

MEDICINE IN DENIAL[©]

Lawrence L. Weed
lweed@earthlink.net
(802) 899-3272

Lincoln Weed
ldweed@cox.net
(202) 728-4318

Table of Contents

Abstract

- I. Introduction**
- II. Disorder and denial in medical practice**
- III. The concept of defined inputs**
- IV. The foundation: coupling patient data with medical knowledge**
- V. "Idols of the mind" in medical practice**
- VI. Building on the foundation: the medical record as a tool to organize the processes of care**
- VII. The nature of medical knowledge as applied to individual patients**
- VIII. Medical education and credentialing as barriers to progress**

I. Introduction

It is in vain to expect any great progress in the sciences by the superinducing or engrafting new matters upon old. An instauration must be made from the very foundations, if we do not wish to revolve forever in a circle, making only some slight and contemptible progress.

— Francis Bacon¹

A culture of denial subverts the health care system from its foundation. The foundation—the basis for deciding what care each patient individually needs—is connecting patient data to medical knowledge. That foundation, and the processes of care built upon it, require standards for managing clinical information analogous to accounting standards for financial information. If businesses were permitted to operate without accounting standards, the entire economy would be crippled. That is the condition in which the \$2 trillion U.S. health care system finds itself—crippled by lack of standards for managing clinical information. The outcome is a continuing state of denial about the disorder that would be exposed if clinical information were managed with order and transparency.

Contrary to what the public is asked to believe, physicians are not trained to connect patient data with medical knowledge safely and effectively. Rather than building that foundation for decisions, autonomous physicians traditionally rely on personal knowledge and judgment, in denial of the need for external standards and tools. Medical decision making thus lacks the order, transparency and power that external standards and tools would bring to it. Physicians are left to carry a prohibitive burden. Acting under severe time constraints, they must connect intricate patient data with crucial details from vast and growing medical knowledge. The outcome is that the entire health care enterprise lacks a secure foundation.

Equally insecure are the complex processes built on that foundation: decision making, execution, observation, and corrective action over time. Responsibility for all these processes falls on physicians. Yet, here again medicine lacks accounting standards to manage clinical information. Inputs are undefined. The outcome is that physicians are not equipped to fulfill their immense responsibility safely and effectively. Other caregivers are not equipped to share that responsibility with physicians. Patients are not

¹ Bacon F. *Novum Organon* (1620), Summary of the Second Part, Aphorisms Concerning the Interpretation of Nature and the Kingdom of Man, Aphorism No. 31 (Montague, trans., 1854); at <http://history.hanover.edu/texts/Bacon/APHOR.html>.

equipped to work effectively with multiple caregivers, nor to assume the ultimate burden of decision making over their own bodies and minds.

In short, essential standards of care, information tools and feedback mechanisms are missing from the marketplace. And the underlying medical culture does not even recognize their absence. This does not prevent some caregivers from becoming virtuoso performers in narrow specialties. But that virtuosity is personal, not systemic, and limited, not comprehensive. Missing is a secure system for enforcing care of high quality by all caregivers for all patients.

Graduate medical education fills this vacuum with harmful habits and illusions that physicians find difficult to let go. Credentialing confers a legal monopoly on physicians, insulating them from competition. The resulting state of denial blocks development of a secure, orderly system of defined inputs, tight feedback and continuous improvement in patient care.

This state of affairs can be transformed. The necessary standards of care and information tools, as described below, are in large part already developed. Their widespread adoption could permit inputs to be defined, new roles for caregivers and patients to emerge, individualized decision making to occur, new feedback loops to develop and a self-correcting, continuously improving system of care to evolve.

Some may disagree, arguing instead that it is enough simply to reduce third party involvement in funding and the doctor-patient relationship. Others argue that the necessary transformation is already underway, as "health information technology" is used to bring "evidence-based medicine" to "patient-centered" care. Yet, neither point of view comprehends the wide gap between what current medical practice delivers and what patient care truly requires. This gap persists regardless of whether health care spending is provider-driven (traditional fee-for-service medicine), payer-driven (managed care) or now "consumer-driven." Until the gap is closed, attempts at cost control and universal coverage will continue to revolve in a circle, without sustainable progress.

Closing the gap could transform how medicine is personally experienced by caregivers and patients. Supported by a system of order and transparency, caregivers could find their work to be less exhausting and more rewarding, emotionally and intellectually, than what they now undergo. The physician's role could disaggregate into multiple roles, all freed from the impossible burdens of performance that physicians are now expected to bear. The expertise of nurses and other non-physician caregivers could deepen, and their roles could be elevated. All caregivers could follow time-honored standards of care that in the past have been honored more in the breach than the observance. All caregivers and patients could jointly use electronic information tools for

matching data with medical knowledge, radically expanding their capacity to cope with complexity. Inputs by caregivers could be defined and subjected to constant feedback and improvement. In that environment, patients and their families could personally manage unique, individual needs, collaborating with caregivers as necessary. A truly evidence-based medicine could develop, where evidence would be used to individualize care rather than standardize it. Under a system of checks and balances, using a common information infrastructure, patients and caregivers could act on incentives for quality and economy far more effectively than before.

Society could thus find enormous opportunities to harvest resources now going to waste. These wasted resources include a vast body of medical knowledge that all patients and caregivers could use more effectively, simple tests and observations that in combination could uncover solutions to patient problems, patients who could become better equipped and motivated to improve their own health behaviors, routine patient care that could become a fertile source of medical knowledge, and the first hand insights of caregivers who could participate in harvesting that new knowledge for the benefit of their own patients.

Were such a transformation to occur, the health care delivery system would become a transparent network that patients/consumers navigate to manage their own health. As an analogy, consider the transportation system – like health care, a system where public safety is at stake. Travelers rely on expert service providers when needed (pilots, auto mechanics, travel agents, for example), but the primary decision makers are travelers themselves. They determine the destination, the route, and the mode of travel for a journey. And their decisions are highly individualized. Two different people driving across the country might choose completely different routes, depending on whom and what they wish to visit and what they encounter along the way. Because such factors are variable, the choice of routes among different travelers is variable. No one would regard such variation as inappropriate. No one would expect travelers to conform to some "evidence-based" determination by experts of the "best" route across the country. Similarly, in medicine, no one should think that two different people labeled with the "same" disease necessarily have comparable medical needs. Nor should we think that the care of unique individuals must conform to "evidence-based" guidelines derived from large population studies. Rather, high quality efficient care would emerge case-by-case, in a progression of many small steps, each one carefully chosen and reliably executed.

Like the transportation system, the health care system should be usable by ordinary consumers when feasible. In travel, rather than relying on professional taxi drivers, we

learn to drive and we buy our own cars. Instead of hiring professional engineers to tell us what cars to buy, we read *Consumer Reports* and judge our personal needs and preferences for ourselves. Instead of hiring professional guides, we read maps and road signs. In some modes of travel (rail and air) we depend on expert service providers, but we as consumers, not those experts, choose the mode of travel. By comparison, if pilots had the authority that physicians have in the health care system, then pilots would decide when consumers travel by air, and unnecessary flights for short distances would be routine.

Essential to the functioning of consumers and expert service providers in the transportation system are reliability of the infrastructure and transparency in the rules for its use. Consider auto transportation. Roads and bridges are maintained in drivable condition. Maps, road signs and electronic systems are provided for navigation. Drivers are licensed and cars are inspected. Traffic laws are defined and enforced. As a result, drivers can find their destinations and arrive safely. The primary risk to safety is the behaviors of other drivers, not breakdowns in the transportation system. In contrast, breakdowns in the safety, quality and economy of health care are endemic.

Consider also airline safety regulation. A key characteristic is that inputs by workers with specialized expertise are carefully defined within an integrated system, every component of which is subject to strict scrutiny and control. Airline mechanics, for example, are subject to careful recordkeeping and inspection requirements. Pilot credentials are based not on formal education but on periodic demonstration of actual competence in flying specific planes. Air traffic control systems, sophisticated cockpit instrumentation and detailed standards of care govern the actions of expert pilots. Pilots do not have professional autonomy. They function within a protective system that is meticulously monitored. As a result, airline travel is so safe that consumer decisions about using it are made virtually without regard to safety concerns. No one chooses among airlines by comparing crash rates or pilot credentials.

Any complex system depends on the quality of its parts and their connections. The functioning of all parts must be reliable and oriented towards a common purpose. In these terms, the transportation system is superior to the health care system. The health care system will never be trustworthy or affordable until its parts and their connections match the transportation system in three key respects. Indeed, its recent evolution is turning in this direction:

- *Inputs by caregivers must be carefully defined and controlled.* During the last decade, this development has begun to take root. The patient safety movement has demonstrated over and over again the need to define and control inputs from

fallible human beings. But this development has focused largely on *execution* of decisions. Decision making itself equally needs definition and control of inputs, because each human mind, left to its own devices, is unreliable and not well connected to other participants in the system. Medical education and credentialing block the necessary changes in this regard. To reform, they must become focused on instilling in practitioners a core of behavior, not a core of knowledge.

- *A trustworthy and transparent intellectual infrastructure for care must be established.* During the last decade, the Internet has revolutionized access to expanding medical knowledge. But the human mind cannot apply complex knowledge effectively without external aids. Caregivers and patients trying to navigate the medical landscape need two information tools: a map of the landscape and a communication system for the journey. The map (tools for coupling medical knowledge with patient-specific data) reveals the landscapes so that individuals can find routes that serve their personal needs (unlike "evidence-based" travel directions dictated in advance). The communication system (structured medical records) enables the patient and multiple caregivers to coordinate their actions by recording and planning each step of the journey, informed by feedback from each other and from what happens along the way. With this infrastructure, all caregivers and consumers can apply complex knowledge to detailed data, and readily understand how their efforts interrelate.
- *The central role of the patient/consumer must be recognized.* During the last decade, this recognition has become increasingly evident in two areas: consumer-driven health care and management of chronic illness. But these developments are incomplete. The consumer-driven care movement focuses more on spending than care. In management of chronic illness, many organizations have developed approaches for helping patients manage their own conditions, but these disparate efforts are not unified by common tools and standards applicable in all medical contexts. Common tools and standards exploit basic principles of orderly problem-solving that everyone grasps. With that simplicity and unity, the health system becomes transparent and usable for all.

Transforming medicine in these respects would reorient patient care towards a single purpose: individualized medical problem solving for each unique patient. That orientation differs fundamentally from both payer-driven managed care and traditional, provider-driven medicine. These are both variations on the same vendor-driven system.

Both variations are disconnected from patient needs because a truly consumer-driven system of care has yet to be built.

To present these concepts, we begin with a detailed case study. Then we describe two information tools² and the standards of care they implement. Finally, we analyze the implications of these tools for reform of medical practice, the development of medical knowledge and medical education. The following outlines Parts II to VIII below.

- II *A detailed case study of a missed diagnosis:* This part examines disorder in medical decision making, explains the need to change medicine's division of labor, and analyzes implications for reform of medical practice.
- III *The concept of defined inputs:* This part explains the necessity for defined inputs by caregivers. This does not mean dictating medical decisions (the approach taken by managed care and evidence-based medicine). Rather, defined inputs mean assurance of reliable functioning, which is a prerequisite to meaningfully comparing outcomes. The concept of defined inputs presented in this section underlies the information tools for medical decision making described in parts IV and VI, and the approach to medical education and credentialing described in part VIII.
- IV *Coupling patient data with medical knowledge:* This part describes software tools designed to elicit relevant patient data and match it with applicable medical knowledge for specific medical problem situations. Using these software tools effectively imposes standards of care that are more rigorous than accepted medical practice. Physicians object to these standards on various grounds, which this part discusses in detail.
- V *Historical and philosophical background:* This part describes how medical practice lacks scientific rigor due to its unwarranted reliance on the unaided human mind – an issue that Francis Bacon identified 400 years ago at the birth of modern science. F. A. Hayek analyzed related issues in market economies, where people avail

² Specifically, the information tools are (1) decision support software designed for coupling medical knowledge with patient data, and (2) electronic medical records designed to organize care around patient problems instead of provider habits. Known respectively as knowledge coupling software and the problem-oriented medical record (POMR), these tools implement standards of care for managing medical information, as discussed in parts IV and VI below. Some basic references are Weed, LL., et al., *Knowledge Coupling: New Premises and New Tools for Medical Care Education*, New York: Springer-Verlag, 1991 (chapter 13 of this volume, authored by Dr. Ken Bartholomew, provides a detailed discussion of use of knowledge coupling software in conjunction with the POMR in a primary care practice); Weed, LL et al., *Medical Records, Medical Education and Patient Care*, Cleveland: Case Western Reserve University Press (1969); Burger, Charles S., "The Use of Problem Knowledge Couplers in a Primary Care Practice", *Healthcare Information Management*, vol. 11, no. 4, Winter 1997, available at www.pkc.com; C.C. Weed. *The Philosophy, Use and Interpretation of Knowledge Couplers*. PKC Corporation, 1982-2008, available at www.pkc.com; Weed LL, Weed I. Opening the black box of clinical judgment, *British Medical Journal, eBMJ Edition*, Vol 319, issue 7220, 13 November 1999, available at <http://bmj.bmjournals.com/cgi/content/full/319/7220/1279/DC2>

themselves of knowledge which individually they do not possess. These principles are directly applicable to medicine, which is both a scientific and commercial endeavor.

- VI *The medical record:* This part examines how the medical record can serve as the essential tool to manage complex cases (typically cases involving chronic disease and multiple problems). When the medical record is structured to reflect the basic steps of orderly problem solving, the complex processes of care become transparent to patients, caregivers and third parties.
- VII *The nature of medical knowledge as applied to patient care:* This part examines how the information tools described in parts IV and VI overcome the limitations of evidence-based medicine by taking into account the medical uniqueness of individuals.
- VIII *Medical education and credentialing:* This part argues that medical education and credentialing must change from a knowledge-based to a skills-based approach. A knowledge-based approach fails to exploit the power information technology for accessing knowledge and completely undermines definition and control of provider inputs to care. Moreover, medical education must be conceived in terms of not only caregivers but also patients/consumers, because they are central to medical decision making and execution.

IV. The foundation: coupling patient data with medical knowledge

A. Importance of defining initial inputs

For the want of a nail the shoe was lost, For the want of a shoe the horse was lost, For the want of a horse a rider was lost, For the want of a rider the battle was lost, For the want of a battle the kingdom was lost. And all for the want of a horse-shoe nail.

— Benjamin Franklin

Orderly problem solving begins with gathering relevant information. Relevance is a function of purpose. In medicine, information is gathered for three basic purposes: (1) maintaining wellness, (2) identifying medical problems at an early stage, and (3) solving identified problems. The output of this information gathering process is an initial database for decision making. The database elements relevant to the first two purposes comprise the *screening database*. The database elements relevant to the third purpose comprise the *initial workup* of each problem identified by the screening database plus any additional problems identified by the patient and provider. (The initial workup is sometimes referred to as the "present illness.")

Gathering a database should involve (1) careful *selection* of data relevant to the three purposes just stated, and (2) accurate *analysis* of the data to determine implications for the patient. Both selection and analysis require matching data with medical knowledge. The results of the matching process need to be organized for problem solving purposes. This total process of matching general medical knowledge with patient-specific data and organizing the results we refer to as *knowledge coupling*. Regardless of whether it is carried out by the unaided mind, software tools or both, knowledge coupling is inherent in developing a database for decision making.

Knowledge coupling is employed not only for the initial database but also for defining problems, formulating diagnostic or treatment plans, evaluating the results of those plans and modifying the plans as needed, all of which involve matching medical knowledge with patient data. Knowledge coupling is thus fundamental to all of medical decision making. But its importance is greatest, and most amenable to improvement,

when developing the initial database. Indeed, undefined, uncontrolled inputs to the initial database are a root cause of harm and waste in patient care.

The pivotal importance of the initial database should come as no surprise. Complex activities rest on a foundation laid by initial choices. Initial choices commit scarce time and resources to a chosen course of inquiry or action. Yet, initial choices are very often made without considering or comprehending relevant information. The best choice may thus be delayed, rejected or overlooked. Initial errors of this kind, like Franklin's missing horseshoe nail, trigger chain reactions that often cannot be stopped or even perceived until too late. Known as "cascade effects" when they occur at the biological level, these harmful chain reactions also occur at the organizational level. They are a recurrent phenomenon in clinical care at both levels.^{3 4} And events at both levels often interact to increase the risk and degree of harm to the patient. Characteristic elements in these chain reactions include multiplying of disorganized data collection and risky medical interventions, delay of beneficial treatment, escalating complexity and increased likelihood of error as a result.

Because initial choices have such great significance, complex activities are usefully conceived in two categories: (1) *threshold* processes where initial information is considered and initial decisions are made; (2) subsequent *follow-up* processes where decisions are executed, feedback is received and new decisions are made. The element of feedback further highlights the importance of threshold processes, because careful initial planning is often needed to establish feedback loops.

