

Four-part series “Topics in Healthcare Innovation” – *Forbes.com*

By David A. Shaywitz, MD, PhD - <http://blogs.forbes.com/davidshaywitz/>

Contents: (full text follows on subsequent pages)

Part I: What silicon valley doesn't understand about medicine

Link: <http://blogs.forbes.com/davidshaywitz/2011/06/17/what-silicon-valley-doesnt-understand-about-medicine/>

Part II: Standards and medical practices: a cautionary prognosis for algorithm-based care

Link: <http://blogs.forbes.com/davidshaywitz/2011/06/21/standards-and-medical-practices-a-cautionary-prognosis-for-algorithm-based-care/>

Part III: Why medicine's killer app is an e-textbook

Link: <http://blogs.forbes.com/davidshaywitz/2011/06/22/why-is-medicines-killer-app-an-e-textbook/>

Part IV: Our metrics fetish – and what to do about it

Link: <http://blogs.forbes.com/davidshaywitz/2011/06/23/our-metrics-fetish-and-what-to-do-about-it/>

See below for full-text.

Part I: What silicon valley doesn't understand about medicine

Link: <http://blogs.forbes.com/davidshaywitz/2011/06/17/what-silicon-valley-doesnt-understand-about-medicine/>

Where's Silicon Valley when we need it? The inexorable rise of healthcare costs has created not only a crisis waiting to happen, but also an urgent need for innovation – exactly the sort of thing for which Silicon Valley is justly famous. So where is it?

While entrepreneurs in the Valley and elsewhere have achieved remarkable success (at least historically) developing new drugs and devices, more fundamental innovation in healthcare delivery – particularly the sort designed to take costs out of the system – has proved far more elusive.

It's not for lack of trying. Almost every legendary Silicon Valley entrepreneur (see [here](#); also [here](#)) has recognized the opportunity: our healthcare system is legendary for its inefficiencies, complexity, and opacity – all areas that seem utterly amenable to a technology-driven solution. Yet, to date, the results have been middling at best. Why hasn't there been a Google or a Facebook in the healthcare space?

One reason may be that existing stakeholders are more entrenched than generally appreciated; we may also be more reluctant to tinker with our health than we are to try out the latest app. It's also likely that at least until recently, entrepreneurs were more incentivized to create expensive solutions than to engineer cost-effective ones (as [this](#) McKinsey report suggests).

Yet I'd argue that the most significant disconnect I've found here in Silicon Valley is a failure to understand the nature of the problem to be solved. Most engineers and computer scientists with whom I've spoken conceptualize medicine as primarily a rational, evidence-based, problem-solving enterprise focused on well-defined conditions, rather than a discipline that in my experience owes more to scientism than science, is far more ambiguous than most engineers tend to recognize, and is founded on relationships, connectedness, trusted advice, reassurance, and frequently the off-loading of significant responsibilities from patient to doctor. Perhaps the physician's greatest responsibility is helping to compassionately mediate the patient's experience of illness.

Simply put, medicine is so much more than simply a diagnosis and treatment (although regrettably in today's world of rapid-fire outpatient medicine, it occasionally seems to be reduced to this cartoon), and patients — especially those with chronic illnesses — value more than just the rigid outcome measures traditional metrics tend to capture. A novel technology platform that overlooks the integrated needs of patients or underestimates or fails to account for the complexity and messiness of illness as it actually occurs and is experienced by patients (and those closest to them) will inevitably fall short.

Former Google CEO Eric Schmidt represents the pervasive Silicon Valley perspective in [this](#) dialog with surgeon and author Atul Gawande. Stanford Medical School professor (and author) Abraham Verghese presents a contrasting perspective in [this](#) recent NYT Op-Ed.

This disconnect between the reality of medicine and its Silicon Valley conceptualization is especially important because technological innovation — the sort of ideas the valley was founded upon — is urgently needed in medicine, and offers the potential to radically disrupt the way we think about health and disease, and can make it easier for us to stay healthy and understand what may be ailing us when we're sick. There's also an urgent need for improved communication between providers so that successful approaches can be shared, and less effective approaches identified and called-out (highlighted in [this](#) BCG white paper). There's also a growing need to adopt what might be called a network or systems view of medicine, and develop ways of integrating complex data over multiple domains to uncover new patterns and suggest new approaches to diagnosis or treatment.

