*Note: this is the penultimate draft; please cite from the published version.

Probability and informed consent

Nir Ben-Moshe, Benjamin A. Levinstein, Jonathan Livengood

nir@illinois.edu; benlevin@illinois.edu; jlive@illinois.edu

Abstract

In this paper, we illustrate some serious difficulties involved in conveying information about uncertain risks and securing informed consent for risky interventions in a clinical setting. We argue that in order to secure informed consent for a medical intervention, physicians often need to do more than report a bare, numerical probability value. When probabilities are given, securing informed consent generally requires communicating how probability expressions are to be interpreted and communicating something about the quality and quantity of the evidence for the probabilities reported. Patients may also require guidance on how probability claims may or may not be relevant to their decisions, and physicians should be ready to help patients understand these issues.

Keywords Bayesianism, Decisions, Frequentism, Informed Consent, Medical Ethics, Philosophy of Medicine, Probability

Introduction

Medical interventions do not always succeed. And sometimes they cause serious health problems or even death. Since the outcome of any given medical intervention is uncertain, the patient's decision whether to undergo an intervention is necessarily made under conditions of uncertainty. Physicians therefore have an obligation to their patients to help them understand both the risks they face and the uncertainty of those risks. Our goal in this paper is to illustrate some serious difficulties involved in conveying information about risks and securing informed consent for risky interventions in a clinical setting. We argue that in order to secure informed consent for a medical intervention, physicians often need to do more than report a bare, numerical probability value. When probabilities are given, securing informed consent generally requires communicating how probability expressions are to be interpreted and communicating something about the quality and quantity of the evidence for the probabilities reported. Patients may also require guidance on how probability claims may or may not be relevant to their decisions, and physicians should be ready to help patients understand these issues.

We take our paper to raise challenges for anyone who thinks that probabilities need to be reported in order to secure informed consent. Minimally, informed consent theorists need to consider whether patients can come to understand the difficult concepts and issues that we discuss in this paper. If patients cannot understand the relevant concepts and issues, it might not make sense to report probabilities, since doing so would not seem to accord with any of the leading rationales for requiring informed consent in the first place. Accordingly, the informed consent requirement may need to be relaxed. Alternatively, if patients can come to understand the relevant concepts and issues when properly advised — such that the concepts and issues we discuss are not a principled barrier to obtaining informed consent — then the medical community should train physicians to provide that advice and to help patients understand it (potentially through third-party patient activist decision theorists who are members of broader patient activist groups). Of course, this need not be an either-or choice: it might very well be the case that the informed consent requirement would need to be relaxed *and* that physicians would need to be further trained, so as to satisfy the existing rationales for informed consent in a wide array of situations in which reporting raw probabilities is insufficient.

Based on the challenges we highlight — and drawing inspiration from the medical ethics literature on informed consent — we make some tentative and, we hope, common-sense suggestions about what ethical practice could be with respect to informed consent for uncertain interventions. However, we fully intend our negative argument and our positive proposal to be a starting point for further discussion of probability and informed consent, not the last word.

Before proceeding, we would like to clarify several important aspects of the dialectic of this paper. First, the aim of this paper is not to defend the idea that physicians have a moral obligation to obtain informed consent or the idea that the law ought to require physicians to obtain informed consent. Indeed, we are attempting to engage people who accept either that physicians have some moral obligation to obtain informed consent or that the law ought to require physicians to obtain informed consent. Accordingly, we will often talk about an informed consent requirement without further specifying whether we are talking about a moral requirement or a legal requirement. We take our arguments to bear on both, but we are primarily concerned with reasonable procedural, legal requirements. Among those who accept that there is or ought to be some moral or legal informed consent requirement, we are especially interested in

those who accept an informed consent requirement that either explicitly endorses informing people about their risks or implicitly does so in virtue of the way the requirement is understood [1-2].