A clear example of the importance of threshold processes medicine is the case study in part **Error! Reference source not found.** Months of unnecessary care and avoidable suffering were set in motion when "classic manifestations" went unrecognized in the initial workup. Other examples appear frequently in the media. Consider the following:

- A Pulitzer-prize winning *Wall Street Journal* article reports on "a deadly discrepancy between the available medical knowledge about aortic aneurysms and the ignorance of many front-line physicians." Headlined "Medical Ignorance Contributes to Toll From Aortic Illness,"⁵ the article explains that aortic disease "kills an estimated 25,000 Americans a year ... a larger toll than that of AIDS and

³ Deyo R. Cascade Effects of Medical Technology. *Annu. Rev. Public Health* 2002. 23:23-44

⁴ Mold, JW, Stein JF. The cascade effect in the clinical care of patients. *New England Journal of Medicine*, 1986. 314:512-14.

⁵ Helliker K, Burton T. Medical Ignorance Contributes to Toll From Aortic Illness. *Wall Street Journal*. Nov. 4, 2003, p. A1, available at <http://www.pulitzer.org/year/2004/explanatory-reporting/works/story7.html>.

most kinds of cancer.” Yet, many physicians are not aware of the prevalence of aortic disease, its risk factors, its presenting signs and symptoms, and important diagnostic and therapeutic advances in caring for the disease. The article suggests that front-line physicians are to blame for their “ignorance.” Yet, such blame misses the point, because ignorance is inevitable. The article suggests that a new medical specialty for aortic illness is needed, but that too misses the point, because the illness, and thus the need for the specialist, are difficult to recognize. The point is that an orderly, structured investigation is essential from the outset of care, because aortic illness is only one of hundreds of possible disease conditions suggested by its various presenting signs and symptoms. If those diagnostic possibilities are systematically investigated from the outset, then prompt diagnosis and treatment of the underlying disease (whether it turns out to be aortic illness or something else) become readily achievable in many cases. The difficulty is that all physicians inevitably will be ignorant of some of those diagnostic possibilities, their manifestations and their treatments. It is thus crucial to minimize dependence on the personal knowledge of physicians or any other caregivers. Only then does it become possible to optimize diagnostic workups of presenting signs and symptoms, whatever their cause. Seen in this light, the *Wall Street Journal* headline radically understates the tragedy the article describes. Rather than “Medical Ignorance Contributes to Toll From Aortic Illness,” a more accurate headline would have been, “Medical Disorganization Multiplies Toll From All Illness.” And rather than blaming “the ignorance of many front-line physicians,” who deserve help rather than blame for being less than omniscient, it would be more accurate to blame the leaders behind the front lines, “the best and the brightest” who are positioned to transform the disorganization faced by all caregivers and patients.

- A 1998 ABC News report⁶ described the case of a patient who experienced fatigue and joint pain, progressing to muscle pain, worse fatigue, disorientation and double vision. She “consulted more than 20 different doctors in search of an explanation,” without success. Finally, she turned to Dr. Charles Burger, who first had her spend “about 30 minutes completing a detailed computer-generated questionnaire that posed virtually every question that medical science suggests is relevant to her symptoms of fatigue—more than 550 questions in all.” Dr. Burger’s computer software then coupled the questionnaire responses with a

⁶ <http://www.pkc.com/videos.aspx>. The software described is from PKC Corporation.

database of possible diagnoses associated with the responses. The patient's signs and symptoms "turned out to be a nearly perfect match for 'hyperventilation syndrome,' a shallow-breathing disorder that can gradually change the body's chemistry." The patient learned corrective breathing techniques, with the outcome that her health is "steadily improving."

Such examples suggest that optimizing the initial workup could pay enormous dividends.

Optimizing threshold processes in medicine must begin with defining inputs to the process of coupling general knowledge with patient-specific data. This knowledge coupling process is employed for both components of the initial database – the screening database and initial workup of any identified medical problems. The following discussion is limited to the initial workup. In part **Error! Reference source not found.** below we address the screening database.

The knowledge coupling function may be carried out by the unaided human mind, or with the aid of software tools. Because the unaided mind is so unreliable and so inefficient at knowledge coupling, software tools are essential. The question thus arises – which should be used first when conducting the initial workup? Should the caregiver first apply personal knowledge and judgment, and then use software tools to help with unresolved problems? Or should the caregiver first use software tools and then turn to personal knowledge and judgment?

A central theme of this paper is that software tools should be employed first. Stated differently, human thought should supplement but not substitute for software tools when the knowledge coupling function is performed. This principle constitutes a new standard of care for medical practice. It is completely contrary to the basic premises of medical education and credentialing. To understand these points, we first need to examine the underlying logical structure of the initial workup, and then compare alternative approaches to conducting it.

B. The structure of the initial workup

We all know what happens when a physician first examines a new patient. The physician's clinical judgment largely determines what questions are asked about the patient's "chief complaint" and medical history, what points are checked in the physical examination and what laboratory tests are ordered. Similarly, the physician's personal judgment is the primary vehicle for analyzing the data collected. Then the physician judges what data and analysis to include in the medical record. Cognitive inputs to the

initial workup are thus not predefined but determined during the patient encounter. Similarly, as to manual inputs when performing the physical examination, the quality of performance by physicians and other caregivers is not verified in advance. Caregivers vary considerably in the skill they bring to performing physical examinations, and they are rarely subject to corrective feedback on the quality of their performance. In short, the initial workup is unreliable at many levels.

Optimizing this initial workup requires breaking it down into three steps—choice, collection and analysis of patient data. Specifically:

1. Initial data must be chosen that is known to be cost-effective for identifying diagnostic or therapeutic options. This requires linking one datum—the patient’s complaint—with comprehensive medical knowledge about what cost-effective initial data.
2. All the chosen data must be collected by conducting the patient/family history, the physical examination and laboratory tests) without error or omission.
3. Once collected, the chosen data must be linked with comprehensive medical knowledge to determine, based on those initial data, what diagnostic or treatment options should be considered, the initial pros and cons of each option, and what steps should be considered next.

Combining the initial data and knowledge inputs should yield the following outputs: a set of options (diagnostic or therapeutic), plus, for each option, evidence (patient data) for and against the option, with proposals for additional data to collect or therapeutic measures to initiate. Equipped with this information, the patient and caregiver apply judgment and preferences to make an initial choice among the options identified.

This basic structure applies to both diagnosis and treatment decisions. That both types of decision have the same logical structure is not surprising. Organized problem solving of any kind involves processing information to identify options, plus evidence for and against each option. This information processing provides the *basis* for decisions. The decision is made when the patient or caregiver applies judgment to choose among the options identified, based on the evidence, in light of personal preferences and values. The objective quality of the decision, and its subjective acceptability to the patient, both depend on optimizing the *basis* for that decision.

C. Two contrasting approaches to the initial workup

The traditional approach to the three steps of the initial workup depends on the personal knowledge and judgment of the physician. The physician chooses and collects

very limited initial data (steps 1-2) for purposes of quickly formulating initial conclusions or hypotheses (step 3). Applying clinical judgment throughout, the physician is highly selective with the first two steps, proceeding to the third step as soon as possible. If the first iteration of this sequence is not successful, the physician keeps repeating the sequence with new data. This can be termed a *judgmental* approach, because it relies heavily on the trained physician's sophisticated clinical judgment in selecting the right initial data on the patient and then analyzing that data accurately (steps 1 and 3). External information tools are not employed for either step.

A contrasting approach begins with defining a detailed, standardized database for a given diagnostic or management problem (step 1), taking all medical specialties into account, based on research of the medical literature. Constructed *before* any patient encounter, this database is entered into software designed to be used *during* the patient encounter. In step 2, the patient and caregiver use the software as guidance for collecting all of the defined data for the given problem, whether or not the caregiver judges the data to be useful for an individual patient. In step 3, the software automatically links all of the data with comprehensive medical knowledge about the significance of the data points and their interrelationships. This is a simple process of association, readily accomplished with software tools and without reliance on the caregiver's knowledge or judgment.

This alternative to the judgmental approach we term a *combinatorial* approach. It uses computer software to combine data with knowledge, in order to identify medically significant combinations of data points. The combinatorial approach contrasts with the judgmental approach not only in its reliance on an external tool to determine decision making inputs, but also in the nature of these inputs:

- The level of detail, in both the data collected and knowledge taken into account, are much greater with the combinatorial approach.
- This detailed information is gathered and presented up front and all at once, rather than in a gradual, piecemeal fashion.
- Inputs are highly standardized (the data collection and identification of linkages to medical knowledge are determined uniformly, regardless of which provider the patient sees), unlike the highly variable inputs of the judgmental approach, where inputs depend on the caregiver's personal knowledge, experience, specialty orientation, reimbursement expectations, time constraints and other factors.

Driven by an external tool, not personal judgment, the combinatorial approach creates a minimum standard for the information taken into account, but not a limit on that information. Caregivers (and patients) are free to supplement the defined initial workup with any additional data they judge useful, and are free to make their own

connections between the patient data and medical knowledge in addition to the connections identified by the software. In contrast, the judgmental approach limits the initial workup to the data and connections suggested by the caregiver's personal judgment, with no assurance of satisfying minimum standards of completeness or accuracy.

The tool-driven, combinatorial approach immediately brings to bear the best current expertise, accumulated from the experience of countless patients and caregivers over time, filtered through peer-reviewed medical literature. In contrast, the judgmental approach relegates the patient largely to the limited expertise possessed by the caregivers available to the patient.

The judgmental and combinatorial approaches contrast in terms of not only their inputs but their outputs. Specifically:

- The output of a combinatorial approach is not a clinical decision, but merely the informational basis for a decision: a set of options with patient-specific evidence for and against each option. The initial options and evidence are generated by the software tool, not by the physician's judgment. The decision to be made based on this output is left to the patient and caregiver. In contrast, the output of a judgmental approach naturally takes the form of the physician's recommendations plus a selective presentation of options and evidence, generated by exercise of the physician's clinical judgment.
- The combinatorial approach automatically generates complete documentation, which is available over time to the patient and multiple caregivers and to clinical researchers. This documentation includes positive, negative and uncertain responses (whether or not the caregiver judges them to be significant). In contrast, typical physician notes are ambiguous as to whether an unmentioned finding was found to be negative or uncertain or was never checked at all.

The judgmental approach, not the combinatorial approach, is the accepted standard of care in medicine. For example, standard coding guidance states that the physician "uses the presenting illness as a guiding factor and his or her clinical judgment about the patient's condition to determine the extent of" the history and physical examination.⁷

The judgmental approach, with its dependence on the unaided mind, is the accepted standard not only in routine primary care but also in advanced specialty care at leading medical centers. With respect to the latter, this conclusion might seem surprising because

⁷ Center for Medicare and Medicaid Services, Evaluation and Management Services Guide (Rev. July 2007), pp. 12-13, at http://www.cms.hhs.gov/MLNProducts/Downloads/eval_mgmt_serv_guide.pdf.

specialists frequently pursue detailed data collection, which is characteristic of the combinatorial approach. Such investigation, however, usually occurs *after* the initial workup, and after a crucial threshold judgment is made—the judgment of which specialist should be consulted first. Different primary care physicians or other gatekeepers vary in their judgments of which specialist should be consulted first, and physicians in different specialties vary enormously in how they approach a given patient problem (not to mention the variation that exists even among physicians within the same specialty). In contrast, the tool-driven combinatorial approach takes into account data and knowledge from all potentially relevant medical specialties at the outset of care, before any physicians exercise judgment. This deferral of judgment is critical, because patients often have multiple problems, and because even a single problem often implicates multiple body systems, each with its own medical specialty.

D. Optimizing the initial workup – evaluating the two approaches

Recognizing the differences between the judgmental and combinatorial approaches is essential to optimizing the initial workup. Optimization requires both selecting the optimal approach and optimizing its execution.

Comparing the judgmental and combinatorial approaches faces significant obstacles. Very few physicians have even considered the combinatorial approach, much less actually used the software tools required. Moreover, rigorously implementing the combinatorial approach is disruptive. Ultimately, using the required software tools entails large changes in mental habits, established procedures, occupational roles, with significant reimbursement implications. Unless those changes are made, the two approaches cannot be accurately compared. Moreover, the basis for comparison is subject to dispute. In particular, comparison on the basis of medical outcomes is problematic at best. The effect of the initial workup cannot be isolated by controlling or adjusting for all the other inputs over time from providers and patients and external forces. That reality makes it impractical to design a clinical trial providing an accurate outcome comparison of judgmental and combinatorial approaches to the initial workup. Innovation must be pursued without the security blanket of definitive outcome comparisons.

What can be done in the face of the above difficulties is to analyze the basic premises of the judgmental and combinatorial approaches in relation to what patients need and modern information tools permit. In doing so we must take into account the realities of

medical practice, including the reactions of physicians, other caregivers and patients. Yet, their beliefs, existing occupational roles, current uses of technology, and immediate institutional needs must not cloud or restrict our long-term vision of what is optimal for individual patients and the health care system as a whole.⁸ At the same time, visions of what could be optimal with greater resources must not divert us from making optimal use of existing resources.

The basic premise of the judgmental approach is that firsthand clinical judgments of highly trained physicians interacting with patients should govern the initial workup. Physician judgments should not be compromised by second-hand, abstract, “evidence-based” generalities incorporated in clinical guidelines or software tools. Those external tools may have some utility as references to consult, but physicians should primarily rely on their own clinical judgment to determine the contents, and assess the results, of each initial workup for each unique patient.

In contrast, the basic logic of the combinatorial approach is that detailed initial data to be collected for a given medical problem should be defined in advance and collected without fail at each patient encounter. For this to happen, software tools, not clinical judgment during the encounter, should govern collection and analysis of the data. Ongoing clinical judgments of physicians may have some utility as a *supplement* to the combinatorial minimum standard, but should not be permitted to lower that standard.

E. Basic objections to the combinatorial approach.

Physicians naturally view the judgmental approach, and the elaborate training needed for the unaided mind to apply it, as inherent in scientifically advanced medical care. By comparison, a tool-driven combinatorial approach seems to impose both crude standardization and excessive detail—“cookbook medicine” taken to a compulsive extreme. These general reactions can be broken down into the following five specific points:

- A combinatorial approach seems prohibitively time-consuming and expensive to physicians, because the knowledge coupling software employed in the combinatorial approach requires routinely collecting data (step 2 of the initial workup) in greater detail than is customary or feasible in real-world medical practice, where rapid throughput of patients is an economic necessity.

⁸ Weed LL, “Physicians of the Future,” *New England Journal of Medicine*, 304: 903-907, April 9, 1981.

- Whatever its feasibility, detailed data collection at the outset of care seems unproductive to physicians. On this view, detailed initial data collection results from premature concern with unlikely possibilities (rare diagnoses, non-standard treatments). Limited time and resources logically should be directed first at the most common and likely possibilities. Knowledgeable physicians can frequently avoid the time and expense of detailed data collection, as illustrated by the physician who correctly diagnosed the Addison's disease case based on very limited data from the initial examination (see part **Error! Reference source not found.** above).
- Physicians view the *standardized* data collection involved in a combinatorial approach as mere cook-book medicine. Every patient and every practice situation are different. Selecting which data are clinically useful for each patient in any particular context (step 1 of the initial workup) must be judged by the physician during each patient encounter.
- Physicians object that the combinatorial approach harms the doctor-patient relationship. Because the patient interacts initially more with a computer and non-physician personnel than with the physician, the combinatorial approach gives the physician less opportunity to observe the patient's condition firsthand and establish a personal relationship.
- Physicians believe that mere information processing by external software tools can never substitute for their informed clinical judgment. Like experts in other advanced fields of knowledge, physicians solve complex problems by applying first principles in a subtle and discriminating manner that no software can replicate. Nor can software replicate physicians' intuitive, instinctive judgments arising from their personal interactions with patients.

Much of the above critique originates in the culture of graduate medical education. That culture teaches reliance on personal intellect. Medical students are first selected for academic proficiency. Then they learn that they must acquire vast medical knowledge. Then they learn to rely heavily on their own intellectual powers when applying that knowledge to detailed patient data. Then they submit to knowledge-based standards for licensure and board certification. Their ordeal indoctrinates physicians with faith in the efficacy of intellect. Reinforcing this faith are the high compensation and status physicians receive for exercising their clinical judgment. Moreover, they are legally insulated from competition by other caregivers who might deliver superior care with less faith in personal intellect.

Not only physicians but their patients acquire this faith. We are all socialized to believe, as the great clinician Herman Blumgart once wrote, that “application of knowledge at the bedside is largely the function of the sagacity inherent in or personally developed by the individual physician.”⁹ Sherwin Nuland has further described this ideal of personal sagacity:

To understand pathophysiology is to hold the key to diagnosis, without which there can be no cure. The quest of every doctor in approaching serious disease is to make the diagnosis and to design and carry out the specific cure. This quest I call The Riddle, and I capitalize it so there will be no mistaking its dominance over every other consideration. The satisfaction of solving The Riddle is its own reward, and the fuel that drives the clinical engines of medicine’s most highly trained specialists. It is every doctor’s measure of his own abilities; it is the most important ingredient in his professional self-image. ... Our most rewarding moments of healing derive not from the works of our hearts but from those of our intellects -- it is there that the passion is most intense.¹⁰

Yet, the works of our intellects cannot be trusted. Those rewarding moments when our intellects solve The Riddle for some patients are inseparable from terrible moments when our intellects fall short. And the problem is not just fallibility. The problem is also that the best intellects can serve only a few. One physician's solution of The Riddle for one patient provides no foundation for improving the care of all patients by all caregivers. Those realities are the Achilles heel of the judgmental approach to the initial workup. A commitment to overcome those realities leads inevitably to the combinatorial approach, implemented with external information tools designed to make it practical.