An especially useful starting point for an aspiring healthcare company would be to think about ways to deliver a “positive good.” Most contemporary medicine is generally regarded, economically, as a negative good — something we use or engage in because we have to, not because we especially want to. While this creates a ready market, it's also results in a tense and uncomfortable relationship between the “producers” of healthcare — who have often invested a significant amount of resources with the intention of making a profit — and consumers, who have little choice but to pay (or co-pay) for expensive care, and consequently tend to resent it.

Back when I was in the clinic, I remember feeling confused by patients who would complain about the cost of care while seeming to have few qualms about the going on pricey vacations or buying expensive gadgets. On reflection, however, I think the key

difference is that patients at least felt they had a choice about buying a gadget, and liked selecting something they thought they might enjoy. The opportunity here for technologies companies is to develop an attractive and engaging platform that will warmly incentivize positive behaviors, and help patients nudge themselves in healthier directions. The business case hinges upon its voluntary adoption, rather than obligatory usage.

Consequently, I'd argue there's a significant opportunity for a platform focused relentlessly (think Jeff Bezos) on the patient, providing patients with significant value and benefits from engagement that are both immediate and accretive. Jamie Heywood's [PatientsLikeMe](#) has achieved much deserved recognition for the ALS patient community it has developed, and may be one model to consider; the question at this point is whether there's it's possible to achieve this level of community benefit and participation for less devastating diseases, especially considering that most patients want to live life, not spend their days fixated on their particular medical condition.

It will also be important to ensure that even as we recognize – and seek to capture, leverage, and ultimately monetize – the value associated with the collection of an ever-increasing amount of data, we also recognize that most people don't want to be perpetually monitored (at least not intrusively). While there's a much-discussed movement called [quantified self](#), focused on capturing and sharing vast quantities of physiological data using sensors and other devices, this sort of excessive monitoring is almost certainly not something most of us want. One challenge will be figuring out how to capture useful physiological information in a way that offers benefit while also remaining unobtrusive and respecting privacy concerns.

Steve Jobs famously recruited Pepsi's John Sculley by asking him, "Do you want sell sugar water for the rest of your life or come with me and change the world?" Perhaps it's time for the next generation of technology entrepreneurs to revisit this challenge, and discover that there are some things in the world even more rewarding than the latest iPad – though I suspect my two year old daughter would emphatically disagree.

Part II: Standards and medical practices: a cautionary prognosis for algorithm-based care

Link: <http://blogs.forbes.com/davidshaywitz/2011/06/21/standards-and-medical-practices-a-cautionary-prognosis-for-algorithm-based-care/>

A distinguishing feature of the U.S. healthcare system – let us suspend judgment, for the moment, on whether this is an asset or a liability – would seem to be its variability, the fact that two patients with a similar disease can receive very different care. This may have its historical routes in the development of medicine as a profession, where physicians traditionally functioned as independent agents, practitioners who could essentially put up a shingle and go to work, striving to provide what each doctor individually perceived as the best care to each patient. However, as extensive recent research has demonstrated, the practice patterns of physicians turn out to vary widely.

This variation is just the sort of thing modern business management (think Six Sigma) seeks to reduce in other industries, and replace with a standardized, optimized process that ensures efficient and effective production. Rather than have ten factories building cars in ten different ways, according to the whim of the local manager, standard operating procedures (SOPs) typically will be created based on perceived best practices, and then utilized in all ten locations.

In medicine, a much-discussed alternative to the current, expert-based system is an approach that relies on SOPs in the form of treatment guidelines or more explicitly, algorithms; the idea is that patients with particular diagnoses are placed on a pathway and treated in a relatively uniform fashion. These can vary in complexity, from some of the basic checklists advocated and implemented by [Peter Pronovost](#) and others, and [described eloquently](#) by Atul Gawande, to very sophisticated treatment algorithms explicitly incorporating a range of contingencies and variables.

