or on-demand reports of the audit trail of accesses to their records. They would then be responsible for detecting and reporting inappropriate uses of their records in the same way that consumers are responsible for reviewing their credit card statements for fraudulent uses of their credit.

The HIPAA regulations require technological, policy, and educational interventions. They affect all PHIs, whether electronic or paper. They also affect how research is conducted, including data mining of medical records, which is considered by the HIPAA regulations to be a form of human-subjects research. The HIPAA regulations essentially extend the concept of harm for research participants to include breaches of confidentiality. Specific procedures for obtaining approval for such research are outlined in the regulations.

Electronic health records are changing the way health information is managed. Especially with interoperability, and possibly centralization, EHRs will allow the application of modern information management tools to health care data. Implementation of EHRs, however, is a difficult task. When implementing an EHR, social and cultural issues must be addressed, and expectation management is critical. Acceptance of an EHR is dependent on acceptance of the underlying goals of the implementation.

2.2 Clinical Alerts and Decision Support

One of the promises of EHRs is that the information they contain can be used to provide automatic alerts such as drug-drug interactions and suggestions for treatment or diagnosis. This naturally raises the question of who is in charge of making medical decisions, the clinician or the computer?

The "standard view" (Miller, 1990) is that human clinicians should retain the ultimate authority to make decisions and that computers should provide advice only. There are two reasons for this standard view. One is simply that computer decision support systems have so far not been shown to be clinically useful, especially in general diagnostic situations (Berner 1994). This does not mean, however, that such systems will never be useful, only that the construction of useful systems is complex. Whether computers will ever provide powerful enough decision support to supplant human clinicians in at least some situations is an empirical question (Moor 1979). The second reason for the standard view is that medical decisions are more than the simple "mapping from patient data to a nosology of disease states" (Mazoué 1990). In other words, many medical decisions cannot be made on entirely scientific grounds. Rather, they require the careful consideration of the underlying goals and values of health care in the context of the individual patient and society at large. This sort of judgment can only be made by those who understand these values and have the skills required to make decisions based on them—namely, humans.

Even if computers remain in an advice-only mode, however, there may still be powerful reasons for following that advice. For example, EHRs and other prescription-writing and dosing programs (e.g. http://www2.epocrates.com, http://www.pdr.net) routinely perform medication interaction checking. If a clinician ignores a warning provided by one of these programs, it is clear that he runs the risk of providing inferior medical care, not to mention of being subject to legal action. Even though the computer has provided "only" advice, the clinician ignores it at his peril. It is easy to imagine that it would be even more difficult to ignore diagnosis or treatment advice.

Because computer alerts and decision support systems can have such power, it is essential to be sure that these systems are properly designed, evaluated, and maintained (Anderson 1994). "Properly" in this context means that clinical decision support systems must adhere to the underlying goals of medicine, which may be different from the underlying goals of a commercial systems designer. Commercial systems must adhere to such health care goals as standard of care, primacy of the best interest of the patient, and informed consent (Goodman 2001). They cannot operate under the usual free-market ethic of *caveat emptor*.

Computer decision support systems that are designed to provide prognosis information are particularly problematic. Obtaining an accurate prognosis can be a difficult task for clinicians, but it is an important one because many treatment decisions are based on prognosis. In addition, the prognosis is a critical piece of information for patients, particularly in cases of life-threatening illness. Prognostic scoring systems can potentially be used

for several purposes (Sasse 1993): quality assessment, resource management (including triage and rationing), and individual patient care decisions.

Quality assessment and improvement is an important goal for any health care institution. Prognostic scoring systems could be used, for example, to compare actual to expected outcomes. This is a reasonable use of scoring systems, assuming they have been properly evaluated, in that it is aligned with, and furthers, the goals of health care. Whether such use will actually improve the quality of care has yet to be demonstrated.

