

# Informed Consent and Medical Artificial Intelligence: What to Tell the Patient?

I. GLENN COHEN\*

## TABLE OF CONTENTS

INTRODUCTION . . . . .	1426
I. A BRIEF PRIMER ON MEDICAL AI/ML . . . . .	1429
II. THE DOCTRINAL QUESTION: HOW DOES THE CURRENT CASE LAW ON INFORMED CONSENT IN THE UNITED STATES APPLY TO MEDICAL AI/ML? . . . . .	1432
A. GENERAL BACKGROUND ON U.S. INFORMED CONSENT LAW . . . . .	1432
B. THREE KINDS OF “PENUMBRAL” INFORMED CONSENT CASES THAT MAY PROVE USEFUL FOR AI/ML . . . . .	1434
1. Provider Experience and Qualification . . . . .	1435
2. Substitute Physicians: “Ghost,” Concurrent, and Overlapping Surgery . . . . .	1436
3. Financial Conflicts of Interest . . . . .	1439
D. WHAT TO TELL THE PATIENTS: APPLYING INFORMED CONSENT CASE LAW TO MEDICAL AI/ML . . . . .	1441
1. Starting With the Standards . . . . .	1442
2. Deeper Into the Analogies . . . . .	1444
III. IS THAT THE RIGHT ANSWER? BEYOND THE DOCTRINAL APPROACH . . . . .	1449
A. AN EMPIRICAL APPROACH . . . . .	1449
B. A NORMATIVE APPROACH . . . . .	1452
IV. SOME REMAINING ISSUES . . . . .	1462

---

\* James A. Attwood and Leslie Williams Professor of Law, Harvard Law School. Faculty Director, Petrie-Flom Center for Health Law Policy, Biotechnology & Bioethics. J.D. © 2020, I. Glenn Cohen. I thank Sara Gerke, Nadia Sawicki, W. Nicholson Price II, and participants at *The Georgetown Law Journal’s* Symposium, Law and the Nation’s Health, and the University of Texas at Austin School of Law’s and the Ohio State University Moritz College of Law’s faculty workshops, and the ETH Zurich Workshop and Lecture Series in Law & Economics for comments. Max Bloom, Benjamin Sanchez, and Priya Sundaresan provided excellent research assistance. This research was supported by the Collaborative Research Program for Biomedical Innovation Law, which is a scientifically independent collaborative research program supported by the Novo Nordisk Foundation (grant NNF17SA0027784).

A. WHAT ABOUT CONSENT? . . . . .	1462
B. THE SPECIAL CASE (?) OF BIAS AND INDIVIDUALIZING CONSENT . . . .	1464
C. SOME (MORE) COLD WATER ON YOUR WAY OUT . . . . .	1465
CONCLUSION . . . . .	1467

## INTRODUCTION

Imagine you are a patient who has been diagnosed with prostate cancer. There are two main approaches to treating it in the United States: active surveillance or the surgical option of radical prostatectomy.<sup>1</sup> Your physician recommends the surgical option and spends considerable time explaining the steps in the surgery, the benefits of (among other things) eliminating the tumor, and the risks of (among other things) erectile dysfunction and urinary incontinence after the surgery.<sup>2</sup> What your physician does not tell you is that she has arrived at her recommendation of prostatectomy over active surveillance based on the analysis of an Artificial Intelligence (AI)/Machine Learning (ML) system, which recommended this treatment plan based on analysis of your age, tumor size, and other personal characteristics found in your electronic health record. Has the doctor secured informed consent from a legal perspective? From an ethical perspective? If the doctor actually chose to “overrule” the AI system<sup>3</sup> and failed to tell you, has she violated your legal or ethical right to informed consent? If you were to find out that the AI/ML system was used to make recommendations on your care and no one told you, how would you feel? Well, come to think of it, do you know whether an AI/ML system *was* used the last time you saw a physician?

This example is hypothetical, but in the real world many physicians are racing toward integrating AI/ML into diagnostics, prognostics, allocation of resources, and treatment itself.

For example, in July 2019 the *Lancet HIV* published research that used electronic health record data and machine learning to predict whether a particular man was at high risk of contracting HIV, thus enabling his physician to determine whether to put him on Pre-Exposure Prophylaxis (PrEP).<sup>4</sup> Should a physician that

1. E.g., Belén López-Pérez et al., *Predicting Prostate Cancer Treatment Choices: The Role of Numeracy, Time Discounting, and Risk Attitudes*, 22 J. HEALTH PSYCHOL. 788, 788 (2017).

2. *Id.* at 788–89.

3. I put “overrule” in scare quotes to capture the ambiguity of the underlying choice architecture here. One could imagine an implementation where the AI/ML recommendation was merely “informational” as opposed to one where the recommendation was treated as a strong or weak default in the sense of what the physician must do to adopt the opposite recommendation (for example, record that they ignored it, consult with a colleague, present to a committee of experts, etc.). There is also a separate level of complexity regarding the multiple players in a healthcare system. Even if a physician maintains full discretion to ignore an AI/ML recommendation, one could imagine pressures from payers (for example, you may do X, but if the recommendation was Y, we will cover only Y or not cover the difference in cost between X and Y).

4. See Julia L. Marcus et al., *Use of Electronic Health Record Data and Machine Learning to Identify Candidates for HIV Pre-exposure Prophylaxis: A Modelling Study*, LANCET: LANCET HIV (July 5,

uses that tool and, at least partially as a result of doing so, recommends PrEP need to disclose that he or she did so to the patient?

Or consider what Doctor Daniel A. Hashimoto and his colleagues call the “[i]ntegration of multimodal data with AI [to] augment surgical decision-making across all phases of care” to achieve a “collective surgical consciousness,”<sup>5</sup> which they describe using the example of bariatric surgery:

Preoperatively, a patient undergoing evaluation for bariatric surgery may be tracking weight, glucose, meals, and activity through mobile applications and fitness trackers, with the data feeding into their EMR. Automated analysis of all preoperative mobile and clinical data could provide a more patient-specific risk score for operative planning and yield valuable predictors for postoperative care. The surgeon could then augment their decision-making intraoperatively based on real-time analysis of intraoperative progress that integrates EMR data with operative video, vital signs, instrument/hand tracking, and electrosurgical energy usage. Intraoperative monitoring of such different types of data could lead to real-time prediction and avoidance of adverse events. Integration of pre-, intra-, and postoperative data could help to monitor recovery and predict complications. After discharge, postoperative data from personal devices could continue to be integrated with data from their hospitalization to maximize weight loss and resolution of obesity-related comorbidities. Such an example could be applied to any type of surgical care with the potential for truly patient-specific, patient-centered care.

AI could be utilized to augment sharing of knowledge through the collection of massive amounts of operative video and EMR data across many surgeons around the world to generate a database of practices and techniques that can be assessed against outcomes.<sup>6</sup>

Many companies and healthcare providers are currently investing heavily in developing medical AI/ML systems including AI-driven X-ray image analysis<sup>7</sup> systems and AI-driven monitoring systems used to identify elderly patients at risk of falling.<sup>8</sup> One recent forecast suggests that the market for these technologies

---

2019), [https://www.thelancet.com/journals/lanhiv/article/PIIS2352-3018\(19\)30137-7/fulltext](https://www.thelancet.com/journals/lanhiv/article/PIIS2352-3018(19)30137-7/fulltext); see also Gina Kolata, *Would You Want a Computer to Judge Your Risk of H.I.V. Infection?*, N.Y. TIMES (July 30, 2019), <https://www.nytimes.com/2019/07/30/health/hiv-aids-prep.html> (discussing additional models to predict a patient’s risk of contracting HIV).

5. Daniel A. Hashimoto et al., *Artificial Intelligence in Surgery: Promises and Perils*, 268 ANNALS SURGERY 70, 74 (2018).

6. *Id.* at 73. As discussed below, in some states, the law of informed consent is limited to surgical settings such that this type of example might prove particularly important.

7. See Ahmed Hosny et al., *Artificial Intelligence in Radiology*, 18 NATURE REVIEWS. CANCER 500, 507 (2018).

8. See E. Ramanujam & S. Padmavathi, *A Vision-Based Posture Monitoring System for the Elderly Using Intelligent Fall Detection Technique*, in GUIDE TO AMBIENT INTELLIGENCE IN THE IoT ENVIRONMENT 249, 250 (Zaigham Mahmood ed., 2019).

will be larger than \$34 billion worldwide by 2025.<sup>9</sup> Perhaps most startlingly, in 2018 the FDA gave permission to market IDx-DR, an AI/ML-based medical device intended to be used in a primary care doctor's office for the disease diabetic retinopathy, which provides a screening decision: either

(1) 'more than mild diabetic retinopathy detected: refer to an eye care professional' or (2) 'negative for more than mild diabetic retinopathy; rescreen in 12 months' . . . *without the need for a clinician* to also interpret the image or results, which makes it usable by health care providers who may not normally be involved in eye care.<sup>10</sup>

The FDA's review included the results of a clinical trial, involving 900 participants, showing that IDx-DR "performed well" in detecting the presence of more than mild diabetic retinopathy.<sup>11</sup>

This Article is the first to examine in-depth how medical AI/ML intersects with our concept of informed consent. To be clear, this is just one of a number of issues raised by medical AI/ML—which includes data privacy, bias, and the optimal regulatory pathway<sup>12</sup>—but it is one that has received surprisingly little attention. I hope to begin to remedy that with this Article. Part I provides a brief primer on medical Artificial Intelligence and Machine Learning. Part II sets out the core and penumbra of U.S. informed consent law and then seeks to determine to what extent AI/ML involvement in a patient's health should be disclosed under the current doctrine. Part III examines whether the current doctrine "has it right," while examining more openly empirical and normative approaches to the question.

To forefront my conclusions: although there is some play in the joints, my best reading of the existing legal doctrine is that in general, liability will not lie for failing to inform patients about the use of medical AI/ML to help formulate

---

9. TRACTICA, ARTIFICIAL INTELLIGENCE FOR HEALTHCARE APPLICATIONS (2018), <https://tractica.omdia.com/research/artificial-intelligence-for-healthcare-applications/>.

10. *FDA Permits Marketing of Artificial Intelligence-Based Device to Detect Certain Diabetes-Related Eye Problems*, U.S. FOOD & DRUG ADMIN. (Apr. 11, 2018) (emphasis added), <https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-artificial-intelligence-based-device-detect-certain-diabetes-related-eye> [<https://perma.cc/XRS6-PNUM>]; see U.S. FOOD & DRUG ADMIN., DE NOVO CLASSIFICATION REQUEST FOR IDx-DR: DE NOVO SUMMARY (DEN180001), at 1 (2018) [hereinafter DE NOVO CLASSIFICATION REQUEST], [https://www.accessdata.fda.gov/cdrh\\_docs/reviews/DEN180001.pdf](https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN180001.pdf) [<https://perma.cc/7ZMJ-KE9C>].

11. DE NOVO CLASSIFICATION REQUEST, *supra* note 10, at 12.

12. See, e.g., I. Glenn Cohen & Michelle M. Mello, *Big Data, Big Tech, and Protecting Patient Privacy*, 322 JAMA 1141 (2019) (discussing patient privacy); Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 SCIENCE 447 (2019) (discussing gender and racial bias); W. Nicholson Price II, *Medical AI and Contextual Bias*, 33 HARV. J.L. & TECH. 65 (2019) (discussing regulatory requirements). See generally BIG DATA, HEALTH LAW, AND BIOETHICS (I. Glenn Cohen et al. eds., 2018) (discussing big data's impact on health law and bioethics and other big-data-related considerations).

treatment recommendations. There are a few situations where the doctrine may be more capacious, which I try to draw out (such as when patients inquire about the involvement of AI/ML, when the medical AI/ML is more opaque, when it is given an outsized role in the final decisionmaking, or when the AI/ML is used to reduce costs rather than improve patient health), though extending it even here is not certain. I also offer some thoughts on the question: if there is room in the doctrine (either via common law or legislative action), what would a desirable doctrine look like when it comes to medical AI/ML? I also briefly touch on the question of how the doctrine of informed consent should interact with concerns about biased training data for AI/ML.

### I. A BRIEF PRIMER ON MEDICAL AI/ML

This is a brief symposium Article, so keep your expectations in check—only the briefest tour will be provided of the basics of how some forms of medical AI/ML work, but hopefully enough to get your head around the way it intersects with informed consent.

Although the exact meaning of “artificial intelligence” (AI) remains contested, the FDA provides a good working definition as “the science and engineering of making intelligent machines, especially intelligent computer programs.”<sup>13</sup>

AI encompasses a series of different approaches. Professors Yu, Beam, and Kohane give an excellent summary of the various kinds of AI that are being used in the medical space with their illustrative example being related to the IDx-Dr technology mentioned above. Although earlier attempts at AI systems “relied on the curation of medical knowledge by experts and on the formulation of robust decision rules, recent AI research has leveraged machine-learning methods, which can account for complex interactions, to identify patterns from the data.”<sup>14</sup> These machine-learning algorithms may be categorized as either supervised or unsupervised.<sup>15</sup>

Supervised methods “work by collecting a large number of ‘training’ cases, which contain inputs,” and by analyzing the patterns between these inputs and their corresponding outputs.<sup>16</sup> In this way, the algorithm “learns to produce the correct output for a given input on new cases,” while identifying parameters that “minimize the deviations between their predictions for the training cases and the observed outcomes in these cases.”<sup>17</sup> In so doing, the theory goes, the algorithms

---

13. *Artificial Intelligence and Machine Learning in Software as a Medical Device*, U.S. FOOD & DRUG ADMIN. (Apr. 2, 2019), <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device> [<https://perma.cc/47HV-7RRQ>] (citing John McCarthy, *What Is Artificial Intelligence?* (Nov. 12, 2007) (unpublished manuscript), <http://jmc.stanford.edu/articles/whatisai/whatisai.pdf> [<https://perma.cc/TD5N-WLHG>]).

14. Kun-Hsing Yu et al., *Artificial Intelligence in Healthcare*, 2 NATURE BIOMEDICAL ENGINEERING 719, 719 (2018).

15. *Id.*

16. *Id.*

17. *Id.* at 719–20.

will develop predictive power for novel cases beyond the dataset on which they have trained.<sup>18</sup>

Under unsupervised methods, on the other hand, the algorithms are presented with unlabeled data and tasked with discerning underlying patterns in the data, identifying outliers, or “produc[ing] low-dimensional representations of the data.”<sup>19</sup> Yu, Beam, and Kohane note that advances in machine learning have been driven by developments in the application of neural network approaches to medical AI, in particular “deep learning—which involves training an artificial neural network with many layers (that is, a ‘deep’ neural network) on huge datasets—to large sources of labelled data.”<sup>20</sup>

The growth in AI/ML has been married with the development of decision support systems, which since the 1970s, have collected relevant information to provide suggestions to medical providers. Today, “decision support systems can actively gather information from patients and [Electronic Health Records (EHRs)], present suggestions to clinicians and store the system outputs in EHRs,” by taking advantage of the EHR’s quite detailed information about the patient, “including clinical notes and laboratory values, enabling the application of natural-language-processing methods to extract codified vocabularies.”<sup>21</sup> In the last few years, the results have been startling in that AI systems have achieved “specialist-level performance in many diagnostic tasks [and] can better predict patient prognosis than clinicians” in many areas.<sup>22</sup> As the many studies Professors Yu, Beam, and Kohane collect show, these results have been impressive—not only in the “low hanging fruit” of image-based diagnosis (facilitated by huge amounts of training data and well-structured tasks) in areas such as radiology, cardiology, dermatology, ophthalmology, and pathology, but also in predicting clinical outcomes. For example, “[d]ata from health insurance claims can be used to predict mortality in elderly patients, patient attributes in the medical notes can be

---

18. *Id.* at 720.