Second, we are not claiming that patients need a perfect understanding of probability and related concepts in order to give informed consent, nor are we suggesting they need all the information required to make an *optimal* decision. However, insofar as the information we point to is frequently especially relevant to rational decision-making, patients will often need *substantial* understanding of these concepts to make informed decisions.[1] We recognize that there is something paradoxical about our position. On the one hand, it seems that people cannot give informed consent without being informed of their risks, which requires understanding the chances of those risks. On the other hand, it seems that people routinely *do* give their informed consent and that they generally do *not* have substantial understanding of the chances of the risks they face, for reasons we provide in this paper. We will emphasize the difficulty of informing people about the chances of their risks, since such difficulties have largely been ignored in the medical ethics literature. One might worry that our proposal is far too strong and that informed consent can be obtained without patients understanding the features of risk that we discuss. We think the right thing to say here is that, despite appearances, people typically do not give informed consent. A determined critic might conclude optimistically that informed consent is not required for *genuine* consent or pessimistically that all consent requirements should be given up. But we are not ready to draw those conclusions. We will return to the difficulty here in the concluding section of the paper.

Third, we are not suggesting that extant justifications for informed consent requirements are wrong or need to be revised or reinterpreted. Rather, we are arguing that, in many situations, it is far from clear that providing patients with anything less than fully-interpreted probability claims along with the relevant evidence supporting these claims and an explanation of how the claims pertain to their decision-making, satisfies the justifications provided for informed consent in the literature. Finally, while we think it is advisable for physicians to use probability and statistics to communicate to their patients what is known about the risks and benefits associated with available procedures, we do not take any position with respect to *best practices* for successfully communicating first-order uncertainty as such.[2]

We proceed as follows. In the second (following) section, we briefly present the current state of the literature on informed patient consent with emphasis on disclosing probabilities. In the third section, which is composed of three subsections, we argue that securing informed consent often requires more than simply reporting bare, numerical probability values. We argue in the

[1] Here we follow authors who have argued that comprehension of the information disclosed is an important component of informed consent, and that medical practitioners should help their patients comprehend the information they disclose through effective communication [3; 4, ch. 4]. However, some authors have raised worries about aspects of the requirement to ensure understanding as part of informed consent [5-7].

[2] We are neutral with respect to whether one should use verbal probability expressions (such as 'likely' or 'probable') or numerical probability expressions (such as '20-percent chance'). For recent literature on the use of verbal and numerical probability expressions, see [8-17]. We are also neutral with respect to several related issues having to do with the presentation of probabilistic information. We are neutral with respect to whether one should use a point estimate or a probability range, for which see [18-21]. We are neutral with respect to whether one should use symbolic-algebraic representations or iconic-geometric representations, for which see [22-24]. And we are neutral with respect to what features of patients, such as numeracy [25], affect their understanding of probability reports. The issues we raise in this paper arise regardless of how these incredibly important and interesting debates are ultimately settled. In this connection, this literature suggests that there is no agreed-upon clinical guidance or accepted standard for how physicians are required to communicate the probabilities of different outcomes or risks involved in procedures.

first subsection that when physicians report probabilities, they generally need to explain how the concept of probability that they are using works. In other words, they need to specify the way in which they interpret their probability claims — whether in terms of degrees of belief, relative frequency, or something else. We argue in the second subsection that patients should often be informed as to the quantity and quality of the evidence for reported probabilities. And we argue in the third subsection that physicians should be prepared to help patients understand how probabilities might or might not be relevant to their decisions. Finally, in the fourth section, we discuss the consequences of our arguments for medical practice. We suggest a dialectical interaction between physicians and patients wherein patients are afforded opportunities to learn about their risks in detail.

Informed patient consent and the disclosure of probabilities

Many researchers in biomedical ethics recognize the importance of disclosing uncertainty in obtaining informed consent. According to Dan Brock, any attempt to obtain valid informed consent should include the presentation of the following information to the patient: the patient's current medical condition, including a prognosis if no treatment is pursued; treatment alternatives that might improve the patient's prognosis, including explanation of the procedures; the significant risks and benefits of each of the alternatives, with their associated probabilities; and a recommendation of the best alternative [26, p. 121]. Similarly, Jay Katz argues that patients need to know "the nature of the ailment, the nature of the proposed treatment, the probability of success or of alternatives, and perhaps the risks of unfortunate results and unforeseen conditions within the body" [27, p. 2].[3] Robert Young notes that "patients do need to know about the kind of risks they face, how likely it is that those risks will eventuate, and, if they do, what the effects will be and when they will occur" [29, p. 536]. And Jessica Berg et al. argue that there are four elements of the risk that should be disclosed: the nature of the risk, the magnitude of the risk, the probability that the risk might materialize, and the imminence of risk materialization [30, pp. 56-57]. All these authors take the disclosure of the uncertainty regarding prospective procedures to be very important, and most are explicit about including the *probability* of the risks and benefits of various treatment options in order to secure informed consent. Before discussing some problems related to the disclosure of probabilities, we wish to commence by explaining why disclosing probabilities is *prima facie* essential for obtaining informed consent. We do so by appealing to various justifications for requiring that physicians obtain informed consent from their patients.