Resource management is a much more difficult problem. We currently do not have a societal consensus on how to manage our health care resources. It is generally accepted that we are no longer able to do everything for everyone but, nonetheless, the health care system continues to function on that premise. Clinicians often find themselves in ethically problematic positions where their traditional role of patient advocate is in conflict with their duty to manage society's health care resources. No computer system can solve this problem. It is possible, however, that accurate information provided by computer systems may assist in the process of making these kinds of difficult decisions.

The most problematic use of prognostic scoring systems is in making individual care decisions. As noted above, decisions based on prognosis are not entirely scientific but are also value-driven (Knaus 1993), thus putting them outside the realm of computers. In addition, prognostic systems are by their very nature based on, and provide, a statistical score. Applying population statistics to individuals is fraught with problems (Thomasma 1988). At the same time prognostic scores can have an aura of certainty and objectivity that they do not warrant. They can also be self-fulfilling: if care is withdrawn due to a poor prognostic score, the patient will certainly die, thus apparently confirming the score. On the other hand, good prognostic statistics are a key to good medical decision making. What is critical is that the data provided by a prognostic scoring system be properly interpreted and applied, which means, in turn, that the users of these systems must be properly trained and qualified (Goodman 2001).

The development of computer-based diagnostic programs has received a great deal of attention in the field of medical informatics. Diagnosis programs are designed to process information about a patient and produce a differential diagnosis list, usually rank-ordered by probability. Most of these programs (c.f. http://www.lcs.mgh.harvard.edu/dxplain.htm) use the Bayes theorem (c.f. Fletcher 1996) to calculate the probability of a diagnosis based on the probabilities that the input signs and symptoms are associated with the diagnosis (there are some notable exceptions to this approach, for example, http://www.isabel.org.uk/). However, these programs have not so far shown a great deal of promise, especially for general diagnosis (Berner

1994). One problem is that Bayesian calculations are strongly dependent on the underlying population statistics. For example, there is a high likelihood that a child with a fever has a viral syndrome, but that is not a useful piece of information for a computer to communicate to a physician. Computer-based diagnostic programs are needed most to remind clinicians about rare or unusual diagnoses, not to determine common diseases. The nature of a Bayesian calculation does not lend itself to detection of rare events. There may be other knowledge discovery tools that are more sensitive to unusual events and that clearly warrant further development in this area.

A similar problem occurs with some processes designed to help with medical decision making. Decision trees (Detsky 1997) can be constructed for some clinical situations, with branches representing outcomes and intermediate states with their associated likelihoods. By simple Bayesian calculation, the likelihood of an outcome can be determined from these trees. By giving each outcome a value ("utility," typically on a scale of 0 to 1, where 0 represents death and 1 represents healthy life), a patient or clinician can get some indication of the most desirable course of action (i.e. the optimal combination of probability and value). Unfortunately, this process has at least three major pitfalls. First, if any of the branch points depend on population statistics (i.e. likelihood of a disease), then the Bayesian calculation is generally overwhelmed by that point, making the rest of the tree irrelevant. Second, a utility scale of 0 to 1 does not capture the full range of possible values. In particular, it is certainly possible to imagine states worse than death (i.e. with negative utilities) (Patrick 1994). Lastly, utilities may be very individual (for example, palliative chemotherapy may be intolerable to some, worthwhile to others). Individual utilities are at least burdensome, and perhaps impossible, to determine accurately. Substituting population-based utilities (averages of utilities chosen by many people given the same situation) (Bell 2001) erases any ability of a patient or clinician to adjust the decisions produced by the tree to reflect personal values.

Clinical alerting and decision making systems can, without question, improve the quality of health care, but they must be implemented properly. Users need training and education about the abilities and limitations of the systems. Systems must be evaluated and maintained. System designers and vendors must understand that their systems will be held to the high standards of medical care. Most importantly, the advice produced by these systems must be understood and acted upon in the context of the overall goals and values of health care.

2.3 Internet-based Consumer Health Information

The dramatic increase in accessibility of information provided by the internet, and especially the World Wide Web protocols, has of course extended to the field of health care. This development is fundamentally good for health care because information is the lifeblood of evidence-based medicine (see Section 2.4 below). Medical information designed for patients can also strengthen patients' ability to make informed judgments about their own care. But the open nature of the web also brings with it the danger of inaccurate or misleading information, both by omission and by comission.

