Is Healthcare Information Technology Based on Evidence?

R. Koppel

University of Pennsylvania, Pennsylvania, PA, USA

Summary

Is healthcare information technology (HIT) based on evidence of efficacy? Are the trillions of dollars already devoted and in the pipeline for HIT implementations based on systematic evaluations? If evaluated, would those evaluations focus on patient safety, return on investment, clinical efficiency, improved clinician satisfaction, and/or workflow integration? Do we have reliable evidence of usable interfaces, of successful implementations, of data standards allowing interoperability, of continuous improvement, of responsiveness to clinician feedback?

While measurement of HIT's efficacy is extraordinarily difficult—complicated by a myriad of other factors involved in providing healthcare and in organizational dynamics—it is not impossible. But is such evidence required before most implementations? Any implementation? Or are the goals of patient safety and efficiency so self-evident, profoundly desired, and laudable that HIT's beneficence is accepted without rigorous data? Note that lack of systematic evidence does not mean HIT is ineffective. HIT may provide untold benefits even if there is no hard proof of those benefits.

We find that HIT is seldom objectively measured, and that evidence of its efficacy is at best spotty, and often influenced by self-promotion. Most measures, especially those associated with cost-benefit analyses, are aspirational or hubris transubstantiated into numbers.

Keywords

Healthcare IT, EHR, EMR, evidence-based, ROI, CPOE, HIT

Yearb Med Inform 2013;7-12

The question posed for this keynote chapter is essential and intriguing, but usually avoided, ignored, or assumed. It asks: Is healthcare information technology (HIT) evidence-based? To what extent is HIT designed, built /or implemented on a foundation of evidence derived from systematic evaluation of what works and what fails to work, of where HIT may be more or less effective, and of the necessary conditions for HIT to work at all?

Some introductory parameters: By HIT we refer in this paper to: Computerized provider order entry systems (CPOE), electronic health records (also electronic medical records) (EHRs/EMRs), electronic medication administration records (eMARs), bar-coded medication administration system (BCMA), and electronic prescribing (eRx). We do not include systems for storing and sending radiology and other graphic documents (picture archiving and communication systems—PACS). Even though these systems are useful and venerable, they are not included because they involve less of the kind of active integration and two-way interactions of the other systems.

We also acknowledge the vast heterogeneity of the types of HIT systems. There are at least a dozen or two commercial enterprise systems (i.e., for hospitals or healthcare systems), a few score of home-built enterprise systems, and over a thousand software programs for clinics and clinicians.

Structure: The chapter has four parts: First, we ask what types of evidence would be needed for HIT to be “evidence-based.” Second, we briefly note the absence of evidence and provide a list of reasons and rationales frequently offered for or assumed when purchasing and implementing systems. We also note that a general *absence of hard evidence* of efficacy does not mean HIT is without support. There are many studies by advocates and others attesting to its value, and some devotees insist HIT's benefits are self-evident. One major

problem is that HIT's effectiveness is wickedly hard to measure. Indeed, the third section enumerates the many difficulties in trying to measure HIT's impact and costs. Last, we offer ten recommendations to encourage systematic evaluation of HIT as a way of improving HIT, medical care, and patient safety.

1 To Answer the Question We Must Unpack It: Evidence of What?

- Is the *design* or *construction* of HIT based on evidence of what is most efficient and effective in practice?
  - Is navigation within and among software screens evaluated and modified to be the least burdensome and most intuitive? Are clicking and scrolling reduced to the minimum? Do clinicians know “where they are” and how to get to where they are going, or at least, how to get back?
  - Is usability of the systems—as measured via human factors research—tested and progressively improved?
  - Are displays of information in graphic and textual form carefully assessed? Are they tested on a wide range of users in many real-life settings?
  - Are observations of clinicians' use of HIT systematically analyzed as to how information can be amassed and analyzed with minimum distractions and minimum unnecessary cognitive burdens?
- Are differing systems evaluated when *implemented in situ* across similar settings? Or are implementations of the same system evaluated in dozens or hundreds of various medical settings, also, by definition, *in situ*? If such comparisons are

available, are they routinely provided to potential purchasers?