The dominant justification given in the literature for both moral and legal requirements to obtain informed consent is respect for autonomy [4, ch. 4; 31-33].[4] However, some authors have considered alternative justifications. For example, Neil Manson and Onora O'Neill point to the role consent plays in regulating how patients may waive their ethical and legal rights [3, pp. 72-77], and Joseph Millum and Danielle Bromwich argue that consent governs the way normative

[3] Katz is quoting from [28].

[4] Patient autonomy as a justification for informed consent has been endorsed by the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research [34], and the President's Commission for the Study of Ethical Problems in Medicine and Biomedical and Behavioral Research [35]. For additional discussion of the relations between autonomy and informed consent, see [29; 36, ch. 7].

boundaries between patient and medical practitioner may be drawn [2]. Other justifications for the practice of requiring informed patient consent include preventing bodily trespass and promoting self-ownership [37], preventing abusive conduct [38], and building trust in the medical profession [1, ch. 7].[5]

All the various justifications for moral and legal informed consent requirements also support informing patients about the probabilities of the risks they face. The justification from autonomy requires that medical practitioners provide the patient with information that allows the patient to deliberate and genuinely choose a course of action for herself. Respect for patient autonomy also requires practitioners to do their best to provide information that allows the patient to choose a *good* course of action based on her rational deliberation. Therefore, medical practitioners must often provide information about the probabilities of the major benefits and risks of plausible courses of action. Without information about those probabilities, patients cannot, in most circumstances, reasonably weigh the alternatives or consider tradeoffs between costs and benefits. Hence, respect for autonomy frequently requires medical practitioners to provide information about probabilities in order to obtain informed consent.

Proponents of alternative justifications of the requirement to obtain informed consent should also endorse the claim that patients must often be provided with information about probabilities of good and bad outcomes in order to secure informed consent. First, information about the likelihood of good and bad outcomes is frequently essential to evaluating whether *prima facie* trespasses on one's body and *prima facie* violations of one's self-ownership are, all-things-considered, justified. Second, providing the patient with statistical information about outcomes helps in preventing abusive conduct and building trust within the medical community; patients would be 'armed' with all the relevant information, thus safeguarding them against abusive conduct (primarily by protecting patients from paternalistic practitioners who are uninformed or under-informed), and medical practitioners would be as truthful as possible, thus promoting trust between patients and the medical community. Finally, since the risks and benefits (and their probabilities) are essential elements of nearly any medical procedure, it is often only possible for patients to legitimately waive their right to bodily integrity or to redraw the normative boundaries between them and their physicians if they are informed as to the probabilities of the risks and benefits. For, in most cases, if patients are not informed as to the probabilities of the various risks and benefits of a proposed course of action, they will not understand either what rights they are giving up or the extent to which they are giving up their rights. Hence, providing patients with all the relevant information about the risks and benefits of a procedure — including the probabilities of those risks and benefits — should generally be part of any morally sound practice of redrawing the normative boundaries between patient and medical practitioner. In sum, providing patients with information about the probabilities of the risks and benefits of a procedure is generally required in order to secure informed consent not only according to the dominant justification for the informed consent requirement, but also according to several prominent alternative justifications for that requirement.

[5] For similar suggestions, see also [39-40]. For a paper that lays out an argument for informed consent based on building trust, but that attacks that argument and the idea more generally, see [41].