F. A software implementation of the combinatorial approach

... a whole calling may have unduly lagged in the adoption of new and available devices. It may never set its own tests. There are precautions so imperative that even their universal disregard will not excuse their omission.

—Judge Learned Hand¹¹

Physicians vary in their innate and acquired intellectual abilities. Moreover, even the most gifted and well-schooled intellects are not reliable when processing large

⁹ Blumgart H., "Medicine: The Art and the Science," *Hippocrates Revisited*, R. Bulger ed. (New York: MEDCOM Press, 1973), p. 34.

¹⁰ Nuland, Sherwin B. *How We Die: Reflections of Life's Final Chapter* (New York: Alfred A. Knopf, 1994), pp. 248-49.

¹¹ *The T. J. Hooper*, 60 F.2d 737, 740 (2d Cir. 1932).

amounts of information on the fly – which is precisely what steps 1 and 3 of the initial workup require. Medicine thus needs a new division of intellectual labor – a division between electronic tools that process information, and users who, based on that information and personal values, apply judgment to arrive at decisions.

This division of labor exploits the fact that human judgment becomes increasingly accurate, efficient and powerful when it receives relevant information and is equipped to take that information into account – provided, however, that the information is actually relevant to the problem at hand and not so voluminous or disorganized as to escape comprehension. The proviso is critical. An overload of extraneous or disorganized information brings to the fore the very cognitive weaknesses that must be overcome. Yet, physician credentialing pretend these cognitive weaknesses do not exist.

The culture of medicine must drop the pretense and instead enforce the consistent use of external information tools. To understand what this means, it is useful first to step back and identify three distinct functions that information technology may perform:

1. storage and retrieval of general knowledge (e.g., Internet access to medical texts);
2. storage, retrieval and transmission of patient data (e.g., electronic medical records, telemedicine tools); and
3. linkage of patient-specific data with general knowledge for decision making purposes (e.g., computerized prescription order-entry systems that use patient data inputs to provide individualized guidance on medication selection and dosing, or the knowledge coupling software discussed below).

The difficulty of each of these three functions has been escalating dramatically for a century, with the extraordinary growth of medical knowledge and corresponding patient data. Performance of the first function has successfully kept pace with this growth, thanks to electronic tools for knowledge storage and retrieval. Performance of the second function, however, has not kept pace, in part because of lack of interoperability among different technologies, and in part because of the disorganized state of medical records as discussed in part **Error! Reference source not found.** below. Here, our concern is the third function – linking data with knowledge. Paradoxically, improvement in the first two functions makes this third function all the more difficult. That is, new tools for the first two functions increase the volume of accessible knowledge and data to combine and comprehend.

That third function is at stake in steps 1 and 3 of the initial workup. Recall that step 1 involves linking one datum – the patient’s problem – with comprehensive medical knowledge about the additional data needed to investigate that problem. Once collected

(step 2), the chosen data must be linked with all relevant knowledge (step 3) in order to “connect the dots” – to comprehend the pieces of data that turn out to be useful, while filtering out extraneous data. The unaided human mind cannot perform this linkage function reliably or efficiently. The difficulty is not just that the volume of medical knowledge exceeds what anyone can learn. (That problem diminishes with tools for efficient knowledge retrieval.) The deeper problem is that no one is able to link complex knowledge reliably with detailed patient data, especially when operating under real-world time constraints.

This is a familiar phenomenon. It occurred, for example, in the Addison's disease case described in part **Error! Reference source not found.** The inherent intellectual difficulty is often exacerbated by situation-specific or provider-specific constraints. A doctor may be so busy and distracted, or so fixated on possibilities within his specialty, that he fails to consider or follow through on possibilities that another physician might pursue immediately. Moreover, existing knowledge is imperfect. Actual patients often do not fit neatly into the patterns that education and experience lead physicians to expect. Every patient is unique in the combination of those characteristics and circumstances that bear on solving the patient's problem. Thus, the subset of data and knowledge that turn out to be useful differs for every patient, even patients labeled with the same disease. And that subset cuts across specialty boundaries. Medical specialties artificially restrict analysis, because actual patient problems turn out to implicate multiple specialties. It is thus not possible to determine in advance the limited subset of data and knowledge that will prove to be relevant to each patient's unique needs.

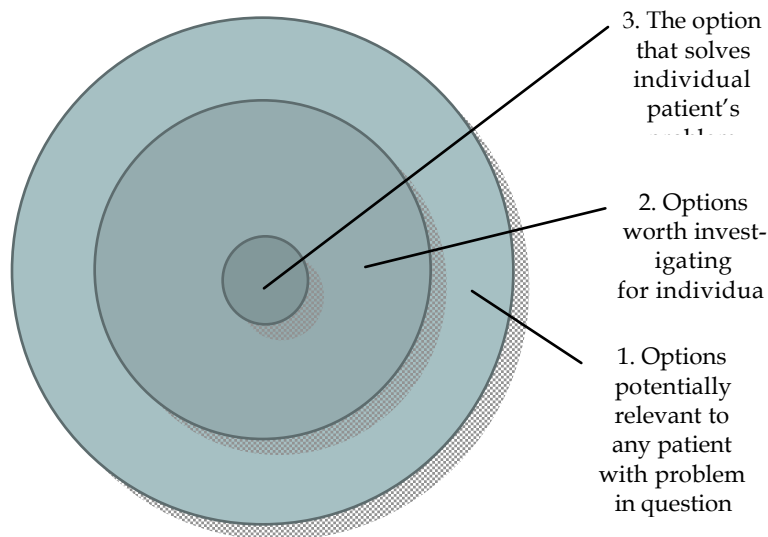
One context in which the mind's limitations have become especially obvious is medication ordering in hospitals. In that context, most observers now recognize that judgmental and manual processes are no substitute for computerized prescription order entry (CPOE) systems. More than two decades of studies have documented compelling quality and cost improvements at institutions that enforce use of well-designed systems. If this is true for the limited function of medication ordering, then it should come as no surprise to find that electronic tools are essential in any setting that requires linkage of patient data with medical knowledge – especially the initial workup.

Electronic tools should filter and organize all information potentially relevant to the problem at hand for the purpose of “connecting the dots” among the limited items of information actually relevant to the individual patient. To illustrate, consider diagnosis of hypertension. Scores of possible causes (and other diagnostic associations) are usefully taken into account for safe and cost-effective diagnosis of hypertension. These diagnostic possibilities implicate numerous medical specialties. Cost-effectively identifying which

of the diagnostic possibilities are relevant to an individual patient requires making hundreds of distinct clinical findings from history, physical and laboratory data. Yet, most physicians consider only a fraction of all this information when conducting initial workups of hypertension.

Another example is management of diabetes. Again, the medical literature shows that managing diabetes requires taking into account scores of different therapeutic options. Identifying the options relevant to an individual patient and initially assessing the pros and cons of the relevant options involves making hundreds of distinct findings at the outset of care. Again, most physicians consider only a fraction of all this information during initial workups of diabetic patients.¹²

Once made, all of the initial findings are coupled with a database of medical knowledge built into the knowledge coupling software. The software's output is a list of diagnostic or therapeutic possibilities (i.e. possibilities for which at least one expected finding appears in the patient), plus the evidence for and against each possibility (i.e. the expected findings that are present and those that are not present). The output also includes comments on findings and options for which further explanation is useful, and supporting citations to the medical literature. The software thus links data with general knowledge, filters out information not potentially relevant to the patient and organizes the rest, arranging it by the option to which the information relates. Complex information is presented in a way that is maximally useful for solving the problem at hand. This total process is termed "problem-knowledge coupling."



¹² The examples given are based on research conducted by PKC Corporation for purposes of building the knowledge coupling software described here.

The coupling process can be conceived as navigating through three concentric circles of knowledge and corresponding patient data. These three circles relate to diagnostic or therapeutic options for the individual patient, as shown below.

The outer circle consists of all known options that are potentially relevant to the problem in question, plus, for each option, medical knowledge about the most useful initial data to collect for purposes of determining which options are worth investigating. This outer circle is completely standardized. The same data should be collected uniformly for all patients with the same presenting problem.

The middle circle is the subset of options, with corresponding knowledge and data, that are worth considering for an individual patient, based on data suggested by the outer circle and collected for that patient. The options worth considering are options for which at least one positive finding is made in that patient. Unlike the outer circle, the information in the middle circle is not standardized; it varies among different patients with the same presenting problem. The initial consideration of these options is simply comparing them to determine if one option stands out as the solution to the patient's problem, or, if not, which options are worth further investigation.

The inner circle is the option ultimately chosen as the best solution for the patient's problem. In some cases, the initial workup will be sufficient to arrive at this inner circle; in other cases, further investigation will be needed before a solution can be identified. In cases of genuine uncertainty, the inner circle is never located, because no clear solution may be ascertainable from coupling existing knowledge with available data.

Sound decision making requires navigating reliably and efficiently from an initial position of ignorance, through the first two concentric circles, to the inner circle. Reliability and efficiency in this process are crucial. All possibilities relevant to the medical problem must be considered, and all possibilities irrelevant to the individual patient must then be excluded, without unnecessary trial and error.

Another way to visualize these concepts is as a grid, with columns for all options and rows for all findings applicable to the problem in question. The grid corresponds to the outer circle depicted above. The data for any individual patient will fill only a small fraction of the cells in the grid, and the pattern formed by those cells (corresponding to the second circle) will differ for almost every patient. When those cells are concentrated in one column, that option may be chosen for that patient (corresponding to the inner circle), but even then different patients in the same column will likely vary—their patterns will overlap but not be identical.

Once the knowledge coupling process identifies possibilities relevant to the patient, the next step is to prioritize those possibilities for further consideration. Knowledge coupling software does not establish priorities but groups the possibilities in a way that facilitates making a judgment about priority. The grouping varies depending on the medical problem involved. For example, the Coupler for diagnosis of hematuria classifies possible diagnoses into the following groups:

- Rapidly Progressing Disorders: May Need Immediate Attention
- Causes for Which Just One Finding Makes Consideration Mandatory
- Other Causes of Hematuria
- Approach for Isolated Hematuria When No Findings Suggest a Diagnosis

Within each one of these groups, the Coupler output lists primary options (i.e. diagnostic possibilities) for which one or more positive findings appear in the patient.

For each option, the listing shows the number of positive findings in the patient and the total number of possible findings for that option. That information gives the user an immediate sense of which options seem to best match the patient. For example, one option might show that 5 of 7 findings were positive, while all other options might show only one or two positive findings. But the apparent best match should not necessarily be the highest priority option. Instead, the highest priority options to consider are rapidly progressing disorders that could cause harm if not treated quickly. The next highest priority group is options for which just one positive finding makes consideration mandatory, regardless of how many positive findings appear for other options. The next group lists other possible options, which can easily be ranked by the positive findings as a proportion of all findings for each option. That numerical ranking, however, is not a sufficient basis for judgment. Needed is further information about each option, including a description of typical manifestations, known variations from what is typical, epidemiological information about prevalence in various populations, possible tests to order and other information useful for diagnostic assessment. This kind of information appears when the user clicks on each option in the list.

Sometimes the description for one of the options matches the patient closely, and all other options match poorly. That description might identify a single test to confirm or rule out the option, or, in the absence of such a test, the user may judge whether the option can be accepted as clearly the correct diagnosis. In other cases, several options might be worth considering based on positive findings and the descriptions provided, while in other cases, none of the descriptions match the patient well and no diagnosis seems plausible. In these two situations, the Hematuria Coupler provides detailed further guidance for each option. For example, under the heading "Approach for Isolated

Hematuria When No Findings Suggest a Diagnosis," the following subheadings appear: (1) Isolated hematuria with no plausible diagnosis, (2) Left renal vein hypertension, and (3) Prostatic venous rupture. Under the first subheading the user finds epidemiological information (e.g. cancer and urological disease prevalence among men above and below age 40 with microscopic hematuria) and detailed explanation and references about possible further testing and monitoring. Similar guidance is provided for left renal vein hypertension and prostatic venous rupture.

Another example is the Coupler for management of Type I and II diabetes in adults (including gestational diabetes, disease complications and related conditions). The options are classified into the following groups (only those relevant to an individual patient will be displayed):

- Glucose-Related Crises Requiring Hospitalization
- Related Conditions Requiring Identification and Management
- Monitoring Diabetes Control
- Controlling and Monitoring Diet
- Exercise
- Overview of Drugs for Diabetes
- Oral Drugs That Increase Insulin Secretion by the Pancreas
- Oral Drugs that Decrease Insulin Resistance & Increase Glucose Use
- Oral Drugs that Slow Digestion of Complex Carbohydrates after Meals
- Insulin Therapy
- Before-Meal Insulin Analogs and Inhaled Insulin
- Basal Insulin Analogs
- Insulin Therapy: Preventing and Monitoring For Common Problems
- Foot Care
- Travel Considerations
- Cholesterol and Triglyceride Management
- Blood Pressure Management
- Preventing, Monitoring For, and Managing Associated Conditions
- Pregnancy and Diabetes
- Monitoring for and Managing Complications of Diabetes
- Other Couplers That Might Provide Further Guidance
- Complementary and Alternative Medicine (CAM)

- Emerging Therapies and Therapies of Limited Availability
- Options For Which Only Cautions Are Present*

* Guidance options that are unlikely to apply to this particular patient, but for which information about cautions might be useful, are included under this heading.

A useful comparison with the combinatorial approach and knowledge coupling software is the concept of simply entering findings into a search engine such as Google.¹³ That approach falls far short of knowledge coupling at three levels. First, the search results depend entirely on what terms the physician enters in the search engine, that is, on the initial findings that the physician judges to be significant. That exercise of judgment completely undermines the initial workup. Different combinations of findings on a patient may point in very different directions. No one can be confident that the right combination of findings will emerge unless the initial workup takes into account, without omission, *all* the findings needed to elicit *all* the options worth considering for the individual patient. Second, Google searches an unfiltered, unstructured body of information – the entire world wide web – completely unlike the distilled and structured body of precisely relevant information in which knowledge coupling "searches" take place. Third, the output of a Google search lacks the structure and precise relevance of the output generated by knowledge coupling software. Accordingly, the recent *BMJ* article on use of Google as a diagnostic aid acknowledged serious limitations in that approach:

We suspect that using Google to search for a diagnosis is likely to be more effective for conditions with unique symptoms and signs that can easily be used as search terms ... Searches are less likely to be successful in complex diseases with non-specific symptoms ... or common diseases with rare presentations The efficiency of the search and the usefulness of the retrieved information also depend on the searchers' knowledge base.

Some observers have pointed out that the Internet has created an unprecedented excess of available information. In the past, much of that information would have been forgotten rather than preserved. Now it seems as if "the art of forgetting" must be restored.¹⁴ This new excess of information magnifies a dilemma already inherent in any

¹³ Tang H, Ng J. Googling for a diagnosis—use of Google as a diagnostic aid: internet based study. *BMJ*, doi:10.1136/bmj.39003.640567.AE (published 10 November 2006), <http://www.bmj.com/cgi/rapidpdf/bmj.39003.640567.AEv1>.

¹⁴ Viktor Mayer-Schönberger, "Useful Void: The Art of Forgetting in the Age of Ubiquitous Computing" (April 2007), [http://ksgnotes1.harvard.edu/Research/wpaper.nsf/rwp/RWP07-022/\\$File/rwp_07_022_mayer-schoenberger.pdf](http://ksgnotes1.harvard.edu/Research/wpaper.nsf/rwp/RWP07-022/$File/rwp_07_022_mayer-schoenberger.pdf); Seth Lloyd, "You Know Too Much" (4/28/07), <http://discovermagazine.com/2007/apr/you-know-too-much>.

attempt to apply complex general knowledge to specific problem situations. Sometimes expressed as "connecting the dots," or "finding the needle in the haystack," this dilemma is effectively addressed with knowledge coupling software.

Use of knowledge coupling software does more than locate relevant information; it also changes modes of expression, which in turn changes perception. In traditional practice, a complex patient encounter is typically reduced to a few words or phrases in the record stating the physician's "impression." That impression is bound to be both incomplete and subjective, and the language chosen to express it may further distort the reality of the patient's condition. It is always the case that language appropriates and alters reality by representing it in terms of the speaker's or writer's perspective, using words and concepts that the audience may understand differently than the speaker or writer.¹⁵ But with the use of knowledge coupling software, this process becomes transparent, defined and subject to organized improvement (see the discussion in part **Error! Reference source not found.** at note **Error! Bookmark not defined.** below). Rather than have the caregiver compose a selective, minimalist statement of impressions from a limited initial workup (or cut and paste some other caregiver's previous text in an electronic record), the combinatorial approach has the patient work with the caregiver to select from pre-defined, careful descriptions of myriad details that might describe the patient's condition. The descriptions from which the patient selects may be pictorial as well as textual. And the patient and caregiver remain free to add their own additional descriptions on specific points.