The appeal of algorithms is simple: it guarantees a basic standard of care for all patients, and ensures that vitally important elements are not forgotten – making sure certain surgical patients receive medicine to prevent blood clots is a classic example; requiring providers to wash hands prior to examining a patient is another. The point is that these should be “automatics” – the kind of interventions that truly almost every patient should receive, minimal reflection required. One much-discussed tragedy of contemporary medicine is the recognition of just how many patients are not receiving these basic fundamentals of care, often due to reluctance of physicians to change how

they've always done things, and their disinclination to be told by anyone else what to do. (A classic example from my training was a transplant surgeon who refused to administer the flu vaccine to his patients after his wife came down with the flu despite receiving the vaccine, leading him to conclude, erroneously, that the vaccine was universally ineffective.) The legendary difficulties experienced by Pronovost and others in trying to implement the most basic checklists speak to the very real, and very worrisome, reluctance of many physicians to modify their behavior, and least when this is presented as an externally imposed requirement.

Another appeal of algorithms is their scalability – you don't need experts to administer them, and in fact that's the point. Patients who satisfy basic criteria can be tracked, essentially put on autopilot (or at least aspects of their care go on autopilot), and a series of prescribed actions is performed. A great example of the success of this approach is the worldwide adoption of highly-active anti-retroviral protocols, which has enabled millions of patients across the globe to benefit from current medications. Offloading the requirement to administer such care from experts to low-cost providers (as discussed in Clay Christensen's fascinating book, [The Innovator's Prescription](#)) also offers the potential to deliver care at a profoundly reduced cost. Not surprisingly, such offloading is also a source of profound tension, particularly among the highly-paid "experts."

There are many obvious and not-so-obvious concerns about algorithm-based care. Clearly, one worry is that such an approach often fails to acknowledge and account for individual differences between patients – it threatens to lump together patients who might benefit from more customized treatment – and who would almost certainly benefit from the perception of receiving personalized, tailored care. Patients with the same formal diagnosis are still different, and most seek and value a sense of connection and recognition for the unique way they are experiencing their illness (captured with particular eloquence [here](#)).

Algorithmic medicine can also represent a significant challenge to physicians, calling into question many of the central tenants of a physician's faith. For starters, most doctors have been taught from the first day of medical school to treat each patient as a unique individual who must be carefully understood and thoughtfully treated. Beyond that, many physicians feel threatened or challenged by the notion that a protocol administered by a non-physician could provide the same level of care; in this sense, algorithms can be perceived as a real double-whammy, potentially threatening not just a doctor's job, but also his or her privileged status (and [authority](#)).

Yet another worry relates to the quality of the algorithms. A frequently critique of US healthcare is the amount of care that doesn't follow standard treatment recommendations. The problem, however, is that all treatment recommendations are not created equal (see [this](#) controversial WSJ op-ed by Jerry Groopman and Pam Hartzband); while some are based on very rigorous science and thus can be considered robust and reliable, many others are not, and often represent simply a political consensus following a conversation among putative academic experts (many of whom have staked their professional careers on one particular therapeutic approach or another). While treatment recommendations (such as those listed in the National Guideline Clearinghouse, www.guidelines.gov) typically include a measure of the degree of confidence behind each suggestion, in practice, treatment recommendations necessarily tend to be presented and, more importantly, implemented in a summarized and simplified form, making it very difficult to discern which are truly reliable.

As discussed in my [last post](#), medicine is generally far more hazy and imprecise than popular perception might suggest (and than Silicon Valley engineers generally appreciate). In particular, for physicians and patients, figuring out the relationship between care provided and outcome is a difficult and often uncertain process, depending on many more factors than what the doctor may have recommended at a given visit. An apt analogy might be to the observation that only a fraction of most diseases are due to genetic causes, and external factors often account for a significant amount of the risk (for example, if one identical twin has type 1 diabetes, the chance of the other twin being afflicted is less than 50%). On top of this, the benefits of providing the right evidence-based care can be difficult if not impossible to see; many patients often need to be treated in order for a benefit (preventing heart attacks through use of statins, say) to be realized in one, and many patients receiving appropriate treatment (such as statins) will get sick (e.g. suffer a heart attack) nevertheless. The tenuous relationship between treatment and outcome for an individual patient interferes with the opportunity for conditioning (positive or aversive) of physicians, and can reduce their incentive to embrace new guidelines or constraints, especially those that are perceived to be of dubious intrinsic value.