Informed consent and statistical decisions

Suppose that Jack goes to his physician seeking a vasectomy. Jack's physician tells him about the risks. For example, she tells Jack that there is a small chance of about 1% that his vasectomy will fail and that his wife Jill will become pregnant after the procedure [42]. Suppose Jack's vasectomy does fail, and Jill does become pregnant. Jack and Jill decide not to terminate the pregnancy. During labor, Jill experiences a lot of pain and asks for an epidural (central neuraxial block). Her physician tells Jill about the risks. For example, she tells Jill that there is a very small chance, of about 0.0006%, that an epidural will result in permanent injury [43, table 3].[6]

Were Jack and Jill adequately informed about their risks (assuming that an adequate *variety* of risks was discussed in ways similar to the selected examples)? We think that the answer is 'no' for three related reasons. First, as we argue in the first subsection below, the claim that some outcome has a given probability is often ambiguous unless an interpretation of probability is specified, and consequently, one cannot secure informed consent using a bare probability claim. Second, as we discuss in in the second subsection, not all probability claims are supported equally well by the evidence. In order for patients to be informed about their risks, they often need to understand the quantity and quality of the evidence for the claims at stake. But even interpreted probability claims do not necessarily say anything about the quantity and quality of the evidence in their support. Hence, in many cases one cannot secure informed consent merely by stating probabilities, even if those probabilities are properly interpreted. And third, as we argue in the third subsection, patients ultimately need to understand how probabilities may or may not be relevant to decision-making in order to be able to give genuinely informed consent. Patients need to understand the conditions under which probabilities are good guides to action. Among other things, even interpreted probability claims that are strongly supported by the evidence may be seriously misleading with respect to effective action. Hence, in many cases one cannot secure informed consent merely by stating probabilities, even if those probabilities are properly interpreted and the evidence for them is clearly described.[7]

Why think, as we do, that Jack and Jill were not adequately informed about their risks? After all, they were told what the risks were, and they were given the probabilities of those risks. Of course, we agree that something is better than nothing, and, in this respect, the physicians we have imagined in our story about Jack and Jill did better than many physicians in real life (at least, if our own experiences are a representative sample). However, "better than nothing" might still be inadequate. What, then, would be adequate? We do not think there is any specifiable threshold or set of necessary and sufficient conditions. Rather, there is a default obligation on physicians to inform their patients, and that obligation is satisfied by good-faith dialogue with the patient, attempting to disclose risks in a way that satisfies the patient without generating an

[6] One might think that physicians should provide only numerically precise frequency statements to patients because statements about natural frequencies are better understood (because more ecologically valid) than decimal percentages. For a recent meta-analysis of work on the natural frequency facilitation effect, see [14]. As with other presentation format issues, we are neutral with respect to whether probabilistic information should be presented as a natural frequency or as a decimal. We maintain that using natural frequencies does not solve all of the problems. We will return to this point later.

[7] Paul Han argues that different types of probability are important to communicate to patients. He focuses primarily on the difference between what he calls "epistemic" and "aleatory" uncertainty (subjective confidence versus known risk) rather than on the difference between Bayesian and Frequentist interpretations. Like us, he concludes that greater conceptual clarity is required for adequate communication [44].

illusion of understanding.[8] Moreover, we think that *if* a physician discloses a probability value, *then* the physician is automatically further obligated to interpret their probability statement. We take it that there is an upper limit to required disclosures, which is set (somewhat vaguely) by a reasonable person standard. When a patient requests information that no reasonable person would want, then the physician is permitted to refuse to provide the information. But we think that none of the information we consider in this paper falls into that category.

Interpreting probability claims

In the long history of probability and statistics, there have been two main approaches to the interpretation of the mathematical machinery of probability: the Bayesian approach and the Frequentist approach.[9] The approach one takes to interpretation usually goes together with the approach to statistical inference one endorses and the approach to decision-making under risk that one adopts. If one is a Bayesian (or Frequentist) about the interpretation of probability, then one is likely to be a Bayesian (or Frequentist) about statistical inference and decision. In this subsection, we review the core ideas of the Bayesian and Frequentist approaches to the interpretation of probability, and we argue that reporting a bare, uninterpreted probability of success with respect to a proposed procedure or intervention is not sufficient to secure informed consent. We return to issues about statistical inference and decision in the second and third subsections.