- Are we assessing various methods of implementation, incorporating facility design, number of clinicians, number of intersecting offices, expertise of the IT team, etc?
- At what point is an implementation fairly evaluated? Six months after the “go live” date? A year after the “go live” date? After the third upgrade? After each patch or version is installed?
- In light of HIT vendors’ reasonable statements that each implementation is an essential element of that system’s safety and effectiveness, is it sensible to evaluate HIT systems independent of their implementations? And given that concern, is “certification” of an “isolated” HIT system a futile exercise?
- Are evaluations based on random controlled trials (RCT) or double blind RCTs? Is RCT possible with “evidence” from HIT use?
- Are clinical decision support (CDS), order sets, disease protocols, or dosage alerts built on the latest medical knowledge? Is CDS examined to ensure it achieves:
  - Reduced alert fatigue by careful titration of alerts to only the most essential?
  - Presentation of alerts in ways that align with clinical workflow and thought-flow?
  - Presentation of information in alerts that is relevant to the user at the point-of-care and the time of decision?
  - Easy access to additional information on how the alerts are determined?
- Are there evaluations of the connections among the many IT systems that are linked to the CPOE or EHR systems? Not only is each connection a vulnerability to the overall system, but because many of the other IT systems interact with each other separately from the EHR, there are factorial interactions and vulnerabilities. The core HIT systems are embedded in a network of other systems, each of which potentially affects many other elements of the network. Is the network evaluated synergistically? Is there evidence of one of HIT’s most basic promises: Interoperability, or proof that HIT is capable of interoperability—sharing information in usable formats with interpretable data? Do HIT systems permit sharing data across a region, across town,

across the hallway, or across the room?

- Is there a reliable calculation of the return-on-investment (ROI) in HIT? Are there savings in time, in staff, in avoided errors, and in fewer repeated tests and laboratory orders? Is the ROI the basis of the decision to use HIT or is it a *post hoc* justification? Implicit in this question is that one knows the cost of HIT and its implementation. As we shall see, this is a wickedly difficult figure to determine.
- Is there evidence of improved patient safety from HIT? Improving patient safety is one of the central claims of HIT; are there consistent and systematic data to support this claim?
  - In the evaluations, are there statistical controls for:
    - Training of clinicians—a critical issue, because of the role of teaching hospitals, “voluntary” attendings, clinicians who practice in different hospitals or offices with differing EHRs and CPOE systems?
    - Patient loads and acuity?
    - History of technology use in each institution and by each clinician user?

2 Alas, Usually Missing:

Answers to all of these questions are usually missing. Why they are missing is discussed below, in section III. But understand, “*missing*” or “*not provided*” data does not mean that HIT programs and implementations are without value or that purchase decisions are uninformed or wrong. “*Missing*” also does not mean that compelling evidence is not forthcoming. What “missing” means, rather, is that much of HIT’s development, selection, and implementations are based on:

- Legacies of earlier systems
- What others tell us may have worked elsewhere
- What we think makes sense
- What we can afford
- What vendors recommend
- What many—usually differing—clinicians and leaders within our institutions have compromised upon
- What customers have told vendors
- What we’ve read
- Subsidies and incentives by governments
- Regulations
- Decisions of larger institutions with

which we are affiliated

- How much time and energy we have to customize the systems
- The legal and technical limits of customization
- Other IT systems with which we must currently connect, and
- On our best judgments as problems emerge.

Many of the development, selection, and implementation decisions are made with great thought and consideration, with earnest debate, with careful reading of the available literature, and with the advice of consultants. But few are evidenced-based in the way we conceptualize serious evaluation or the scientific method.

Also, the lack of evidence in building or implementing HIT does not mean these systems are ineffective. Although there is little systematic research, many HIT systems appear to work for several functions: EHRs can enable physicians and patients to maintain a complete, and omnipresent medical record. CPOE systems allow physicians and other health care professionals to enter medication orders directly into a computer system, avoiding handwriting or transcription errors, and speeding orders to pharmacies and laboratories. CDS provides information to physicians or nurses when they order or administer medications, for example, warning that the proposed dose exceeds the normal range or that the patient is listed as being allergic to a proposed drug. These systems help physicians and nurses to order and administer medications in a timely fashion. Many of these technologies, also, may reduce redundant tests and procedures.