There may also be another, more subtle concern about algorithms, which perhaps speaks to why Americans in particular may be resistant to this sort of care (as I suspect is true). In the most extreme form of algorithm-based care, every patient would receive the "standard of care" treatment for their particular ailment. This has the potential to

benefit the population, since many would otherwise receive sub-optimal care. At the same time, it also means that the patient's expectations are necessarily limited – they know they will receive the same relative good care everyone else is getting, but seems to foreclose the possibility that they will receive exceptional care, that their doctor will construct a treatment that uniquely benefits them. In other words, if someone has a particular form of cancer, they know they are likely to do about as well as everyone else (say survive another few months), but may lose the hope of an individually-constructed cure. My own sense (and perhaps this reflects my own medical training at a major referral center, though I actually believe it's more universal) is that most patients come to a doctor with belief that their situation will be different, that their physician will at least uniquely understand them, and ideally be able to come up with an innovative, personalized approach. A formalized algorithm delivered by an interchangeable provider undercuts this hope and relationship.

A final challenge associated with algorithms is the concern that they lock in a particular standard of care; in the same way that evolution would stall without an adequate background mutation rate, so too might medical progress be inhibited by the lack of therapeutic variability, and the opportunity for an inspired physician to hit upon a novel approach that works better. Remove variation and you remove the opportunity for serendipity. While deliberate studies remain possible (and presumably required for any change in a particular algorithm), many ideas for these studies originate in the sort of unexpected observations that standardized algorithms would seem to preclude. Of course, this must be balanced with the presumed therapeutic harm that comes from doctors deciding arbitrarily what an ideal regimen should be, and winding up with a result that's significantly worse than what the relevant treatment guideline might have produced. (See [here](#) for a previous discussion of these issues – one of my favorite pieces.) Moreover, since the relevant treatment information – what was done, what was the result – often isn't captured in a useful way, important opportunities for learning, and identifying both best practices and not-so-good practices are typically missed, representing a significant lost opportunity.

This defines the challenge: developing approaches that ensure a basic standard of care, while explicitly leaving adequate space – especially in areas where treatment guidelines are not especially robust – to permit customization and individualization of care, and to enable, and ultimately capture, innovation. Doctors, as many have argued, need to be more receptive to enlightened change; yet, equally important, those who are looking to

reform medicine could use a healthy dose of humility as well, and recognize that standardization for the sake of standardization may not be a good thing for patients, physicians, or the future of medicine.

Part III: Why medicine's killer app is an e-textbook

Link: <http://blogs.forbes.com/davidshaywitz/2011/06/22/why-is-medicines-killer-app-an-e-textbook/>

If you walk the floors the nation's leading hospitals, chances are you'll see a number of physicians hunched over computers, busily pecking away. Look closer, and you'll see that many are probably using an application called UpToDate, a for-profit platform that provides doctors timely (up-to-date – get it?), authoritative information on a wide range of medical conditions, written and edited by an expansive team of experienced clinicians, many associated with top-tier academic medical centers.

Fundamentally, UpToDate is an electronic textbook, incorporating a minimal amount of technology that enables rapid cross-referencing and frequent online updates. These two features, together with consistently high-quality content, have enabled UpToDate to render the classic medicine textbooks effectively obsolete; if you are a physician, and you want to look up a medical condition, odds are the first place you'll look is UpToDate. Many drug developers, investors, and industry analysts also subscribe. The company was founded in 1992 by Burton Rose, a Harvard nephrologist, and was acquired by Wolters Kluwer in 2008 for an undisclosed price. (FWIW, I have no relationship with UpToDate, other than as a paid subscriber.)

The popularity of UpToDate offers a window into how physicians currently view and value information. Most obviously, UpToDate suggests physicians prize authority; they want to know the latest thinking associated with a disease area, and are interested in a distillation of the latest published data and expert consensus. Significantly, UpToDate generally offers not just a summary of recommendations but also an explanation of the associated physiology and a review of relevant clinical studies. Thus, it provides doctors not so much a recipe for treatment, but a sense of knowledge, empowering them to make informed therapeutic decisions without explicitly telling them what to do.