Traditionally, Frequentists understand probability claims to be statements about how often some event occurs within a specific collection of events, called the *reference class*. When a Frequentist says that some event type has a 1% chance of occurring, she is saying that in some specific reference class, 1 in every 100 members is a token of that event type. The reference class might be infinite, in which case the Frequentist takes the probability of an event type to be the *limiting* frequency with which it occurs. Put briefly: Probability is a measure of the (limiting) relative frequency of some type of event in a (possibly infinite) collection of events. By contrast, a Bayesian understands a probability claim to be a statement about her *credence*, by which we mean either her personal degree of belief that some proposition is true or the degree to which her evidence supports some proposition. So, Bayesians — but not Frequentists — take probability claims to be agent-relative.

So what is it that the physician is telling Jack when she says that there is a 1% chance that his vasectomy will fail? If she is a Bayesian, she is reporting her credence that Jack's vasectomy will fail. If Jack's physician is a Frequentist, she is making a claim about the frequency with which vasectomies fail in some reference class. Clearly these interpretations are not equivalent. The Bayesian interpretation — but not the Frequentist interpretation — licenses inferences regarding the physician's attitude toward a proposition. The Frequentist interpretation — but not the Bayesian interpretation — licenses inferences regarding repeated sampling from a specific reference class. Since the two interpretations license at least some different inferences, the question becomes pressing as to whether Jack's decision depends on the interpretation given to

[8] Patients might waive the default obligation if, like Han Solo, they do not want anyone to tell them the odds. But patients need to be afforded the opportunity to learn about their risks and to understand what they are giving up when they waive the physician's obligation.

[9] There are many ways to work out the details of each main approach. For examples of varieties of Frequentism, see [45-46]. For examples of varieties of Bayesianism, see [47].

the chance claim.

More concretely, we can imagine a few different interpretations of the physician's claim. She might mean that around 1% of vasectomies performed by Jack's physician have failed. She might mean that 1% of vasectomies performed in Jack's country in the last 10 years have failed. Or she might mean she is personally 1% confident that Jack's vasectomy will fail. The first two interpretations are Frequentist. In the first interpretation, the reference class is patients Jack's physician has personally vasectomized. On the second, it is vasectomy patients in Jack's country in the last ten years. The third is Bayesian. On either Frequentist interpretation, Jack may think he differs in important ways from the relevant reference class. For instance, he may suspect he is younger or older, or more or less fertile, than the patient population. Therefore, his personal risk could be higher or lower than the 1% figure quoted. On the Bayesian interpretation, Jack may wonder whether the physician's personal views would be widely shared amongst experts. If, for instance, other experts would assess his risk at 3%, he may very well opt not to undergo the operation.

Accordingly, the physician has a *prima facie* duty to inform Jack about her interpretation of probability as part of obtaining his consent, since Jack's decision plausibly depends on the interpretation. In such cases, knowing only a bare probability value does not provide an adequate basis for rational decision-making, including waiving normative requirements. If so, then knowing only a bare probability value is not sufficient to secure informed consent.

Physicians might be tempted to say that there is no major ambiguity in their probability claims because *obviously* they are reporting actual frequencies. But we are not so sure. Suppose a physician has access to several related, but only *partially informative* studies. For example, a physician might be tasked with treating an elderly Canadian of Inuit descent but only have evidence from a study on African Americans living in the United States and a study on elderly people living in Australia. When a physician has access to one or more studies involving participants who are similar to the patient in many respects but also different from the patient in potentially important ways, frequency information does not directly bear on the question at issue. In such cases, the physician's chance claim is best interpreted as a report of credence with respect to the proposition that the intervention will be successful. Similarly, suppose a physician is trying to communicate the risks with respect to a procedure that has never been tried or that has been tried only a very small number of times. In such cases, if the physician says that the procedure has some probability of success, she is not making any claim about *actual* frequencies.[10] She might be making a claim about *hypothetical* frequencies, but we think it is more likely that she is simply expressing her credence that the procedure will be successful. Moreover, even if physicians *intend* to report actual frequencies in all cases where they report