So, lack of scientific evidence [1-7] and the mountain of conflicting evidence [1, 8-15] for the many HIT features and elements noted above do not negate HIT’s benefits to patients and healthcare providers. Humans built bridges before trigonometry and calculus. We covered wounds before the germ theory of disease. Moreover, just as much of all-but-modern medical practices were based on theories we now hold as absurd (e.g., humors, blood surfeit, demonic possession) and were “supported” by dubious evidence of efficacy, physicians nonetheless often achieved successful outcomes. We shall continue to invest trillions in HIT systems because they

are usually better than paper, because we so dread medical errors, and because the complexity of medical care is so daunting.

3 Why Do We not yet Have Evidenced-based HIT?

There are many reasons we lack consistent and valid evidence of HIT's efficacy. Measurement of HIT's efficacy is hard; its measurement *in situ* is even harder. Consider:

The messy reality of medical practice: In general, the real world applicability of evidence-based medicine (EBM) is frequently overstated. Our research model is the clinical trial, where studies are conducted with carefully selected samples of patients to observe the effects of the medicine or treatment without additional interference from other conditions. Unfortunately, the clinical trial model differs from actual medical practice because hospitals and doctors' waiting rooms are full of elderly patients suffering from several co-morbidities and taking about a dozen or more medications, (some unknown to us). It is often a great leap to apply findings from a study under "ideal conditions" to the fragile patient. So physicians must then balance the "scientific findings" with the several vulnerabilities of real patients. Clinicians are obliged to constantly deal with these messy tradeoffs, and the utility of evidence-based findings is mitigated by the complex challenges of the sick patients, multiple medications taken, and massive unknowns. This mix of research with the messy reality of medical and hospital practice means that evidence, even if available, is often not fully applicable. *When applied to HIT,* very similar factors come into play. No two hospitals or practices are the same; every CPOE system must be installed into an existing network of other IT systems, workflows, clinician experience, etc.

Evidence for HIT's Effects: Added to these limitations of EBM just discussed, are additional challenges of measuring HIT's *impact* independent of the many other factors that co-vary with it. The number and complexity of variables involved in medical settings and in medical care are staggering and each HIT implementation is unique. (As the saying goes, "you've seen one EHR implementation, you've seen one EHR implantation.") HIT

training among interns, senior physicians, nurses, and the myriad of other clinicians varies dramatically. There are constant flows of professionals and students; medical practice is also a teaching practicum. The number and criticality of interactions with other systems (within an institution and across the globe) is usually unknown and often far more nuanced than anyone suspects *de novo*. A list of factors that affect HIT's influence in any institution, while not infinite, is beyond most users' comprehension [10-13]. Moreover, this list does not begin to include the promised outcome measures (e.g., reduced errors, efficient billing, ward acuity measures, nursing personnel needs, cost and time savings, inventory control), which are certainly as complex and varied as the context, patient types, and users.

Random Controlled Trials: The gold standard of research is the random controlled trial method (RCT). Even better is the double-blind study, where neither subjects nor researchers know which is the test intervention and which is the placebo or standard intervention. Such research designs are almost impossible to imagine for HIT. Double-blind, or even single-blind research designs would require all of the participants' active involvement in not examining the type of system they would be obliged to use for a year or two. Also, can anyone believe that a hospital and its clinicians can spend 400 million Euros and 4 years implementing an EHR about which the clinicians are not fully informed? Then one also must assume the implementation is followed by a few years of evaluation of the EHR-in-use while also keeping everyone in the dark? More basic, for the reasons enumerated in the two paragraphs directly above (on obtaining similar settings, linkages, staff, IT, etc), the research design requirements for RCT are extremely difficult to enact in the real world.

Acts of faith: We assume more HIT generates better and safer healthcare. That assumption is assisted by massive efforts by vendors, vendor trade associations, lobbyists, legislators, policy "experts", business groups, governmental agencies charged with encouraging HIT use, insurance carriers hoping HIT will reduce costs, academics, and many healthcare providers [15-22]. The lobbying and marketing budget of vendors is many millions of dollars [17, 18, 20]. As with any faith, expressing doubts about HIT's benefits is viewed as heretical [22,23].

