2018 University of Virginia Institute for Practical Ethics & Public Life
Patricia Hollingsworth Prize in Ethics

DANGERS OF THE DATA-DRIVEN EQUILIBRIUM SHIFT IN MEDICINE

Philip H. Schroeder

March 14th, 2018

We are entering an age in which everything and everyone generates data, with each human life increasingly reducible to a perpetual stream of characters and digits forever embedded within the fabric of binary code, the ensemble of ones and zeroes, that has come to define the world we call home. The resulting trails of digital exhaust will provide a roadmap for our past, our present, and even our future. Such roadmaps harness unprecedented potential to revolutionize the fight against humanity's greatest challenges, including and especially those within medicine. However, they also present unprecedented dangers to patients, physicians, and the practice of health care as a whole. In this paper, I strive to shed light on these dangers. Specifically, I discuss concerns of ethics with the increasing dominance of data in patient care. Although focused in medicine, the themes of this paper are of great relevance to the data-driven cultural changes occurring throughout science and society at large.

Equilibrium Shift in Decision-Making

Equilibrium of Decision-Making

The exponential growth in data and computational capabilities is giving rise to what I will refer to as "data-driven decision-making." Decisions of this type are informed by quantifiable patterns found in large sets of data. From stockbrokers to epidemiologists, data-driven decision-making can be found in many forms and is becoming increasingly prevalent in nearly every field, with great success. The insights used to inform data-driven decision-making generally derive from complex analyses of multi-dimensional relationships between variables and outcomes of interest. An example of this in medicine is prognostic modeling, in which machine-learning algorithms use a patient's biomedical data (e.g., genomic makeup, blood culture results, vital signs) to determine their risk for developing certain morbidities.[1] Diagnostic and treatment decisions that derive from such analyses are data-driven and are revolutionizing medicine, increasing the precision and speed with which therapies can be personalized for individual patients.[1]

Due to the limited cognitive capacity of humans to keep pace with complex computational analysis, the reasoning behind data-driven decisions often exceeds human comprehension. That is, humans may not be able to understand the *why* behind data-driven decisions. For example, a physician may not fully understand *why* a patient is at high risk of sepsis, as determined by a data-driven analysis of blood culture results, but may still decide to treat the patient accordingly. In this way, data-driven decision-making can generate friction with the more traditional "knowledge-driven decision-making," which is founded in human intuition or established knowledge. The dichotomy between knowledge-driven and data-driven decision-making in health care is becoming increasingly recognized in medical and legal literature. For example, data-centered versus patient-centered care and implicit versus explicit personalized medicine have arisen as popular characterizations of data-driven versus knowledge-driven decision-making in medicine. [1,2]

Of course, the classifications of knowledge-driven and data-driven are not absolute, nor are they mutually exclusive. To use the previous example, biomedical knowledge of sepsis can help to inform physicians' decisions to treat patients determined to be at high risk by data-driven analyses. I describe this interaction between data-driven and knowledge-driven decision-making as a dynamic equilibrium, in which data can inform knowledge-driven decision-making and knowledge can inform data-driven decision-making. This interaction is subtle and its implications are complex. For this reason, the precision with which the interaction is characterized is of great importance. I choose to describe it as an equilibrium for three reasons. First, although situated in a state of balance between forces, an equilibrium often does not involve equal contribution from such forces and can shift depending on how such contributions change over time. Second, equilibriums function on a continuous spectrum, with little change between any two consecutive moments, allowing for incremental shifting to occur almost without notice by a single generation of spectators. Finally, the location at which an equilibrium is situated is entirely self-regulated, controlled solely by the forces of which it is comprised, not by the quality of its output or the vitality of agents that rely upon such output. These three dynamics highlight the key characteristics and the potential dangers of the equilibrium of data-driven and knowledge-driven decision-making. In doing so, they guide the remainder of my discussion. I start with the first.

Equilibrium Shift

Due to limits of human cognition, the extent to which knowledge and intuition can explain data-driven insights is inversely related to the amount of factors considered in the analysis. That is, data-driven decision-making becomes increasingly immiscible with knowledge-driven decision-making as the decisions increase in complexity. Therefore, as we continue expanding the amount of variables that can be quantified and analyzed, an equilibrium shift is occurring, in which decision-making is becoming increasingly data-driven and decreasingly knowledge-driven. This equilibrium shift will undoubtedly revolutionize our ability to solve problems in medicine and beyond. However, the shift also presents significant dangers and threatens to expose vulnerabilities we have yet to fully realize.