[10] A referee wondered whether our claim here is really correct and suggested that the physician could be making a claim about the frequency of success in a range of related cases. At a first pass, we want to resist this suggestion and say that the imagined physician is not making a claim about the actual frequencies *with respect to that specific kind of procedure*. Our interlocutor could respond by saying that the physician includes the new procedure in a wider reference class. But if that response works, then the procedure is not best thought of as a new one that has never been tried. Going a bit further, suppose the physician says, "There is a 20% chance of death with this procedure." If the physician means that in some wider reference class, death results 20% of the time, it seems to us to be important that the patient knows both *that* the claim is about a wider reference class (since Gricean conversational maxims would ordinarily suggest that the physician's claim is maximally specific) and also *what* the reference class really is. Moreover, if the physician's claim is about the frequency of success for some different but related procedures, then, in addition to ambiguity in the physician's claim, there is an important inferential problem, which we discuss in more detail in the subsection below.

probabilities, they need to present information in a way that is intelligible to their patients.[11]

On the evidence for probability claims

In the previous subsection, we reviewed the core ideas of the Bayesian and Frequentist approaches to the interpretation of probability, and we argued that reporting a bare, uninterpreted probability of success for a proposed medical procedure or intervention is generally not sufficient to secure informed consent. In this subsection, we argue that even an *interpreted* probability claim is often not sufficient to secure informed consent. In order for patients to be informed about their risks, they frequently need to understand the quantity and quality of the evidence for the probabilities that their physicians report.[12] T.M. Cook et al. suggest a particularly strong version of such a requirement. Lamenting the lack of reliable estimates for the frequency of injuries from neuraxial block, they say that physicians need to know and accurately report the frequency with which complications arise in order to obtain genuinely informed consent. They write:

> Knowledge of the incidence of such complications [as permanent injury from neuraxial block] should be an essential component of the clinical decision-making and consent processes, but there are few good data which can be quoted to support such discussions, leaving both patient and clinician in a quandary. Figures (ranging from 1:1000 to 1:100 000) are quoted, but their doubtful validity questions the ability to obtain genuinely informed consent from patients offered these procedures [43, p. 180].

Cook et al. do not provide an argument for the claim that physicians need to know the frequency with which complications arise in order to obtain genuinely informed consent from their patients, but we take the following to be a charitable attempt to lay out an argument on their behalf.

> [C1] If physicians do not know the frequency with which complications occur for a medical intervention, then physicians cannot non-accidentally report a true and reliable estimate of the frequency with which complications occur for that intervention.
>
> [C2] Physicians can satisfy their obligation to secure informed consent regarding a medical intervention only if they can non-accidentally report a reliable estimate of the

[11] One might think of the issue here as a consequence of the fact that consent is a propositional attitude with intensional context, for which see [48]. A patient might assent to undergoing a procedure that has a 70% probability of success but not a procedure that succeeds 70 times out of 100.

[12] What we have in mind here is similar to what is sometimes called "ambiguity" [44]. However, we think there is an important distinction to be made between evaluating the quantity and quality of evidence for a probability claim and the complete uncertainty about probability (what we would count as genuine ambiguity) that figures in Ellsburg's paradox and related puzzles in decision theory. While some worry that ambiguity aversion will lead patients to avoid decision making if they are presented with higher-order probabilities [44], other studies suggest that in cases where the quantity and quality of evidence is good, "including weight of evidence content … attenuates perceived information uncertainty" [49, p. 1302].

frequency with which complications occur for that intervention.[13]

[C3] If physicians do not know the frequency with which complications occur for a medical intervention, then they cannot satisfy their obligation to secure informed consent with regard to that intervention.

The standard advanced by Cook et al. is very strict. The physician must *know* the frequency of complications. Cook et al. do not specify in any detail how much evidence is required to know such a thing. However, their procedure for obtaining an estimate gives some indication. They attempted to "identify both numerator (number of major complications) and denominator (number of CNB) information for a 12 month period by a review across the breadth of anesthetic and pain management practice in the UK National Health Service (NHS)" [43, p. 180]. Something that Cook et al. do not address is whether and to what extent physicians need to *communicate* the quantity and quality of their evidence. We think that communicating the quantity and quality of the evidence is indispensable for securing informed consent. Consider the following example.