It is important to understand that use of Couplers breaks down specialty boundaries. For significant medical problems the outer and middle circles of information encompass multiple specialties. Specialized knowledge inevitably is incomplete and fragmented relative to actual patient needs. As a result, physicians cannot know all the tests and observations that might be relevant to complex medical problems. Nor can they know how to interpret and interrelate all the results. Nor do they have time for the research needed to fill in the gaps in their knowledge. Even when their knowledge is sufficient, physicians are frequently unable to apply their knowledge to detailed data in an organized, reliable manner. For those reasons, failures to collect, comprehend or even keep track of potentially useful data are endemic in medicine. These are not failures of individual physicians. Rather they are failures of a non-system that imposes burdens of information processing too great for physicians to bear.

¹⁵ *Knowledge Coupling*, note 2 above, p. 52.

Despite the relief from these burdens that knowledge coupling software provides, many physicians initially find that its use is disruptive and disturbing. The disruption is in part external, because using Couplers effectively involves changing office procedures and making greater use of non-physician personnel. But the disruption is also internal, because a change in self-image is involved. The software constantly confronts physicians with options and evidence going beyond what they would take into account if left to their own devices. Physicians soon realize that all their hard-won knowledge is radically incomplete; indeed, it is often misleading. Entrenched mental habits must change.

It becomes apparent to physicians, their co-workers and their patients that their role is not to *learn* medical knowledge but rather to *access and apply* the limited knowledge relevant to each patient's individual needs. This new role cannot be performed reliably if the physician exercises clinical judgment about what options are worth considering or what initial data are worth collecting for each patient. Rather, the caregiver's role is to collect, without omission, all the data made relevant by the outer circle of knowledge (which does not preclude collecting additional data that the caregiver or patient believe are relevant).

G. Answering objections to the combinatorial approach

In light of this basic description of the combinatorial approach and the external tool needed apply it, let us revisit the five objections outlined in part I.E above. In reading this discussion, keep in mind the three steps of the initial workup (see part I.B) and the three circles of information (see part I.F).

1. Feasibility of detailed initial data collection

As we have seen, the tool-driven combinatorial approach requires collecting initial data in far greater detail than is customary. Physicians object that the level of detail involved is prohibitively time-consuming. In many practice settings, physicians are expected to see several patients every hour. That pace leaves no time even to collect detailed data, much less time to analyze and review it all with the patient. Even in settings where the pace is slower, most physicians view the combinatorial approach to the initial workup as unacceptably time-consuming.

This objection is not meaningful unless it takes into account the *utility* of the combinatorial approach. If it is sufficiently valuable in terms of quality and long-run cost-effectiveness, then the combinatorial approach should be enforced without regard

to the prior habits or expectations or economic demands of providers and third party payers. The utility of the combinatorial approach is the subject of sections 2-5 below. In this first section, however, we examine its feasibility.

The amount of time required by the combinatorial approach is less than it may appear at first glance. Only data that are simple, non-invasive, quick and inexpensive to obtain are selected for the initial workup. Non-physician caregivers and patients themselves, guided by knowledge coupling software, can gather those data with great efficiency and reliability.^{16 17} Entering history and symptom data in the software requires the patient to do nothing more than click once (to say “yes”) on those findings that the patient recognizes as describing his or her condition (positive findings). The patient clicks twice to say “not sure.” Negative findings require no action at all. And most findings are negative for most patients. Most patients thus can rapidly make hundreds of findings simply by paging through computer displays, clicking on the occasional positive or uncertain item. Some findings require the patient to click on a button to review an explanatory graphic or text. That extra step for the patient saves time, because the patient relies on the software rather than the caregiver for initial help.

Moreover, positive and uncertain findings can be annotated with free text when the patient or caregiver wish to elaborate. That extra step also becomes a timesaver, because it elicits and captures significant patient data in context for later use. One practitioner estimates that gathering data in this manner is 5-10 times faster than a verbal discussion covering the same amount of data.¹⁸

The efficiency of the combinatorial approach becomes most obvious at step 3 of the initial workup: linkage of detailed patient data with comprehensive medical knowledge bearing on the significance of the data. With knowledge coupling software, that step occurs instantaneously. Equally important, the software’s output organizes all this information for rapid comprehension, while filtering out extraneous information that can safely be ignored. Specifically, the software displays a list of diagnostic or therapeutic options suggested by the initial positive and uncertain findings on the patient (the middle circle of information), while omitting options for which not a single positive finding is made (the outer circle). The options displayed are logically grouped, and

¹⁶ Bartholomew K. The Perspective of a Practitioner, in Weed LL. et al., *Knowledge Coupling: New Premises and New Tools for Medical Care Education*, New York: Springer-Verlag, 1991, p. 240.

¹⁷ Burger, C., “The Use of Problem Knowledge Couplers in a Primary Care Practice”, *Healthcare Information Management*, vol. 11, no. 4, Winter 1997, available at www.pkc.com

¹⁸ Personal communication with Harold D. Cross, M.D.

clicking on each option displays all of its positive, uncertain and negative findings as evidence for or against that option in that patient, along with additional description, commentary and citations about the options and findings.

This output is generated automatically, without the time-consuming process of drafting or dictating/typing/proofing a narrative summary. The software's output is more complete, organized and precise than a narrative summary and more compatible with coding systems. The output can be printed or exported to electronic medical record systems, and reports can be generated in rich text format for printing or electronic transmission to the patient and other providers. In short, using software tools to couple medical knowledge with patient data generates enormous direct administrative efficiencies. The more important source of efficiency, however, is the indirect effect on productivity.

Using knowledge coupling software for the initial workup resembles use of a map. A map is a highly efficient information tool for navigating in unfamiliar territory. Its efficiency results from its communicating carefully selected geographic information in a distilled, visual form that is maximally usable at the point of need. Similarly, knowledge coupling software provides a map to the landscape of medical knowledge. Patient-specific data items entered into the software, like coordinates on the X and Y axes of a map, enable the user to find immediately the patient's location in the landscape of medical knowledge. Collecting detailed data is, of course, more time consuming than using coordinates on a map. But this time investment pays large dividends.

The caregiver and patient can rapidly compare the patient's condition with more-or-less similar situations described in the software's output, including not only the options and data it suggests but also the commentary it provides and the medical literature it cites. Detailed information becomes manageable because it is filtered and organized based on individual relevance to each patient's unique combination of needs. A careful initial workup is thus more feasible than it may first appear. As described by a pioneering user of knowledge coupling software, Dr. Ken Bartholomew:

... let me dispel the notion that [Couplers] are extremely time-consuming if used properly. Certainly, the first few times using a Coupler will be time consuming, but no more so than reading a textbook chapter on a given problem. ... Because of the timeliness of the data that is built into the couplers, and since they are so pertinent by being problem-oriented [i.e. designed analyze is presenting problems, not medical specialties], I do not need to spend 20 or 30 minutes going through indices of textbooks to find what may or may not be appropriate information.¹⁹

¹⁹ Bartholomew K. The Perspective of a Practitioner, in Weed LL et al, *Knowledge Coupling: New*
06/27/08

The whole question of feasibility is thus transformed. Once individualized problem solving becomes the goal, then nothing less than detailed data collection seems feasible for achieving that goal.

Moreover, to the extent that the goal is efficient use of personnel resources, knowledge coupling software makes it feasible to increase reliance on inexpensive non-physician caregivers. As described by one user of knowledge coupling software, Dr. Charles Burger:

We have trained medical assistants to a high level of skill in information gathering and physical examination. They serve as the main information gatherers, entering historical and physical examination findings into the coupler. Since couplers define the universe of what needs to be done for each problem, the medical assistants become extremely skillful in these basic but important tasks. This use of medical assistants allows the physician or nurse practitioner to spend most of his or her time clarifying and annotating the history, checking certain physical findings, and, most important, reviewing the results of the coupler session with the patient to make decisions regarding possible diagnoses or management options.²⁰

Similarly, as described by Dr. Bartholomew:

In my clinic we have experimented with nurses ... doing the "pre-workup." This is the bulk of the time-consuming process. When I enter the examining room, I have a coupler [output] that is largely done. With a good nurse, a large portion of the common physical findings can be entered in the computer and a note left if something is in question. The physician then rechecks any physical findings that are positive or questionable. This is, in fact, extremely time saving and allows you to ... "become a consultant in your own practice."

Physicians who organize their practices to take advantage of knowledge coupling software may find significant economic benefits.²¹ And the rewards are more than economic. Use of knowledge coupling software not only lessens the burdens on the physician but makes practice more interesting and satisfying. As described by Dr. Bartholomew:

The physician in this setting enters the equation at a higher level of expertise, and, instead of spending the whole day gathering mundane data, spends much more time reviewing the complexities of the cases that need that extra caution to the patient's benefit. Furthermore, I must admit that the practice of medicine in this setting is simply more fun. It is more fun because it is more intellectually rewarding, and this is itself an excellent reason to be using couplers. By having this

Premises and New Tools for Medical Care Education, New York: Springer-Verlag, 1991, p. 240.

²⁰ Burger, Charles S., "The Use of Problem Knowledge Couplers in a Primary Care Practice", *Healthcare Information Management*, vol. 11, no. 4, Winter 1997, p. 19, available at www.PKC.com

²¹ Burger, C. "A Coupler Centered Practice: Business Case Analysis for Couplers," available at <http://www.pkc.com/papers/ccp3.pdf>.

extra time to spend on more complex cases, the physician can then begin to use the couplers to function at a higher intellectual level than a busy practice usually affords. ...

... The couplers are full of information; I have never failed to learn something new each time I have used one. ... Using couplers begins to become a reinforcing loop—because they are fun to use, you use them more. The more you realize that valuable information, beyond your own personal store of knowledge, is being brought to bear on each of the patient's problems, the more secure you feel, the more patient gratification you generate, and the more gratification you have from your own practice.²²

The foregoing addresses the feasibility and benefits of the combinatorial approach from the physician's perspective. But in the long run that is the wrong perspective. Medicine's division of labor needs to shift away from the physician and towards external information tools, non-physician caregivers, and patients themselves. That conclusion becomes increasingly obvious as we further address physician objections to the combinatorial approach.

2. Utility of detailed initial data

Physicians, even if persuaded that detailed data collection is feasible, question whether it is truly productive at the outset of care. The only advantage of detailed initial data, it seems, is identifying as many diagnostic or therapeutic options as possible. For any individual patient, most of those possibilities will turn out to be inapplicable (the outer circle of information, as described above). Physicians believe that their expertise enables them to leapfrog over the outer circle. They can rapidly identify the options of probable relevance to the individual patient. Those are the options worth investigating (the middle circle), and only limited data are needed to identify them, physicians believe. The important question, on this view, is whether objective evidence of probability, rather than subjective, variable physician judgments, should be the basis for identifying the options worth investigating. Whatever the answer to that question may be, expert physicians collect initial data selectively rather than exhaustively. If and when it becomes necessary to investigate improbable options (rare diagnoses, non-standard treatments), then more detailed data collection may be needed and specialists may be consulted. An iterative process of successive elimination thus occurs. On this view, collecting detailed initial data at the outset of care defeats a primary purpose of expert functioning in the initial workup—to avoid unnecessary data collection.²³

²² Ibid.

²³ This section draws on "The Database," in *Knowledge Coupling*, and on *The Philosophy, Use and*
06/27/08

This entire point of view rests on a mistaken premise—that detailed initial data collection is not needed to identify the options worth investigating. The reality, however, is the opposite. The only reliable way to determine the options worth investigating for an individual patient is first to collect patient-specific data in great detail.

Much of the detailed initial data will turn out to be irrelevant. But the relevant and irrelevant data cannot be distinguished in advance. They vary for each patient with the same presenting problem. Detailed data must therefore be gathered at the outset. In no other way is it possible to take into account all options of potential relevance, and narrow them down to options of actual relevance for an individual patient. These threshold inquiries are compromised when detailed data collection and analysis are deferred.

“The matter of time is essential in all estimates of the value of information,” Norbert Wiener once observed.²⁴ In medicine, the value of detailed information is greater at the outset of care than the value of the *same* information gathered piecemeal, over time. That reality explains the traditional precept (honored more in the breach than the observance) that a detailed patient history “is still the most essential part of clinical decision making. ... Extra time spent gathering data almost always saves much more time in the long run.”²⁵

This observation from 20 years ago is now true more than ever, because growing medical knowledge increases the utility of data available from a well-designed patient history (as well as physical examination and basic laboratory tests). As Dr. Ami Schattner has written: “Diagnostic difficulties can often be resolved by simple means. ... close attention to the safest, least expensive, and most informative test of all—the history—can help resolve diagnostic and therapeutic dilemmas weeks and sometimes months before they might otherwise be resolved.”²⁶ Moreover, the patient history, physical exam and basic laboratory tests are often needed to use advanced diagnostic tests productively, in both formulating inquiries and interpreting results. For example, as Jerome Groopman points out, the settings for multidetector CT scans (which enable rapid scans of large

Interpretation of Knowledge Couplers, note 2 above, co-authored and authored, respectively, by Chris Weed.

²⁴ N. Wiener. *The Human Use of Human Beings*. Boston: Houghton Mifflin, 1950, 1954; New York: Avon Books, 1967, p. 168.

²⁵ Mold, JW, Stein JF. The cascade effect in the clinical care of patients. *New England Journal of Medicine*, 1986. 314:512-14 (p. 513).

²⁶ Schattner A. “Simple Is Beautiful: The Neglected Power of Simple Tests.” *Arch Intern Med* 164: 2198-2200 at 2199 (Nov. 8, 2004).

volumes of tissue) can be adjusted to take into account not just the specific question posed by the referring physician but also the patient's history.²⁷

One would hope that economic pressures induce more selective use of costly diagnostic technologies and more effective use of traditional approaches. But, as Dr. Schattner elsewhere writes, precisely the opposite has occurred: "physicians have become 'fascinated', 'preoccupied' and 'obsessed' with their new instruments," which "are vigorously promoted by the large companies that supply them." This state of affairs is destructive at many levels:

With the increasing availability of powerful diagnostic instruments, physicians have become distanced from both their patients and the basic clinical data. Today, tests and procedures are considered infallible and ordered in increasing numbers – often almost blindly, repeatedly and sometimes even without examining the patient. Thus, many are redundant, inconclusive or misleading, in addition to being unnecessarily expensive. Uncertainty, false positive findings and fear of lawsuits often beget more tests or procedures and may trigger dangerous cascades. This testing-dominated approach undermines the value of clinical skills, which tend to become underestimated, underused and finally lost.²⁸

A primary reason for neglect of basic clinical observations and tests is their multiplicity. That makes it difficult to know or recall what tests are available for any given problem or which tests should be used when. Equally difficult is comprehending all the data generated. Accordingly, Dr. Schattner argues, fully exploiting information from the patient history "mandates a closely linked and thoughtful use of large, preferably electronic, databases." But Dr. Schattner does not address the core issues of how and when external databases are to be used. In traditional medical practice, the physician's unaided mind largely determines the content of the initial history during the patient encounter; afterwards the physician may go to medical libraries and electronic databases for external guidance, if time permits. This sequence is backwards. External guidance should be organized *before* and used *during* the three steps of the initial workup – choice, collection and analysis of patient data. That is what knowledge coupling software accomplishes. *Then* the practitioner and patient may supplement the software's output with additional

²⁷ Groopman J., *How Doctors Think*, note **Error! Bookmark not defined.** above, pp. 193-94.

²⁸ Schattner, A. "Clinical paradigms revisited," *Medical Journal of Australia*, 2006; 185 (5): 273-275, available at http://www.mja.com.au/public/issues/185_05_040906/sch10143_fm.html. See also Dr. Schattner's response to letters to the editor, criticizing "imaging without forethought." Ibid, 2006; 185 (11 / 12) : 671 - 672 [Letters], available at http://www.mja.com.au/public/issues/185_11_041206/arn11029_letters_fm-2.html It is remarkable that these issues are still the subject of debate. By now it should be obvious that caregivers need information tools to use clinical imaging tools effectively as much as they need the imaging tools to reveal internal organs. In both contexts, ignoring modern technology is unacceptable.

observations suggested by their personal knowledge, experience and judgment – that of the practitioner who may have seen many other similar patients, and that of the patient whose intimate personal knowledge and experience of her own condition may reveal crucial variations from what is expected.

The preceding discussion is framed in terms of diagnostic decision making, but a similar analysis applies to treatment decision making, especially management of chronic disease. Many problems require considering a multiplicity of treatment options, including expensive technologies, a variety of medications, other interventions and changes in patient behaviors. Evaluating the possibilities for each individual patient requires detailed data collection and analysis. That does not occur without the necessary standards and tools. As a result, patient needs are easily sacrificed to the short-term economic interests of other parties.

This notion that routinely collecting detailed data is useful may seem to contradict research in the psychology of decision making. Summarizing this research, Malcolm Gladwell states that in some situations "you need to know very little to find the underlying signature of a complex phenomenon." In such cases, Gladwell concludes, collecting and considering detailed data bury the few data points that matter in a mass of extraneous information.²⁹ This point is valid as far as it goes, but it overlooks two crucial elements. First, when many possibilities must be considered and each one has a different signature (e.g., multiple diagnostic possibilities, each with its own cluster of findings), then detailed data must be collected in order to consider all the possibilities. Second, finding the few data points that matter buried within detailed information is readily accomplished with software tools. Breaking our dependence on the unaided mind solves the needle in a haystack problem.