One very consequential example of the assumptions and unexamined faith that can distort perceptions and research is the seminal article by Jha et al [24] that found US hospitals required subsidies for HIT purchases, incentives to continue HIT use, certification of HIT systems, better IT departments, and fewer technophobic physicians. A review of the survey instrument on which that New England Journal of Medicine paper was written, however, reveals they asked no questions about usability, data standards, interoperability, friendly user interfaces, clunky software, irrational choices in menus, lost laboratory results, and non-responsive vendors. The only options the hospital respondents were offered were about needs for subsidies, incentives, certifications, better IT departments, and fewer technophobic physicians. While not intentionally biased, it nevertheless assumes HIT is beneficial, and any impediments must be found among users and their institutions.

Evidence for CDS: While few doubt the eventual value of CDS, most CDS alerts are generally ignored or overridden because they are viewed as useless or just annoying. Override rates are as high as 97% [25, 26] - a reality that makes the evaluation of CDS inherently problematic. Does one count only actions based on the small percent of alerts that generate change, and ignore the vast bulk of alerts that generate only annoyance or rage? And even then, does one count only those actions that are associated with beneficial change, and discount those CDS-inspired actions that result in new or additional errors (a non-trivial occurrence). Any evaluation, moreover, should reflect the distractions, interruptions in work flow, cognitive burden, and errors associated with the many overridden/ignored alerts and recommendations. CDS alerts are in fact, the most hated feature of HIT. Measuring their efficacy is therefore challenging and efforts at systematic assessments are disappointing. In one major study, Metzger and colleagues [27] found that CDS detected only 53 percent of all medication orders that would have resulted in fatalities and caught from 10 to 82 percent of orders that would have caused serious adverse events. Drugs prescribed for a wrong diagnosis were caught only 15 percent of the time (that is, in ideal cases in which the computer already had the patient's record and could "know" that the drug was inappropriate), and drugs

that were wrong for a patient of a given age were intercepted only 14.1 percent of the time. Nevertheless, CDS remains on everyone's list as an essential feature of HIT's benefits.

Data Standards: Without unified data standards and data formats, creation of evidence of HIT's efficacy across differing facilities and HIT systems is almost impossible. Proprietary interests, legacy systems, and existing capital investments make agreement on standards difficult. But without unified data standards and data formats, we create a tower of Babel within each medical facility, and we severely attenuate the utility of HIT. Perhaps worse, without unified data standards we cannot share information across systems; we fail to achieve real interoperability. These towers of Babel become isolated from each other; a noisy but deaf city.

Interoperability, Evidence Not Relevant: Interoperability, an underpinning of useful HIT, remains elusive and is often sidestepped by HIT vendors because those vendors benefit from proprietary systems that communicate primarily within their own brand—encouraging the purchase of suites of products. Regulatory capture is common, and vendors insist regulation will retard innovation, a refrain reinforced with lobbying, sycophant researchers, and true believers. As Professor Kalra wrote in a recent keynote chapter of this yearbook series [28]:

“Achieving...interoperability across the breadth of health and healthcare is the challenge that needs urgently to be addressed. It is hard to believe that the origins of the EHR archetype are now almost 15 years old. This was conceived as the most basic level of knowledge representation for the EHR: standardised clinical models to provide a common representation to which data from heterogeneous systems can speed up transfer of new medical knowledge”

Yet, HIT systems with no or minimal interoperability continue to be sold even though interoperability represents a foundation of HIT's value. One might argue evaluators cannot demonstrate interoperability's value because there is so little interoperability on which to collect evidence—a self-perpetuating condition.

Marketing Strategy: Another factor limiting evidence is marketing and market share capture strategies. Because these large HIT systems are so costly, and because they take

several years to implement, there is no room for buyer's remorse. Once an organization buys them, it is generally wed to the system for a decade. Organizations do not easily junk a system that cost one to five hundred million dollars in direct cost and four times that figure in personnel and other costs for implementation, linkages, and training. The vendors thus seek to capture market share as soon as possible, and are encouraged to rush HIT products to market before they are sufficiently tested. Aggressive marketing and subsidies also obscured or prevent objective discussion of HIT's merits and challenges. The vast funds involved, and the consequential career implications of those participating in HIT purchases enhance intimidation of critics and those who report problems with the technology [14]. The general faith in technology and the sincere desire to see HIT improve medical workflow encourages so many to define critics as technophobes, incompetents, and non-team players.