19. *Id.*

20. *Id.* As they further explicate:

The recent renaissance in AI has to a large extent been driven by the successful application of deep learning . . . . The basic architecture of deep neural networks consists of an input layer and an output layer, and a number of hidden layers in between. . . . Many modern neural networks have more than 100 layers. Neural networks with many layers can model complex relations between the input and output but may require more data, computation time or advanced architecture designs so as to achieve optimal performance. . . . However, deep-learning algorithms are extremely ‘data hungry’ for labelled cases. Only recently, large sources of medical data that can be fed into these algorithms have become widely available, owing to the establishment of many large-scale studies . . . .

*Id.* at 720–21.

21. *Id.* at 722. Of course, the data contained in EHRs are far from perfect, especially because in the United States, EHRs have primarily been developed for billing purposes and reflect some of that genesis in what they capture. For a good discussion of the problems, see Sharona Hoffman & Andy Podgurski, *The Use and Misuse of Biomedical Data: Is Bigger Really Better?*, 39 AM. J.L. & MED. 497, 515–21 (2013) (discussing input errors, incomplete or fragmented records, and flaws in data coding or standardization).

22. Yu et al., *supra* note 14, at 722.

employed to classify cancer patients with different responses to chemotherapy, and clinical predictors for the prognosis of patients receiving thoracic organ transplantation can be identified.”<sup>23</sup> Among other examples, there have also been AI/ML systems for predicting cardiac arrest, need for ICU transfer, and hypoxemia during surgery.<sup>24</sup>

One further distinction is useful: some would group medical AI/ML based on its level of “opacity,” usually a metaphor for how difficult or easily the way it draws conclusions from data can be understood by human beings, with the most opaque forms sometimes referred to as “black box.” Professor W. Nicholson Price II offers a good account:

Algorithms can be opaque for multiple reasons. Sometimes, algorithms are nontransparent because, while they may rely on explicit rules, those rules are too complex for us to explicitly understand—for example, patients whose measurements place them in a particular region of  $n$ -dimensional (where  $n$  is large) characteristic-space are at a higher risk of stroke. In particular, these rules may be impossible to explain or to understand by following the process of scientific/medical discovery: mechanistic lab experiments followed by confirmatory clinical trials. Other times, the relationships used in a black-box algorithm are literally unknowable because of the machine-learning techniques employed—that is, no one, not even those who programmed the machine-learning process, knows exactly what factors go into the ultimate decisions. A key distinguishing feature of black-box algorithms, as the term is used here, is that it refers to algorithms that are inherently black box (i.e., their developers *cannot* share the details of how the algorithm works in practice)—rather than to algorithms that are deliberately black box (i.e., their developers *will not* share the details of how the algorithm works). Black-box algorithms are especially likely to evolve over time as they incorporate new data into an integrated process of learning-and-applying.<sup>25</sup>

Thus, in some but certainly not all instances, the rules used by algorithms used in medical AI may be not just *not explained* but *not possible to explain*. When we think about how informed consent interfaces with AI/ML, it will be useful to

---

23. *Id.* at 726 (citing Dursen Delen et al., *A Machine Learning-Based Approach to Prognostic Analysis of Thoracic Transplantations*, 49 *ARTIFICIAL INTELLIGENCE MED.* 33, 33–42 (2010); Maggie Makar et al., *Short-Term Mortality Prediction for Elderly Patients Using Medicare Claims Data*, 5 *INT’L J. MACHINE LEARNING & COMPUTING* 192, 192–97 (2015); Terence Ng et al., *A Clinical Decision Support Tool to Predict Survival in Cancer Patients Beyond 120 Days After Palliative Chemotherapy*, 15 *J. PALLIATIVE MED.* 863, 863–69 (2012)).

24. *Id.* (citing Matthew M. Churpek et al., *Multicenter Development and Validation of a Risk Stratification Tool for Ward Patients*, 190 *AM. J. RESPIRATORY & CRITICAL CARE MED.* 649, 649–55 (2014); Matthew M. Churpek et al., *Predicting Cardiac Arrest on the Wards: A Nested Case-Control Study*, 141 *CHEST* 1170, 1170–76 (2012); Scott M. Lundberg et al., *Explainable Machine-Learning Predictions for the Prevention of Hypoxaemia During Surgery*, 2 *NATURE BIOMEDICAL ENGINEERING* 749 (2018)).

25. W. Nicholson Price II, *Regulating Black-Box Medicine*, 116 *MICH. L. REV.* 421, 430 (2017).

keep this idea of opacity in mind even though there has been some pushback on the underlying concept and its importance in the legal and bioethics literature.<sup>26</sup>

## II. THE DOCTRINAL QUESTION: HOW DOES THE CURRENT CASE LAW ON INFORMED CONSENT IN THE UNITED STATES APPLY TO MEDICAL AI/ML?

This Part examines what current case law on informed consent might say about the obligation to disclose medical AI/ML. It begins with a brief primer on the basics of U.S. informed consent law. Finding that the typical case—involving the failure to adequately disclose risks and benefits of a procedure—is not particularly helpful in thinking about AI/ML, it considers three more penumbral types of informed consent cases that may be better analogies. It then interrogates those analogies in greater depth.

### A. GENERAL BACKGROUND ON U.S. INFORMED CONSENT LAW

In the United States, it was not until the early 1960s that the law began to recognize that liability could attach if a physician did not inform the patient of the risks and benefits of the proposed treatment or nontreatment.<sup>27</sup>

[T]o recover on a claim for breach of the duty to secure informed consent, plaintiffs must demonstrate four elements: (1) failure to disclose a specific risk in violation of the governing standard; (2) materialization of that risk; (3) “causation”—that is, if the risk been disclosed, the patient, or a prudent person in the patient’s position, would not have proceeded as she did; and (4) that no exception, like emergency, excuses the failure to disclose.<sup>28</sup>

The main disagreement at the level of doctrine is on the first factor as to the scope of the disclosure duty—what must be disclosed. U.S. jurisdictions are roughly evenly divided between a physician-based standard and a patient-based

---

26. See, e.g., Alex John London, *Artificial Intelligence and Black-Box Medical Decisions: Accuracy Versus Explainability*, 49 HASTINGS CTR. REP. 15 (2019) (arguing that opaque reasoning is actually common in non-AI/ML medicine, and that for medicine confidence in the accuracy of results may sometimes be more important than the ability to explain how results are produced); Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 FORDHAM L. REV. 1085, 1090–91 (2018) (suggesting it is more useful to discuss properties of “inscrutability” and “non-intuitiveness” separately rather than opacity or nonexplainability).

27. E.g., Nadia N. Sawicki, *Modernizing Informed Consent: Expanding the Boundaries of Materiality*, 2016 U. ILL. L. REV. 821, 822–23. Today the tort is considered a species of negligence; although grounded in negligence principles, it actually originated as a cause of action for battery. E.g., JESSICA W. BERG ET AL., *INFORMED CONSENT: LEGAL THEORY AND CLINICAL PRACTICE* 41–44 (2d ed. 2001); Sawicki, *supra*, at 827 n.18. My focus in this short Article is what the law requires in terms of informed consent. There is also a robust literature on informed consent in clinical medicine as a freestanding ethical requirement. The two do intersect, and although I briefly touch on some of the more ethical writing in this area in Part III, for the most part I focus on the legal requirements in the United States.

28. Robin Fretwell Wilson, *The Promise of Informed Consent*, in THE OXFORD HANDBOOK OF U.S. HEALTH CARE LAW 213, 217 (I. Glenn Cohen, Allison K. Hoffman & William M. Sage eds., 2015–2016 ed. 2016).

standard of disclosure,<sup>29</sup> although in practice the distinction between the two standards may blur.<sup>30</sup> The nomenclature is pretty accurate—the physician-based standard, which was more dominant in the early days of the doctrine, answers what a physician must disclose by reference to what a reasonable physician would customarily disclose—or in the words of the 1960 Kansas Supreme Court decision in *Natanson v. Kline*<sup>31</sup> only “those disclosures which a reasonable medical practitioner would make under the same or similar circumstances.”<sup>32</sup>

By contrast, the patient-based standard is well captured by the seminal case of *Canterbury v. Spence*<sup>33</sup> from the D.C. Circuit in 1972, involving the failure to disclose to a young man undergoing surgery for back pain that there was a risk of paralysis:

In our view, the patient’s right of self-decision shapes the boundaries of the duty to reveal. That right can be effectively exercised only if the patient possesses enough information to enable an intelligent choice. The scope of the physician’s communications to the patient, then, must be measured by the patient’s need, and that need is the information material to the decision. Thus the test for determining whether a particular peril must be divulged is its materiality to the patient’s decision: all risks potentially affecting the decision must be unmasked. And to safeguard the patient’s interest in achieving his own determination on treatment, the law must itself set the standard for adequate disclosure.

...

From these considerations we derive the breadth of the disclosure of risks legally to be required. The scope of the standard is not subjective as to either the physician or the patient; it remains objective with due regard for the patient’s informational needs and with suitable leeway for the physician’s situation. In broad outline, we agree that “[a] risk is thus material when a reasonable person, in what the physician knows or should know to be the patient’s position, would be likely to attach significance to the risk or cluster of risks in deciding whether or not to forego the proposed therapy.”<sup>34</sup>

On the patient-based standard the key idea is “materiality,” but that idea has proved somewhat vague at best and slippery at worst in the reported decisions.<sup>35</sup> To use one good formulation from the Supreme Court of South Dakota: “Material information is information which the physician knows or should know

---

29. See, e.g., *id.* at 218; Sawicki, *supra* note 27, at 829.

30. See, e.g., BERG ET AL., *supra* note 27, at 51.

31. 350 P.2d 1093 (Kan.), *reh’g denied*, 354 P.2d 670 (Kan. 1960) (clarifying the prior decision).

32. *Id.* at 1106; see Sawicki, *supra* note 27, at 830.

33. 464 F.2d 772, 776 (D.C. Cir. 1972).

34. *Id.* at 786–87.

35. See Sawicki, *supra* note 27, at 833.

would be regarded as significant by a reasonable person in the patient's position when deciding to accept or reject a recommended medical procedure."<sup>36</sup>

Although phrased broadly, most courts tend to focus on information "about the patient's diagnosis and proposed treatment; the treatment's risks and benefits; alternative procedures and their risks and benefits; and the risks and benefits of taking no action," what Sawicki calls the "standard risk-and-benefit-disclosure."<sup>37</sup>

Unfortunately for our purposes, in thinking about how the use of AI/ML fits into this picture, we need to move from these core examples of what is material into more penumbral examples. Without attempting to be comprehensive, let me give some examples of the kinds of cases the courts have dealt with, which we will use as scaffolding in the next Part.

At the start, I also want to emphasize a more general doctrinal limitation: some U.S. states have limited the legal obligation to get informed consent to surgical or other invasive procedures, and therefore do not require it in other medical contexts.<sup>38</sup> For such states, the use of medical AI/ML in surgical contexts may impose informed consent obligations but not, for example, in the PrEP case mentioned above.

B. THREE KINDS OF "PENUMBRA" INFORMED CONSENT CASES THAT MAY PROVE USEFUL FOR AI/ML

We now move from the core to what, I think it is fair to say, are more penumbral situations for U.S. informed consent. In my view these three categories of cases, more than the core ones, provide some useful analogies to medical AI/ML: provider experience and qualification cases, cases involving substitute providers (including "ghost" surgery), and cases about financial conflicts of interests. But because we are drawn to the shadows of these penumbral cases for our analogies is a general harbinger for what I will claim later in this Article: it may be hard to ground legal obligations of informed consent for medical AI/ML in the current doctrine.

---

36. *Wheeldon v. Madison*, 374 N.W.2d 367, 371 (S.D. 1985). There are many other slightly different formulations. For example: "A material risk is one which a physician knows or ought to know would be significant to a reasonable person in the patient's position in deciding whether or not to submit to a particular medical treatment or procedure." *Sard v. Hardy*, 379 A.2d 1014, 1022 (Md. 1977). "Materiality may be said to be the significance a reasonable person, in what the physician knows or should know is his patient's position, would attach to the disclosed risk or risks in deciding whether to submit or not to submit to surgery or treatment." *Wilkinson v. Vesey*, 295 A.2d 676, 689 (R.I. 1972) (citation omitted). It is not clear what difference, if any, the small differences in the formulation might mean for informed consent law in general, and there is no reason to think the differences matter for AI/ML cases specifically.

37. Sawicki, *supra* note 27, at 831 (collecting cases).

38. Steven G. Rado, *A Patient's Right to Know: A Case for Mandating Disclosure of Physician Success Rate as an Element of Informed Consent*, 18 HEALTH MATRIX 501, 504 (2008); *see, e.g., Morgan v. MacPhail*, 704 A.2d 617, 619 (Pa. 1997).

## 1. Provider Experience and Qualification

In *Johnson v. Kokemoor*,<sup>39</sup> a patient became a quadriplegic after basilar bifurcation aneurysm surgery and brought an informed consent claim arguing that it was a violation for the physician to fail “to divulge the extent of his experience in performing this type of operation” and that information about the doctor’s lack of experience was relevant because the surgeon’s relative inexperience went to the comparative risk of having the surgery.<sup>40</sup> The Supreme Court of Wisconsin decided that “[a] reasonable person in the plaintiff’s position would have considered such information material in making an intelligent and informed decision about the surgery.”<sup>41</sup> Sawicki, collecting cases, observes that a “handful of courts in other states have also held that information about a provider’s credentials or experience with a given procedure may need to be disclosed, particularly where those facts suggest there might be an increased risk of injury.”<sup>42</sup> Another handful of states “have held that providers who misrepresent their credentials in response to patient inquiries” might be liable for breach of informed consent.<sup>43</sup>

But these courts are, overall, in the minority. The substantial majority of courts have rejected the notion that the failure to disclose the physician’s experience or qualification breaches the duty of informed consent, on the theory that only information about the procedure itself is material.<sup>44</sup> And at least one court has rejected the application of tort even when a patient specifically asked about qualification and was misled, though this court suggested a fraud action might be viable.<sup>45</sup>

As doctrinal scaffolding, perhaps one might think there are cases where the failure to disclose AI/ML involvement in decisionmaking or actual procedures, especially when AI/ML is meant to enable a regular doctor to perform at the level

39. 545 N.W.2d 495 (Wis. 1996).

40. *Id.* at 497.

41. *Id.* at 505. The court notes at one point that “[a]ccording to the record the plaintiff had made inquiry of the defendant’s experience with surgery like hers” and “[i]n response to her direct question about his experience he said that he had operated on aneurysms comparable to her aneurysm ‘dozens’ of times,” but in fact this type of aneurysm is considered more difficult than the ones the defendant had handled in the past. *Id.* From the opinion itself, it is a little hard to tell whether the court considered that the patient had asked this question to be a *necessary* element of its finding liability, or whether a failure to affirmatively share lack of experience even when not asked directly would also allow a finding of liability. Other decisions do, though, put more emphasis on the patient having directly asked.