The fundamental basis of the doctor-patient relationship (or the relationship between any reputable health care provider and their client) is the primacy of the patient's best interests. When a provider suggests a course of treatment, the patient can reasonably expect that the provider is suggesting what in the provider's judgment is best for the patient. Sometimes there are several reasonable courses of action that may be appropriate in a given situation and the provider will assist the patient in making choices through the process of informed consent. The concepts of best interest and informed consent derive directly from the principles of autonomy, beneficence, and nonmaleficence.

The relationship between a salesman and a customer is quite different. There, the suggestions by the salesman of choices that might be made by the customer are largely based on the self-interest of the salesman, not the best interest of the customer. A good salesman has the ability to make the customer believe, however, that he has the customer's interests in mind. Experienced consumers are well aware of this and understand the nature of the relationship, namely *caveat emptor*.

The concepts of informed consent and *caveat emptor* may come into direct conflict on the World Wide Web. Because authorship and, more importantly, the author's intent can be difficult to determine on web sites, it is difficult for patients to know what ethical construct to apply to a particular site. Some sites provide authoritative medical information that is designed to enhance the ability of patients to understand their conditions and make appropriate health care choices. Other sites are designed to sell something to those same patients while masquerading as sources of information. Patients may be unable to distinguish between the two types of sites, either because the latter sites have purposely been made to appear like the former or simply because patients may not be aware of the purposes of commercial medical web sites. To confuse matters further, there may be a mix of commercial information and authoritative information on the same site, either well or poorly distinguished.

Several certification programs purport to assist consumers in determining the quality of a medical web site (e.g. http://www.hiethics.org, http://www.truste.org, http://www.hon.ch). Such programs are only as good as the awareness they generate among consumers and the quality of their underlying requirements for certification. The effectiveness of certification programs has yet to be demonstrated. Requirements for certification vary widely, but there are several themes that appear consistently. These include:

- Clear mission and appropriate use statements
- Clear attribution and dating of medical material and claims
- Clear indication of advertising material
- Commitment to use of scientifically supported medical information
- Clear contact and complaint resolution information
- Appropriate security for PHI
- Ability to opt-out or opt-in to sharing of PHI
- Ability to amend PHI
- Agreement to bind business partners to policies of site
- Clear and timely notifications of any changes to policies
- Extensive disclosures, including
  - privacy practices
    - data sharing with third parties
    - aggregation and re-identification
    - use of tracking technology
  - financial
    - ownership
    - sponsors
    - third party revenues from data sharing

Many of the requirements are disclosure-dependent, meaning that the commitment is to inform the consumer about the site practices, rather than to eschew certain practices, such as preferentially including information from a financial sponsor, altogether. As a result, the burden of determining the quality of a site's information rests squarely on the shoulders of the consumer, even though the site bears a logo of certification.

Empowering health care consumers through readily-available health information is a valuable use of the internet. The nature of the internet environment, however, raises the spectre of abuse of vulnerable patients. Reputable web sites with health information must be careful to inform users about the nature of the site's information, both through extensive disclosures and avoidance of deceptive marketing practices. Patients must be educated to approach commercial web sites with *caveat emptor* firmly in mind.

2.4 Evidence-based Medicine, Outcome Measures, and Practice Guidelines

Modern medicine is defined in part by its use of therapies demonstrated to work by scientific evidence. Outcome studies and practice guidelines purport to provide such evidence. Outcome studies data must therefore be accepted unless convincing reasons to discard them exist. Outcomes research, however, faces a number of practical and philosophical problems which raise important ethical questions about the proper use of their results.

For this discussion, I define outcomes research as the statistical examination of outcomes as a function of diagnostic or treatment strategies using large numbers of subjects. Outcomes research often utilizes multiple studies combined via meta-analysis. Results from outcome studies are usually descriptive and use the language of statistics and probability. Practice guidelines typically combine outcomes research results with "expert" or "consensus" panels to produce prescriptive recommendations for clinical practice, often in the form of algorithms.