While the idea of using a medical textbook to guide clinical decision making does not strike most physicians as a particularly radical idea, it's useful to consider alternative approaches that are not generally used. For instance, you might imagine an open source textbook, perhaps analogous to Wikipedia, that could have emerged as the reference of choice; or you could envision (as many have) an information system that told physicians how their peers were actually treating other patients in similar situations, or how other patients actually fared in response to different treatment approaches. You could also imagine the regular use of an online physician community to crowdsource information

and leverage experience; while this sort of activity does occur today (mediated by companies such as SERMO) and seems to provide an occasionally useful source of orthogonal input, it certainly has not emerged as the go-to reference for most physicians.

There's one other important feature of UpToDate that I've discovered while working on the commercial side of things: it turns out that the guidelines offered by these learned authors frequently do not correlate with what the vast majority of practicing physicians are actually doing. Even as many physicians seek out the latest advice and information, they don't necessarily translate this into practice.

Consider, for example, sitagliptin, a DPP-IV inhibitor developed for the treatment of diabetes. Read UpToDate, and you might assume that the drug would be utilized only in highly specialized situations, and would consequently achieve marginal sales at best, essentially as a niche product. Yet drug usage and sales data tell a far different story – by these criteria, sitagliptin has achieved solid blockbuster status.

Industry critics will pounce on this contradiction as offering evidence of the unhealthy impact of pharma marketing, while proponents might highlight the benefit realized by patients and argue that most medical guidelines, including those articulated in UpToDate, are scripted by self-righteous academic elitists who reflexively shun both novelty and industry.

I'm not sure I'd endorse either of these perspectives, and would instead suggest that in general, what drives treatment decisions for individual patients can be so much more complex or obscure than most outsiders would suspect.

Here's a case in point from my days in management consulting: we were looking at a market where there were several very similar products, essentially indistinguishable clinically. Product A was about to go off-patent and become available as a generic, and we were trying to anticipate the effect this would have on sales of Product B, still on patent. I remember one conversation in particular – we asked a physician who tended to prescribe Product B whether he'd immediately transition his patients to the generic form of Product A. Probably not, he told us, explaining that if he did, his patients would want to know why they hadn't been put on Product A to begin with, and the time required to adequately explain this to his patients would chew up too much of the allotted fifteen minutes. Thus, until his office started to get "hassled" by pharmacists

and payors to switch to generics, he planned to keep his patients on Product B. While this particular rationale for a therapeutic choice may be idiosyncratic, factors such as ease of explanation and likelihood of being hassled play a surprisingly large role in physician decision making, even if absent from traditional guidelines such as UpToDate.

In college, I remember that a particularly useful resource was an unofficial guide to classes produced anonymously by students, offering a colorful and highly opinionated companion to the infinitely more sober course catalogue; this guide purported to tell you how things “really were,” and often but not always enhanced decision-making. While I’m sure there’d be a range of legal obstacles to overcome, it’s intriguing to speculate whether physicians might benefit from a similar sort of guide, one that strived less for political correctness and more towards offering the sort of information that busy physicians in practice tend to value. Meaningfully optimizing patient outcomes would naturally remain a central focus, of course, but optimizing reimbursement would be discussed as well.

Such a guide (which I’d envision as a companion to, not a replacement for, UpToDate) might also provide patients, policymakers, and entrepreneurs with a useful view into how physicians actually think, and the range of considerations that matter to them (spoiler alert: therapies that reduce patient calls to the office nurse will be highly valued). I suspect most of us would be appalled, amazed, and ultimately enlightened by what such an approach would reveal about how medicine, in the trenches, is actually practiced.

Part IV: Our metrics fetish – and what to do about it

Link: <http://blogs.forbes.com/davidshaywitz/2011/06/23/our-metrics-fetish-and-what-to-do-about-it/>

“How’m I doing?”

It’s not just former NYC mayor [Ed Koch](#) who’s asking. At some level, we are all preoccupied with this question, and ask it not just of ourselves (are we [winning?](#)), but also of those with whom we work, and of the businesses we choose to patronize.

We’re fairly obsessed with rankings, when you think about it – what restaurant scored highest on Yelp? Which college (or, for those living in NY or the bay area, pre-K program) was ranked highest? How’s my book faring on Amazon? Find me an author who doesn’t monitor his ranking compulsively (or for that matter, a blogger who doesn’t keep track of her page views).

Presumably, there’s something in our nature that’s hardwired to be comparative.