Usability: Until very recently, vendors generally rejected usability testing—and thus evidence thereof—as “subjective,” “immeasurable,” “already achieved,” “already measured,” and sufficiently “reflected in user feedback.” [29–32] Vendors also insist usability is so affected by local implementation decisions that original designs are only the canvas on which purchasers paint new user interfaces. Thus, they argue, usability is not subject to their control—a claim with some validity in that local decisions do significantly alter usability. But overall, it is a disingenuous argument, analogous to US automobile manufacturers in the 1960s and 70s insisting safety was entirely in the hands of drivers, with manufacturers resisting demands for better brakes, collapsible steering wheel columns, and even fewer sharp knobs in the cabin. Bad usability of HIT harms patients. Clinicians' ability to find needed test results, to view coherent listings of current medications, or to review “active” problems is a patient safety necessity and central to efficiency and effectiveness. Insisting usability is immeasurable, uncontrollable by vendors, et cetera helps vendors evade responsibility for usability as a key requirement, and certainly does not encourage evidenced-based usability testing within HIT. Vendors' positions on usability—which are beginning to change—ignored 50 years of usability research by human factors experts and industrial engineers.

Customization: Customization is a double-edged sword. On the one hand, each institution and each clinician is promised with (and flattered by) the ability to make significant modifications to the system(s). And indeed there are excellent reasons why customization is needed. On the other hand, customization is used as a marketing ploy, and creates massive problems. It makes upgrades and patches far more difficult than they need be, it attenuates systematic collection of evidence of HIT's efficacy, and it endangers patients and healthcare when clinicians must practice across differing HIT systems, each with different ways of viewing data, ordering medications, arranging problem lists, etc. Customization is thus often a perverse benefit: a workaround for integration of IT systems that should be but are not designed for interaction; perpetuating a laissez faire digital environment of autonomous silos when interoperability is the absolute requirement for better care and patient safety.

Return on Investment: To calculate return on investment (ROI), one needs to know the cost of the investment and the economic value of the “return.” With most of HIT, we have neither, or at least neither with reliable estimates. Recently (2013), the RAND corporation published a report that seriously questioned the economic savings from HIT they previously estimated. [14, 20] In addition, the new publication brought to light the fact that their previous RAND report of 2005, so effective for HIT promoters, was in fact supported by HIT vendors Cerner and GE. Vendors and system developers have long provided favorable analyses of ROI [15, 19, 20]. But these are generally self-justification, incomplete arithmetic, marketing, or acts of faith transubstantiated into numbers [2, 4] because: 1. One has little idea how much these systems cost in total; 2. There will be unknown additional needed personnel in IT, in medical, nursing, or pharmacy informatics; 3. Each hospital HIT upgrade requires about 40,000 person hours to implement and test [34]; Few if any organizations predict such costs; 4. As outlined above, we seldom know the outcomes of HIT's use [34]; 5. The cost of the systems is not disclosed and often obscured or modified via joint marketing agreements, reductions on future upgrades or add-ons, fees for demonstrations to other potential clients; 6. Healthcare providers suffer dramatic productivity losses when the HIT is

newly implemented; and many of those losses continue for years if not indefinitely. [35]

4 What Should Healthcare Institutions and Clinicians Do?

Clinicians and healthcare IT staff should not accept HIT vendors' assessments of their own HIT without independent investigation. Even though systematic evaluation of HIT is remarkably difficult, it is not impossible. A recent study by Duke, Li and Dexter [36] compared the efficacy of two types of CDS alerts using a randomized study of physicians. Although the findings showed no improvements with either method, the work illustrates rigorous evaluation methodology in a circumscribed but potentially valuable area of EHR functionality.

Needless to say, evaluation could be made far more manageable if there were transparent pricing and clear reporting of implementation and training costs. Insisting also on data standards would permit both interoperability and meaningful system comparisons. The vast funds and personnel involved and the patient safety consequences demand nothing less. Equally important, *improvement* of HIT will only be achieved if it is based on careful and unbiased evaluation of HIT design, implementation, training, and use. While acknowledging evaluations cannot reach the level of RCTs, they will still be far better than the current slate of marketing hype, *ad hoc* testimonies, and self-analyses. We must listen to the frustrations of clinicians and of local IT personnel, and then act constructively to address those problems. Denying those frustrations, or failure to examine HIT's flaws, is counterproductive and will condemn patients to unsafe care and condemn clinicians to unnecessary burdens and stress.