42. Sawicki, *supra* note 27, at 839 (collecting cases).

43. *Id.*

44. Sawicki cites the following cases as representing the “vast majority” view:

*Shock v. United States*, 689 F. Supp. 1424, 1425–26 (D. Md. 1988); *Duffy v. Flagg*, 905 A.2d 15, 21 (Conn. 2006); *Ditto v. McCurdy*, 947 P.2d 952, 958–59 (Haw. 1997); *Foard v. Jarman*, 387 S.E.2d 162, 166–67 (N.C. 1990); *Duttry v. Patterson*, 771 A.2d 1255, 1259 (Pa. 2001); *Wlosinski [v. Cohn]*, 713 N.W. 2d [16.] 20 [(Mich. Ct. App. 2006)]; *Johnson v. [.] Jacobowitz*, 884 N.Y.S.2d 158, 162 (N.Y. App. Div. 2009); *Zimmerman v. N.Y.C. Health & Hosp. Corp.*, 458 N.Y.S.2d 552, 554 (N.Y. App. Div. 1983); *Avila v. Flangas*, No. 04-95-00106-CV, 1996 WL 63036, at \*2 (Tex. App. Feb. 14, 1996); *Thomas v. Wilcaf*, 828 P.2d 597, 600–01 (Wash. Ct. App. 1992).

*Id.* at 840 n.81.

45. *Duttry*, 771 A.2d at 1259.

of a specialist, goes to the lack of experience/qualification of a doctor that ought to be disclosed.

## 2. Substitute Physicians: “Ghost,” Concurrent, and Overlapping Surgery

There are two related phenomena that some colloquially call “ghost” surgery: in overlapping surgery, “operations performed by the same primary surgeon such that the start of one surgery overlaps with the end of another” and such that “[a] qualified practitioner finishes noncritical aspects of the first operation while the primary surgeon moves to the next operation.”<sup>46</sup> By contrast, in concurrent surgery, the primary surgeon may be responsible for “critical parts” of operations that occur at the same time.<sup>47</sup> The latter practice received wide attention after the *Boston Globe*’s “Spotlight” team published an investigative report on concurrent surgeries taking place at the well-regarded Massachusetts General Hospital.<sup>48</sup>

There have been a few reported cases on breaches of informed consent in these kinds of situations. In *Hurley v. Kirk*,<sup>49</sup> a case involving a laparoscopic hysterectomy, the Supreme Court of Oklahoma held that “the doctrine of informed consent requires a physician to obtain the patient’s consent before using a non-doctor to perform significant portions of a surgery for which the physician was engaged to perform thereby subjecting the patient to a heightened risk of injury.”<sup>50</sup>

In *Perna v. Pirozzi*,<sup>51</sup> Perna consulted Dr. Pirozzi, a specialist in urology, who examined Perna and recommended surgery for the removal of kidney stones.<sup>52</sup> Pirozzi was part of a medical group that also included Drs. Del Gaizo and Ciccone and allegedly operated with no doctor having an individual patient but instead having all physicians in the practice treat all patients.<sup>53</sup> Indeed,

[I]t was not the practice of the group to inform patients which member would operate; the physicians operated as a “team,” and their regular practice was to decide just prior to the operation who was to operate. If, however, a patient requested a specific member of the group as his surgeon, that surgeon would perform the operation. Nothing indicated that Mr. Perna was aware of the group’s custom of sharing patients or of their methods for assigning surgical duties.<sup>54</sup>

---

46. Michelle M. Mello & Edward H. Livingston, *The Evolving Story of Overlapping Surgery*, 318 JAMA 233, 233 (2017).

47. *Id.*

48. Jenn Abelson et al., *Clash in the Name of Care*, BOS. GLOBE (2015) (describing controversies arising from scheduling concurrent surgeries, during which a surgeon may move back and forth between two surgical procedures without disclosure to patients).

49. 398 P.3d 7 (Okla. 2017).

50. *Id.* at 8.

51. 457 A.2d 431 (N.J. 1983).

52. *Id.* at 433.

53. *Id.*

54. *Id.*

Perna claimed he specifically requested Pirozzi for the surgery, a claim that none of the defendants contradicted.<sup>55</sup> The operation was performed by Dr. Del Gaizo, assisted by Dr. Ciccone, and Dr. Pirozzi was not present (in fact he was off duty on that day).<sup>56</sup> Del Gaizo and Ciccone were not aware that Dr. Pirozzi's name appeared on the consent form.<sup>57</sup> When Perna experienced postsurgical complications, he and his wife sued claiming, *inter alia*, "that Mr. Perna's consent to the operation was conditioned upon his belief that Dr. Pirozzi would be the surgeon."<sup>58</sup>

The New Jersey Supreme Court found that liability could lie, although it parceled that liability into different categories. As against Drs. Del Gaizo and Ciccone, the ones who *actually* performed the surgery when the patient thought it would be Dr. Pirozzi, the court held that this should be pursued as the tort of battery, not informed consent.<sup>59</sup> As the court wrote:

The medical profession itself recognizes that it is unethical to mislead a patient as to the identity of the doctor who performs the operation. American College of Surgeons, Statements on Principles, § I.A. (June 1981). Participation in such a deception is a recognized cause for discipline by the medical profession. See American College of Surgeons, Bylaws, art. VII, § 1(c) (as amended June 1976). By statute, the State Board of Medical Examiners is empowered to prevent the professional certification or future professional practice of a person who "[h]as engaged in the use or employment of dishonesty, fraud, deception, misrepresentation, false promise or false pretense . . ." N.J.S.A. 45:1–21. Consequently, a statutory, as well as a moral, imperative compels doctors to be honest with their patients.<sup>60</sup>

In the claim against Dr. Pirozzi, who was supposed to but did not actually perform the surgery, the court held that:

[T]he action follows from the alleged breach of his agreement to operate and the fiduciary duty he owed his patient. With respect to that allegation, the Judicial Council of the American Medical Association has decried the substitution of one surgeon for another without the consent of the patient, describing that practice as a "deceit." A patient has the right to choose the surgeon who will operate on him and to refuse to accept a substitute. Correlative to that right is the duty of the doctor to provide his or her personal services in accordance with the agreement with the patient. Judicial Council of the American Medical Ass'n, Op. 8.12 (1982).

---

55. *Id.*

56. *Id.* at 434.

57. *Id.*

58. *Id.*

59. *Id.* at 439.

60. *Id.* at 439–40.

Few decisions bespeak greater trust and confidence than the decision of a patient to proceed with surgery. Implicit in that decision is a willingness of the patient to put his or her life in the hands of a known and trusted medical doctor. Sometimes circumstances will arise in which, because of an emergency, the limited capacity of the patient, or some other valid reason, the doctor cannot obtain the express consent of the patient to a surrogate surgeon. Other times, doctors who practice in a medical group may explain to a patient that any one of them may perform a medical procedure. In that situation, the patient may accept any or all the members of the group as his surgeon. In still other instances, the patient may consent to an operation performed by a resident under the supervision of the attending physician. The point is that a patient has the right to know who will operate and the consent form should reflect the patient's decision. Where a competent patient consents to surgery by a specific surgeon of his choice, the patient has every right to expect that surgeon, not another, to operate.

The failure of a surgeon to perform a medical procedure after soliciting a patient's consent, like the failure to operate on the appropriate part of a patient's body, is a deviation from standard medical care. It is malpractice whether the right surgeon operates on the wrong part or the wrong surgeon operates on the right part of the patient. In each instance, the surgeon has breached his duty to care for the patient. Where damages are the proximate result of a deviation from standard medical care, a patient has a cause of action for malpractice. Although an alternative cause of action could be framed as a breach of the contract between the surgeon and the patient, generally the more appropriate characterization of the cause will be for breach of the duty of care owed by the doctor to the patient. The absence of damages may render any action deficient, but the doctor who, without the consent of the patient, permits another surgeon to operate violates not only a fundamental tenet of the medical profession, but also a legal obligation.<sup>61</sup>

The court remanded the case for further proceedings.<sup>62</sup> While confident that tort law permits redress for facts like these, the court admittedly struggled with how to characterize the claim in terms of the right tort category.<sup>63</sup> It seems plausible to me that a court today, faced with the facts of *Perna* and a more mature informed consent jurisprudence, might characterize the case more as a straightforward breach of informed consent tort. It is clear, though, from both cases, that liability can lie when someone other than the promised physician performs a critical part of a surgery.

Indeed, there is at least one case where the court allowed a patient to sue where the patient consented to a particular physician performing the entire surgery, but another physician performed the beginning of the surgery before the consented-to

---

61. *Id.* at 440–41 (footnotes omitted).

62. *Id.* at 441.

63. *See id.* at 438–39.

physician took over.<sup>64</sup> This might seem like a particularly good analogy to the AI/ML world, because patients may be unaware of the role of AI/ML in only a *particular part* of their treatment. For example, the physician performs the surgery herself, but it is based on an AI/ML recommendation of which surgical technique to use in this specific case.

### 3. Financial Conflicts of Interest

A third useful penumbral case, because it arguably goes to the *reasons* why a physician adopted a particular course, is informed consent and financial conflict of interest. Some courts have been receptive to this theory. The most famous one, read by most law students at some point in their law school studies, is *Moore v. Regents of California*,<sup>65</sup> involving the failure of a physician to disclose that he intended to derive a cell line from the patient's spleen leukocytes after performing a splenectomy on the patient.<sup>66</sup> The California Supreme Court held that “(1) a physician must disclose personal interests unrelated to the patient’s health, whether research or economic, that may affect the physician’s professional judgment; and (2) a physician’s failure to disclose such interests may give rise to a cause of action for performing medical procedures without informed consent or breach of fiduciary duty.”<sup>67</sup> In language that might be helpful in thinking about the AI case, the court went on:

[A] physician who treats a patient in whom he also has a research interest has potentially conflicting loyalties. This is because medical treatment decisions are made on the basis of proportionality—weighing the benefits *to the patient* against the risks *to the patient*. As another court has said, “the determination as to whether the burdens of treatment are worth enduring for any individual patient depends upon the facts unique in each case,” and “the patient’s

---

64. See *Grabowski v. Quigley*, 684 A.2d 610 (Pa. Super. Ct. 1996). The patient consented to a back surgery to be performed by Dr. Quigley but Dr. Bailes instead began the surgery when Quigley was not available, with Quigley taking over only midway through the procedure. *Id.* at 612–13. The court found that this could support a cause of action, though the court couched it in terms of battery rather than informed consent. *Id.* at 615. Intriguingly, the court seems to suggest it was framing the tort in this way to make it *easier* for plaintiffs to recover than in breach of informed consent:

Unlike an informed consent case where it must be shown that “as a result of the recommended treatment, the patient actually suffers an injury the risk of which was undisclosed, or the patient actually suffers an injury that would not have occurred had the patient opted for one of the undisclosed methods of treatment[,]” it is not necessary for a plaintiff to prove such specific medical findings under a theory of battery.

...

Where it is proven that an unauthorized invasion of a person’s personal integrity has occurred, even if harmless, the person is entitled to nominal damages.

*Id.* (citations omitted). For present purposes, I am more interested in the question of when the failure to inform the patient as to who would be doing the surgery is actionable than whether it is more correct to treat the claim as battery rather than breach of informed consent.

65. 793 P.2d 479 (Cal. 1990).

66. *Id.* at 481.

67. *Id.* at 483.

interests and desires are the key ingredients of the decision-making process.” A physician who adds his own research interests to this balance may be tempted to order a scientifically useful procedure or test that offers marginal, or no, benefits to the patient. The possibility that an interest extraneous to the patient’s health has affected the physician’s judgment is something that a reasonable patient would want to know in deciding whether to consent to a proposed course of treatment. It is material to the patient’s decision and, thus, a prerequisite to informed consent.<sup>68</sup>

The key analogy for AI/ML one might argue (we will evaluate it below) is that just as a physician who fails to disclose a financial conflict of interest deprives the patient of the ability to evaluate how an “extraneous” interest might be affecting judgment, a patient to whom the use of AI/ML is not disclosed is deprived of the ability to evaluate how that is affecting the physician’s judgment. As I will discuss in greater depth below, one can imagine some important distinctions between the two contexts, though they may just limit the cases relating to AI/ML for which the analogy holds (for example, to cases where the AI/ML may be optimized to something other than this particular patient’s health).

A related but separate strand of doctrine has to do with the pressure during the era of managed care for physicians to limit care based on financial incentives. The American Medical Association adopted an ethical opinion requiring physicians to disclose “any financial incentives that may limit appropriate diagnostic and therapeutic alternatives that are offered to patients or that may limit patients’ overall access to care,” while noting that this disclosure could come from the health plan itself and satisfy the obligation.<sup>69</sup> As Sawicki summarizes the lay of the land:

In a few cases, courts recognized the validity of claims that a physician’s failure to disclose financial incentives constituted a breach of duty, allowing them to proceed under theories of informed consent or malpractice. A Minnesota appellate court in 1997, for example, stated that a physician’s failure to disclose a kickback scheme “presents a classic informed consent issue.”<sup>70</sup>

Although these cases exist, it would be an exaggeration to suggest they ever fully bloomed.

This section has sought to set out the basic elements of the law of informed consent in the United States. As we saw, when it comes to what must be disclosed, the states are roughly evenly divided between a physician-based and a patient-based standard of disclosure. Although the run-of-the-mill informed consent case fact pattern involves something like the failure to disclose a potential

---

68. *Id.* at 484 (citations omitted).

69. Sawicki, *supra* note 27, at 842 (quoting AM. MED. ASS’N: COUNCIL ON ETHICAL AND JUD. AFF., OPINION E-8.132: REFERRAL OF PATIENTS: DISCLOSURE OF LIMITATION (2007)).

70. *Id.* (quoting *DAB v. Brown*, 570 N.W.2d 168, 171 (Minn. Ct. App. 1997)) (citing *Shea v. Esensten*, 208 F.3d 712 (8th Cir. 2000); *Neade v. Portes*, 739 N.E.2d 496 (Ill. 2000)). For a discussion of the law and ethics surrounding informed consent in these settings, see Mark A. Hall, *A Theory of Economic Informed Consent*, 31 GA. L. REV. 511 (1997).

risk, I have suggested that there are three more penumbral recurring fact patterns that will prove more useful in thinking about the medical AI/ML context. These cases involved: provider experience or qualifications; substitute providers and “ghost” surgeries; and provider conflicts of interest. In the next section, we will consider, among other things, what those cases might have to say about medical AI/ML.

D. WHAT TO TELL THE PATIENTS: APPLYING INFORMED CONSENT CASE LAW TO  
MEDICAL AI/ML

Now that we have a grasp on both the integration of AI/ML into medicine, especially as to making recommendations for treatments, and the core and penumbra of modern U.S. informed consent law, we can see our animating question is actually two related questions. First, when, if ever, under the law of informed consent do physicians need to tell patients that AI/ML was involved in helping to guide the treatment decision the physician ultimately adopted? Second, conditional on believing such disclosure is sometimes appropriate, how much detail must they share about the AI/ML recommendation and the AI/ML system itself? For example, if the system is more of an opaque one with relatively little explainability is *that* something that must be disclosed to the patient? If the system is less opaque, does the physician have to explain the reasons why (in the sense of what characteristics of the patient were at play) the system reached its decision?

The existing doctrine does not address, and indeed academics and lawyers have not really started thinking about, this issue. This is not unusual when a new technology emerges—we struggle with whether it can be assimilated into existing doctrines or whether it requires something new. One familiar approach to this recurring phenomenon is what one might call the “doctrinal—analogue approach” or perhaps simply the “common law” approach because it mirrors the way common law judges have for centuries dealt with new issues—drawing analogies to cases that have come before or at least more familiar hypothetical examples. The common law nature of this approach is perhaps quite fitting for a doctrine such as informed consent that has developed in fits and starts in the common law setting.