The Addison's disease case discussed in Part **Error! Reference source not found.** powerfully illustrates these points. Recall that the physicians began by considering only a limited range of diagnostic options, based on probabilities. As they explained it: "the clinician usually begins a diagnostic investigation by considering (and excluding) the *most common* diagnoses. As those most common diagnoses become less likely, many *less common* diagnoses are considered" (emphasis added).³⁰ Consistent with this accepted practice, an evidence-based ranking of diagnostic possibilities for fatigue would assign a low rank to Addison's disease, because that condition is rare, both in the general population and in the population of patients with severe fatigue.

²⁹ Gladwell, M. *Blink* (New York: Little Brown and Co., 2005), pp. 136-45.

³⁰ Keljo D, Squires R. Clinical Problem Solving: Just in Time. *N Engl J Med* 1996; 334:46-48 (p. 48).

But those broad populations were not relevant for this particular patient. Instead, relevant to her was the limited group of individuals with combinations of findings similar to hers. The initial findings on her included fatigue, shortness of breath, hypotension, weight loss and numerous, deeply pigmented moles—each of which are documented manifestations of Addison's disease. Soon after the initial encounter, additional known symptoms of Addison's disease appeared. In the limited group of patients with such a combination of findings, Addison's disease should be ranked high as a probable diagnosis. But in this case it was not even considered as a possibility until the patient was near death.

The reason for this delay emerges from two elements of the situation. These two elements occur routinely in both primary and specialty care:

- First, none of the findings recited above are specific to Addison's disease. On the contrary, the various findings each suggested numerous diagnostic possibilities. Collectively, these possibilities implicated almost every specialty in medicine.
- Second, faced with this complexity, the patient's physicians began their consideration with diagnostic possibilities that they viewed as "the most common." Addison's disease was apparently not considered because it is rare. This approach is accepted medical practice.

In beginning with the most common possibilities, physicians commit a basic conceptual error. Indeed, this approach is backwards. What is common or rare in the general population should be the last information to consider, not the first. The first information to consider *for an individual patient* should be all the possible diagnoses suggested by the patient's particular combination of findings, regardless of the probability of those diagnoses occurring in larger populations. Addison's disease would have immediately emerged as a highly probable diagnosis if the non-specific findings on this girl were *combined*. The underlying principle is clear: findings that are non-specific when viewed in isolation often become highly specific when viewed in combination. A corollary principle is that judgments of probability are highly misleading as applied to individual patients, because those judgments are derived from large population studies where the few variables examined become isolated from the detailed patient data needed for combinatorial analysis.

These principles give the combinatorial approach enormous power. But that power depends on gathering data in sufficient detail and then reliably matching it with medical knowledge. The reliability of both tasks depends on using software tools rather than the unaided mind to guide initial data collection and analysis.

Data collection guided by software tools is highly standardized. Why mandate that departure from accepted practice?

3. Utility of standardized initial data

Physicians object that the combinatorial approach does not allow for case-by-case discretionary judgments in selecting initial data. This view rejects a central element of the combinatorial approach—that *standardized* initial data must be determined in advance for a given presenting problem. In other words, for every patient who has that problem, the caregiver must habitually collect all data specified in advance, without any omissions based on the caregiver's clinical judgment. To many physicians, this standardization is mere "cookbook medicine."

This view is backwards. In reality, cookbook medicine results from our human propensity to process only information that supports our preconceptions. Each physician has a personal set of preconceptions, generating enormous variability. The only escape from this dilemma is to standardize initial data collection and analysis.

Paradoxically, *standardized* data collection is essential to capturing the *uniqueness* of individual patients. Their medical individuality emerges when data are collected from a predefined universe of standardized data with sufficient consistency and detail to show individual variations. In contrast, when physicians have discretion to shorten a standardized initial workup, key individual variations may be lost. Indeed, this loss of information is to be expected.

Consider, for example, the problem of diabetes management. An initial workup of a diabetic patient should cover some 330 findings. Each patient's set of positive findings will differ. And those variations among diabetic patients may be crucial to their care, because every finding is carefully selected in advance for its direct relevance to diabetes management. If physicians substitute their own judgments for the predefined initial workup, some of those judgments will result in omissions. And some of those omissions will reflect physician idiosyncrasies (time constraints, clinical prejudices, unsupported hypotheses) rather than accurate judgments of which findings can safely be ignored. In short, allowing judgmental *omissions* from a predefined initial workup reintroduces the very idiosyncrasies that patients need protection against. (Judgmental *additions* to the initial workup are a quite different matter, as discussed below.)

Some experienced physicians may be confident that they can judge when it is safe to abbreviate an initial workup. Their talent and experience, they believe, often enable them to recognize a clear diagnosis or treatment without performing a complete workup.

Those physicians will readily acknowledge, however, that their judgments are superior to the judgments of some other physicians with less talent or experience, and certainly superior to the judgments of non-physicians. Permitting those *other* practitioners to omit portions of the initial workup at their discretion will therefore leave patients at risk. Moreover, even the most self-confident physicians must concede that their superior judgments will not always be accepted on faith by patients or others. Indeed, patients or colleagues or payers or regulators might prefer to rely on initial workups conducted by non-physician caregivers, who would not presume to cut corners during the initial workup based on their personal knowledge or judgment.

In short, everyone involved needs some objective standard for identifying trustworthy judgments. The combinatorial approach provides the necessary objective standard—that is, detailed patient data selected in advance based on review of the medical literature. Patients or others who are unwilling to accept physician judgments on faith can simply demand completion of the predefined initial workup, even if the physician views it as unnecessary overkill. A self-confident physician should welcome that demand, because a complete initial workup presumably will vindicate his superior judgment—that is, the physician will have already considered anything significant the complete initial workup reveals.

Actual experience with the combinatorial approach reveals that the self-confidence of even the best physicians is often misplaced. The human mind, no matter how gifted and well-schooled, simply cannot be trusted with the intricate information processing that individualized decision making entails. When the medical literature shows that 300+ initial findings are needed to manage diabetes, for example, the physician's case-by-case judgments of when to dispense with some of those findings are bound to be fallible. Every patient's combination of findings will vary. Any individual variation may turn out to be significant. And this differentiation increases with chronic conditions, as each patient's evolving disease interacts with his or her unique physiology, psyche and circumstances over time.

Moreover, two other factors mandate rigorous enforcement of a combinatorial minimum standard. First, however justifiable some omissions from the initial workup might seem, they create an unexplained gap in the patient's record of care. Caregivers other than the initial examiner cannot distinguish between findings that were checked and found to be negative and findings that were simply omitted and never checked. The resulting uncertainty invites either wasteful duplication or uninformed follow-up.

Second, omissions contaminate the patient's record of care as a source of data for discovering new knowledge. The recorded care of many thousands of patients should

provide reliable data from which patterns can be identified, as we shall discuss further in Parts **Error! Reference source not found.** and 0.

These points are a reminder that medicine needs something like standards of accounting in the business world. In a business, every deposit and every payment of funds must be recorded. No one would attempt to judge when it is unnecessary to do so, because everyone understands that recording each deposit and payment is essential to maintaining controls over financial operations. Similarly, in medicine, rigorously collecting and recording every positive finding of a predefined data set is necessary to maintain reliable systems for decision making and feedback. A combinatorial approach to the initial workup makes this point obvious, because it means that each item of data has been carefully chosen in advance for its potential utility.

Here users may reasonably object that they need some flexibility in the timing of data collection for some items (particularly lab tests). It may not always be practical in some practice settings to immediately obtain each and every test result that the knowledge coupling software identifies as needed for the initial workup. The software accommodates this need by designating certain items as data to be collected "if available." When necessary, the user can defer collection of these items in the hope that other data will be enough to arrive at a solution. The key point to recognize is that the software should make this exercise of judgment by the user transparent. The missing data items are recorded as uncertain findings, not as negative findings. Anyone reviewing the Coupler output is thus alerted to the possible need to obtain the missing data.

4. Effect on the doctor-patient relationship

Whatever utility the combinatorial approach to the initial workup may have, physicians still object that it harms their relationship with patients. The combinatorial approach means that the patient interacts initially not with the physician but with a computer and with non-physician personnel who conduct the physical examination. This division of labor is unacceptable to some physicians and patients. They contrast an idealized scenario where the physician sympathetically inquires about the patient's history, personally performs the physical exam, observes the patient's demeanor and immediately orders tests and procedures he judges to be relevant. Through this interaction, it is believed, the physician develops personal rapport, observes firsthand the patient's physical and emotional condition, and establishes his immediate command of the situation.

Physicians are right that their personal relationships with patients are critical. At its best, the human interaction between caregiver and patient is itself therapeutic. But physicians' idealized image of their relationship with patients is hardly consistent with how the initial workup is usually conducted. Time constraints often prevent a sympathetic discussion of the patient's condition, cultural barriers may hinder communication, graduate medical education fosters behaviors and attitudes that interfere with communication, and not all physicians have good communication skills to begin with. In short, the personal relationship between the physician and patient is not necessarily a positive element. Moreover, that personal relationship provides no comfort to patients if not backed by professional expertise the patient can trust.

Patients need to be able to trust the physician to gather and correctly take into account the right information. An objective basis for trust arises when physicians employ a combinatorial approach to the initial workup implemented with software tools. In contrast, a judgmental approach to the initial workup hardly inspires patient trust. "The average patient visiting a doctor in the United States gets 22 seconds for his initial statement, then the doctor takes the lead,"³¹ according to one study. Other observers conclude that the physicians are deficient in bringing out patients' own concerns.³² This state of affairs fosters cynicism among patients. As expressed in the headline of a *New York Times* article on the subject: "Tell the Doctor All Your Problems, But Keep It to Less Than a Minute." The hurried, judgmental approach remains accepted practice even though research has "linked poor communication to misdiagnoses, the ordering of unnecessary tests, and the failure of patients to follow treatment plans."³³

Critics of current practice advocate improved interviewing techniques by physicians.³⁴ Yet, no improvements in interviewing will ever bring initial workups to an acceptable level of quality. So long as physician judgment determines the content of the workup, the inevitable outcome is enormous variation from one physician to another. Drs. John Bjorn and Harold Cross documented this phenomenon more than 35 years

³¹ Langevitz W, Denz M, Keller A, et al. "Spontaneous talking time at start of consultation in outpatient clinic: cohort study." *BMJ* 2002. 325:682-693.

³² Ibid.; Marvel MK, Epstein RM, Flowers K, Beckman HB. Soliciting the patient's agenda: have we improved? *JAMA* 1999. 281:283-287.

³³ M. Levine, "Tell the Doctor All Your Problems, but Keep It to Less Than a Minute." *New York Times*, June 1, 2004.

³⁴ Ibid.; Marvel MK, Epstein RM, Flowers K, Beckman HB. Soliciting the patient's agenda: have we improved? *JAMA* 1999. 281:283-287.

ago³⁵ when they recruited one of their patients to take her case to a number of physicians in their community. This patient found that each physician elicited different information and drew different conclusions, even though they were examining the same symptoms in the same person.

Since then, innumerable studies have further documented that variation among providers is the norm in numerous medical contexts (most of these studies do not examine the context of the initial workup, the importance of which is not generally recognized). These studies, however, typically do not compare responses of different providers to a single patient but rather compare provider responses to a single disease condition. Variation is conceived as departure from "evidence-based" guidelines for a particular disease condition. This concept does not account for the possibility that some variations by physicians justifiably reflect the varying needs of individual patients labeled with the "same" disease. In contrast, the Bjorn and Cross study, by comparing responses of multiple physicians to a single patient, clearly shows that variation reflects provider idiosyncrasies. This conclusion comes as no surprise to patients. In fact, the only consistency some patients experience is that multiple physicians repeatedly ask them for the same basic information—which is hardly the consistency patients want.

Patients and doctors are adrift. For them, the system fails even to define the optimal initial workup, much less disseminate and enforce it. This state of affairs is what harms the doctor-patient relationship. Patients see that physicians' initial workups are variable, ad hoc and incomplete. The patient is too often left wondering whether the physician considered all the relevant diagnostic or treatment possibilities, gathered the right data, correctly interpreted the data obtained, and carefully recorded what was done. Physicians themselves may have the same doubts about their own work, and that of their colleagues (which is one reason why patients find themselves repeatedly asked for the same data). The result is that patients shuttle from one costly specialist to another, because no one involved can tell whether the uncertainty they face reflects gaps in personal knowledge (which the next specialist might know enough to remedy), or whether that uncertainty reflects gaps in current medical science (which makes the next specialist consultation an exercise in futility).

This disorder undermines the therapeutic benefit inherent in caregiving itself. Medical attention satisfies a basic human need for relief and sympathy. But if patients

³⁵ Bjorn J, Cross H. *The Problem-Oriented Private Practice of Medicine*. 1970. Chicago: Modern Healthcare Press, pp. 24-28.

perceive disorder and lack of care in what is done to them, that perception alone may destroy the therapeutic benefit that medical attention itself potentially confers.

Unlike a judgmental approach to the initial workup, an orderly, combinatorial approach lays a secure foundation for a system. Both parties feel protected against the limits and idiosyncrasies of the unaided mind. The burden of uncertainty is reduced, and open communication is enhanced. The combinatorial approach brings both greater effectiveness and greater honesty to the physician's work. As Dr. Richard Rockefeller has written of his experiences in patient care using knowledge coupling software: "our sense of worth and competence is better served by improved outcomes in the realm of the possible than by compensatory fantasies of omniscience."³⁶ Similarly, Dr. Jerome Groopman (not a user of knowledge coupling software) has written:

Does acknowledging uncertainty undermine a patient's sense of hope and confidence in his physician and the proposed therapy? Paradoxically, taking uncertainty into account can enhance a physician's therapeutic effectiveness, because it demonstrates his honesty, his willingness to be the more engaged with his patients, his commitment to the reality of the situation rather than resorting to evasion, half-truth and even lies. And it makes it easier for the doctor to change course if the first strategy fails, to keep trying. Uncertainty sometimes is essential for success.³⁷

But without knowledge coupling software, it is difficult to distinguish between genuine uncertainty and mere personal unawareness of applicable knowledge.

A sociologist has done a survey of patient attitudes towards use of PKC knowledge coupling software in a primary care practice.³⁸ He found that a majority viewed their experience favorably, a significant minority were neutral and a significant minority viewed their experience unfavorably. Elements of the experience that were viewed favorably included thoroughness and depth of data collection, inclusion of personal lifestyle details bearing on diagnostic and treatment decisions, lessened reliance on the doctor's personal knowledge, objective presentation of decision options and evidence, and printouts of the detailed knowledge coupling results (which facilitate recall, understanding, follow-up, and discussion with family members and other providers). Unfavorable reactions appeared to reflect antipathy to computers and a preference for the familiar, personal questioning by an authoritative physician. The sociologist who conducted the survey wrote that some patients view the computer as "fostering an

³⁶ *Knowledge Coupling*, note 2 above, p. ix.

³⁷ Groopman J., note **Error! Bookmark not defined.** above, p. 155.

³⁸ Weaver R. "Informatics Tools and Medical Communication: Patient Perspectives of "Knowledge Coupling" in Primary Care," *Health Communication*, 15(1), 59-78 (2003).

'impersonal' environment" even though the software elicited detailed findings about the patient's condition and personal circumstances. Also problematic for a few patients is the fact that use of computers protects against the mind's fallibility: "Whereas most respondents applaud this, a few, ironically, seemed disheartened by the tacit admission of the mind's limitations, [which] seemed to diminish their confidence in the care and advice being delivered."³⁹ A contrasting view was expressed by one of the survey respondents, who commented:

I have had a couple of problems that a series of doctors failed miserably to diagnose. None of them ever picked up a book while I was in the office nor hinted that they consulted any reference. Perhaps their performance wouldn't have been so pitiful if they had used references. Computers can be a quick way to find and check information.⁴⁰

Dr. Ken Bartholomew, who pioneered use of PKC knowledge coupling software to implement the combinatorial approach, described the reactions of his patients as follows:

Not only do patients see the thoroughness involved in the use of couplers, but they sense that we care enough to give them the kind of thoroughness that they feel entitled to. With the coupler's systematic review of details in the patient's life that could be relevant to the current problem, the patient feels that his or her individual situation has been thoroughly examined and all possible conclusions have been taken into account. In management couplers, they further see the many different combinations of therapy and understand that the care of a complex, long term problem requires a detailed understanding of the patient's unique situation, followed by a careful monitoring of the options chosen. Even when a diagnosis is still in question, they have, in my experience, been completely satisfied with the outcome of the encounter. In addition, by receiving a printout of the findings and possible causes, they feel empowered to review the situation at home and to watch for signs and symptoms that may aid the diagnostic process in the days or weeks to come. The use of couplers teaches them that there is a time course to disease and that not all signs and symptoms necessarily occur "by the book" or simultaneously. By thus empowering our patients with information, as opposed to leaving them in a void, we reinforce their collaborative role as part of a team working toward an understood goal. ... it is only when this occurs that the optimum physician/patient relationship is built.⁴¹

Another physician user of knowledge coupling software, Dr. Richard Rockefeller, has further described the importance to patients of detailed data collection about their personal situation and experience:

³⁹ Ibid., p. 75.

⁴⁰ Ibid. p. 73.

⁴¹ Bartholomew K, in *Knowledge Coupling*, note 2 above, pp. 238-39.