Recommendations

These recommendations are offered with reservations about their probability of implementation. That is, the forces that limit evidenced-based HIT, will also limit the ability to enact these recommendations. Nevertheless, we should at least try.

1. Establish clear metrics for HIT's core functions. This will require separate ex-

aminations for CPOE, EHRs, pharmacy IT, etc. The tests must include time-at-task as well as completion of the task. In the USA, the HIT certification process has been an open-book test with unlimited time making it remarkably dissimilar to clinical environments. Unlike testers, clinicians do not have an hour to enter each medication order.

2. If, as vendors insist, HIT is dramatically altered once implemented, then the testing must be conducted *both in vitro* (in laboratory settings) and *in situ* (in actual field conditions—a hospital, clinic or office).
  - 2.1 If we need interim test-bed settings (e.g., a “standard” set of hospital linkages to other IT systems, a “standard” workflow environment), than this should be established. At least a realistic test bed is a better guide than only a laboratory test.
3. Recognize the need for multi-method testing procedures: Discovering how HIT works in reality requires the full range of research techniques and data sources. Koppel et al's [37] study of CPOE employed surveys, observations, focus groups, shadowing of physicians and nurses, one-on-one interviews with many different kind of staff, expert interviews (with IT and hospital leaders), and shadowing of pharmacy personnel as they used the system. Koppel et al's study of medication barcoding administration used all of these methods plus analysis of almost half a million scans of patient IDs and medication IDs, vendor interviews, review of vendor specifications, and interviews with dozens of hospital and IT leaders from throughout the nation.[38]
4. Require data standards. Without unified data standards and data formats, achieving interoperability is nearly impossible. Without both, HIT's utility—and the ability to evaluate HIT's utility—is profoundly attenuated. Note that several organizations and groups are actively involved in providing data standards and semantic interoperability. The need for data standards does not mean there is a need for new standards. Some of those standard-setting organizations include the association publishing this chapter, The International Medical Informatics Association, along with others: HL7, the International Health Terminology Standards Development Organization, which owns and administers the rights for SNOWMED CT, etc.
5. Establish consistent usability tests for every major screen and function. These should include careful examinations of system navigation, way finding, and ability to determine where in the system one is.
6. Evaluate graphic presentations of data. HIT offers extraordinary abilities to convert numeric data to easily viewed formats. But confusing and poorly annotated graphic displays are worse than none at all.
7. Use the tests to help vendors improve their products and to help healthcare providers select the best products for their needs. Do not allow proprietary interests to influence the assessment process or the distribution of findings. However, vendors should have the ability to annotate and dispute any reports offered by clinicians and testing services.
8. Make these evaluation processes transparent.
9. Publish the findings.
10. Do not allow hidden contractual agreements, (e.g., joint sales agreements, fees for demonstrating a product, or fees for attesting to the excellence of a product) to distort colleagues' judgments. It's permissible to reimburse hospitals and clinicians to talk with potential customers, but those customers must know if money or goods are being provided as compensation [39].

In sum, we readily acknowledge the impact of HIT is often hard to measure, always nuanced, and profoundly complex. Unintended side effects (both good and bad) as well as intended effects are often discovered slowly, and only with vigilance, thoughtful examination, and openness to surprises. Policymakers, health care executives, and clinicians must gain a balanced understanding of the powers, problems, and implications of the technology if they are to assess evidence of its efficacy. But as daunting a challenge as that is, there are no viable alternatives. The often oversold technology—a belief assiduously nurtured by an HIT industry with much to gain—does not negate the significant benefits

HIT offers. And as HIT evolves, it will be of even greater value to patients, clinicians, and budgets. Ironically, the extravagant hype, the rush to market, and the reluctance to measure its problems and effects may be more of a danger to its continued growth than are its multifaceted failures. The continuing, and well-orchestrated, chorus of promises may deafen the industry's ability to hear its customers and to recognize their needs. Patient safety, which has so much to benefit from good HIT, will suffer until the healthcare experts and the HIT industry are willing to carefully evaluate the evidence of these systems, and then use that evidence to improve them.