Concerns of Ethics with the Data-Driven Equilibrium Shift in Medicine

I now seek to introduce concerns of ethics associated with the data-driven equilibrium shift in health care. In doing so, my intention is not to be exhaustive, but instead to raise more questions than answers and to preserve a constructive malleability that, in turn, can foster diverse reflection.

Privacy & Coercion of Patients

The most obvious concerns that arise with the increasing dominance of data in health care are those of patient privacy and coercion. A well understood principle in data science is that the greatest determinant in the degree to which any data-driven analysis can be optimized is the quantity and quality of available data. Thus, as medical practice becomes dominated by data-driven decision-making, the quality of patient care becomes increasingly dependent on the quantity and quality of patient data available to the health care providers.[1] This includes the quality and precision of evaluation, monitoring, diagnosis, and treatment of a patient. With level of care dependent on the availability of data, it is increasingly reasonable and convenient for health systems to assign blame for misdiagnoses, preventable progression of morbidities, or overall unfavorable health outcomes to the lack in quality or quantity of patient data. This, in turn, renders patients vulnerable to a dangerous dynamic, in which hesitation to surrender private data can be leveraged by health providers and insurance companies to deflect liability for poor health outcomes or justify changes in insurance cost or availability. It is already commonplace for insurance companies to reward patients for

behavior that is thought to improve health (e.g., exercising, eating health). As personal data collection and sharing become behaviors deemed beneficial to health, or at least beneficial to the care that supports health, there is increasing incentive to reward patients for sharing private data, whether that be daily step count, genetic makeup, or social media behavior. Unlike those for exercising and eating healthy, reward systems for sharing data can expose vulnerabilities and disproportionally underserve patients based on generational, cultural, and socioeconomic differences. Further, once the data is retrieved by health systems, measures to increase its utility often undermine its security.[1] That is, increased fluidity of data exchange improves the heterogeneity and generalizability of datasets, which, in turn, increases the power and number of insights that can be identified, but also increases the likelihood that anonymized data can be reassociated with individual patients.[3,4] In addition, broader access to datasets increases the amount of researchers able to evaluate the data and improve clinical algorithms, but decreases control over the information.[3] These tradeoffs, again, promote a dynamic in which the level of patient care is directly proportional to the extent of data privacy sacrificed by or for individual patients.

Informed Consent

As discussed, data-driven analyses are not only being leveraged to improve identification of appropriate treatments, but also identification of the appropriate patients for such treatments.[1] As a result, physicians are inevitably becoming less knowledgeable about why they are pursuing certain therapies or interventions. This raises a simple question: How can informed consent be achieved when even the physician does not entirely understand how an intervention works or why it is indicated for the patient? Many would argue the first issue of this question, of the physician not understanding how the intervention works, is neither novel nor ethically problematic. That is, since the very beginning of medical practice, it has been commonplace for physicians to suggest therapies and prescribe medications without complete knowledge of how they lead to favorable outcomes. To provide some current examples; acetaminophen for pain, lithium for bipolar disorder and the keto-diet for epilepsy. Since the correlation between interventions and positive outcomes is most often discovered long before the causation, it would be excessively consequential to patients if we were to require a biomedical understanding of the intervention to be elucidated prior to its approval. In these cases, informed consent is achieved through the physician's communication of the

performance of the therapy with the patient's pathology in previous trials, allowing patients to gain a relative understanding of potential benefits and risks. However, the second issue in this question, of the physician not understanding why the therapy is indicated for the patient, is in fact both novel and ethically problematic. Under these circumstances, the patient is robbed of any opportunity to consider their values and priorities, both medical and nonmedical, within the grater context of the risks they face. Regardless of how its current definition is interpreted, informed consent simply cannot be achieved with such lack of clarity.

Moral Accountability

At its most fundamental level, the practice of medicine is one in which a human commits to helping a fellow human. In this commitment, there is an oath of care to which the caregiver is held morally accountable. In turn, this moral accountability motivates the caregiver to make purposeful decisions that are thought to help heal the patient. However, this influence is reciprocal, not unidirectional as easily perceived. That is, the moral accountability motivates purposeful decision-making and, in turn, the purposeful decision-making motivates the caregiver to both recognize and embrace the moral accountability to care for the patient. With protocol development increasingly dominated by data-driven decision-making, caregivers progressively lose understanding of why certain medications or procedures are pursued for their patients. With this equilibrium shift comes a growing disconnect between what caregivers understand and what caregivers are obligated to do based on protocol. With this disconnect, caregivers lose ability to make purposeful decisions for their patients and are left vulnerable to the moral distress of executing therapies with which they may not agree. This, in turn, may push caregivers to grow numb to or even reject the moral accountability to care for the patient.