Outcomes research would not be possible without computational power. It is through data mining, knowledge discovery, and meta-analysis that results are obtained and all of these endeavors are impractical without computers. The internet is also critical in that it serves as the primary medium through which the results of outcomes studies and practice guidelines are distributed (c.f. http://www.cochrane.org, http://www.guideline.gov).

Outcomes studies have significant practical problems. Potential methodological flaws include the inherent difficulties of meta-analysis, inaccurate description of variables, and potentially inadequate sample sizes to detect small effects. Results may generalize poorly due to limited study populations (for example, males only or outdated therapeutic regimens) (Lagasse 1996, Gifford 1996). Most significantly, the design of outcomes research contains inherent value judgments which may not be apparent in the reporting of the results.

These value judgments are evident in various aspects of study design. For example, the outcome measure of "cost effectiveness" is value-driven because the answer to the question of what constitutes a reasonable expenditure of health care dollars is not scientific, but rather social and political. Also, complex systems do not lend themselves to simple measures, so a choice of measure must be made. This choice often involves the values of the researcher or the funding agency. Some important outcomes (for example pain, quality of life, reassurance, or justice) may be difficult, or even impossible, to measure and may therefore be inappropriately ignored

(Kerridge 1998). Fundamentally, the choice of outcome measure and its use is value-driven.

If the results of outcome studies are linked to resource allocation, then several additional problems occur. First, it is often impossible to compare studies of the outcomes of treatments for different conditions, but such comparisons must be made if treatments for different conditions are to compete for health care resources. Second, even if we use outcome studies only to compare treatments for a single condition, we must know which treatment is better and also by how much in order to understand how to manage expenditures (Shiell 1997). Third, the agenda of treatments studied may itself be driven by considerations of cost rather than health (Tanenbaum 1994). It is all too tempting for those who manage health care resources to assume that a lack of evidence for the efficacy of a treatment implies that the treatment has no efficacy and should therefore receive little support.

The purported ability of outcomes research to improve health care rests on the single assumption that data about probabilities of outcomes is valuable in making optimal decisions about health care delivery. That assumption is open to criticism on several fronts. First, as in the case of prognostic scoring systems, it can be very difficult to apply probabilistic, population-based data to individual patients, especially if the study population poorly matches the individual's background. Second, as was demonstrated by the poor performance of general diagnostic systems (Berner 1994), the technique of probability-based decision making itself is of limited value, due in part to the sensitivity of the Bayes theorem to population statistics. Finally, the use of probability-based rather than causal-based reasoning as the primary method of health care decision making is a significant departure from historical precedent. Probability-based reasoning implies an acceptance of a utilitarian philosophy which is not consistent with much of the moral philosophy of medicine. In addition, the use of probability-based reasoning implies an acceptance of induction (the expectation that one event will follow another from past experience of such sequences) as a reasoning model, as opposed to the more familiar use of causal models such as pathophysiology (Goodman 1996). Given that computer diagnostic systems do not demonstrate the efficacy of such reasoning, and given the inherently value-driven nature of outcome studies, it is clear that the results of outcome studies must be used with great caution.

2.5 Data Mining

The use of computers in health care has engendered an explosion of the quantity of electronically encoded data. A central theme of medical informatics is the use of this data for knowledge discovery, a process

commonly termed “data mining.” The purpose of data mining is to identify significant data patterns that would otherwise go undetected. Data mining is used in outcomes research, epidemiology, drug and genome discovery, biomedical literature searching, and many other areas. Data mining can also be used to detect unusual data patterns which might be indicative of disease outbreaks or fraudulent activities.

One of the great promises of the EHR is that clinical data will become available for data mining. As increasing amounts of health data become computerized, it is easy to imagine that data mining of EHRs will become the primary form of clinical research. Electronic records containing PHI need proper protections. The HIPAA regulations incorporate procedures for such research, recognizing that a breach of privacy is a form of harm.