In the prevailing medical paradigm, orientated as it is toward general knowledge, patients feel appreciated and well cared for to the extent that their problems are generic, that is, match population-based classifications of pathology. When the dimensions of their suffering and needs extravasate beyond these borders, as commonly happens, patients discover their idiosyncrasies to be sources of frustration and anger. They often find themselves alienated not only from a system which fails to meet their needs but also from themselves, to the extent they identify with the systems implicit devaluation of their uniqueness. [In contrast, the combinatorial] approach ... takes the set of attributes, historical circumstances and preferences which differentiate one individual from the next as central to, rather than as a troublesome distraction from, the important work of the therapeutic encounter. Patients' satisfaction is enhanced as their aptitudes and contributions – self-knowledge, willingness and ability to gather data pertaining to the problem, among others – are appropriately valued. Finally, as a reward for relinquishing the comfortable (but also dangerous, increasingly untenable and ultimately unfulfilling) illusion of being wholly provided for by an omnipotent parent, patients are afforded the human rewards of collegiality, including equal stature in the patient/doctor relationship and control over decisions affecting their health.

Written in 1991, long before current discussion of "consumer-driven care," "information therapy" and the like, Dr. Bartholomew and Dr. Rockefeller presciently described an ideal that the marketplace is beginning to recognize. But the recognition is incomplete, and so the necessary tools and standards of care are yet to be accepted.

5. Information processing, clinical judgment and the two stages of decision making

Our only remaining hope and salvation is to begin the whole labour of the mind again; not leaving it to itself, but directing it perpetually from the very first, and attaining our end as it were by mechanical aid.

– Francis Bacon⁴²

We have seen that clinical judgment is unreliable in both selection and analysis of initial patient data, and we have examined *selection* of initial data at some length. Here we further examine *analysis* of patient data.

a. Analysis as information processing

By "analysis" of patient data we mean a simple process of association between data items and corresponding medical knowledge – for example, the association between a

⁴² Bacon F. *Novum Organon* (1620), Preface to Second Part, available at <http://history.hanover.edu/texts/Bacon/novpref.htm>.

cluster of patient findings and a diagnosis explaining those findings, or between possible treatments for the diagnosis and patient findings bearing on suitability of each treatment.

Conceived in this way, data analysis involves raw information processing – establishing linkages between patient findings and a database of medical knowledge. That function is readily carried out with software tools used by all caregivers and patients themselves. This simple concept of analysis is to be distinguished from more complex processes involving logic and inference. In turn, both simple and complex analytical processes of decision making may be distinguished from instinctive, intuitive processes.

Physicians object that a simple process of association is crude and incomplete. The analytical sophistication involved in understanding pathophysiology is lacking. Moreover, even if software tools with that sophistication were designed, those tools could never capture the subtleties of personal interactions with patients and the resulting instinctive element of clinical judgment.

The reality, however, is that giving free rein to clinical judgment degrades analysis of initial data. Inevitably, physicians jump to conclusions based on skimpy data and limited personal knowledge. As Dr. Jerome Groopman has written of diagnostic analysis, for example, "research shows that most doctors quickly come up with two or three possible diagnoses within minutes of meeting a patient ... All develop their hypotheses from a very incomplete body of information."⁴³ Their analyses are similarly incomplete in the treatment context, where physicians often fail to consider much of the information needed for identifying and choosing among available treatment options.

The exercise of judgment naturally leads physicians to shortcut organized threshold processes and plunge into follow-up processes prematurely. Without a secure foundation, those follow-up processes are piecemeal, disorderly, risky and full of waste. Acting in a hit-or-miss fashion, physicians order tests and treatments as soon as possible. If they do not hit on a solution, more tests and treatments are tried. The escalating volume of data to ponder often becomes too large and disorganized to take it all into account, while the volume of medical knowledge relevant to interpreting all that data becomes too large to recall or comprehend. Coordination, follow-through and feedback frequently fall apart. Risk of mishap lurks every step of the way.

This quagmire traps both primary care physicians and the specialists they consult. A vivid illustration is the Addison's disease case discussed in part **Error! Reference source not found.** In that case, what was needed was not sophisticated clinical reasoning but simply pattern recognition – the association between the initial findings and the

⁴³ Groopman J. *How Doctors Think*, note **Error! Bookmark not defined.** above, p. 35.

correct diagnosis. Software tools would have been superior to expert judgment for that limited task.

Autonomous clinical judgment is not capable of performing that task reliably, no matter how well physicians are educated. Constraining clinical judgment is the mind's capacity for information processing. To compensate, the mind's normal modes of operation include various simplifying approaches (heuristics) that limit the information taken into account, as decades of research in cognitive psychology has demonstrated. In the last decade, thanks to the patient safety movement, medicine has woken up to the implications of this research for explaining why execution of medical decisions so often goes awry. More recently, implications for medical decision making itself have become the focus of attention.⁴⁴

But the weaknesses of clinical judgment go beyond the mental heuristics on which the mind relies to cope with complexity. The problem is also that judgment is idiosyncratic and personal. Except in the simplest matters, no two people take into account the same information in the same way. Idiosyncratic variations among physicians arise from varying individual abilities, varying medical backgrounds (training, experience, specialty orientation), other influences (emotional, cultural and financial), and other contingencies (the time available at the patient encounter, the stage of the patient's condition at that point, the patient's recall of needed information, the interpersonal dynamics between patient and caregiver, the sequence in which different specialists happen to be consulted).

Most of us try to use our judgment to recognize and overcome internal cognitive weaknesses and external influences. But our capacity to do so is just as limited as our capacity for raw information processing. The mental heuristics identified by cognitive psychologists "appear to be integral components of human information processing. As with visual illusions, awareness does not prevent us from being susceptible to their effects ..." ⁴⁵ Much the same can be said of other influences (financial interests, emotional needs, cultural preconceptions) that further distort judgment. Inevitably, reliance on judgment compromises analysis of initial data.

⁴⁴ Well before the patient safety movement made the issue prominent, however, the implications of cognitive psychology for medical decision making were discussed in the literature. See Weed LL., *Physicians of the Future*, note 8 above; Weed LL. *Knowledge Coupling*, note 2 above, pp. 8, 37-42, 212, 226-27.

⁴⁵ A Elstein. Heuristics and Biases: Selected Errors in Clinical Reasoning. *Acad. Med.* 74:791-793 (1999), p. 793.

b. The two stages of decision making and the proper role of judgment

What, then, is the role of judgment, if any, under the combinatorial approach? Recall from Part III.A our discussion of the two stages of decision making: (1) building an informational foundation designed to identify options for decision with the pros and cons of each option, and (2) choosing among the options. The first stage requires processing information—making associations between general knowledge and patient-specific data. The second stage involves applying judgment to choose among the options in light of the evidence. Judgment at this second stage involves logical analysis, intuition and personal values.

The role of human judgment is to *supplement* the first stage of decision making and *govern* the second. In the first stage, the predefined initial workup may be supplemented with additional information judged by the patient or caregiver to be relevant. But judgment should not be permitted to cut short the initial workup. The first stage must be completed, so that, in the second stage, judgment is informed and its basis is transparent. And in the second stage, it is the patient's judgment, not the physician's, that should be decisive.

If patient data are carefully chosen in advance, collected without omission and reliably coupled with medical knowledge at the initial workup, then the need for judgment in some cases is almost eliminated. That happens when the first stage reveals a clear solution to the patient's problem. A clear solution is one that anyone's judgment would accept. The Addison's disease case study discussed above provides an example. There, assembling the right information would have revealed a clearly correct diagnosis and treatment at the outset of care, with minimal judgment required. In such cases, the locus of decision making authority should not matter, because the right decision almost makes itself once the right information is assembled.

Significant judgment is needed only if uncertainty remains and opinions vary after rigorous information gathering and processing at the initial workup are completed. At that point, the patient and caregiver are usually faced with a range of possible diagnostic or treatment options to investigate further and choose among. Uncertainty means that no one option stands out as superior to all the others. Once a choice is made among the options, continuing feedback and adjustment over time may be needed (particularly in cases of chronic disease), taking into account new patient data coupled with relevant medical knowledge. These follow-up processes involve a series of choices and the exercise of judgment at each point, but judgment should be highly structured, not open-

ended. That is, ongoing judgments should be structured within processes that are rigorously organized and documented (the subject of part **Error! Reference source not found.** below) as well as continuously informed by reliable knowledge coupling.

c. **Physician objections to separate stages of decision making**

Physicians tend to reject compartmentalizing of the decision making process into separate stages. On their view, as data and impressions emerge from the very beginning of the initial workup, physicians believe they can and should employ their trained judgment to quickly form hypotheses and thereby determine what further data to obtain from the history, physical examination and laboratory tests, which then provide feedback on the hypotheses, and so forth. In this iterative process, physicians believe, continual exercise of informed judgment by gifted, highly trained and experienced experts is superior to a rigid combinatorial approach where pre-determined data are collected uniformly and linked to medical knowledge without resort to expert judgment. Moreover, on this view, it is an illusion to think that compartmentalizing the two stages of decision making protects against the pitfalls of human judgment. Judgment is inescapable throughout.

In the first stage of decision making, preexisting medical judgments by *someone* determine what options should be considered, what counts as evidence for and against those options, and what initial data should be collected as evidence. That someone, on this view, should be the treating physician—not the patient, not third party payers, not clinical researchers, and not the authors of texts, clinical guidelines or software tools. Only the treating physician's judgment, on this view, combines medical science, firsthand contact with the individual patient and firsthand experience with similar medical problems in many other patients. Only the treating physician has the expertise to integrate advanced medical knowledge and experience with each patient's individual needs.

Implicit in this point of view are different concepts of expertise. Before responding to this point of view, we should unravel these differing concepts of expertise in general and physician expertise in particular. It should be understood that these traditional concepts of expertise all assume that the human mind is the primary vehicle for information processing — for combining general knowledge with specific data. When external tools assume that function, they directly reveal new patterns and relationships that the human mind may only infer indirectly, if at all (see part 0 below). Thus, traditional expertise as discussed below is superseded, just as the stethoscope is superseded for diagnosing what chest X-rays reveal directly. (Accordingly, the following

discussion of traditional concepts of expertise is a digression, which may be of little interest to readers who accept the need for external tools to process complex information.)

Three schools of thought about expert decision making. Physicians naturally believe that their informed judgment involves a set of cognitive abilities (knowledge, analytical skill and developed intuition) that only they can acquire and apply. This belief in "the art of medicine" seems supported by a distinction recognized by psychologists and philosophers. They distinguish between explicit knowledge (associated with conscious, deliberate judgment) and tacit knowledge (associated with intuitive or instinctive judgment), expressed in the phrases "knowing that" and "knowing how," respectively.⁴⁶ Explicit knowledge involves factual information, principles and logical relationships that can be articulated. Tacit knowledge involves skills and perceptions that cannot be articulated ("we know more than we can tell"). Experts thus appear to acquire tacit knowledge that texts and software tools cannot reproduce. Indeed, expertise is sometimes defined as possession of specialized tacit knowledge. On this basis, it may seem plausible to believe that physician expertise cannot be reproduced in software tools, and that physician judgment should not be subordinated to the combinatorial standard.

This school of thought seems to be further supported by research showing the impressive powers of instinctive judgment. These powers have been studied in many contexts, not limited to professional expertise. Relying on this school of thought, Malcolm Gladwell's bestseller *Blink* argues that conscious, deliberate judgments are often less trustworthy than first impressions and snap judgments ("rapid cognition").⁴⁷ Gladwell suggests two reasons for this phenomenon. One is that deliberate decision making often buries the few key factors that matter in excessive information (this point misses the potential for external tools to uncover the few key factors that matter). Another reason is that deliberate decision making lacks the power of our "adaptive unconscious" to comprehend intangible, unarticulated factors of relevance (of course, this point misses the risk that unconscious heuristics cause crucial information to be overlooked).

By comparison with this school of thought favoring instinctive judgment, the culture of graduate medical education favors deliberate over instinctive judgment. Scientific

⁴⁶ Barbiero, D. "Tacit knowledge." *Dictionary of Philosophy of Mind*, <http://philosophy.uwaterloo.ca/MindDict/tacitknowledge.html>.

⁴⁷ Gladwell M. *Blink*, pp. 13-14.

understanding of pathophysiology is viewed as the core element of physician expertise. The acquired knowledge and analytical capacity involved in that expertise⁴⁸ is what Dr. Sherwin Nuland described as "every doctor's measure of his own abilities ... the most important ingredient in his professional self-image (see note 10 above). Some non-academic practitioners may disagree, viewing their developed intuitions as no less important than scientific knowledge and analysis, but they still view the latter as essential to their expertise.

A completely different school of thought points in just the opposite direction. This school discounts the value of expert decision making, regardless of whether expertise is attributed to deliberate or instinctive judgment. As summarized in a meta-analysis by professors William Grove and Paul Meehl, decades of studies in various fields (including medicine) have compared subjective, impressionistic, expert judgments with mechanical, algorithmic procedures such as multiple regression, weighted sums of predictive factors, and actuarial tables. These procedures combine items of data (e.g. findings on a patient) to arrive at predictive conclusions (e.g. diagnoses) in a formulaic manner. Grove and Meehl conclude that the mechanical procedures perform as well or marginally better than the judgments of expensive expert professionals. Based on this conclusion, they argue that the legal authority conferred on expensive experts in many contexts is not justified.⁴⁹ This school of thought suggests that human judgment should be minimized not only in the first stage of decision making, as argued here, but also in the second stage, contrary to what is argued here.

To summarize, we have described three alternative bases for medical decision making – (1) deliberate judgments based on explicit knowledge, (2) instinctive judgments based on tacit knowledge, and (3) mechanical substitutes for judgment – and three schools of thought about these alternatives. In academic medicine the prevailing school of thought is that deliberate judgment is the core element of physician expertise. The

⁴⁸ Some would argue that knowledge, not analytic capacity, is the primary basis for an expert's deliberate judgment: "Studies of expertise have repeatedly demonstrated that the expert is distinguished, not by the possession of any general skills, but by the ready availability from memory of appropriate knowledge to resolve the problem." Norman G R. Problem-solving skills, solving problems and problem-based learning. *Med Educ* 1988;22: 279-86. If knowledge is indeed the primary basis of expertise, then the claim of experts that they are irreplaceable is that much weaker. As a vehicle for knowledge retrieval, software tools are superior to the expert's mind, which cannot readily retain or retrieve from memory the variable bodies of knowledge relevant to real world problems.

⁴⁹ Grove W, Meehl P. Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: the clinical-statistical controversy. *Psychology, Public Policy and Law* 1996; 2:293-323, p. 316. See also Dawes R. *Rational Choice in an Uncertain World* (New York: Harcourt Brace Jovanovich, Inc. 1988), pp. 202-19.

second school of thought (not specific to medicine) is that the power of instinctive judgment is undervalued and deserves increased acceptance; some view this power as a defining characteristic of true expertise. The third school of thought, accepted by many cognitive psychologists, is that both deliberate and instinctive judgments are overvalued and very often should be replaced with formulaic, non-judgmental alternatives.

The necessity of distinguishing between the two stages of decision making. These various schools of thought about the nature of expertise, and physician experts who object to compartmentalizing the process of decision into separate stages, miss two essential points.

- First, different forms of human judgment and substitutes for judgment are each unreliable in the second stage of decision making (making choices) because important information is so often overlooked in the first stage (identifying available choices and the pros and cons of each). That explains why the studies cited by Grove and Meehl show the accuracy of human judgment to be only marginally different from non-judgmental substitutes, and why both are highly fallible. Their fallibility diminishes, however, when the right tools and standards of care are brought to the task of assembling information in the first stage of decision making.
- Second, what matters is the division of labor between judgment and external tools. At both stages of decision making, the mind's limited capacity to recall and process complex information in the first stage constrains expert judgment. Bypassing that constraint with simple software tools enhances medical decision making far more than sophisticated software that simulates the expert's flawed functioning within that constraint.⁵⁰

The ideal of continuous improvement. In our efforts to improve decision making, the guiding principle should be to continuously improve the *basis* for decisions in the first stage as a foundation for exercising judgment in the second stage. Improving the basis for decisions not only informs the exercise of judgment but also clarifies when to employ deliberate or instinctive judgment or a substitute for both. Above all, improving the first stage of decision making involves obtaining and acting on feedback over time—the key to continuous improvement of any complex activity.

⁵⁰ "Reengineering Medicine," note **Error! Bookmark not defined.** above, p. __. A similar point was recently made in Fenstermacher K. "The tyranny of tacit knowledge: What artificial intelligence tells us about knowledge representation," *Proceedings of the 38th Hawaii International Conference on System Sciences - 2005*, available at http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?tp=&arnumber=1385736&isnumber=30166.

In medicine, continuous improvement in the basis for decision making is not attainable unless a combinatorial approach is enabled with knowledge coupling software. Improvement may then occur at many levels. This is easiest to understand at the level of deliberate judgment, which can be articulated and captured in knowledge coupling software. But improvement may also occur at the level of instinctive judgment. Continual use of knowledge coupling software with many patients over time develops and enriches physicians' tacit knowledge. At every patient encounter, the combinatorial approach brings to bear a wealth of significant information going beyond the physician's personal store of knowledge and experience. Elicited with knowledge coupling tools, this broader information is not abstract or general but rather is directly relevant to the immediate problems of an individual human being. Caregivers who have this experience with thousands of their patients over time may develop instincts and intuitions that are unattainable in current medical practice.