Before stepping on to other approaches—more openly empirical or normative—in the next Part, it is worth determining just how far this approach can take us. At the risk of showing my cards too early: my tentative conclusion is that following this approach suggests that in most instances, the failure to inform the patient of the use of medical AI/ML will not violate the law of informed consent. I offer the best, or perhaps more accurately “my best” or “most creative,” arguments for the contrary result nonetheless, but I am not sure if I were to recast myself in the role of judge that I would find them completely persuasive. There are, though, a few instances where a claim of a violation seems more plausible, even if not plausible, as I try to adumbrate in the remainder of this Part.

## 1. Starting With the Standards

Although recognizing that the doctrinal distinction between physician-based and patient-based standards for disclosure can be overstated, it is nonetheless useful to start there. Under the physician-based standard, a physician must make “those disclosures which a reasonable medical practitioner would make under the same or similar circumstances.”<sup>71</sup> Here, we run into a bootstrap problem—given a brand new technology, and what is more a new way of using a technology in medical practice, how can we say what reasonable medical practitioners in fact do?<sup>72</sup> One possible way to cut the Gordian knot is to reach for analogies, or “similar circumstances.” But, for better or worse, we find several candidate analogies.

One way of conceiving medical AI/ML is as simply as another input to a reasoning process. If we were able to open up the thought process of the typical physician deciding what surgical technique to use or whether to recommend a particular patient to go on PrEP, we would find a lot of potential inputs. The physician may be drawing on a varied assortment of vague memories from a medical school lecture, what the other doctors during residency did in such cases, the latest research in leading medical journals, the experience with and outcomes of the last 30 patients the physician saw, etc. It is beyond cavil that a physician who fails to describe each of these steps of the reasoning does not violate the law of informed consent; indeed, a physician who did so would strike most patients as odd. For this reason, a duty to disclose use of medical AI/ML seems like an unlikely fit for the physician-centered standard.

Now perhaps the situation might be different if a patient were to specifically ask a question—one could imagine an oncologist in the patient role asking her treating oncologist whether she had considered a recent article in the field in recommending X over Y treatment, and indeed in some of the cases discussed above this is exactly what happens. If a patient were to ask a physician “is this what the AI/ML recommended?” or “did you rely on an AI/ML?” and the physician were to mislead the patient by falsely denying that they did so, then that might be a good basis for a breach of informed consent claim (assuming, as always, that all the other elements of the tort are satisfied), analogous to the cases about credential misrepresentation cases discussed above. But few patients will know to ask, although that might change in a future where the technology is more ubiquitous and well-known.

---

71. *Natanson v. Kline*, 350 P.2d 1093, 1106 (Kan. 1960).

72. This also highlights one of the disadvantages of the physician-based standard in informed consent (indeed, there is a similar problem for medical malpractice). *See* Sawicki, *supra* note 27, at 830. Newly adopted technologies, even if they are better, take a long time to get integrated as the standard of care. Therefore, when the law focuses on what a reasonable medical practitioner would say (or do), it is likely to disfavor disclosing the use of (or using) new technologies.

On the patient-based standard, the question is what a reasonable patient would find material. Even on this standard, if the analogue is all the inputs to the decisionmaking of the physician, including the most recent journal articles, it seems implausible that a court would treat this as material and force disclosure.

Perhaps things look a little different when it comes to more opaque forms of AI/ML. To use a fanciful example, suppose in deciding whether to operate or which drug to prescribe, a physician consulted a Magic 8-Ball or Astrology. Surely, a reasonable patient would want to know that *that* was the reason for a medical decision and would find that material in deciding whether to adopt the recommended course from this physician or instead seek out a second opinion. The argument is that relying on opaque medical AI/ML is unlike relying on journal articles or medical school teaching and more like the 8-Ball or Astrology.

But is that really a fair analogy? With the Magic 8-Ball or Astrology (with apologies to Susan Miller!) not only can the physician not explain *why* it works, but she also has no epistemic warrant *that* it works. Conversely, with medical AI/ML (especially its more opaque forms), the argument is that it is more like Aspirin, in that the former may be true but the latter is not—assuming that physicians have good reason to believe the AI/ML is likely to lead to better decisions.<sup>73</sup> The epistemic warrant for that proposition need not be firsthand knowledge—we might think of medical AI/ML as more like a credence good, where the epistemic warrant is trust in someone else. That might seem strange, but upon reflection this is true of most FDA-approved pharmaceuticals. The physician is likely quite ignorant of the underlying trial design or results that led FDA to believe that the drug was safe and effective, but her knowledge that it has been FDA-approved supplies the necessary epistemic warrant. That might suggest that part of what will matter is whether there are similar indicia of reliability or preclearance processes for medical AI/ML.

This would lead to a rule whereby physicians would need to disclose reliance when such indicia of reliability (or other epistemic warrants) are absent, but not when they are present. Such a distinction is interesting, but does it really reflect the distinction “a reasonable patient” would draw? That partially depends on whether we think the answer to that question is an empirical one (what do actual patients care about?) or an at least partially normative one (how can we construct the reasonable patient to satisfy our normative goals of having an informed consent doctrine?). Later in this paper, I wade into such waters. But one thing such a rule has going for it is that through the doctrine of informed consent, we might force the purchasers, users, and makers of medical AI/ML into what we think of as a desirable social policy—some form of preclearance review or other indicia of reliability and improvement in medical practice. But if *that* is what we are aiming for, why not impose the requirement directly rather than try to back into it in such a roundabout way via the doctrine of informed consent?

---

73. For discussions of this point and a comparison to other medical knowledge, see London, *supra* note 26.

## 2. Deeper Into the Analogies

Let us take a moment to catch our breath after dashing through that rabbit hole. Although it seemed quite natural to start there, I am not sure the distinction between physician- and patient-based standards has proven so useful in helping us understand the medical AI/ML case that it should serve as the key branching point for the analysis. Indeed, it seemed that to make use of the standards we had to jump quickly into the contested analogies. For that reason, from this point on, I will instead dive directly into the possible analogies without trying to bisect their treatment on the two standards, especially because in many instances I am not confident the two standards would suggest different outcomes.

I have suggested above that if we analogize medical AI/ML to other inputs of medical decisionmaking (training, journal articles, etc.) it seems difficult to conclude that the failure to disclose reliance on AI/ML would violate the law of informed consent. But, that may just mean we need to think about different analogies. We could instead view AI/ML, anthropomorphized, as another personage. If in *Hurley v. Kirk*,<sup>74</sup> discussed above, it was a breach of the duty of informed consent to have a non-doctor perform part of a surgery,<sup>75</sup> then, the argument goes, it is a similar breach to make a treatment decision based on the decision of non-doctor in the form of the AI/ML. Similarly, if in *Perna v. Pirozzi*,<sup>76</sup> it was a violation of the duty of informed consent to have someone other than Dr. Pirozzi, the doctor agreed upon, to perform the surgery,<sup>77</sup> here, it is a violation to have the AI/ML play an undisclosed role in formulating a recommendation.

How useful is this analogy? It seems strongest when, analogous to *Perna*, the substitution of AI/ML for a doctor is complete. That said, cases like *Hurley* suggest even when the consented-to doctor is present, but not exclusively performing the operation, there may be an informed consent problem if that is not disclosed. On the other hand, both cases involve surgery and carrying out a treatment, where there is arguably a stronger expectation of exclusivity, rather than formulating a recommendation for treatment. Indeed, one striking thing about the mass of informed consent cases is how often the subject is surgery or something else performed on the patient and how seldom it is diagnosis—and as discussed above, in some states the doctrine is officially limited to such cases.

To test the persuasiveness of the analogy, one should imagine a variant of *Perna* involving diagnosis and recommendation: The patient goes to see Dr. House, a renowned diagnostician, but the diagnosis is instead carried out by one of his junior colleagues without disclosing that to the patient. It seems possible to me that in a jurisdiction that does not limit the legal obligation of informed

---

74. 398 P.3d 7 (Okla. 2017).

75. *Id.* at 8–9.

76. 457 A.2d 431 (N.J. 1983).

77. *Id.* at 433, 441.

consent to surgery, an informed consent cause of action could lie in such a case, though the intuition is helped in the hypo by House's stature as a well-known diagnostician. Just as a patient consents to surgeon X not Y to do the surgery, so the patient consents to physician X not Y to make a diagnosis. To push it to the extreme, it seems to me that if one found out that a urologist was making diagnostic recommendations as to breast cancer, and the patient was not informed, that would set up a plausible claim for lack of informed consent. The analogy seems stronger to me the more the AI/ML recommendation is automatically adopted, where it becomes more of a true substitute, and weaker to the extent it ends up as just one input in the recommendation.

If I were forced to argue for a breach of informed consent in this case, this is an analogy I would use. But, stepping back, how persuasive is it? I am not sure. I worry that part of the intuition pump in my Dr. House hypo or the urologist diagnosing breast cancer is that there is likely medical malpractice in the activity; that is, the "intuition pump" hypo is picking up the wrong thing—malpractice rather than breach of informed consent. This is possible, but if I negate that unstated element of the hypothetical, by saying that this urologist, it just so happens, renders his diagnosis in comportment with the prevailing standard of care, it is not clear to me that the intuition of wrongdoing goes away.

A third analogy is the financial conflict of interest cases. What is nice about such cases as an analogy is that they get at the notion of compromised *reasoning* better than do the analogies to substitute providers. In these cases, the financial motive (to commercialize a cell line, to satisfy the financial incentives of managed care) is viewed as problematic precisely because it compromises the reasoning of the physician. Can the AI/ML case be viewed in that light? Possibly—just as financial incentives, undisclosed, might be viewed as material to a patient because they compromise reasoning, so the presence of an AI/ML system does the same. But the analogy is a bit procrustean. The AI/ML does not so much compromise reasoning as to replace (or at least supplement it) it with a different kind of reasoning.

It seems to me that the analogy gets stronger the more opaque the machine learning is—having a recommendation whose basis one does not understand might be argued to disrupt reasoning about an ultimate recommendation the way the effects of financial conflict of interests on reasoning might be thought to disrupt reasoning. But the analogy is far from perfect. Moreover, in both *Moore* and the managed-care cases, the thing affecting reasoning was extraneous, and at least potentially contrary, to the patient's interests. Medical AI/ML, by contrast, is ideally being adopted because it is meant to improve patient care, which *is* in the patient's interest. I say "ideally" because one could imagine cases where the adoption story looks more like the managed-care one—for example, a hospital system adopts medical AI/ML to reduce costs after a study that shows it does not affect patient care one way or the other or improves some patient care and worsens other patient care or (most cynically) leads to a small diminution in the

quality of patient care that is cost-justified.<sup>78</sup> Such cases look more like *Moore* and especially like the managed-care cases, though once again the latter are already on weak footing in terms of what courts entertain.<sup>79</sup>

Finally, in some instances, the provider qualification cases might be helpful analogies. One goal of AI/ML, as evinced by IDx-DR, is the democratization of expertise—to enable the average physician to achieve the level of sophistication in diagnosis or treatment of the specialist.<sup>80</sup> If one squints hard enough, perhaps we can see the physician as unqualified but for the use of the AI in some of these instances, and thus it is necessary to disclose the use of AI/ML just as it would be to disclose the lack of experience in *Johnson v. Kokemoor*.<sup>81</sup> Similarly, in a more recent case involving a patient who underwent a complicated heart surgery, the Iowa Supreme Court held that, given that the “heart procedure is a very complicated procedure . . . harder to perform than a heart transplant,” the information that a physician had no experience with the procedure and no specific training could be material to a reasonable person deciding whether to consent to the procedure, thus giving rise to a jury-triable issue.<sup>82</sup> The complexity of the AI/ML

---

78. This hypothetical case may be contrasted with other more difficult ones where the goal is one in keeping with the patient’s interest, but problem may lie in “problem formulation,” or more specifically “the label on which the algorithm is trained.” Obermeyer et al., *supra* note 12, at 450–51. Without delving too deeply into the details, in a recent paper, Obermeyer and colleagues examine the reason why an algorithm commonly used by hospitals to predict patients who would benefit from “high-risk care management” programs that provide extra resources and attention to certain patients shows racial biases. *Id.* at 452.

They show that:

The bias arises because the algorithm predicts health care costs rather than illness, but unequal access to care means that we spend less money caring for Black patients than for White patients. Thus, despite health care cost appearing to be an effective proxy for health by some measures of predictive accuracy, large racial biases arise. We suggest that the choice of convenient, seemingly effective proxies for ground truth can be an important source of algorithmic bias in many contexts.

*Id.* at 447. Part of what is so interesting about this case is that the authors believe that “the algorithm manufacturer’s choice to predict future costs is reasonable,” but not “the only reasonable choice,” and it is one that produces the race differentiated results while trying to act in the patient interest. *Id.* at 450–51. The bias and the mechanism by which it arose were largely unknown before the leaders in the field investigated, and not the kind of thing that the average physician or even hospital system could know ahead of time and disclose to a patient (let alone for the patient to understand it and its overall implication for their particular care in the way that the informed consent doctrine ideally aims at). As with questions of malpractice, I think we face a question of division of labor. One can believe that this *is* a real problem, and yet also believe the law of informed consent is an unlikely place to solve it, preferring, for example, FDA or other forms of premarket review as more plausible. This might be particularly clear if one has doubts about the deterrent force of informed consent liability, as discussed in Part IV below.

79. There would also be more difficult questions in what we might call “mixed motive” cases. For example, imagine a case where the AI/ML systems are adopted because they both improve care *and* reduce costs, but there are two other AI/ML systems on the market that are not adopted, for among other reasons, because they reduce cost *less* than this one does.

80. See Price, *supra* note 12, at 17–18.

81. 545 N.W.2d 495, 498 (Wis. 1996).

82. *Andersen v. Khanna*, 913 N.W.2d 526, 542 (Iowa 2018).

engine in recommending a particular course of action might be loosely analogized to this.

These analogies might also tell us something about what must be disclosed. If we will tolerate the anthropomorphizing of the AI/ML as another “member of the care team,” then perhaps it is necessary to give the patient information about *its* expertise. This could be information such as whether the AI/ML was reviewed by FDA or another regulator, questions about false positives and negatives, etc. Perhaps one might argue that physicians using the AI/ML might be required to explain how they use it (follow the advice always, overrule it in certain situations, etc.) and perhaps how much they do or do not understand about the AI/ML due to its opacity.<sup>83</sup>

Here, there are a handful of cases that might be useful in building such an argument. In *DeGennaro v. Tandon*,<sup>84</sup> involving an accident during dental surgery, the Connecticut Appellate Court found that a dentist’s use of

equipment with which she was unfamiliar and . . . an office that was not ready for business, [was] the type of provider specific information that a reasonable person in the plaintiff’s position would consider material in weighing the risks of this dental procedure and in deciding whether a viable alternative was to seek a different provider to perform the procedure.<sup>85</sup>

---

83. A recent article in the *AMA Journal of Ethics* captures well the challenges from a physician’s perspective of meeting the ethical (not the legal) requirements of informed consent when the AI/ML is opaquer:

The opacity of an AI system can make it difficult for health care professionals to ascertain how the system arrived at a decision and how an error might occur. For instance, can physicians or others understand why the AI system made the prediction or decision that led to an error, or is the answer buried under unintelligible layers of complexity? Will physicians be able to assess whether the AI system was trained on a data set that is representative of a particular patient population? And will physicians have information about comparative predictive accuracy and error rates of the AI system across patient subgroups? In short, if physicians do not fully understand (yet) how to explain an AI system’s predictions or errors, how could this knowledge deficit impact the quality of an informed consent process and medical care more generally?