References

- Karsh B-T, Matthew B Weinger MB, Abbott PA, Wears RL et al. Health Information Technology: Fallacies and Sober Realities. *J Am Med Inform Assoc* 2010;17:617-23.
- Chaudhry B Wang J, Wu S, Maglione M, Mojica W, Roth E, et al. A Systematic Review: Impact of Health Information Technology on Quality, Efficiency, and Costs of Medical Care. *Ann Intern Med* 2006;144(10):742-52.
- Koppel R, Localio AR, Cohen A, Strom BL. Neither panacea nor black box: responding to three Journal of Biomedical Informatics papers on computerized physician order entry systems. *J Biomed Inform* Aug 2005;38(4):267-9.
- Garg A, Adhikari N, McDonald H, Rosas-Arellano MP, Devereaux PJ, Beyene J, et al. Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: a systematic review. *JAMA* 2005;293(10):1223-38.
- Wears R, Berg M. Computer technology and clinical work: still waiting for Godot. *JAMA* 2005;293(10):1261-3.
- Nebeker JR, Hoffman JM, Weir CR, Bennett CL, Hurdle JF. High Rates of Adverse Drug Events in a Highly Computerized Hospital. *Arch Intern Med* 2005;165:1111-6.
- Shulman R, Singer M, Goldstone J, Bellington G. Medication Errors: A Prospective Cohort Study of Hand-Written and Computerized Physician Order Entry in the Intensive Care Unit. *Crit Care* 2005;9(5):R516-R521.
- Ash JS, Sittig DF, Poon EG, Guappone K, Campbell E, Dykstra RH. The Extent and Importance of Unintended Consequences of Computerized Physician Order Entry. *J Am Med Inform Assoc* 2007;14:415-23.
- Campbell EM, Sittig DF, Ash JS, Guappone KP, Dykstra RH. Types of Unintended Consequences Related to Computerized Provider Order Entry. *J Am Med Inform Assoc* 2006;13(5):547-56.
- Aarts J, Ash J, Berg M. Extending the Understanding of Computerized Physician Order Entry: Implications for Professional Collaboration, Workflow and Quality of Care. *Int J Med Inf* 2007;76:S4-13.
- Han Y, Carcillo J, Venkataraman S, Clark R, Watson S, Nguyen T, et al. Unexpected increased mortality after implementation of a commercially sold computerized physician order entry system. *Pediatrics* 2005;116:1506-12.
- Harrison M, Koppel R, Bar-Lev S. Unintended consequences of information technologies in health care: an interactive socio-technical analysis. *J Am Med Inform Assoc* 2007;14:542-9.
- Spencer J, Koppel R, Ridgely MS. AHRQ Guide to reducing unintended consequences of HIT Implementation and Use; 2011 www.ucguide.org (accessed May 13, 2011).
- Kellermann AL, Jones SS. What It Will Take To Achieve The As-Yet-Unfulfilled Promises Of Health Information Technology. *Health Aff* 2013;32(1):63-8.
- Madara JL. Open Letter to Office of the National Coordinator of HIT Chair, Dr. Farzad Mostashari on the American Medical Association's position on "Meaningful Use" regulations. <http://www.ama-assn.org/resources/doc/washington/stage-3-meaningful-use-electronic-health-records-comment-letter-14jan2013.pdf> (accessed January 21, 2013)
- Silverstein S. Contemporary Issues in Medical Informatics: Common Examples of Healthcare Information Technology Difficulties. <http://www.ischool.drexel.edu/faculty/ssilverstein/failurecases> (accessed May 13, 2011).
- Koppel R, Gordon S. First Do Less Harm: Confronting The Inconvenient Problems of Patient Safety. Ithaca, NY: Cornell University Press; 2012.
- Koppel R, Majumdar SR, Soumerai SB. Electronic Health Records and Quality of Diabetes Care. Editor's Correspondence. *N Engl J Med* 2011;365:2338-9.
- Berger R, Kichak J. Computerized Physician Order Entry: Helpful or Harmful? *J Am Med Inform Assoc* 2004;11:100-3.
- Hillestad R, Bigelow J, Bower A, Girosi F, Meili R, Scoville R, et al. Can Electronic Medical Record Systems Transform Health Care? Potential Health Benefits, Savings, and Costs. *Health Aff* 2005;24(5):1103-17.