Some physicians may choose to believe that their exceptional minds are capable of selecting and analyzing initial data as well or better than lesser minds equipped with external tools. These physicians, however, must admit that their expertise is not available to most of those who need it, nor is it transparent and subject to organized, continuous, verifiable improvement. And those physicians must admit that even their judgments may on occasion be wrong. Their exceptional ability may then cause them to take the patient faster and farther in the wrong direction than would otherwise occur.

Moreover, regardless of ability, a physician's personal experiences in patient care inevitably affect his or her judgments of what patient data should be collected and what the data means. (This becomes explicit whenever physicians justify their conclusions with the statement, "in my experience ... "). Yet any physician's personal experience is inevitably limited and randomly different from that of other physicians. Judgment should be informed and tested by evidence that may lie outside personal experience and that may contradict beliefs derived from experience. Merely personal judgment should not compromise either stage of decision making.

Enforcing this principle was the original goal of evidence-based medicine. But that goal is not achievable by evidence-based medicine in its current form. Evidence-based medicine compromises decisions by misplaced use of population-based knowledge for unique individual patients. The only way to extract some utility from population-based knowledge for patient care, and above all the only way to develop a more individualized body of medical knowledge, is to enforce detailed data collection on all patients, without the case-by-case exercise of judgment as to what data are necessary.

To permit judgmental departures from a combinatorial minimum standard increases each patient's exposure to wrong judgments. And a wrong judgment of initial data can be disastrous. Consider again the Addison's disease case. There, the initial data were sufficient to suggest the correct diagnosis. Yet, the physicians misjudged the data and thereby overlooked the correct diagnosis for months. Their patient almost died as they wandered down the various blind alleys their judgments suggested.

What is needed is a system designed to minimize erroneous judgments while incorporating prior, accurate judgments applicable to the problem at hand. This is traditionally a central function of experts – not just to exercise judgment but first to apply established knowledge – to filter out extraneous prior judgments and identify relevant prior judgments for solving the problem at hand – that is, moving from the outer to the middle circle of knowledge (see part I.F above). This research function is not feasible in settings where time is short, economic pressures are intense and information is in a state of disorder. Studies of professional expertise thus suggest that "an individual's ability to 'bring order to the informational chaos that characterizes one's everyday environment' determines whether that professional continues to perform competently."⁵¹

Bringing order to informational chaos is precisely what knowledge coupling software accomplishes. It does so far more efficiently than physicians could ever do. And once these basic information processing tasks are carried out, human judgment becomes far more accurate, efficient and powerful. Stated differently, the right software tools enable ordinary human judgment to *accomplish* what costly expert judgment can merely attempt.

To reiterate, judgment is empowered when (1) directly relevant information is presented in an organized form with extraneous information filtered out, and (2) time is available to consider that information with care. Knowledge coupling software is thus designed to filter and organize information, freeing up time for the caregiver and patient to consider it thoroughly. Moreover, the software incorporates pre-existing medical judgments that were reached under ideal conditions. Software builders have more opportunity than practicing physicians to exercise careful judgment in evaluating the medical literature. The literature in turn offers peer-reviewed judgments of leading authorities who similarly have more opportunity than ordinary practitioners to deliberate with care.

⁵¹ Pew Health Professions Commission, *Reforming Health Care Workforce Regulation*, San Francisco: The Commission, 1995, p. 26, quoting Pottinger, *Competence Testing As a Basis for Licensing: Problems and Prospects*, Washington, D.C., National Center for the Study of the Professions, 1977.

In some cases, that threshold step is itself enough to reveal a clear solution (the inner circle), as discussed. These cases are the low-hanging fruit. In more difficult cases, however, that threshold step reveals uncertainty and the need for more investigation.

That investigation begins a series of follow-up processes calling for repeated exercises of judgment. These judgments must be reached in a scientifically rigorous manner. What that means is the subject of part **Error! Reference source not found.** below. As we shall see, just as knowledge coupling software minimizes the role of expert clinical judgment at the initial workup, so rigorous scientific practices lessen the need for uncertain clinical judgments during these follow-up processes after the initial workup.

This notion of lessening reliance on expert judgment contradicts the expectations of highly educated people in general and physicians in particular. Advanced education teaches its recipients to rely heavily on their own minds in applying established knowledge. In no field is this more true than in medicine. Recall Dr. Nuland's words: "Our most rewarding moments of healing derive not from the works of our hearts but from those of our intellects."⁵² The activity of developing new knowledge, where the intellect has an essential creative role, is taken as the model for the activity of applying established knowledge. Professional education offers rewards of money and status for this reliance on personal intellect. But this state of affairs runs directly counter to the development of both modern science and market economies. The next part examines those developments in relation to medical practice.

⁵² See note 10 above.

VII. The nature of medical knowledge as applied to individual patients

Much of this paper has focused on the human mind's limited ability to recall and apply medical knowledge under real-world conditions. This constraint means that physicians process only a fraction of the medical knowledge relevant to patient needs. Beyond this gap between the human mind and medical knowledge is an equally important gap between medical knowledge and patient realities: our knowledge comprehends only a fraction of the complexity faced by caregivers and patients. Indeed, our "knowledge" is not just incomplete but in part is simply wrong.

As in other areas of science, medical knowledge is only a provisional approximation of reality. Medical knowledge must constantly be tested against that reality, in both patient care and clinical research. In patient care that testing process demands taking into account all potentially relevant knowledge and patient-specific data at the outset of care, and then carefully monitoring and adjusting whatever course of action is chosen, as we have discussed. In clinical research, that testing process demands continuously harvesting feedback on knowledge by carefully examining meticulous records of the application of knowledge to individual patients.

From this point of view, we consider different forms of medical knowledge, its relation to patient-specific data and how these issues are understood in current medical practice. We begin with the concept of evidence-based medicine.

A. The gap between evidence-based medicine and individual patient needs

... we must bring men to particulars, and their regular series and order, and they must for a while renounce their notions and begin to form an acquaintance with things.

— Francis Bacon⁵³

For centuries an unresolved debate has been carried on in medicine about the meaning of an individual patient's risk⁵⁴ and the concepts of induction and probability as applied in patient care.⁵⁵ During the past two decades the debate has provided the

⁵³ Bacon F. *Novum Organon* (1620), see note 1 above, Aphorism No. 36.

⁵⁴ Goodman S. Probability at the bedside: the knowing of chances or the chances of knowing? *Ann Intern Med* 1999;130:604-606.

⁵⁵ Hanckel F. The problem of induction in clinical decision making. *Medical Decision Making*. 1984;4:59-06/27/08

framework for critiques of managed care (in the U.S.) and evidence-based medicine (EBM).

"EBM and managed care share a common ethical and epistemologic focus on outcomes measured across populations."⁵⁶ Both epistemologically and ethically, this focus is misplaced in the context of patient care. As Chris Weed has observed, statistical information about population outcomes answers questions that patients and their caregivers do not ask:

Statistical answers are rationally useful only when one is interested not in individual cases but in regularities occurring in large numbers of cases. In other words, the patient is not in the position of a gambling casino. He or she is not in a position to say that if money is lost to the customers one day (the patient's illness is misdiagnosed or mismanaged), the odds are that it will be recovered several times over within a week. The patient may in a very real sense be "open for business" only for that day.⁵⁷

Epistemologically, outcomes measured across populations do not permit us to know the expected outcome for an individual who differs in relevant respects from the measured population—and relevant differences are the norm, not the exception. Ethically, population-based decision making rests on an "assumption that outcomes faced by individual persons can offset each other," as Asch and Hershey observe. That assumption "effaces the moral distinction between those persons." Outcomes among population members can be permitted to offset each other in veterinary care for livestock herds, but not in health care for human beings.⁵⁸

Therefore, in optimizing decisions for individuals, the uniqueness of each individual must be rigorously taken into account. This reality is now invariably acknowledged by evidence-based medicine proponents. But they underestimate the extent of individual uniqueness within populations, and fail to define a rigorous, systematic approach for taking that uniqueness into account. For example, one authoritative statement of EBM

68.

⁵⁶ Tonelli M. The Philosophical Limits of Evidence-Based Medicine. *Academic Medicine*, 1998; 73: 1234-1240 (p. 1238).

⁵⁷ Weed CC. *The Philosophy, Use and Interpretation of Knowledge Couplers*, note 2 above, p. 4. The quoted statement, written in 1982, predated the rise of managed care and evidence-based medicine, but it articulated the reason why physicians and patients distrusted those developments.

⁵⁸ Asch D, Hershey J. Why some health policies don't make sense at the bedside. *Ann Intern Med* 1995; 222:846-850 (p. 848). One reason for the population perspective is that so much of medical care is paid by third parties. A third party payer views cost-effectiveness in relation to the entire covered population. But individuals within that population view cost-effectiveness in relation to their individual needs, and health coverage (both group and individual) is purchased on that basis. The benefit of that bargain is denied to its intended beneficiaries when their individual needs are not even taken into account.

principles explains: "we must remember that recommendations can be made only for average patients, and the circumstances and values of the patient before us differ."⁵⁹ This seems to imply that average patients are medically similar, differing only in non-medical circumstances and values, and therefore that medical recommendations can apply across a population, subject only to individual adjustments for non-medical differences. The reality, however, is that virtually all individuals differ *medically* from the average. That is, they differ not just in their circumstances and values but in their medical characteristics bearing on the problem presented. The details of those differences, not population averages, must be the start and end points of individual decision making. Averages are useful, if at all, only for clues about which details are the highest priority for analysis.

Misconceiving individual variations as exceptions to population-based rules perverts the core element of evidence-based medicine, its "hierarchy of evidence." At the top of this hierarchy are randomized clinical trials (RCTs). Randomization deliberately excludes most individual variations, even those (such as co-morbidities) manifestly relevant to individual decision making. Other forms of population-based evidence, such as large observational studies and systematic reviews, similarly exclude or overlook potentially relevant data about individual variation. Evidence-based medicine thus directs clinicians to look first to population-based general rules and then to consider individual variations merely as exceptions to those rules.

This hierarchy is upside down. Knowledge about large populations is useless, indeed misleading, until other, more individually applicable knowledge is first taken into account. Thus the *sequence* in which evidence is considered is crucial. Recall from the Addison's disease case study in Part II.A how the physicians first looked for diagnoses thought to be common in the population of people suffering from severe fatigue, despite details about their patient revealing her to be quite unlike that large population and very much like the small subpopulation of people with Addison's disease. Recall further how the article characterized the patient's presentation of the disease as atypical, when in reality her presentation should have been seen as one of many possible expected variations within recognizable patterns of the disease. This example illustrates how the population-based perspective suppresses awareness of individual variation, and inhibits clinicians from taking into account patient-specific details.

⁵⁹ Guyon G. et al. (for the Evidence-Based Medicine Working Group). "Users' Guides to the Medical Literature. XXV. Evidence-Based Medicine: Principles for Applying the Users' Guides to Patient Care. " *JAMA* 284:10; 1290-1296 (Sep. 13, 2000, p. 1294)

More complete than this population-based perspective is a perspective that begins and ends with awareness of each patient's uniqueness. Continuing with the Addison's disease example, that disease can be defined in terms of a single abnormality – deficiency in adrenal-cortical hormones. Yet that one element can interact with unique individual physiologies in enormously variable ways. We thus can conceive Addison's disease in either population-based or individualized terms, that is, either (1) the basic element common to the entire population with the disease (the hormone deficiency), or (2) the variable ways in which unique individuals respond to the hormone deficiency and manifest the disease. The second, individualized conception is too detailed to be captured in textbook terms or comprehended by the unaided mind.⁶⁰

These principles apply in the therapeutic as well as diagnostic context. Consider clinical trial results for drug therapy. Such results are typically expressed as "average treatment effects" in the trial population relative to a placebo or alternative therapies. Yet, the average obscures heterogeneous treatment effects within the trial population, not to mention differences between that population and the populations that may later use a drug therapy.⁶¹ Thus, for an individual patient considering use of the drug, the population average is much less relevant than a detailed comparison of the patient's medical characteristics with those of two subpopulations: individuals who experienced a favorable outcome and those who experienced an unfavorable outcome. Suppose detailed comparison of careful medical records on the patient and careful records on the clinical trial population reveals that the patient closely resembles the subpopulation with a favorable outcome. In that case, the population average is essentially irrelevant to that patient. Decision making based on large population studies is becoming increasingly

⁶⁰ The utility of these two perspectives varies with the context. For example, in the therapeutic context, the population-based concept of Addison's disease is useful, because most patients with the disease can be successfully treated with simple hormone replacement therapy. In the diagnostic context, the individualized perspective becomes somewhat more important, because Addison's disease manifests itself in variable ways. Yet, Addison's disease should be relatively simple to diagnose, because recognizable patterns of the disease can be ascertained from organized data collection and analysis. Many conditions are much more complex to diagnose or treat, and even simple conditions can become difficult in patients who have co-morbid conditions or other complicating factors.

⁶¹ See McMurray J, O'Meara E. Treatment of Heart Failure with Spironolactone – Trial and Tribulations, 2004. *New Eng J Med* 351:6; 526-529; Juurlink D et al., Rates of Hyperkalemia After Publication of the Randomized Aldactone Evaluation Study. 2004. *New Eng J Med* 351:6; 543-551. These articles describe a large discrepancy between a favorable clinical trial of a drug for advanced congestive heart failure and unfavorable outcomes associated with increased use of the drug after the trial results were published and disseminated in practice guidelines. The articles analyze a number of foreseeable explanations for this discrepancy, including clear differences between the real-world population receiving the drug and the randomized trial population.

pointless now that we can analyze individual variation and small subpopulations at the level of genes and protein molecules.

The primacy of patient-specific evidence does not mean that population-based evidence is useless. On the contrary, once detailed initial data is taken into account, population-based evidence can then provide useful guidance to prioritize investigation. In the diagnostic context, where detailed patient data initially suggest several plausible diagnoses, it is then useful to know which of those possibilities are most common. Other things being equal, a common condition should be a higher priority for investigation than a rare condition. But using population-based knowledge of prevalence to select the highest priority diagnostic possibility is only an intermediate step. The next step is collecting new patient-specific data to confirm or rule out that possibility. Once collected, that new data supersedes population-based evidence as a basis for decision.

Similarly, in the treatment context, patient-specific data bearing on suitability of a treatment for an individual patient should normally supersede population-based guidelines about aggregate therapeutic efficacy. Efficacy often varies substantially among individuals within a population, including a randomized population in a clinical trial. That heterogeneity is obscured when population-based therapeutic outcomes are expressed as population averages. Individualized decision making requires taking into account patient-specific information bearing on four dimensions of individual variability:

the patient's (1) baseline probability of incurring a disease-related adverse event ("risk without treatment" or "susceptibility/prognosis"), (2) responsiveness to the treatment, (3) vulnerability to the adverse effects of the treatment, and (4) utilities [personal values and preferences] for different outcomes.⁶²

The above discussion suggests that randomized clinical trials in their present form are a poor substitute for rigorous medical practice as a source of evidence to inform decision making. In particular, analysis of drug safety and efficacy could become more accurate, timely, and affordable, if rigor was brought to medical record keeping and coupling of medical knowledge with patient data. In contrast, as Dr. Scott Gottlieb has observed, current regulatory approaches are futile:

The fundamental problem inside the FDA is ... the quality of information on which the FDA can base its evaluations. Today, the data that medical reviewers receive in conjunction with the process for approving new products are from highly structured clinical trials, carried out on homogenous populations of patients that

⁶² Kravitz R, Duan N, Braslow J. Evidence-Based Medicine, Heterogeneity of Treatment Effects and the Trouble with Averages. *The Milbank Quarterly*, 2004, 82:4; 661-687 (p. 662); Erratum at <http://www.milbank.org/erratum4-4.pdf>.

are carefully screened and preselected and then given new drugs under special protocols. There is little chance that such trials will ever provide a complete review of how a new treatment will perform when it is used in much broader populations of patients in real-world clinical settings, where patients do not always take their medicines on time or at all; where patients might have other medical problems or be of advanced age or in frail health; and where they have comorbidities or unusual diets, or they fill prescriptions for medications or dietary supplements that interact with one another, subtly or otherwise.⁶³

Dr. Gottlieb goes on to point out that clinical trials large enough for validity are extraordinarily time-consuming and costly, and it is sometimes impossible to recruit patients with the desired characteristics for such trials. Moreover, as few as one percent of all adverse drug events are estimated to be reported by physicians.⁶⁴ Thus, passive reporting systems that are dependent on physician initiative and not integrated with medical practice are insufficient for harvesting reliable information. Instead, practitioners and patients must function with information tools from which researchers can harvest new medical knowledge as a byproduct. This dual purpose for well-designed information tools can create enormous synergies.

The population-based perspective distorts not only clinical decision making and research but also concepts of quality control. Consider, for example, the numerous studies documenting large geographic variation in utilization rates of health care services without any corresponding variations in medical needs in the respective geographic areas.⁶⁵ These studies raise two basic questions: (1) what utilization rate of the services in question is appropriate, and (2) how can appropriate usage be enforced uniformly? At first glance, these questions may seem readily answerable by population-based analysis. For example, in some cases the answer to the first question is that variations may seem clearly attributable to provider-induced demand (because of correlation between utilization rates and the supply of providers). The answer to the second question might be to counteract provider-induced demand by using evidence-based practice guidelines to restrict provider discretion.