Daniel Schiff & Jason Borenstein, *How Should Clinicians Communicate with Patients About the Roles of Artificially Intelligent Team Members?*, 21 *AMA J. ETHICS*. E138, E140 (2019), [https://journalofethics.ama-assn.org/sites/journalofethics.ama-assn.org/files/2019-01/cscm3-1902\\_0.pdf](https://journalofethics.ama-assn.org/sites/journalofethics.ama-assn.org/files/2019-01/cscm3-1902_0.pdf) [<https://perma.cc/K4DN-U5U4>]. Unfortunately, the main recommendations they make are for professional communities to grapple with this, for physicians to better educate themselves, and perhaps push for more “transparent” algorithms. *Id.*

84. 873 A.2d 191 (Conn. App. Ct. 2005).

85. *Id.* at 200. One interesting, albeit unusual case, somewhat on point involves breach of informed consent by a hospital. In *Nguyen v. IHC Medical Services*, 288 P.3d 1084 (Utah Ct. App. 2012), the plaintiff’s son was severely injured in a car accident and was admitted to a pediatric intensive care unit. *Id.* at 1086. The child needed constant ventilation, but his physician wanted to perform a CT scan which required the child be moved to a different floor. *Id.* at 1087. The hospital did not have equipment that could provide mobile ventilation, but a piece of equipment had been brought to the hospital for use in a sales demo that did have this capacity. *Id.* The child’s physician was on the committee evaluating the equipment and requested it be used to allow for transport of the child between floors. *Id.* The committee allowed the use, but during transport the machine failed, and the child died shortly thereafter. *Id.* The

On the other hand, I have some doubts about the analogy. When a physician recommends Aspirin, they do not tell the patient they have no idea how it works.<sup>86</sup> Yet we do not think that is an instance of breach of informed consent, so it is hard to see the failure to disclose a lack of understanding of an AI/ML or even that a mechanism is inexplicable by human reasoning at the current time as problematic. More generally, given that many courts have not imposed duties to disclose information about lack of experience or qualification unless a patient asks,<sup>87</sup> it may be that this is not a strong foundation on which to build an analogical house.

Where does that leave us? The doctrinal–analogical path, in the common law way, has left us with competing possible analogies. In the end though, on the general question of whether the failure to disclose the use of AI/ML violates the law of informed consent, if one was scoring the match, I think that the skeptics probably win the day. Perhaps there are certain cases where the argument for a violation is more plausible—when the medical AI/ML is more opaque, when it is given an outsized role in the final decisionmaking, when it is used to reduce costs or optimize for a goal other than improving patient health, and cases involving

---

Court of Appeals of Utah held that “[w]hen a hospital authorizes the use of an unfamiliar piece of equipment under the hospital’s control, the hospital should have an independent duty to obtain informed consent—a duty that may well be typically discharged, as a practical matter, by the treating physician on the hospital’s behalf.” *Id.* at 1092. The case has not particularly borne fruit, but it does raise at least the possibility that a *hospital system* might itself also face informed consent duties when it adopts AI/ML, or permits individual physicians to do so, especially when those systems have not been well-tested in the hospital itself. Nicolas Terry, in an article about the possible use of a Google Glass visor in surgical settings notes the *Nguyen* case as well as makes the point that the case is an outlier:

There is a good argument that using Glass during diagnosis or surgery may not pose a material risk that requires any additional informed consent. First, in contrast to the major innovations that occur daily in the health care industry, Glass may seem a relatively modest addition. Second, Glass could replace a potentially unsafe technique, such as walking away from the patient to view an image on a computer screen or a light table, with a less risky one. The use of any new technology in the operating room (OR) as a surgical adjunct (e.g., the use of a new surgical technique) may, however, require additional risk disclosure. Another argument could be made: that using a wearable device for any purpose in the OR has the potential to transform a procedure, process, or workflow so significantly that it creates the potential for outcomes or risks that should be communicated to a patient to secure informed consent. In practical terms, this issue will reduce to the persuasiveness of expert testimony on the materiality of the risk of using Glass, whether the jurisdiction follows the patient expectations test or the customary standard for reasonable disclosure. Courts have held consistently that the physician, not the health care institution, owes the patient the duty of informed consent despite cogent arguments that this duty should have migrated to the institution. An area of weakness, however, may be where the institution has better information about risks than the physician, which could be the case if the institution adopted and implemented an experimental program, such as equipping a department’s professional staff with Glass.

Nicolas P. Terry et al., *Google Glass and Health Care: Initial Legal and Ethical Questions*, 8 J. HEALTH & LIFE SCI. L. 93, 99–100 (2015).

86. See, e.g., London, *supra* note 26, at 17 (using Aspirin as an example of a medical intervention for which we cannot explain its mechanism).

87. See Sawicki, *supra* note 27, at 840–41.

surgery.<sup>88</sup> But overall, I do not think the existing doctrine strongly supports applying the tort here.

### III. IS THAT THE RIGHT ANSWER? BEYOND THE DOCTRINAL APPROACH

We have now seen where the doctrinal approach is likely to take us, but of course, it is not the only way to think about the problem. Even if the courts came to the same conclusion as I do as to the obligations to disclose AI/ML usage under the *existing* case law—that is, largely *against* obliging disclosure—a legislature could come in and change the state of play by specifically setting out when the use of AI/ML must be disclosed or else informed consent liability might lie. Courts might also look beyond the existing doctrine in thinking about this problem.

Should they do so? Here, I consider two approaches to answering the question from outside the case law. The goal is to begin a conversation, not definitively answer it.

#### A. AN EMPIRICAL APPROACH

Using both the physician-based and patient-based standard, determining what should be disclosed can be seen as resting, at least in part, on a set of empirical questions.

Under the physician-based standard, a physician must make “those disclosures which a reasonable medical practitioner would make under the same or similar circumstances.”<sup>89</sup> Well, one way of figuring out whether it should be disclosed is to ask what do physicians do in practice?

In any trial, there might be expert testimony on the matter. But we need not wait for litigation, where shopping for experts to support one’s case is rampant.<sup>90</sup> We could go out in the field and try to measure what physicians who use AI/ML currently do in terms of disclosure.

Is this a satisfying way of approaching the issue? One problem is that we are at the dawn of the practice, and few currently use AI/ML at all, so our sample is likely to be small. Indeed, one might worry about bootstrapping from the practice in this moment of uncertainty to establish a standard of what we *ought* to make physicians do. Now perhaps we could supplement this a little by also asking physicians in the same practice areas who do not yet use AI/ML what they think

---

88. One can imagine a possible future where more and more of the surgery, perhaps all of it, is turned over to an AI/ML system as a sort of “autopilot” with the surgeon there to intervene if something unexpected happens. Such a future may not be all that far off. Consider, as a first step toward this, the *da Vinci Surgical System*. It “facilitates minimally invasive robotic surgery by translating the surgeon’s natural hand movements into corresponding micro-movements of instruments positioned inside the patient through small incisions, or ports,” and has already given rise to several lawsuits. Catherine Mullaley, *Washington Supreme Court Holds that Medical Device Manufacturers Have a Duty to Warn Hospitals* - Taylor v. Intuitive Surgical, Inc., 43 AM. J.L. & MED. 165, 165 (2017).

89. Natanson v. Kline, 350 P.2d 1093, 1106 (Kan. 1960); Sawicki, *supra* note 27, at 835.

90. See, e.g., Christopher Tarver Robertson, *Blind Expertise*, 85 N.Y.U. L. REV. 174, 177 (2010).

the standard of disclosure should be, though as I will discuss in the next section, whether that clarifies or muddies the water depends a lot on what one thinks the doctrine of informed consent is supposed to do.

An allied approach—one that blends the empirical with the analogical but might generate a more robust sample size—would empirically examine the prevailing practices for using adjacent technologies. Here I think the use of “dumb” (relatively speaking) computer decision aids by physicians is probably the closest to how AI/ML is likely to be used in the foreseeable future. My own lay impression (but empirical legal studies is all about determining the actual truth of the matter!) is that few physicians explicitly disclose that they have used a computer decision aid “in the background” in deciding on a course of treatment,<sup>91</sup> which on the empirical path, would be a knock against imposing a duty of disclosure here.

On the patient-based standard, the empirical question goes to what patients view as material. The question, to use one formulation, is whether this is “information which the physician knows or should know would be regarded as significant by a reasonable person in the patient’s position when deciding to accept or reject a recommended medical procedure.”<sup>92</sup>

It is possible to empirically measure how significant patients find the use of an AI/ML to help direct their case in deciding whether to accept or reject a recommendation for care. The simplest (but by no means simple, just ask anyone who has run a vignette study) way to do this would be to present participants with vignettes where a piece of information relating to the use of AI/ML is present or absent and see if it makes a difference as to whether the patient would accept or reject the recommendation only when the additional information is included.

I have my own guesses about what such an empirical examination might reveal—that the patients will care a lot when the decision is an important one, and when AI/ML is not just one of many inputs but rather the predominant input or automatically followed—but it is just a guess. The goal of the empirical approach is to actually measure what patients find material.<sup>93</sup>

---

91. I say “in the background” to contrast this with the use of decision aids by physicians with patients aimed at shared decisionmaking, for example in neonatology. See, e.g., Ursula Guillen & Haresh Kirpalani, *Ethical Implications of the Use of Decision Aids for Antenatal Counseling at the Limits of Gestational Viability*, 23 SEMINARS IN FETAL & NEONATAL MED. 25, 26 (2018).

92. *Wheeldon v. Madison*, 374 N.W.2d 367, 371 (S.D. 1985).

93. Indeed, there has been at least one empirical investigation I know of that, although not exactly keyed to the question of interest, get close to it. In a 2017 report, PWC reported the results of a YouGov survey of over 11,000 respondents across the world to gauge comfort with robotic surgery. PWC, *WHAT DOCTOR? WHY AI AND ROBOTICS WILL DEFINE NEW HEALTH* 24 (2017), <https://www.pwc.com/gx/en/news-room/docs/what-doctor-why-ai-and-robotics-will-define-new-health.pdf> [<https://perma.cc/779R-GRYG>]. They:

[A]sked our survey participants if they would be willing for a robot to perform a minor or major surgical procedure instead of a doctor if studies showed that they could do it better than a doctor (e.g. more quickly, more accurately, with a faster recovery time).

We defined a minor surgical procedure as a non-invasive or minimally-invasive surgery, such as cataract surgery or laser eye surgery . . . .

. . .

As I will emphasize below, it is not clear that even well-designed empirical studies would completely get at the law's requirement of materiality on the patient-centered standard. The standard refers back to a "reasonable person in the patient's position," which is inherently a moralized standard and one would have to figure out how to translate from raw data about what *actual* persons cared about to a conception of what a *reasonable* person should care about. But still, raw data on patients' views about AI/ML in medicine and whether it matters to them could be useful, and as far as I know, no one has published precisely on this issue yet.<sup>94</sup>

But even for the raw data, there is an inherent trickiness about the emerging nature of AI/ML: even if one were to undertake the empirical examinations suggested above, there is a question of when. We might get a different view about, in particular, what matters to patient's decisionmaking regarding AI/ML if we conducted the examination today rather than five or ten years from now when AI/ML is more commonly and openly implemented in healthcare. That is not uncommon. Technologies diffuse slowly and our attitudes change. But here is where it gets tricky: whether the use of AI/ML to help direct their care matters to patients depends in part on (a) how commonly it is used and (b) how much patients know about how commonly it is used, which is itself at least somewhat related to (c) whether we require disclosure about the use of AI/ML today.

To use some provocative analogies to drive home the point: consider the way in which the law has historically "exceptionalized" (to use the language of its critics) the treatment of information pertaining to the risks involved with HIV

---

We defined major surgery as invasive surgery—such as replacement of a knee or hip joint, removal of a tumour, or heart surgery.

*Id.* at 35. For minor surgeries, they found that a low of 36% (in the UK) and a high of 73% (in Nigeria) of survey respondents would be willing to undergo *minor* surgery performed by a robot instead of a doctor. *Id.* Unsurprisingly many fewer persons were willing to undergo major surgery performed by a robot instead of a doctor. As they summarize their findings, survey respondents in:

Nigeria, Turkey and South Africa were far more willing (69%, 60%, and 51% respectively) than unwilling to undergo major surgery performed by a robot. Again, the UK was the least willing of all the countries surveyed (27%) and Belgium (60%), Germany (59%), the UK (58%), Sweden (55%), the Netherlands (53%) and Norway (51%) were all much more unwilling than willing to undergo major surgery performed by a robot.

*Id.* at 37 (emphasis omitted). It is not clear how much these results depend on how the question was asked and also how well they would transfer to AI/ML advising a physician as opposed to a robot performing a surgery, but it is helpful in starting to think about the empirical approach and what it might teach us.

94. For the closest data I have seen, see the preceding footnote. There are also more sophisticated questions about how individuated the "reasonable person" should be, for which the data might provide some answers. Suppose we got one response in cancer and another in cardiac, should it be the reasonable cancer patient? Or what about gender or racial or class differences? Of course, this is a classic question that can be asked about the use of reasonable person across law. One is reminded of Alan Patrick's Herbert's description of the "reasonable man" (now, thankfully, we would say "person") as an "excellent but odious character" who "stands like a monument in our Courts of Justice, vainly appealing to his fellow-citizens to order their lives after his own example." A.P. Herbert, Fardell v. Potts: *The Reasonable Man*, in MISLEADING CASES IN THE COMMON LAW 14, 16 (1927).

through, among other things, special rules criminalizing the failure to disclose HIV before sex that are not applied to other medical issues.<sup>95</sup> Or consider the way in which imposing special informed consent requirements relating to abortion are often an attempt to dissuade women from getting abortions.<sup>96</sup> Choosing to exceptionalize treatment of AI/ML as opposed to, for example, other information about why a physician reached the recommendation he or she did (for example, memories of medical school lectures, journals, etc.) is not simply a value-neutral decision that is purely data-driven.

What does all this mean? Empirical analysis of physician or patient reactions can provide some information is useful, but translating this analysis into a recommendation for action requires a set of decisions that seem ineluctably value-laden.<sup>97</sup> In the next section, I try to wade into those normative waters.

#### B. A NORMATIVE APPROACH

In trying to determine whether the existing doctrine should be modified in its treatment of AI/ML, I will argue we cannot avoid being normative to some extent. In asking ourselves about a “reasonable” patient or physician there is an inherent moralization in the construction of that person, and especially in the patient-based standard, an implicit normative judgment that we care about what failure to disclose would aggrieve a “reasonable” patient and will put to one side the reactions of the “unreasonable” one.

But perhaps it is best not to merely be normative “around the edges,” but to instead be more thoroughgoing about it. Perhaps the best way to think about how the law of informed consent should apply to medical AI/ML is to ask, normatively speaking, why do we have informed consent *at all* and work backwards from those reasons to the AI/ML case?

Of course, it is useful to emphasize that what the law requires for informed consent and what is considered ethical as a nonlegal requirement of the medical profession *can* diverge, and indeed the two *do* diverge in practice. That is, in many instances, it is not legally actionable to fail to make the disclosures that medical ethics would impose upon physicians. At the same time, it seems entirely natural in thinking about how to apply (or extend) a legal doctrine to a new area to return to the well, especially with a doctrine like informed consent that is so clearly married to ethics in a way that, say, the statute of frauds or the rule against perpetuities is not.

---

95. *E.g.*, Kim Shayo Buchanan, *When Is HIV a Crime? Sexuality, Gender and Consent*, 99 MINN. L. REV. 1231, 1339–42 (2015); Craig Konnoth, *An Expressive Theory of Privacy Intrusions*, 102 IOWA L. REV. 1533, 1555 (2017).