Jill's physician tells her that there is a 0.0006% probability of permanent injury from an epidural, and she explains that her probability claim is expressing the frequency with which permanent injuries occur in such procedures. That strikes Jack and Jill as an acceptably low chance of injury. But Jack wants to know more. After all, Jack thought that the probability of his vasectomy working was high, but Jill got pregnant anyway. Once bitten, twice shy, Jack asks for more information about the physician's claim that there is a 0.0006% probability of permanent injury from an epidural. The physician says, "Look, Jack, there was a large study of complications from epidurals; and the study was explicitly trying to determine the frequency of permanent injury in these procedures." The physician then refers Jack to Cook et al. who report that there were 161,550 obstetric epidural procedures performed in the UK during the year covered by their study with only one patient suffering permanent injury [43].

How good is the evidence that Jill has a very small chance of suffering a permanent injury from her epidural? Clearly it is much stronger than if Cook et al. had sampled 100 patients. A little bad luck in a sample of that size would produce an estimate more than 1500 times too large (assuming that Cook et al.'s estimate is correct). It is not unusual for medical studies to vary widely in sample size. For example, the estimate of vasectomy failure from the aforementioned study was based on 540 women [42]. But the quantity and quality of the evidence for a probability claim is relevant to whether and how the claim should guide action.[14] Suppose two researchers collect data on the success of a procedure. The first researcher studies 100 people and observes 1 failure. The other researcher studies 10,000 people and observes 100 failures. The

[13] A referee wondered whether a physician pleading ignorance could be the basis for valid consent. Up front, we take pleading ignorance to be incompatible with the demand in [C2]. In a broad range of typical therapeutic, clinical cases — probably in most of them — we think pleading ignorance would be impermissible, since it would either mean that the physician was not competent to advise the patient or that the physician was untruthful and unwilling to show proper concern for the patient's interests. However, not all cases are the same, and in especially difficult or unusual cases or in cases where everything has been tried, even competent and caring physicians may have no good estimate to give. In those cases, disclosing ignorance is probably *required* and is perhaps also *sufficient* to obtain informed consent. Hence, we think that the requirement in [C2] is probably too strong. In experimental research contexts, we are much less sure what to say. We think that uncertainty still needs to be disclosed, but we are not sure that the form of such disclosure will be similar to the therapeutic, clinical context.

[14] By "evidence for a probability claim," we here mean evidence that bears on an estimate of the population mean, that is, limiting frequency. When the sample is small, we have lower quality estimates of the limiting frequency.

point estimates of the probability of success will be 0.99 in both cases, but the likely error in the two estimates will be very different. Patients may not know how and why larger samples tend to yield more secure estimates (though we suspect that most people will understand intuitively that it is true), and we are not suggesting that patients need to be taught about sampling variability and statistical error. However, we do think physicians need to communicate *something* about the quantity and quality of the evidence for the probabilities they report.[15] Reporting the size of the samples might not be the right strategy, since patients might not distinguish for themselves between samples of size 100 and 10,000. If they do not distinguish between such cases, then it would be better to give more qualitative assessments, for example, that this estimate is derived from a very large amount of high-quality data or from a small amount of high-quality data or from a small amount of low-quality data. Can we say more about *what* physicians should communicate about the quantity and quality of their evidence and *how* they should do so?

Physicians could adopt the usual Frequentist strategy and describe the reliability of the statistical method used in producing the estimate. This would fit well with the fact that so much of the medical literature uses Frequentist statistical tools. For example, Jill's physician might point to the 95% confidence interval of (0, 3.4) that Cook et al. report for their point estimate of 0.6 permanent injury events per every 100,000 cases [43]. "You see, Jack," she might say, "if lots of researchers gathered data about the frequency of permanent injuries from neuraxial block, then about 95 out of every 100 intervals constructed according to the method used by Cook et al. would have the true frequency in its range." But Frequentist confidence intervals are difficult to understand, even when they have been carefully explained in detail. Rink Hoekstra et al. provide empirical evidence that many researchers in psychology — a discipline that makes extensive use of Frequentist statistical tools — do not have a solid understanding of confidence intervals [55].[16]

If people who are familiar with statistics — people who learn statistics as part of their professional training and who use statistics every day in their research work — misunderstand confidence intervals at such high rates, then physicians cannot expect typical patients to find confidence intervals intuitive. One might be tempted to think that physicians should avoid inferential statistics altogether and report only raw, observed frequencies. Patients would then have to draw their own conclusions about the limiting frequencies. But we worry that this would put too heavy a burden on physicians, perhaps requiring that they provide all the information needed for patients to draw their own conclusions.[17] For one thing, physicians might need to

---

[15] Peirce used the label "weight of evidence" to refer to what we would call the "quantity" of the evidence [50]. For further discussion, see [51, ch. 6; 52; 53, ch. 14]. For discussion of relatively recent use of the phrase "weight of evidence" in biomedical science, see [54].