⁶³ Gottlieb S., *Opening Pandora's Pillbox: Using Modern Information Tools To Improve Drug Safety*. note **Error! Bookmark not defined.** above, p. 939.

⁶⁴ *Ibid.*, pp. 939-41.

⁶⁵ See generally Congressional Budget Office, *Geographic Variation in Health Care Spending*, Feb. 2008, available at <http://www.cbo.gov/ftpdocs/89xx/doc8972/02-15-GeogHealth.pdf> See also Baker L., Fisher E., Wennberg J. Variations in Hospital Resource Use For Medicare And Privately Insured Populations in California, Feb. 12, 2008 at <http://content.healthaffairs.org/cgi/content/full/hlthaff.27.2.w123/DC1> (evidence that the amount of resources used in the care of chronically ill patients varies widely across hospitals, regardless of the type of insurance coverage).

But such conclusions beg the most important question—how do we improve the basis for decision making, so that patient care decisions will be optimal for each patient, regardless of geographic area? Evidence-based practice guidelines derived from large population studies are no answer to this question because they do not take into account relevant, patient-specific details. The only way to take those details into account is to employ external tools designed for selection and analysis of detailed data in light of medical knowledge. If individual decisions are optimized in that manner, and if the personnel, procedures, drugs, devices and facilities used to execute those decisions are reliable, then whatever level of aggregate usage results from those decisions is appropriate, regardless of whether or how usage varies geographically. Consider again the transportation system analogy discussed in part I. If roads and traffic systems are carefully monitored and-maintained in good working order, and if various routes and modes of transportation are available for individual travelers to choose from, and if travelers are well-informed about options available, then travelers will make whatever choices best fit their situations. The transportation system should provide a range of choices that fits the existing range of individual needs, but the system should not predetermine and dictate some particular pattern of collective choices as the best outcome.

The above discussion has contrasted two points of view: (1) that individual variations are exceptions to the population-based norms, and (2) that individual variations are the norm, and thus population-based knowledge may not be relied upon to guide individual decision making. Rejecting (1) in favor of (2), as this paper does, may seem to some readers to be an unrealistic exaggeration. The following section addresses that issue.

B. General knowledge and individual uniqueness

1. Two forms of medical knowledge

Medical knowledge is often conceived as generalizations applicable to populations of individuals. This population-based concept of medical knowledge is not complete. Medical decision making routinely demands taking into account individual variations from what population-based knowledge leads us to expect.

a. Population-based knowledge

Consider first what is meant by the population-based concept of medical knowledge. Some of the great advances in medical science are applicable to large populations without regard to variations among individuals. Salk's development of a polio vaccine, for example, involved a single therapeutic agent that prevents polio in virtually everyone. In that context, differences among individuals have no significance except in the very few cases where the vaccine may not be safe or effective. Similarly, even in contexts where individual variations have critical and pervasive significance, individual needs cannot be addressed without applying general principles of pathophysiology applicable to large populations. All diabetics, for example, despite their enormous variability, have in common some dysfunction in hormonal regulation of blood glucose levels. Indeed, diabetes is *defined* in those terms. Medical knowledge is thus naturally conceived as applicable to large populations.

This population-based concept of medical knowledge also applies to variations within a large population when multiple individuals with a particular variation or set of variations in common are grouped into a definable subpopulation. For example, the population of diabetics can be grouped into myriad subpopulations, ranging from extremely broad groupings (Type I or Type II diabetics) to narrower ones (Type I diabetics with chronic renal insufficiency and cardiovascular disease). Or the population of individuals with a fatigue condition can be grouped into approximately 160 diagnostic subpopulations (such as individuals with Addison's disease), as explained in part **Error! Reference source not found.** above.

Population-based knowledge may be either explanatory (principles of pathophysiology) or descriptive (statistical information without an explanation for what is described). In either case, population-based knowledge is fallible. Consider, for example, clinical trial results showing statistically significant favorable outcomes for a drug in a randomized population of individuals with a given disease and without comorbidities. The outcome data are descriptive; they do not explain why the drug is successful for some individuals and not others in the clinical trial population. Nor does outcome data explain when the drug is appropriate for individuals with comorbidities or other differences from the clinical trial population. Nor does outcome data explain when off-label use of the drug is appropriate. Pathophysiologic explanations to answer these questions may be offered. But the answers must be tested against the results of actual use with patients. And those results will not yield trustworthy feedback without meticulous medical records providing detailed data about the patients receiving the drug.

b. Knowledge about individual variation

Population-based knowledge as just described has become extraordinarily voluminous and complex. But a further category of medical knowledge exists – medical knowledge about variations at the individual level. This occurs with individuals who seem anomalous in some way, that is, whose characteristics depart unexpectedly from some relevant subpopulation. Such knowledge often takes the form of case reports such as the *New England Journal* article we have discussed, where the patient's presentation of Addison's disease was described as atypical. In addition to such reports, medical knowledge of variations at the individual level could, in principle, take the form of patient data in medical records. But those data are not accessible for research purposes in the way that published case reports are accessible, and therefore cannot be considered part of the body of medical knowledge (although that could change, as discussed below).

By comparison, genomic and proteomic information about individuals may become accessible for research purposes and therefore become part of the body of medical knowledge.

Variations at the individual level are usually viewed as exceptions from the norm. The norm, we are taught, is "the textbook case," the model that other cases are expected to resemble. Yet, "the textbook case is so rare that everyone runs to look at it in the medical center when it is found," as Dr. Ken Bartholomew has observed. In reality, the textbook case is the exception, and variations from it are the norm. "Patient presentations are not one textbook scenario but thousands of similar, yet unique, combinations of presentations that our experience enables us to categorize." The dilemma is that *personal* experience is limited, and the mind categorizes experience simplistically. The textbook case thus becomes "a self-fulfilling prophecy"; unexpected variations from it are not searched for or recognized.⁶⁶ The Addison's disease case study discussed above illustrates this phenomenon.

Individual case reports are important to the progress of medical science. Reports of variations from known patterns can lead to uncovering new diseases, rare manifestations of known diseases, drug side effects, and new understanding of disease mechanisms.⁶⁷

But poor feedback mechanisms block this progress at multiple levels. The unaided mind, following textbook knowledge, may not even recognize variations; or variations may be misclassified; or medical records may too incomplete and unstructured to generate

⁶⁶ Bartholomew K. "The Perspective of a Practitioner," in Knowledge Coupling, note 2 above, p. 240.

⁶⁷ Vandembroucke J. In Defense of Case Reports and Case Series. 2001. *Annals of Internal Medicine*. 134:330-334 (Feb. 20, 2001).

adequate case reports; or published case reports may not be considered by caregivers or patients when those reports are directly relevant.

To reiterate, no caregiver can be aware of all the known patterns to consider, much less variations from those patterns. Pattern recognition in the face of such complexity requires tools external to the mind. Once those tools are employed, the mind develops new perceptions of both the patterns and the individuals with whom the patterns are compared. The next section describes those new perceptions.

2. The concept of individual uniqueness

As the Addison's disease case illustrates, medicine involves pattern recognition. The dilemma is that known patterns are rough generalizations about large populations, and as such are usually an imperfect fit with unique individuals. Every individual is a unique combination of myriad similarities to and differences from other individuals. What constitutes a similarity or difference depends on the particular diagnostic or therapeutic context. The similarities mean that different individuals can be medically classified together in the same category – a trait or set of traits in common with other individuals. The differences mean that various individuals classified in the same category are nevertheless different from each other in various respects that may be key to solving the medical problem they seem to have in common.⁶⁸

The similarities and differences arise initially from each individual's unique genetic heritage and unique developmental history. Each individual is a recombination of pre-existing biological elements, which are built into an enormously complex set of interconnected structures and interacting processes. The recombination of elements is not static but continuously evolving, subject to both internal and external forces. An important internal force is the human body's extraordinary capacity for self-regulation (known as homeostasis) and self-repair. As a result, the normal physiology of healthy persons become increasingly differentiated over time.

This complexity increases by orders of magnitude when normal physiology is disrupted by pathophysiologic processes, psychological processes, the physical environment, the social environment and medical interventions. These interacting elements introduce multiple layers of complexity and disruption. As stated by Dickinson Richards, in words that apply both to human physiology and its external environment:

⁶⁸ The following discussion draws heavily on Weed, CC. *The Philosophy, Use and Interpretation of Knowledge Couplers*, note 2 above.

Man's power became ever greater, but this curiously made matters worse not better, because his power became too great for his understanding, and moved even further beyond his awareness of consequences. ... Man's unbridled use of his technological armament throws whole segments of the natural order out of balance, with the full meaning of this obscure, the outcome unknown.⁶⁹

These realities mean that each person's illness will be a unique course of events, never precisely reproduced in any other person. Chronic illness in particular becomes highly personalized in this way. Consequently, when different individuals are labeled with the same illness, their medical condition and therapeutic needs may in fact differ radically. Diagnosis and treatment of each person's illness must take into account myriad resemblances to and differences from many other persons' experiences of the "same" illness. Doing so far exceeds the capacity of the human mind.

None of these points are surprising. Indeed, they are consistent with intuition, common experience and basic scientific knowledge. But without the tools needed to act on these points, their implications elude us.

C. Some implications of individual uniqueness

She could not eat or sleep, grew visibly thinner, coughed, and, as the doctors made them feel, was in danger. They could not think of anything but how to help her. Doctors came to see her singly and in consultation, talked much in French, German and Latin, blamed one another, and prescribed a great variety of medicines for all the diseases known to them, but the simple idea never occurred to any of them that they could not know the disease Natásha was suffering from, as no disease suffered by a live man can be known, for every living person has his own peculiarities and always has his own peculiar, personal, novel, complicated disease, unknown to medicine – not a disease of the lungs, liver, skin, heart, nerves and so on mentioned in the medical books, but a disease consisting of one of the innumerable combinations of the maladies of those organs.

– Leo Tolstoy⁷⁰

⁶⁹ Richards D. "Hippocrates and History: The Arrogance of Humanism," *Hippocrates Revisited*, note 9 above, pp. 25, 26.

⁷⁰ Tolstoy L. *War and Peace* (1869), Book Nine, Chapter 16 (Maude trans.).

And generally, let this be a rule, that all partitions of knowledge be accepted rather for lines and veins than for sections and separations; and that the continuance and entireness of knowledge be preserved. For the contrary hereof hath made particular sciences to become barren, shallow, and erroneous, while they have not been maintained from the common fountain.

– Francis Bacon⁷¹

To reiterate, each person can be viewed medically as a unique combination of resemblances to and differences from other persons. This reality is significant epistemologically. The unique combination embodied in any individual is "unknown to medicine," in Tolstoy's words, for three reasons. First, "every living person has his own peculiarities" in how a known disease occurs and manifests itself (as the Addison's disease case study illustrates). Second, different persons have different combinations of known diseases (this difficulty was absent from the Addison's disease case study). When two or more known diseases combine in a single person, their effects are likely to interact with each other and with the person's unique physiology so that the disease processes and manifestations to some extent depart from prior knowledge about the interacting elements. Third, disease may involve elements unknown to medical science, such as newly evolved pathogens or unidentified disease processes, interacting with known elements and the person's unique physiology. Thus a person's total medical condition can be regarded as a single, aggregate, new disease entity, described by Tolstoy as that person's "own peculiar, personal, novel, complicated disease, unknown to medicine."

Within that person's complex, unique pathophysiology are elements that medical science is able to understand and/or manipulate therapeutically, to some degree. This scientific knowledge is expressed in terms of elements occurring across a population, abstracted from the unique individuals to whom they apply. But effectively applying that general knowledge to specific individuals very often requires delving into their uniqueness.

How can caregivers and patients apply population-based knowledge to unique individuals in an organized way without being overwhelmed by complexity? And how can clinical researchers and basic scientists systematically test existing knowledge and harvest new knowledge from caregivers' daily encounters with unique patients? The

⁷¹ Bacon F. *The Advancement of Learning*, IX(1), at <http://www.fullbooks.com/The-Advancement-of-Learning2.html>

with those other items. The problem list is thus simultaneously holistic and reductionist, providing the benefits of both perspectives. As described by Dr. Ian Lawson:

the POMR system is affecting nosography, the way in which disease and disability are described. ... the interrelationships of problems are as important as the individual problems themselves. Symptoms and problem profiles, rather than summary diagnostic labels, often prove more sensitive in therapeutic management and may eventually lead to a different kind of care organization and epidemiology.⁷³

With the ongoing revolution in genomics and proteomics, the myriad resemblances and differences among individual human beings are becoming far more sharply defined at the molecular level. The massive scope and intricacy of this knowledge, and its infinitely variable applicability to individuals, make it increasingly obvious that the minds of highly educated physicians cannot be relied upon to recognize the patterns that define unique individuals and their medical needs. In that environment, we will heed Bacon's warning not to "falsely admire and extol the powers of the human mind," and we will embrace the use of external tools to empower the mind.

Both the human mind and external tools use language to reference clinical concepts. Lack of precision and consistency in the use of language has long been recognized as an obstacle to semantic interoperability among disparate health information technologies, particularly electronic health records. Accordingly, major efforts have been underway for many years to develop standardized medical terminology, taxonomies of medical concepts and corresponding coding systems. These efforts, however, valuable as they are, leave unresolved the problem of unstructured clinical judgment by physicians. For example, using standardized terminology to record the results of an initial workup does not assure that the contents of the initial workup will be complete or accurately coupled with medical knowledge. Assuring those goals requires some form of knowledge coupling tools as described above. Standardized terminology and coding is pursued most fruitfully when it is driven by needs that arise in developing knowledge coupling tools and using those tools in medical practice.

Knowledge coupling tools are derived from and linked in detail to an underlying electronic "Knowledge Network," which in turn is derived from and linked to the medical literature. The Knowledge Network is composed of explicitly defined medical entities

⁷³ Lawson I. Comments on the POMR. In Driggs MF ed. *Problem-Directed and Medical Information Systems*. New York: Inter-Continental Medical Book Corporation, 1974, p. 40. Dr. Lawson goes on to observe: "More to the point, this will also create immediate conflicts with "third party" agents and their prototype definitions of "eligible" illness. Indeed, the sooner their computer experience gets wise to (or gets "blown" by) the realities of multiproblem interrelational analysis and management, the better for us all."

(findings, disease conditions, medical procedures) and interrelationships among those entities. Standardized terminology and coding of these entities and relationships, when mapped to the Knowledge Network and applied through knowledge coupling software, has far more utility than the same standardized terminology and coding employed in an unstructured manner by physicians. For example:

- Building a module of knowledge coupling software (a "Coupler") for a given medical problem involves defining relationships among numerous medical entities (for example, a disease entity and the finding entities used to diagnose the disease). As the Knowledge Network is built up, it naturally reveals new entity relationships that may not be apparent from reviewing the medical literature (for example, a diagnostic entity may also be a finding entity in relation to another diagnosis, so that the original finding entities indirectly are related to the second diagnostic entity). When a new Coupler is built for a different medical problem, the existing Knowledge Network is traversed for relevant entities and relationships among them (to continue the example, the indirect relationship between the finding entities and the second diagnostic entity might be relevant to the building of the new Coupler, though not to the first Coupler). Traversing the Knowledge Network may thus reveal connections that medical literature searches would not reveal or reveal only with great difficulty.
- The Knowledge Network content includes not only entity and relationship information but also (1) classifying information that allows entities and relationships to be grouped and retrieved in various ways, and (2) textual explanation of the significance of the entities and relationships. The Knowledge Network thus provides a highly efficient repository of distilled medical knowledge useful for constructing clinical guidance tools—the Couplers.
- Couplers organize information from the Knowledge Network in a problem-oriented manner, that is, in a manner relevant to the specific problem-solving context for which the Coupler is built. The ability to traverse the Knowledge Network for knowledge relevant to the problem context makes it possible to partially automate the process of building Couplers. But this automated element is only the first step. Much thought goes into selecting and editing Knowledge Network content to make it maximally useful in the problem context addressed by the Coupler.
- The volume of medical literature is growing exponentially, much more so than the volume of actionable medical knowledge within the literature. A primary reason is that each article devotes much space to explaining the context and significance

of its subject matter, so that related articles consume much space with overlapping explanation. Yet, that explanation does not usually address all of the various medical specialties to which the subject matter is relevant. In contrast, knowledge coupling software and its underlying Knowledge Network enables each piece of actionable knowledge to be expressed concisely and then viewed as needed in the countless medical contexts and specialties to which it may become relevant. Enabling such access to directly relevant knowledge in specific problem solving contexts is far more efficient and effective than ordinary medical literature as a mechanism for storage, retrieval and transmission of medical knowledge.

- Enormous time, money and talent are currently invested in graduate medical education, publication of medical literature, medical libraries, and conferences. All these are mechanisms for transmitting medical knowledge to caregivers using that knowledge. The voltage drop in these transmissions is enormous. Patients have no assurance that the information residing in the minds of their caregivers corresponds to their individual medical needs. In contrast, knowledge coupling tools and the underlying Knowledge Network make it possible to reallocate scarce resources from futile attempts at transmitting medical knowledge to productive *use* of knowledge.