96. MARK A. HALL, DAVID ORENTLICHER, MARY ANNE BOBINSKI, NICHOLAS BAGLEY & I. GLENN COHEN, *HEALTH CARE LAW AND ETHICS* 161–62, 753–55 (9th ed. 2018) (collecting and discussing cases).

97. There are also institutional role questions about which actors could make use of such information were it available. Courts might be reluctant to rely on this kind of empirical data and instead prefer expert testimony on the issue, but perhaps legislators or other regulators might be keener on letting it influence their own potential interventions.

As befits a Symposium issue, this will be but a brief jaunt into the vast ethics literature on medical informed consent, but hopefully it is enough to show that thinking in normative terms may be useful in thinking about the AI/ML context.

Before we start, it might be beneficial to consider one meta-objection: if there is any doubt as to whether to disclose, why not just disclose? This highlights one of the odd central tensions of the focus on materiality in the patient-centered standard: the obligation to disclose attaches for things that individuals *really* care about (information that would motivate a different decision), but if they do *not* really care, then what is the downside to disclosing even in that case? One answer is that the list of things that people do not really care about is huge. Disclosing all of the things people do not really care about is costly, especially if we take seriously the mantra that informed consent is a process, not a piece of paper, and as such we ought to aim for true interactivity and understanding by the patient rather than a signature on a piece of paper they barely read and understood even less.<sup>98</sup> In particular, if we were to take seriously the need to get true informed consent about the use of AI/ML in a particular treatment recommendation, the amount of complex information that an average patient will have to absorb to make a truly informed decision is quite large and this is no small undertaking.

Of course, it is possible to merely “check the box” by saying or putting into writing something fairly vague like “in reaching the best treatment recommendation I have consulted a large number of inputs including recommendations of supervised and unsupervised machine learning . . .” and be done with it. If the goal was to merely “touch the base” or satisfy an empty ritual, then gold star for you in stating that. But if one believes the goal is for the patient to truly understand the role that AI/ML has played in the recommendation and how that should make the patient more or less confident in the recommendation, that is a much larger undertaking.

Even if one could costlessly and perfectly engage in that disclosure, such disclosure may not always be good for patients. Many believe that “‘overdisclosure’ makes it difficult for patients to distinguish meaningful risks from trivial ones,”<sup>99</sup> which is one of the main reasons for tying the doctrine to materiality.<sup>100</sup>

---

98. This relates to the critique that the law of informed consent fails to achieve what the ethical justification for informed consent demands. As the President’s Commission stated in the early ’80s: “Ethically valid consent is a process of shared decisionmaking based upon mutual respect and participation, not a ritual to be equated with reciting the contents of a form that details the risks of particular treatments.” PRESIDENT’S COMM’N FOR THE STUDY OF ETHICAL PROBLEMS IN MED. & BIOMEDICAL & BEHAVIORAL RESEARCH, MAKING HEALTH CARE DECISIONS: THE ETHICAL AND LEGAL IMPLICATIONS OF INFORMED CONSENT IN THE PATIENT-PRACTITIONER RELATIONSHIP 2 (1982), [https://repository.library.georgetown.edu/bitstream/handle/10822/559354/making\\_health\\_care\\_decisions.pdf?sequence=1&isAllowed=y](https://repository.library.georgetown.edu/bitstream/handle/10822/559354/making_health_care_decisions.pdf?sequence=1&isAllowed=y) [<https://perma.cc/BU59-EL4L>]. For a valuable recent summary of the literature documenting the gap between the ethical goal and real-life practice, see Valerie Gutmann Koch, *Eliminating Liability for Lack of Informed Consent to Medical Treatment*, 53 U. RICH. L. REV. 1211, 1220–27 (2019).

99. Wilson, *supra* note 28, at 229.

100. When I have presented this work in the past, a faculty member asked an intriguing question: whether this problem could be managed with a two-stage solution—asking patients if they want to know

Remember that the main reason to adopt these systems in the first place is because of our belief that over the great run of patients they meaningfully improve patient health. Thus, related to the exceptionalizing discussion above, the more one believes that emphasizing the role of AI/ML may lead patients to *distrust* the recommendation to their detriment and not follow that recommendation, the more one may be worried. Of course, where to set the line may depend on one's (a) prior belief as to (among other things) whether on average following the recommendation of an AI/ML improves patient care, (b) how much so, and (c) how much medical paternalism to tolerate when it leads to better patient outcomes (at least over the great run of patients).

For all these reasons, the simple answer—"always disclose"—does not seem so clearly correct as an omnibus response and, normatively speaking, we need to think more deeply about what the law of informed consent is for. As Nir Eyal writes in the *Stanford Encyclopedia of Philosophy* entry on the subject, from an ethical perspective there are seven "main arguments" for informed consent: (1) protection, (2) autonomy, (3) prevention of abusive conduct, (4) trust, (5) self-ownership, (6) non-domination, and (7) personal integrity.<sup>101</sup>

One approach to the problem would be to consider whether disclosure about the use of medical AI/ML is needed on each of these grounds, and then see which of these grounds are the ones we think are persuasive. Let us work our way through this list to see whether it sheds any light.

Regarding the patient *protection* theory, the idea is that patients have a heterogeneous set of medical and non-medical interests and it is wrong to assume the doctor is better equipped than the patient to know what best protects those interests.

---

how the treatment recommendation was formulated and only then disclosing and explaining the AI/ML role. This proposal nicely tries to bifurcate the heterogeneity of patient preference by having them opt in to disclosure. My skepticism arises because this assumes people know well, *ex ante*, their preferences to have information disclosed to them when they do not yet know what that information is. It would be a little like telling someone "I heard a large number of people talking about you yesterday, would you like to know what they said?" I am not sure I would have a way of knowing *ex ante* whether in fact I really would like to know (and even less benefit from knowing) without knowing the content of the information. Moreover, if it were a good idea for AI/ML, why not apply it for much more run-of-the-mill disclosures like risk and benefit information? Well, in one sense, perhaps we do. "Although there are few court decisions on the issue, the disclosure doctrine's grounding in autonomy suggests that patients should be able to refuse information offered by the physician." HALL ET AL., *supra* note 96, at 178–79 (citing *Spar v. Cha*, 907 N.E.2d 974, 983 (Ind. 2009)). That is a little different from mandating a two-stage approach, but in the ballpark. The idea remains interesting to me, although at some point I do wonder about infinite regress in the sense of what disclosure I might need to determine whether I in fact wanted to have the disclosure that would ordinarily constitute informed consent, and so on.

There is some echo of this idea in the literature about managing deception in research ethics and the question of whether we could rely on "second-order consent" such as for cases involving social psychology research that require deception for validity, by telling all participants that it is possible that the protocol may or may not use deception and having them consent with that disclosure in mind, essentially consenting to be deceived. *See, e.g.*, Dave Wendler, *Deception in Medical and Behavioral Research: Is it Ever Acceptable?*, 74 MILBANK Q. 87, 101–23 (1996).

101. Nir Eyal, *Informed Consent*, in STANFORD ENCYCLOPEDIA OF PHILOSOPHY (Edward N. Zalta ed., rev. ed. 2019), <https://plato.stanford.edu/entries/informed-consent> [<https://perma.cc/ZZ67-MKW9>] (last visited Apr. 30, 2019).

When it comes to medical interests, at least, there is an argument that the heterogeneity may actually favor use of the AI/ML because one of the main goals of this innovation is to offer care that is *more* particularized in its chance of success for this patient than either the patient or physician might choose unaided. And it is unclear how disclosure of the use of AI/ML to reach the recommendation is likely to enable a typical patient to better protect his or her own medical interests, especially when the AI/ML is complex and thus the patient is unlikely to be able to evaluate whether its use in his or her particular case furthers or stymies that patient's interests.

That seems plausible as to patient's medical interests. As to the heterogeneity of patient's nonmedical interests, those are harder to capture by the medical AI/ML itself, though some believe it may be possible to at least some extent.<sup>102</sup> It seems to me, though, that this argument tells us more about what good clinical use of an AI/ML recommendation should look like—the physician's recommendation should be discussed with the patient and a decision made in a shared decisionmaking framework that is also sensitive to nonmedical interests.<sup>103</sup> It tells us

---

102. *E.g.*, Camillo Lamanna & Lauren Byrne, *Should Artificial Intelligence Augment Medical Decision Making? The Case for an Autonomy Algorithm*, 20 *AMA J. ETHICS* E902, E902–05 (2018) (proposing marrying EHR and social media data to help better determine what an incapacitated patient would have wanted).

103. One interesting possibility is for the providers to collaborate with the patients in changing certain facts or parameters in the AI/ML decisionmaking and then together observe how those changes affect the recommendation produced. To use a tangible example, one of the most ethically and personally challenging medical experiences is deciding what to do with a significantly premature baby in the Neonatal Intensive Care Unit (NICU). One thing that the parents really care about in making decisions is the chance that the baby will survive given certain interventions and the level of deficits the child will experience. There are “[o]nline, user-friendly, population-based outcome prediction tools,” including the “National Institute of Child Health and Human Development (NICHD) Neonatal Research Network ‘calculator’” that uses “three patient characteristics—gender, administration of maternal steroids, and multiplicity—that are often known before birth to produce the estimated survival outcome for an infant with a given profile.” Genevieve Allen & Naomi Laventhal, *Should Long-Term Consequences of NICU Care Be Discussed in Terms of Prognostic Uncertainty or Possible Harm?*, 19 *AMA J. ETHICS* 743, 745 (2017).

These tools are “widely used by neonatologists despite uncertainty about their usefulness and impact.” *Id.* AI/ML is in some ways a much more sophisticated version of these calculators. Could we imagine an idealized shared decisionmaking practice where the patient and provider “play” with the AI/ML to see how fragile its recommendation is across certain known unknowns and parameter changes? There is some evidence that allowing individuals to “modify” an algorithm reduces “algorithm aversion.” See generally Berkeley J. Dietvorst, Joseph P. Simmons & Cade Massey, *Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them 2* (Aug. 6, 2016) (unpublished manuscript), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2616787](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2616787) [<https://perma.cc/KYK7-YZL6>] (reporting the results of experiments).

This possibility raises several interesting questions. First, if people trusted the algorithm more because they were allowed to manipulate it more, would that be a good thing? If they have real concerns that are not met with the “play,” (to put it hyperbolically) there is a sense in which this feels like opiating the masses by letting them play with a shiny toy. Second, what if the “modified” algorithms produces recommendations and patient choices that we have reason to believe are *less* likely to improve patient health? Is there still a compelling reason to build that in or should we think about this as a trade-off that at some point tilts away from patient empowerment? Third, how would allowing such modifications impact liability in medical malpractice if the result of the modification is a bad event that would not occur but for the modification? Should the patient internalize that risk or should it remain with the

less about whether the role of AI/ML in helping the physician formulate the initial medical recommendation should be disclosed and in what way. That is, unless one had an exceedingly special case of someone whose non-medical interests depended not on the quality of or confidence in the recommendation, but the fact that the recommendation was arrived at via AI/ML rather than via another method,<sup>104</sup> it does not seem to track.

---

physician? Fourth, to what extent does this possibility depend on the relatively non-opaque AI/ML and the physician's understanding of "what makes it tick," so that it can productively assist the patient in understanding fragility and parameters? Finally, to what extent is such a high-touch, shared decisionmaking approach practicable in a resource-constrained, physician-rushed, healthcare system that most Americans find themselves in?

Now here the answer may be that it could be appropriate for some exceedingly high-risk, high-uncertainty decisions but not for more quotidian decisions, but even there one wonders about certain dynamics if this became required. In particular, if the high-cost shared decisionmaking approach was required by law if AI/ML is used to form part of the recommendation, to what extent would that dissuade hospitals or physicians to adopt the AI/ML? Does the answer depend on whether the time spent with the patient modifying the algorithm to see what happens is also reimbursable by insurers? Will insurers actually be willing to pay for this? Should patients be allowed to pay "extra" for this but it not be part of the baseline care given.

This whole line of questions, I think, exposes larger questions about how much one wants medicine to be "consumer-driven" as opposed to "medical paternalist." Patients used to the experience of high-end, highly-consumer oriented U.S. medicine may have one intuitive answer to these questions, whereas U.S. patients who have difficulty finding a healthcare provider or whose experience is drawn from the Canadian or UK NHS systems, for example, may feel differently.

104. Such cases are difficult to formulate but not altogether impossible to imagine. In the case of deceptive research, Wendler gives the examples of people who think they are playing a game against another human being but in fact are playing a game against a computer, and had they known they would not have participated because playing video games was against their religious beliefs. See Wendler, *supra* note 100, at 98. Perhaps there are people who have similarly strong non-medical anti-AI/ML views: "My brother was killed in an assembly line accident involving artificial intelligence, I sure as hell ain't going to trust an AI to tell me how to treat my cancer!" This again, though, puts pressure on the idea of the normative nature of the reasonableness inquiry in materiality—it is not enough that it actually mattered to you, but instead that it would have to have mattered to a reasonable person like the patient.

A different kind of non-medical interest some might try to marshal might be phrased as a right to a human decision-maker or a demand for algorithmic due process. For discussions of how due process should apply in the AI context, see Danielle Keats Citron, *Technological Due Process*, 85 WASH. U.L. REV. 1249 (2008); Kate Crawford & Jason Schultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C. L. REV. 93 (2014). When it comes to the actual law of due process, the use of AI/ML to inform treatment decisions seems very different from using it to determine when a person merits a particular sentence or whether their welfare benefits should be audited. With the possible exception of the Veterans Affairs health system, the requisite state action seems absent. But what about the animating ideas, if not the Due Process Clause itself? Although a patient may have claim on her physician as to some kinds of reason-giving connected to treatment decisions, that claim strikes me as quite different from those that, for example, a convicted person has as to the judge who sentences them. This is true in the comparison of medicine and law more generally. An appeal of a conviction is quite different from a second opinion from a doctor. A claim to one's liberty is quite different from a claim to a particular treatment option. And the constitutional requirement of notice in the termination of welfare benefits is quite different from the requirements of informed consent.

What the state owes us when its heavy machinery alters our life is quite different from what we may be owed when we seek care. I think we will find the specific fiduciary nature of the physician-patient relationship a surer guide to what is owed in terms of disclosure for informed consent than the Due Process Clause and its jurisprudence (which, for the record, has thus far not been particularly successful in litigative pushes for more algorithmic due process). I do not purport to resolve here whether the

Putting to one side what the precise philosophical foundation of *autonomy* as a basis for informed consent might be,<sup>105</sup> it is not clear that failing to disclose the *method* by which the physician arrived at a recommendation is a violation of the patient's autonomy, understood as a kind of right to self-determination. Consider again the failure to disclose the role in formulating a treatment recommendation of a varied assortment of vague memories from a medical school lecture, what the other doctors during residency did in such cases, the latest research in leading medical journals, the experience with and outcomes of the last thirty patients the physician saw, etc. We do not ordinarily think that it is a violation of the patient's autonomy to fail to disclose all of that and instead to simply present a treatment recommendation and await patient questions to say more. Of course, here again the specter of finding the right analogy may haunt us. If the analogy is *inputs to decisions*, then the violation of the patient's autonomy seems more difficult to conceptualize. If instead the right analogy is to *another personage doing the work* then the concern for autonomy violations seems more legitimate—it is easier to see it violates a patient's autonomy to have the patient agree to Dr. X but then substitute Dr. Y after the patient is unconscious. If that is how we understand the AI/ML case the claim of an autonomy violation seems more persuasive. Here, the ideas discussed above about epistemic warrant and indicia of reliability may be helpful guides—perhaps, instead autonomy is violated when a physician fails to alert patients when indicia of reliability are absent?