Quantifying Illness

The treatment of human illness is an extraordinarily multidisciplinary task, requiring insights from a diverse spectrum of thought and practice. From the physiological to the divine, some influences of illness can be reduced and precisely quantified while others can barely be defined. A fundamental role of the caregiver is to navigate the diverse landscape of human emotional, spiritual, and intellectual experience in pursuit of insights that can help patients. Data-driven decision-making, by its very nature, prioritizes aspects of illness that can be quantified. This poses a threat to all patients who hold priorities and values that cannot be captured by algorithm-conducive inputs.

Furthermore, in recognition of this limitation of data-driven decision-making, the medical community is already advertising ideas to quantify aspects of illness that simply should not be quantified. For example, in their efforts to promote "person-centered" care, Vikki Entwistle and Ian Watt, along with several peer commentaries, suggest developing metrics to measure patients' "personal capabilities," which include "capabilities to reason, to feel and respond to emotion, to intend and initiate action, to be self-aware and self-directing, to experience particular kinds of suffering and so on, and capabilities to participate socially in a group or community of beings that recognizes each other as having significant ethical privileges."[1,5,6] If we were to quantify such aspects of the human experience, we would risk driving the equilibrium shift in a way that does not benefit the quality of health care: in a way that exacerbates the consequences of data-driven care (i.e., depersonalizing the patient) without harnessing its benefits (i.e., optimizing multidimensional decision-making). Not only would the resulting metrics most certainly misrepresent the intended patient qualities, they would provide false assurance that such qualities are being considered while drowning out the patients' subjective experiences by reducing them to another input within the myriad of interconnecting utility functions that fuel data-driven care.[7] Further, it would be poison to our very humanity to regress to such levels of reductionism, to attach calibrated metrics to aspects of illness such as the toll of suffering, the strength of human will, and the presence of mind and spirit. If health care is to continue fulfilling its fundamental responsibility to society of treating illness, in every sense of that word, caregivers must resist any attempt to tame the aspects of the human condition that demand and deserve unadulterated freedom.

Final Word of Caution

To end, I return to the third reason for my characterization of the interaction between knowledge-driven and data-driven decision-making as an equilibrium, which reads: the location at which an equilibrium is situated is entirely self-regulated, controlled solely by the forces of which it is comprised, not by the quality of its output or the vitality of agents that rely upon such output. More than anything, this is a word of caution. This equilibrium shift will undoubtedly progress due to the nature of its forces (i.e., the limitations of knowledge-driven decision-making and increasing potential of data-driven decision-making). However, despite its advantages, this progression may also come with a grave cost, even beyond those discussed in this essay. The pre-shift status of the equilibrium intrinsically designates greater value to human knowledge and, by extension, to human perspective. The vitality of many cherished aspects of health care, and society as a whole, rely upon us continuing to value this human perspective. I fear that the preservation of the most valuable aspect of our humanity, our ability to truly care for one another, is an agent of which the vitality relies on the current status of the equilibrium and, thus, is threatened by the equilibrium shift. As new generations of health care providers and administrators arise during progressively later stages of this shift, it is unclear how they will conceptualize the practice of medicine, and more importantly, the role we share as humans within that practice.

References

1. Price, NW. Black-box medicine. Harvard Journal of Law & Technology. 2015; 28(2):420-67.

2. Entwistle, V. A., and Watt, I. S. Treating patients as persons: a capabilities approach to support delivery of person-centered care. The American Journal of Bioethics. 2013; 13(8):29-39.

3. Malin, B. and Sweeney, L. How (not) to protect genomic data privacy in a distributed network: using trail re-identification to evaluate and design anonymity protection systems. Journal of Biomedical Informatics. 2014.

4. Ohm P. broken promises of privacy: responding to the surprising failure of anonymization. University of California Los Angeles Law Review. 2010.

5. Little, M. A better grounding for person-centered medicine? The American Journal of Bioethics. 2013; 13(8):40-42.

6. Maagt, S. D., and Robeyns, I. Can person-centered care deal with atypical persons? The American Journal of Bioethics. 2013; 13(8):44-46.

7. McMaster, R. Capabilities and patients as persons: Ethical implications for health economics. The American Journal of Bioethics. 2013; 13(8):48-50.