- Healthcare Information Management and Management Systems Society. <http://www.himss.org/ASP/aboutHimssHome.asp> (Accessed 12/25/12)
- EHRA of HIMSS (Electronic Healthcare Record Association of HIMSS) <http://www.himsssehra.org/ASP/index.asp> (Accessed 12/25/12)
- McCormick D, Bor D, Woolhandler S, Himmelstein D. The Effect of Physicians' Electronic Assess to Tests: A response to Farzad Mostashari. *Health Affairs Blog*. March 12, 2012: <http://healthaffairs.org/blog/author/dmdbswdh/> (Accessed 12/25/12)
- Jha AK, Desroches CM, Campbell EG, et al. "Use of Electronic Health Records in U.S. Hospitals," *N Engl J Med* 2009;360:1628-38.
- Ridge MS, Greenberg MD. Too many alerts, too much liability: sorting through the malpractice implications of drug-drug interactions clinical decision support. *Journal of Health Law and Policy* 2012; 15(2):257-96.
- Koppel R. The marginal utility of margin guidance: commentary on Ridgely and Greenberg. *Journal of Health Law and Policy* 2012;15(2):311-8.
- Metzger J, Welebob E, Bates DW, Lipsitz S, Classen DC. Mixed results in the safety performance of computerized physician order entry. *Health Aff (Millwood)* 2010;29(4):655-63.
- Kalra D. "Health Informatics 3.0" Yearb Med Inform 2011;8-14.
- Lawler EK, Alan H, Pavlovic-Veselinovic S. Cognitive ergonomics, socio-technical systems, and the impact of healthcare information technologies. *Int J Ind Ergon* 2011;41:336-44.
- Committee on Patient Safety and Health Information Technology. Institute of Medicine of the National Academies Health IT and Patient Safety: Building Safer Systems for Better Care. Washington DC: National Academies Press; 2011.
- Kannry J, Kushniruk A, Koppel R. Meaningful usability: Health information for the rest of us. In: Ong K, editor. *Medical Informatics: An Executive Primer* (2nd Ed). Chicago: Healthcare Information and Management Systems Society (HIMSS); 2011.
- Koppel R, Kreda DA. Healthcare IT usability and suitability for clinical needs: challenges of design, workflow, and contractual relations. *Stud Health Technol Inform* 2010;157:7-14.
- Middleton B, Bloomrosen M, Dente MA, Hashmat B, Koppel R, Overhage JM. Enhancing patient safety and quality of care by improving the usability of electronic health record systems: recommendations from AMIA. *J Am Med Inform Assoc* 2013 Jun;20(e1):e2-8.
- Spencer J, Koppel R, Ridgely S. AHRQ Guide to reducing unintended consequences of HIT Implementation and Use, 2011. www.ucguide.org (Accessed January 21, 2013)
- Terry K. Doctors' 10 biggest mistakes when using EHRs. *WWW.Medscape.com*, May 1, 2013 <http://www.medscape.com/viewarticle/803188> (accessed May 11, 2013)
- Duke JD, Li X, Dexter P. Adherence to drug-drug interaction alerts in high-risk patients: a trial of context-enhanced alerting. *J Am Med Inform Assoc* 2013 20:494-8.
- Koppel R, Metlay JP, Cohen A, Abaluck B, Localio AR, Kimmel SE, et al. Role of computerized physician order entry systems in facilitating medication errors. *JAMA* 2005;293(10):1197-203.
- Koppel R, Wetterneck T, Telles JL, Karsh BT. Workarounds To Barcode Medication Administration Systems: Their Occurrences, Causes, And Threats To Patient Safety. *J Am Med Inform Assoc* 2008 July/August;14:408-23.
- Goodman KW, Berner ES, Dente MA, Kaplan B, Koppel R, Rucker D, et al. Challenges in ethics, safety, best practices and oversight regarding HIT vendors, their customers, and patients: a report of an AMIA special task force. *J Am Med Inform Assoc* 2011 January;18(1):77-81.

Correspondence to:
 Ross Koppel, Ph.D. FACMI
 Sociology Dept. and School of Medicine
 University of Pennsylvania
 Philadelphia, PA 19104-6299
 USA
 E-mail: rkoppel@sas.upenn.edu