In addition to regulating research on databases containing PHI, HIPAA also provides a mechanism to de-identify data. Once de-identified, data is free from regulation under HIPAA. Two methods of de-identification are allowed. One requires a statistical determination of the level of de-identification necessary to make re-identification unlikely. The other prescribes the removal of a specific set of identifiers (the “safe harbor” method) (http://privacyruleandresearch.nih.gov/research_repositories.asp). The safe harbor method is much simpler and is likely to be the method of choice for most situations. However, the safe harbor method removes a great deal of information that might be critical for answering relevant questions. For example, date elements must be removed except for the year, making determination of age to the accuracy necessary in the pediatric population essentially impossible. There is very little known about how useful or useless data de-identified by the safe harbor method will be.

The goal of de-identification is to make it statistically unlikely that the PHI of an individual patient can be reconstructed from a de-identified data set. Whether the safe harbor method accomplishes that goal has not been verified. There is evidence that some information that can be included in de-identified data may in fact be unique to a particular patient. For example, the ICD-9 diagnosis code, especially when combined with other data such as medications, may map closely to patient identifiers (Clause 2004). It is the clear duty of a researcher using de-identified data to avoid re-identification, but once such data sets become public it will be impossible to limit re-identification activities. Worse, because so much personal information is publicly available, it may be possible to use external sources of data in combination with de-identified medical data to construct fairly complete PHI information that could be used in ethically inappropriate ways (denial of health care, insurance, employment, etc.).

When selecting methods for rule-creation from data mining results, it is important to know what the rules will be used for. For example, we have

already seen that using Bayesian rules are not particularly effective for general diagnosis. A system for general diagnosis would be much more valuable if it accurately detected rare events. This is not because of anything in the nature of medical diagnosis itself, but because in a practical sense computers are not needed to diagnose common illnesses. It is important that the strengths of a particular data mining technique match the intended use of the results.

Similarly, it is important that database design take into account the range of possible queries that might be made of it. At the most basic level, one cannot examine data that is not included in the database at all. On a more complex level, it may well be possible to detect patterns in data that cannot then be adequately explained using that data. For example, some outcome measures such as lung function of patients with cystic fibrosis could be constructed and calculated for a number of different institutions. Given the results, one would then naturally ask what the best institution was doing right and what the worst institution was doing wrong. The answers to those questions could well be impossible to obtain from the original data set. Without those answers, the information about institutional outcomes may be useless, and perhaps damaging (Donaldson 1994). The problem of identifying the information to be measured and recorded is as old as epidemiology itself and has been made more acute by the computational power available today.

A particular concern regarding data mining arises when those results identify new patterns in population subgroups. This can happen when doing population-based or genomic research. It can even occur in research on de-identified data sets. Invidious discrimination requires differentiation between groups. We have experienced the ongoing evil of discrimination along the familiar lines of race, sex, age, and others. It is therefore of concern if data mining creates new subgroups that could then be the target of discrimination. It is easy to imagine discrimination along genetic lines (slow vs. fast drug metabolizers, for example). But any subgroup could be affected. Research where there is the possibility of new subgroup identification should be carried out with great caution, carefully weighing the potential medical benefits against the risks of harm from discrimination.

Data mining will undoubtedly provide important information for epidemiology, clinical decision support, and the practice of evidence-based medicine. It is important to realize, however, that there are ethical and social concerns about the use of the results. As with any health informatics technique, data mining must be used with a clear understanding of, and to further, the underlying goals of health care.

REFERENCES

Anderson, J.G. and Ayden, C.E. (1994). "Evaluating Medical Information Systems: Social Contexts and Ethical Challenges, in K.W. Goodman, Ed., *Ethics, Computing and Medicine: Informatics and the Transformation of Health Care,* Cambridge: Cambridge University Press, 57-74.

Ash, J.S., Stavri, P. Z., and Kuperman, G. J. (2003). "A Consensus Statement on Considerations for a Successful CPOE Implementation," *J Am Med Inform Assoc.*, 10:229-234.