It is hard to see why informed consent in the AI/ML case creates a *protection against abuse*, in the sense of preventing such “offenses as assault, deceit, coercion, and exploitation.”<sup>106</sup> Perhaps there are cases where AI/ML is implemented not in the patient's interest but in tension with that interest—for example, an AI/ML that is designed to recommend cheaper but less efficacious treatments than the physician otherwise would recommend. Even then, I am not sure whether this rises to the level of “deceit,” but in any event, the objection in this hypothetical seems more to be that the AI/ML is making *those* recommendations, not the *failure to disclose* the AI/ML's part in making treatment recommendations. As with the autonomy theory of informed consent, the more the case seems to us like a substitute physician “scrubbing in,” the more plausible it is to view this as “deceit,” and thus activating this normative theory.

The *trust* justification for informed consent has both a forward-looking variation—informed consent protects “ongoing societal trust in caretakers and medical institutions, for example, as a precondition of ongoing compliance with medical advice”—and a backwards-looking variation in which it is “an intrinsically important way to honor the trust that the patient has placed in the physician, and as part of the fiduciary role that the physician has undertaken.”<sup>107</sup>

---

algorithmic due process (or more generally algorithmic accountability) literature will have more to say as to the use or adoption of AI/ML in healthcare *outside* the context of informed consent.

105. See Lamanna & Byrne, *supra* note 102.

106. Eyal, *supra* note 101.

107. *Id.* (citing Steven Joffe & Robert D. Truog, *Consent to Medical Care: The Importance of Fiduciary Context*, in *THE ETHICS OF CONSENT* 347, 352 (F. G. Miller & A. Wertheimer eds., 2010)).

The forward-looking version is well-captured by the Swedish philosopher Torbjörn Tännsjö. He argues that in many cases informed consent can be justified based on the supposition that individuals know best what will make their life go well. But:

[E]ven in cases where they are mistaken, it is crucial that they are treated in the manner they want. Otherwise they will not seek medical advice when they need it, and they must have to fear that they will be treated in a manner they dislike, should they end up in certain situations. This means that their right to informed consent should be granted, even at some cost in the individual case (where they are not taken best care of). This is so because if they—we—could not in this way trust the medical system, the price would be even higher.<sup>108</sup>

Whether failing to disclose AI/ML involvement in formulating a recommendation hampers forward-looking trust in this way is (again) at least in part an empirical question.<sup>109</sup> But it seems to be a tougher argument to make than for some of the other examples Tännsjö uses. Compare it, for example, to his argument as to why we ordinarily provide a Jehovah's Witness the right to informed consent concerning (and thus the right to refuse) a blood transfusion, where he argues "[u]nless this right did exist, the Witness would not dare to seek medical advice and would *all the time* have to fear becoming the victim of some accident where, unbeknownst to her, blood would be provided."<sup>110</sup> It seems unlikely to me that "AI-phobia," even where present,<sup>111</sup> is so great in the medical care sphere that many individuals would choose to abstain from or deleteriously delay all or most beneficial healthcare for fear of undisclosed AI/ML involvement in their healthcare.

On the backwards-looking rationale, the question is whether not disclosing the use of AI/ML violates the nature of the physician–patient relationship, a relationship imbued with trust vulnerability? Of course, if we had a clear answer to this question, the inquiry we are undertaking could be much more cursory. Once again, perhaps reasoning by analogy could help. Even the analogy of AI/ML as inputs in decisionmaking—as opposed to substitute surgeons—one could imagine ways of making patient recommendations that would violate fiduciary duties, such as "I flipped a coin in deciding on your cancer therapeutic rather than researching it" or "I asked my five-year old to pick the drug he thought was prettier." But the most obvious problem with those methods goes to our lack of faith in their accuracy. Doctrinally speaking, the issue seems more about malpractice (failure to meet the diagnostic standard of care) than informed consent. By

---

108. Torbjörn Tännsjö, *Utilitarianism and Informed Consent*, 40 J. MED. ETHICS 445, 445 (2013).

109. Unhelpfully, this may also again raise the bootstrap issue: as it becomes more ubiquitous, the need to disclose will go down on this rationale; but its ubiquity is partially dependent on whether we require informed consent because that will also affect its adoption rate.

110. Tännsjö, *supra* note 108, at 445.

111. For some limited data on the attitudes of the public as to robotic surgery specifically, see *supra* note 93 and accompanying text.

contrast to these outlandish hypotheticals, AI/ML is not likely to be used (and indeed might be objectionable *with or without* informed consent) if we thought it was less accurate than the status quo, or to be a bit more omnibus, if it produced worse outcomes.

The interesting question is what would happen if one was convinced the new input was likely to *both* produce a much better outcome *and* be a cause of concern for the patient? Suppose the young doctor trying to make a difficult diagnosis calls up his much more senior mentor, who tells him his instincts are all wrong, and instead the diagnosis is likely something else. The young doctor acquiesces even though he cannot quite understand why the mentor reached that conclusion. The young doctor presents that diagnosis to the patient but discloses none of this dialogue or his internal thought process. Is this a violation of the fiduciary role, as a normative matter? I am not sure. If the junior doctor has confidence that his senior mentor is more likely to get it right and improve the patient's lot, and has a good epistemic warrant for that belief, then it does seem as though he is working in the patient's interest and not his own. But what do we make of the failure to disclose of the input?

One way to try to answer the question is by reference to convention; is this part of what a reasonable patient would expect from his fiduciary? But that, of course, brings us back to the empirical path and some of the difficulties alluded to above.

A different way is to be "normative all the way down" and to suggest that our theory of informed consent has an answer to this question orthogonal to whatever the empirical analysis tells us—that is, there is a way for a proper fiduciary to behave in this context that is not dependent on what we find through empirical analysis.

One example of this approach might be drawn from a 1989 article, where Howard Brody proposed (indeed as a legal standard!) a "Transparency Standard" for informed consent:

According to this standard, adequate informed consent is obtained when a reasonably informed patient is allowed to participate in the medical decision to the extent that patient wishes. In turn, "reasonably informed" consists of two features: (1) the physician discloses the basis on which the proposed treatment, or alternative possible treatments, have been chosen; and (2) the patient is allowed to ask questions suggested by the disclosure of the physician's reasoning, and those questions are answered to the patient's satisfaction.

According to the transparency model, the key to reasonable disclosure is not adherence to existing standards of other practitioners, nor is it adherence to a list of risks that a hypothetical reasonable patient would want to know. Instead, disclosure is adequate when the physician's basic thinking has been rendered transparent to the patient. If the physician arrives at a recommended therapeutic or diagnostic intervention only after carefully examining a list of risks and benefits, then rendering the physician's thinking transparent requires that those risks and benefits be detailed for the patient. If the physician's thinking has not followed that route but has reached its conclusion by other

considerations, then what needs to be disclosed to the patient is accordingly different. Essentially, the transparency standard requires the physician to engage in the typical patient-management thought process, only to *do it out loud in language understandable to the patient*.<sup>112</sup>

If this is what a doctor owes a patient, it would be problematic to fail to disclose the input of medical AI/ML into a treatment recommendation because it is at least in part a “basis on which the proposed treatment, or alternative possible treatments, have been chosen.”<sup>113</sup>

I think the more difficult question is whether more opaque AI/ML, even if its use is disclosed to the patient, could ever satisfy Brody’s proposed standards, in particular his suggestion that one might be able to explain the “thought process . . . out loud in language understandable to the patient.”<sup>114</sup> The question is whether what is required is to explain “out loud” that one was relying on an AI/ML whose ways of operating are not explainable by the physician, the epistemic warrant one has for relying on that opaque AI/ML (for example, preclinical testing), or whether what is required is the ability to explain how and why the AI reached this conclusion. If it is the last of these, then some forms of medical AI/ML would seem *per se* to violate this standard.

The transparency standard is certainly interesting as a normative vision, but one might wonder—in the mode of the Dworkin of *Law’s Empire*—how much it “fits” into the existing law of informed consent?<sup>115</sup> That is, if *this* is the standard, it seems physicians are often honoring it only in the breach, even in more quotidian cases, by failing to disclose their reliance on colleagues, journal articles, and many other inputs as the basis of their decision. That is not itself determinative of whether the law should go in this direction—perhaps more of the law of informed consent is “wrong” than we have previously thought—but it does suggest that the theory is a bit more radical in its implications than it might seem at first blush.

The *self-ownership* justification for informed consent,<sup>116</sup> often traced back to Locke,<sup>117</sup> is good at tying informed consent to its origin in the tort of battery. But it has much less to say about fairly specific questions of application in a context like this one.

The *personal integrity* rationale is similar, though it may have more to say about the focus on violations of one’s body as opposed to say a physician’s “use of a magic wand to treat an unwilling patient’s ailments without cutting his

---

112. Howard Brody, *Transparency: Informed Consent in Primary Care*, 19 HASTINGS CTR. REP. 5, 7–8 (1989).

113. *Id.* at 7.

114. *Id.* at 8.

115. See generally RONALD DWORKIN, *LAW’S EMPIRE*, at vii (1986) (requiring the new rules created by constructive interpretation fit the “best justification of our legal practice as a whole”).

116. Eyal, *supra* note 101.

117. JOHN LOCKE, *SECOND TREATISE OF GOVERNMENT* § 27, at 19 (C.B. Macpherson ed., Hackett Publishing Co. 1980) (1690) (arguing that “every man has a *property* in his own *person*”).

skin.”<sup>118</sup> It does not seem to say much about this specific context unless there is some other theory that connects personal integrity to the presence of AI/ML as opposed to human inputs.

Finally, we come to the *non-domination* account, an account that is much more prominent as to sexual, rather than medical, consent. It centers on the idea that “no one should be under the arbitrary control of another and that informed consent requirements help to prevent such arbitrary control,”<sup>119</sup> with a particular concern about hierarchy. If one squints, perhaps it has something to say about our context—the patient believes he or she is consenting to the control of the physician, when in fact the physician herself is partially under the “control” of an AI/ML. Perhaps the concern is more pungent where the AI/ML is more opaque—although not arbitrary, the physician cannot explain it, but again perhaps the litmus should be accuracy or epistemic warrant not explainability—and where the physician has fewer degrees of freedom to deviate from the AI/ML recommendation based on institutional set up or the control of an insurer over reimbursement. Notice how this account echoes two different analogies—the substitute-surgeon case and perhaps the managed-care case—where a third party is partially constraining choice and perhaps that should be disclosed.

\*\*\*

If one had hoped that the underlying normative rationales for informed consent doctrine would clearly point us towards an approach to the AI/ML problem, one’s doe-eyed optimism should now be partially occluded by tears. But there are some lessons to be learned. Perhaps the best we can do is identify a few cases where overlap between normative theories of informed consent makes stronger the case for disclosure. In my own view, after surveying all these rationales, these are cases where the physician faces restrictions on disregarding the AI/ML recommendation (be they hospital policy or insurer reimbursement rule) and cases where the physician lacks a good epistemic warrant to believe that the AI/ML recommendation is correct.<sup>120</sup> These cases are where I find the strongest normative reasons to disclose the involvement of AI/ML and also highlight these features for the patient.

If the law of informed consent as to medical AI/ML is up for shaping, my own view is that this would be the direction in which to nudge it.

---

118. Eyal, *supra* note 101.

119. *Id.*

120. To reiterate: that need not necessarily be an understanding of *how* the recommendation was arrived at, in the sense that people typically mean when they think of explainability. It could be enough to have confidence that the recommendation was accurate in the sense that it was more likely to result in a good outcome for a general population of patients. The epistemic warrant for the latter might be a preapproval process, such as FDA review, clinical trial data, or something else.

## IV. SOME REMAINING ISSUES

Part II presents my conclusions about where the doctrine is, and Part III my views, however tentative, about where it should go. This Part discusses a few more severable issues related to consent itself, racial bias, and the role of the non-disclosure elements of the cause of action for breach of informed consent in the project.

## A. WHAT ABOUT CONSENT?

As with much of the doctrine of informed consent, our focus thus far has been almost entirely on the *informed* part—when do patients need to be informed about the use of AI/ML and what facets should be disclosed? The basic question of consent is straightforward—the patient should have a right to accept or reject the treatment recommendation formulated (typically only in part) via AI/ML. The latter could be accomplished by seeking another opinion, for example. The same would be true for AI/ML involvement in a surgical procedure. Once informed, the patient could choose another surgeon.

The more interesting question, though one far away at the moment, is what happens if in the future AI/ML has so pervaded the practice of medicine that no or few physicians would engage in diagnosis without AI/ML? Professors Fromkin, Kerr, and Pineau have recently argued that “existing medical malpractice law will eventually require superior ML-generated medical diagnosis as the standard of care in clinical settings” because “[o]nce computerized diagnosticians demonstrate better success rates than their human trainers, effective machine learning will create legal (and ethical) pressure to delegate much, if not all, of the diagnostic process to the machine.”<sup>121</sup> These authors are worried about this future as one that leads to “overreliance” on ML and what they call a “diagnostic monoculture.”<sup>122</sup>

---

121. A. Michael Fromkin, Ian Kerr & Joelle Pineau, *When AIs Outperform Doctors: Confronting the Challenges of a Tort-Induced Over-Reliance on Machine Learning*, 61 ARIZ. L. REV. 33, 35 (2019).

122. As they put it:

By “diagnostic monoculture” we mean a scenario in which the medical and legal systems standardize on a mechanized approach to diagnosis in a given sub-specialty. Diagnostic monoculture exemplifies a more general problem that arises when society comes to rely, to its detriment, on a dominant mode of thinking to the exclusion of other possible solutions. In this case, a diagnostic monoculture that leads to less input from human physicians could make quality control of diagnostic databases much more difficult. The problem becomes far more serious once reliance on ML goes beyond diagnosis to treatment. The reduction in new data from physicians—that is to say the creation of a loop in which outcomes added to the database are solely or overwhelmingly the result of ML-informed treatment decisions—creates scenarios in which we cannot rule out the risk that sub-optimal conclusions are reached. If a set of symptoms is consistently producing an erroneous ML diagnostic, and physicians act on that erroneous diagnostic, where will ML get the data to suggest a different diagnosis which leads to better treatment? If the answer is “nowhere” then we have a problem. Worse, it is not even clear that either the ML system or an outside observer necessarily would know that the results were sub-optimal. From a human perspective, the challenges associated with understanding and auditing an ML system’s predictive diagnostic process will become

I have my doubts about how likely this is to occur in the foreseeable future—a world where no reasonable physician would fail to use medical AI/ML in making their decision seems unlikely. But let us imagine, perhaps counterfactually, that this is a near- or middle-term likely future. In such a world, the consent part of informed consent would formally be satisfied—the patient could refuse to accept the recommendation of a physician based on AI/ML or the offer of AI/ML-guided surgery—but the alternative in such a world might be none at all, foregoing the medical care because no physician (or none accessible to the patient) does it *without* the AI/ML assistance. Is that a problem? If using AI/ML really produced better patient outcomes across the board, then it seems desirable for it to become the standard of care. As a matter of informed consent, ethically or legally, it is not clear why we should shed tears if in such a world, patients do not have access to the non-AI/ML approach. Many diagnostic tests and approaches have been long superseded by much better ones, yet no patient has a plausible rights claim to the old ways in such instances. If our pretheoretical intuitions as to AI/ML feel different, this might be just a case of exceptionalizing the new rather than a deep difference between the technologies.