Of the many reasons for keeping score, the most important reason – or at least, the one most frequently cited – is to motivate improvement. A restaurant with a poor Zagat score for décor, for instance, might invest in new wallpaper, a student who gets a “C” on a math test may decide to study harder, and a hospital that’s found to rank poorly on infection prevention might institute more stringent practices.

Businesses are especially keen on quantitation in general, and on metrics in particular, as these would seem to offer an objective standard against which improvements can be measured. Process modifications that accelerate inventory turns or reduce days sales outstanding are likely to improve the bottom line. Similarly, if doctors had information telling them that they weren’t providing specific patients with the optimal care, these physicians might modify their practices (see [here](#) for a useful discussion of an effort along these lines in Sweden).

We intuitively use and value metrics, and view metrics as important and powerful tools that allow us to start making sense of a complex world, and permit us to make the many decisions that define our personal and professional lives. Metrics enable us to prioritize time, resources, and mindshare.

The problem with metrics, scoring, and rankings is directly related to their instinctive appeal: once there’s a number affixed to something – a process, a business, and yes, a

person – it can be prohibitively difficult to move beyond that; in our minds, we tend to replace the complexity of something with a number or rank; in fact, that's very much the goal and the point.

This may be reasonable if the number that's been derived effectively captures the aspects and qualities you are most interested in; if you want the fastest car, you should compare the maximum speed each model can demonstrate. If you are trying to lose weight, it makes sense to compare calories.

More often than not, however, we tend to overreach – and in a fairly predictable way. First, we measure things that are easy to count (even if of dubious relevance), and then we roll these up into a grand score and ultimately, a final rank list.

Perniciously, this process offers something for almost everyone: let's say we're talking about a business. In this case, managers have objective parameters to track, workers have objective performance goals towards which to strive, consultants have balanced scorecards to create and deliver, and investors and other stakeholders can compare these results to industry best practices. What's not to like?

The Achilles heel of this whole process is the assumption that the parameters you are tracking are truly relevant to the outcomes you are trying to achieve. I worry – deeply worry – this often is not the case. Far too frequently, we track the factors that most easily lend themselves to counting, rather than thinking deeply about whether they are truly the most important.

I remember vividly an example from the pharmaceutical industry, where a newly-formed division of a very large company, in an effort to gain credibility with senior management, instituted an elaborate series of process metrics, to the point where department members seemed to spend almost as much time recording and tracking these metrics as they did designing clinical protocols. Moreover, the plan worked: senior management was delighted that the new department “got it.” Unfortunately, these metrics didn't speak at all to the most relevant concerns: were the right studies being done, were the right questions being asked, were the new department's efforts ultimately helping the company? In this case, the metrics were able to assess how fast the department was moving, but not whether it was heading in the right direction.

This habit of testing first and asking questions later (if at all) has also impacted the way we think about competency and achievement. Consider that almost every profession

now, from medicine to law to massage therapy to motorcycle repair, has increasing numbers of proficiency exams, tests designed to signify a particular level of achievement and competence. Almost invariably, the professional societies responsible for these assessments make serious bank administering these exams, yet it's often unclear or not known whether these tests have been demonstrated to exhibit content validity (the right things are being measured), appropriate measurement properties (another key factor), and predictive validity. We take comfort from the fact that practitioners of all stripes have been duly "certified," yet rarely ask whether this certification process has either meaning or value. To have been tested, it seems, is enough.

One organization that has given considerable thought into just these sorts of questions is a group within the FDA known as SEALD (Study Endpoint and Label Development); they've put together a much-discussed document summarizing what medical product companies (disclosure: including mine) must do to ensure that the appropriate patient parameters are monitored, and in the right way. The guidance developed by this group – available for download [here](#); Figure 3 on p.7 is the money shot – strikes me as generally relevant to the way most metrics should be developed. Without question, the process (as medical products companies are discovering) is extremely rigorous, and certainly not for the faint-of-heart. Yet, in focusing on the relevance of the questions as well as the measurement qualities of the instrument, the process offers the promise that the metrics produced from this exercise will be something worth counting on.

Moreover, as we strive to improve the quality of health delivery by increasing the amount of data that is collected, analyzed, and disseminated, it's essential that we take the time and care to develop the right measures, and resist both the pressure (particularly from consultants) and the temptation (notably within ourselves) to settle for the metrics that are easiest to create.