Beauchamp, T.L., and Childress, J.F. (1994). *Principles of Biomedical Ethics,* 4th Ed., New York: Oxford University Press.

Bell, C.M., and Chapman, R.H., et al., "An Off-the-shelf Help List: A Comprehensive Catalog of Preference Scores from Published Cost-utility Analyses," *Med Decis Making,* 21(4): 288-294.

Berner, E.S., and Webster, G.D., Et Al. (1994). "Performance of Four Computer-based Diagnostic Systems," *N. Engl J Med.*, 330:1792-1796.

Clause, S.L., and Triller, D.M., Et Al. (2004). "Conforming to HIPAA Regulations and Compilation of Research Data," *Am J Health Syst Pharm.*, 61:1025-1031.

Delbanco, T., and Sands, D.Z., "Electrons in Flight—E-mail Between Doctors and Patients," *N Engl J Med.*, 350:1705-1707.

Detsky, A.S., and Naglie, G., Et Al. (1997). "Primer on Medical Decision Analysis: Part 1-getting Started," *Med Decis Making,* 17(2):123-125.

Donaldson, M.S., Lohr, K.N., and Bulger, R.J. (1994). "Health Data in the Information Age: Use, Disclosure, and Privacy—part 1," *JAMA,* 271(17):1308.

Fletcher, R.H., Fletcher, S.W., and Wagner, E.H. (1996). *Clinical Epidemiology: The Essentials*, 3rd Ed., Baltimore: Williams & Wilkins.

Gifford, F. (1996). "Outcomes Research and Practice Guidelines: Upstream Issues For Downstream Users," *Hastings Cent Rep.*, 26(2):38-44.

Goodman, K.W. (1996). "Outcomes, Confidentiality, and Appropriate Use," *Crit Care Clin.*, 12 (1):109-122.

Goodman, K.W., and Miller, R.A. (2001). "Ethics and Health Informatics: Users, Standards and Outcomes," in Shortliffe, E.H, and Perreault, L.E., Eds., *Medical Informatics: Computer Applications in Health Care and Biomedicine*, New York: Springer, 257-281.

Institute of Medicine (2000). *To Err Is Human*, Washington, D.C., National Academy Press.

Kane, B., and Sands, D.Z. (1998). "Guidelines for the Clinical Use of Electronic Mail with Patients," *J Am Med Inform. Assoc.*, 5:104-111.

Knaus W.A., and Wagner D.P., Et Al. (1991). "The APACHE III Prognostic System. Risk Prediction of Hospital Mortality for Critically Ill Hospitalized Adults," *Chest,* 100(6):1619-1636.

Knaus, W.A. (1993). "Ethics Implications of Risk Stratification in the Acute Care Setting," *Camb Q Healthc Ethics*, 2:193-196.

Kerridge, I., Lowe, M., and Henry, D. (1998). " Ethics and Evidence Based Medicine," *BMJ,* 316:1151-1153.

Lagasse, R.S. (1996). "Monitoring and Analysis of Outcome Studies," *Int Anesthesiol Clin.*, 34(3):263-277.

Mazoué, J.G. (1990). "Diagnosis Without Doctors," *J Med Philos.*, 15:559-579.

Miller, R.A. (1990). "Why the Standard View Is Standard: People, Not Machines, Understand Patients' Problems," *J. Med Philos.*, 15:581-591.

Moor, J.H., 1979). "Are There Decisions Computers Should Never Make?" *Nat Syst*, 1:217-229.

Patrick, D.L., and Starks, H.E., Et Al. (1994). "Measuring Preferences for Health States Worse Than Death," *Med Decis Making*, 14:9-18.

Sasse, K. (1993). "Prognostic Scoring Systems: Facing Difficult Decisions with Objective Data," *Camb Q Healthc Ethics*, 2:185-91.

Shiell, A. (1997). "Health Outcomes Are about Choices and Values: An Economic Perspective on the Health Outcomes Movement," *Health Policy*, 39:5-15.