On the flipside for those like Professors Froomkin, Kerr, and Pineau who are concerned about this possible future for deskilling or other reasons, certain doctrinal decisions about the law of informed consent may have a role to play in forestalling this future. If it turns out, empirically speaking, that many patients have negative reactions when the role of AI/ML is disclosed, then a version of the doctrine that requires that disclosure is more likely slow the adoption of AI/ML because physicians and hospital systems are at least somewhat sensitive to patient reaction. Of course, if one's prior is that patients are in the wrong to fear the use of AI/ML—that is, their reactions do not align with what actually helps improve their health and other interests but instead is driven in part by lack of understanding about AI/ML—then this might be an argument against the doctrine mandating too much disclosure. This is one of the inherent tensions in the doctrine's focus on a "reasonable" patient, which is at least something of a moralized construction.

Thus, it may be that one's attitudes about which direction the doctrine should move in regarding this question is at least in part driven by one's attitude towards AI/ML adoption. If one fears adoption will be too fast and will not improve patient care, then one should favor more robust disclosure requirements. If one thinks the opposite is true, then one should favor nondisclosure. Of course, one might push back and say, "that is not what the law of informed consent is for!"

---

significant. Those challenges become greater if the output of the ML diagnostic system is then fed into a second ML treatment system. In that case, absent personalized medicine, for any given set of symptoms one might get consistent treatment decisions leading to less varied treatment-to-outcome data.

*Id.* at 37–38.

That is, one might normatively defend a view that the law of informed consent is and should be about patients' rights claims and when patients authorize certain things to be done to them. To put it differently, one might claim that the positive or negative effects of adoption or innovation are beside the point when it comes to the shape of the doctrine. I do not pretend to resolve who has the better of this dialogue here.

B. THE SPECIAL CASE (?) OF BIAS AND INDIVIDUALIZING CONSENT

Many are rightly concerned that in medical AI/ML, deficiencies in the representativeness of the data sets used to train the system may result in poor performance for some races, ages, etc.<sup>123</sup> Consider an AI/ML system that tries to improve breast cancer care by making treatment recommendations after analyzing the results of mammograms.<sup>124</sup> Now suppose it turns out that the training data was underinclusive of African-American women, and this was relevant because they tend to have differences in breast density from Caucasian women. As a result, the algorithm is likely to give treatment recommendations that are worse for the African-American population. Should informed consent look different in such a case?

One possibility would be to just say that in such cases, the AI should not be used. We might achieve that result through regulatory premarket approval regimes or on the back end through malpractice law. If that is framing of problem and solution, then informed consent law is beside the point. But one might worry that in some instances this will make the Perfect the enemy of the Good, that even if the AI/ML performs slightly worse in this population it produces better results than the alternative; moreover, in most cases it may be that we *worry* but are not be completely sure about how an AI/ML approach will do for subpopulations. This might carve out some role for informed consent—physicians should disclose to these populations when they use AI/ML involving training-set data that might make medical recommendations for these populations less accurate than they might otherwise be. The analogy might be to the “expanded access” use for an otherwise terminally ill patient of a cancer therapeutic designed for a different kind of cancer: this has worked in cancer X, there are reasons to think it might work for your kind of cancer, but we do not have enough data yet to be sure and there are risks, will you consent to using it?<sup>125</sup>

---

123. *E.g.*, Latice G. Landry & Heidi L. Rehm, *Association of Racial/Ethnic Categories with the Ability of Genetic Tests to Detect a Cause of Cardiomyopathy*, 3 JAMA CARDIOLOGY 341, 342 (2018); Arjun K. Manrai et al., *Genetic Misdiagnoses and the Potential for Health Disparities*, 375 NEW ENG. J. MED. 655, 655 (2016); Price, *supra* note 12, at 5–6, 32–33; see I. Glenn Cohen et al., *The Legal and Ethical Concerns that Arise from Using Complex Predictive Analytics in Health Care*, 33 HEALTH AFF. 1139, 1141 (2014).

124. Adam Yala et al., *A Deep Learning Mammography-Based Model for Improved Breast Cancer Risk Prediction*, 292 RADIOLOGY 60, 60–61 (2019).

125. There is also the question of *as against what* and disclosure of failure to adopt an AI/ML. That is, to use an example above, Obermeyer et al., *supra* note 12, even if a “high-risk patient management” algorithm produces these biased results when it comes to race, how does it compare to the level of bias, or outcomes in general, in hospitals that do not use the algorithm? If the hospitals perform significantly

In theory, such a disclosure regime sounds like a good solution, but in practice it faces several challenges. Professor W. Nicholson Price II has done a good job illuminating those challenges in a paper considering whether FDA labeling should include this kind of information about contextual bias.<sup>126</sup> As he writes, “to really know how to impose labeling requirements that contain enough information to meaningfully inform use, we need to know a lot more about the relevant sources of patient and provider variation than we know now,” meaning that “it will be difficult to get right.”<sup>127</sup> On the flip side, robust requirements of informed consent in this setting might be a cure for one obstacle he sees in FDA-labeling requirements—the rampant use off-label of most approved drugs. As he notes, “drug labels rarely specify that they are principally tested in relatively ancestrally homogeneous populations, we might think of the widespread use of drugs or other treatments in ancestral minorities in whom the treatments were not originally tested as a sort of ersatz off-label use” and likewise “we should expect that medical AI would be used off-label just as other medical treatments are.”<sup>128</sup> If we imposed a robust informed consent basis for the duty to disclose, though, providers might feel the need to more directly disclose the label-restrictions, thus beefing up the impact of a labelling regime.

There is an important difference between what is required for FDA labeling and what is required for informed consent. Nothing in informed consent law per se requires a physician to disclose she is recommending an off-label use of a drug or device, and on the flipside the mere fact that the use is contemplated in the label does not immunize from a possible informed consent tort. Instead, I am using the analogy to labeling as a way of thinking what *might* be done.

### C. SOME (MORE) COLD WATER ON YOUR WAY OUT

We have struggled, perhaps mightily, with how law, empirics, and normative ethics might shape how the law of informed consent ought to handle the use of medical AI/ML, especially when they are helping “in the background” to shape medical recommendations. Hopefully, we have made some headway, but there are two splashes of cold water it would be irresponsible not to subject you to, dear reader, on your way out of this Article.

First, on the matter of liability. Recall there are four elements of a breach of informed consent tort, to use one common formulation:

- (1) failure to disclose a specific risk in violation of the governing standard; (2) materialization of that risk; (3) “causation”—that is, if the risk been disclosed,

---

worse, should that be disclosed by hospitals that do *not* adopt the algorithm? It seems odd to think that any doctrine of informed consent would require discussing all the technologies that might improve care a particular hospital has not bought and implemented, but if the obligations to disclose attach only asymmetrically to those who use the (by hypothesis care improving and bias reducing) AI/ML, doesn't that problematically discourage adoption?

126. See Price, *supra* note 12, at 48.

127. *Id.* at 49.

128. *Id.* at 50.

the patient, or a prudent person in the patient's position, would not have proceeded as she did; and (4) that no exception, like emergency, excuses the failure to disclose.<sup>129</sup>

Our discussion has focused on that first requirement, but although necessary, it is not sufficient. To use the law and economics formulation of the problem, when the tortfeasor considers making changes to his behavior, he considers the chance and cost of liability using all elements of the tort. The tort will only be active in cases where both the undisclosed risk materializes *and* the plaintiff can show that "if the risk been disclosed, the patient, or a prudent person in the patient's position, would not have proceeded as she did." In many cases involving AI/ML it will be hard for the plaintiff to meet those requirements such that the force of liability will be blunted as a curb on behavior. For example, how often will the plaintiff be able to convince the fact finder that had the medical AI/ML been disclosed, he would have opted for a different treatment option? And in what subset of those cases will they be able to show that the relevant risk had materialized?<sup>130</sup>

That said, even if for these reasons the risk of liability may be low in many areas of medicine, even a relatively remote threat of liability has been a powerful spur for the issuing of practice guidelines and the like; so it is not clear that the low likelihood of lawsuits dooms the likelihood of developing informed consent practices around AI/ML.<sup>131</sup>

Second, even if we were to achieve perfect compliance with our ideally designed informed consent regime—the physician discloses the use and details of medical AI/ML involvement in the exact cases and in the exact way we could only heretofore dream of—what follows? A significant strand of research has documented the failures of informed consent—it is hard to communicate complex information accurately to patients. Even after well-designed and well-intentioned efforts, when tested, most patients do not understand the information presented to

---

129. Wilson, *supra* note 28, at 217.

130. There are also a series of interesting subsidiary doctrinal questions that would need to be resolved. One interesting one, that has to do with the gap between the theory and practice of informed consent as Peter Schuck so nicely highlighted a quarter century ago. Peter H. Schuck, *Rethinking Informed Consent*, 103 YALE L.J. 899, 900–05 (1994). One of these interesting doctrinal issues is that the informed consent doctrine typically

presumes that, by submitting oneself to treatment, a patient consents, absent any objection, to the many minor bits and pieces that make up the entire treatment encounter. This "bundling" concept of informed consent arises by legal implication even though the patient receives no risk/benefit disclosure about each of these component parts. For instance, one court ruled that no special informed consent is required for use of forceps during delivery, because this is simply a tool for carrying out the general treatment plan that the patient agrees to upon admission to the hospital.

Hall, *supra* note 70, at 553 (citing *Sinclair v. Block*, 633 A.2d 1137, 1140 (Pa. 1993)). Where does the AI/ML aspect fall within the bundle?

131. Of course, whether that descriptively true phenomenon—changes in practice in medicine based on perceived liability risks even though in reality such risks are quite small—is a positive or negative thing, normatively speaking, is a big question I will not purport to resolve here.

them, and even when understood, many patients show a reluctance to use relevant information to make the decision.<sup>132</sup>

Of course, it would be better if the informed consent processes we invested in developing actually made a difference in patient comprehension and decision-making, but importantly, this is not the only channel by which they may influence behavior. It is possible that the act of disclosure, even if not readily understood by the listener, influences the behavior of the physicians or systems deciding whether/how to implement AI/ML. To put the point differently, the comprehension problem for informed consent as to AI/ML seems similar to that for many other kinds of informed consent. The courts and legislators have spent considerable time and resources developing that law. You may think this has done really important work to protect patients and safeguard patients' rights and interests or you may think the game is not worth the candle. But it would seem to me your attitude on this issue is unlikely specific to AI/ML. So, if you thought informed consent was important enough to get this far into this Article, chances are the AI/ML issue strikes you as important too. If not, well, hopefully it was a pleasant enough distraction of a read.

#### CONCLUSION

Despite the huge financial investment and media and scholarly engagement we are seeing as to medical AI/ML, little has been written on informed consent for these technologies. In this Article, I have tried to remedy that. Here are my central claims:

While there are possible argumentative avenues, drawing on analogies to cases involving qualifications, ghost surgery, conflicts of interest, etc., overall the existing legal doctrine of informed consent does not robustly support an obligation to disclose the use of medical AI/ML.

There are a few fact patterns where obligations to disclose are more easily grounded in existing doctrine: when the patient asks about the basis for decision making (the more specific the better) and is misled by the physician, when the medical AI/ML is really "in charge" or has a quite central role in choosing the decisionmaking, when medical AI/ML is implemented in a way that, by design, may conflict with the patient's best interests (such as to recommend against more costly but more health-promoting options), and when the physician lacks an epistemic warrant that the AI/ML is reaching good decisions (which implicates, but is not the same as, opacity in explanation). And in some states, even in these cases, the doctrine is limited to apply only to surgical interventions. Although these fact

---

132. For good summaries of this literature, see Omri Ben-Shahar & Carl E. Schneider, *The Failure of Mandated Disclosure*, 159 U. PA. L. REV. 647, 668–70 (2011), and see also Sawicki, *supra* note 27, at 871 (noting that "[c]onsumer protection advocates recognize that excessive disclosures may overwhelm patients and leave them struggling to distinguish between relevant and less-relevant facts" and that "some providers may view disclosure and consent requirements as a way to sanitize or excuse inappropriate care—as, for example, where significant financial conflicts of interest cause a physician to provide treatment that jeopardizes patient safety").

patterns strike me as the ones where it is easier to ground an obligation to disclose in the doctrine, it is still far from easy. And even if the disclosure element is required, other elements of the cause of action may defeat liability in the great run of cases.

Does the doctrine get it “right”? That depends, I have argued, on what we mean by “right.” One way of answering is, does the doctrinal result correctly line up with what we would find if we tried to ascertain, through empirical research, “those disclosures which a reasonable medical practitioner would make under the same or similar circumstances” (physician-based standard) or the disclosures that “would be regarded as significant by a reasonable person in the patient’s position when deciding to accept or reject a recommended medical procedure” (patient-based standard)? We could try to find out. It would be difficult, labor-intensive, potentially expensive, and imperfect, but we could try to make progress on these questions.

Whether we ought to do so, I argue, depends on whether one thinks the doctrine *should* track the answer to those empirical questions. To answer that, I argue, we face some unavoidably normative questions of what informed consent is “for.” I have canvassed the oft-articulated normative bases for the obligation to provide informed consent, I think they do not do too badly in lining up with the doctrine. In particular, some of the same “fact patterns” mentioned above where I think the doctrine is most likely (even if still not all that likely) to support disclosure obligations are also areas where much of the normative underlay would also counsel for more robust disclosure.

This has been the crux of my argument. On the side, I have added two more ancillary and more tentative claims. First, should medical AI/ML become widespread (and I have my doubts), the ability of patients to meaningfully consent to its use will be undermined by the difficulty of finding a provider who does *not* rely on it. If that result obtains, I do not think it is a normative worry from within the doctrine of informed consent so long as we believe the technologies are actually improving patient care overall. Second, many are rightly concerned that in medical AI/ML, deficiencies in the representativeness of the data sets used to train the system may result in poor performance for some races, ages, etc. Informed consent is certainly *not* the solution to this problem, but it is worth thinking about whether including required disclosure on known or knowable issues along these lines as part of the informed consent process would be helpful.

The focus of this Article has been when the law should obligate disclosure of information relating to AI/ML in healthcare as part of informed consent, the policing of *underdisclosure*. Beyond the clinical encounter and the doctrine of informed consent, though, one might also worry about the healthcare system promoting AI/ML *too much*. The general worry is that as with stem cell research before it, AI/ML is problematically prone to “science ‘hype’—in which the state of scientific progress, the degree of certainty in models or bench results, or the potential applications of research are exaggerated—is receiving increased

attention from the popular press, the research community, and scientific societies.”<sup>133</sup> The more specific worry is that hospital systems may oversell their use of AI/ML in an attempt to gain business, and that patients may overestimate the amount of AI/ML involvement in their healthcare, or the value derived from it. These are important problems, but to the extent the law polices them (and opinions will differ how successfully) it is not through the law of informed consent, but instead through the law of false advertising.<sup>134</sup>

---

133. Timothy Caulfield, Douglas Sipp, Charles E. Murry, George Q. Daley & Jonathan Kimmelman, *Confronting Stem Cell Hype*, 352 *SCIENCE* 776, 776 (2016). To use the language of the Gartner Group’s iconic hype cycle, AI/ML is almost certainly not yet at the “Peak of Inflated Expectations,” and the “Trough of Disillusionment” is unfortunately still in our future. *Gartner Hype Cycle*, GARTNER, <https://www.gartner.com/en/research/methodologies/gartner-hype-cycle> [<https://perma.cc/D3L3-HC6D>] (last visited Apr. 30, 2019).

134. For a (fun!) introduction to this area of law, see, for example, David A. Hoffman, *The Best Puffery Article Ever*, 91 *IOWA L. REV.* 1395, 1401–06 (2006).