Tanenbaum, S.J. (1994). "Knowing and Acting in Medical Practice: The Epistemological Politics of Outcome Research," *J Health Polit Policy Law*, 19(1):27-44.

Thomasma, D.C. (1998). "Applying General Medical Knowledge to Individuals: A Philosophical Analysis," *Theor Med.*, 9: 187-200.

Yasnoff, W.A., and Humphreys, B.L., Et Al. (2004). "A Consensus Action Agenda for Achieving the National Health Information Infrastructure," *J Am Med Inform Assoc*., 11:332-338.

SUGGESTED READINGS

Goodman, K.W. (ed.), *Ethics, Computing and Medicine: Informatics and the Transformation of Health Care*, Cambridge University Press, Cambridge.
This is the central text for the area of ethics in medical/health informatics.

Howell, J.H., and Sale, W.F. (eds.), *Life Choices: A Hastings Center Introduction to Bioethics*, 2nd ed., Georgetown University Press, Washington DC.
An introductory text to bioethics in general.

Institute of Medicine, *To Err is Human*, National Academy Press, Washington, D.C.
This report sets the agenda for the reduction of medical errors.

ONLINE RESOURCES

Certification organizations for medical web sites
http://www.hiethics.org
http://www.truste.org
http://www.hon.ch

Examples of evidence-based medicine and guideline distribution sites
http://www.cochrane.org
http://www.guideline.gov

Examples of web-based diagnostic systems
http://www.lcs.mgh.harvard.edu/dxplain.htm
http://www.isabel.org.uk

HIPAA Privacy Regulations overview and information on research
http://www.hhs.gov/ocr/hipaa/
http://privacyruleandresearch.nih.gov/research_repositories.asp

QUESTIONS FOR DISCUSSION

1. The mission statement for an effort by a medical organization to collect data from multiple local sources reads: "The mission is to transform what is now a disconnected set of data into a form that is complete for any given patient, no matter where they are seen. It should be available to different groups or health care professionals for different reasons. This will involve an assessment of the current data in terms of its location, accuracy, and accessibility to different parties, identification of these parties with an understanding of the kinds of data they may need, and then a matching of these needs to the restructuring of the database itself." Discuss the technical and political challenges of this mission. How does HIPAA impact this mission? Where is the primary focus of this project, on the patients or providers?

2. You are the CIO of a large hospital system. An eight-year old girl who is a patient at one of your hospitals is in need of a liver transplant. A suitable donor has not yet been found. Many people die while awaiting liver transplants because of a shortage of organs. The girl's parents wish to set up a web site describing their daughter's illness and prognosis, particularly the critical need for a liver, in the hope that this will help them find a donor. They ask you to make this part of your hospital system's public web site because they think that that location will give it more legitimacy and attract more internet traffic. Do you allow this? Why or why not? Discuss in terms of the HIPAA regulations and the principles of Beauchamp and Childress. Include any other factors or reasoning you consider important.

3. Describe three barriers to EHR adoption. What is meant by "minimum necessary" in the HIPAA regulations? Describe what technical and policy measures would be necessary in order to have third-party centralized EHRs. Discuss the advantages and disadvantages of centralized EHRs.

4. You wish to code into an EHR a clinical algorithm for the treatment of hyperbilirubinemia in the newborn (http://aappolicy.aappublications.org/cgi/content/full/pediatrics;114/1/297). "This algorithm depends on the age of the newborn in hours and the level of total serum bilirubin (a blood test). "Turn-around time for this result ranges from 30 minutes to two hours. Describe technical barriers to this project. How would you test the code? What mechanisms do you need to

maintain the code if the algorithm changes? How would you detect and handle missing or inaccurate data?

5. What are the characteristics necessary for a useful general diagnostic support system? What data mining techniques are available besides Bayesian algorithms? Are any of them suitable for a general diagnostic support system?

6. How would you measure the loss of data due to safe harbor de-identification? How would you determine ease of re-identification?