[16] Hoekstra et al. modeled their study on a famous study by Gerd Gigerenzer, which provided empirical evidence that researchers frequently misinterpret p-values, another mainstay of Frequentist statistics [56]. In their study, Hoekstra et al. told a story about a professor who reports a 95% confidence interval of (0.1, 0.4) for a mean value being estimated. They then asked 118 researchers to say whether each of six statements was true or false. The number of researchers reporting the wrong answer ranged from 45 to 102. For example, 70 out of 118 endorsed the false claim that there is a 95% probability that the true mean lies between 0.1 and 0.4, and 68 out of 118 endorsed the false claim that if we were to repeat the experiment over and over, then 95% of the time the true mean falls between 0.1 and 0.4. The first of these mistakes a Frequentist confidence interval for a Bayesian credible interval. The second treats the confidence interval as fixed and the parameter being estimated as variable, whereas the Frequentist thinks of the interval as variable and the parameter as fixed [55].

[17] A referee observes that even the idea of what is needed by patients presents serious difficulties, since what is needed by some patients might not be needed by others. It is not clear, then, whether the requirement is to provide all the information that *this patient* needs or all the information that a reasonable patient would need or all the information

report facts about experimental designs — such as how the data were sampled — since the sampling procedure sometimes matters for the inference drawn.[18]

It may seem, then, that doctors should abandon frequencies and instead report their (Bayesian) credences on the assumption that such reports are easier to understand. However, Bayesian credences are not without problems of their own. Here, we want to illustrate two problems having to do with communicating the evidence that supports a given credence under a Bayesian interpretation. The first problem is that two (or more) very different bodies of evidence may produce the same credence. Jill wants to know the probability that she will suffer a permanent injury. If Jill's doctor were a Bayesian, she would calculate the probability of permanent injury given her evidence. Suppose Jill's doctor started her career thinking that all the possible values for the probability of success were equally likely. That is, she thought it was just as likely that permanent injuries were very probable as that they were extremely unlikely and so on. She then observed some number of trials, all of which she regarded as independent. If Jill's doctor has seen one permanent injury in the course of 98 trials, she calculates the probability that Jill will suffer a permanent injury as 2%.[19] Jill's doctor would *also* have reported a 2% chance of permanent injury if she had observed no permanent injuries in 48 trials or if she had observed 3 permanent injuries in 198 trials. But those bodies of evidence lead to very different credence *distributions*, and hence, they are not all equally informative. In order to communicate the quantity and quality of her evidence, Jill's physician might specify a range — a *credible interval* — such that she has a high credence that the frequency of permanent injuries is in the interval. For example, she might be willing to bet at 20 to 1 odds that the number of permanent injuries per 100,000 is in the interval (0, 4).[20] That would be an expression of her credence regarding the frequency, and it would represent her assessment of the quantity and quality of the underlying evidence that bears on Jill's risk of permanent injury. But it would go well beyond reporting an interpreted probability value.

More commonly, a doctor will not have such readily articulable evidence, which raises problems of transparency and expert disagreement. To illustrate, suppose that Richard is a morbidly obese 66-year-old male considering whether to continue to rely on a urinary catheter or undergo urethral reconstructive surgery because of stricture. He has a long history of urinary tract infections that are resistant to antibiotics and is also on immuno-suppressants because he has received a kidney transplant. He has had Crohn's disease since his twenties and had large parts of his intestines removed decades ago. Richard asks his doctor how likely it is that the operation will be a success. Because the doctor knows that Richard has complications that make him importantly different from the average patient undergoing this operation, she does not report historical frequencies of success. Instead, she takes his history into account and reports her personal confidence level of 70% that the operation will be successful. She might even include a

---

that a suitably large percentage of ordinary patients need or something else. But since here we are being critical of the idea of simply presenting all the needed information, we take the referee's point to be grist for our mill.

[18] Even if one is a Bayesian, the sampling procedure may matter in special cases [57].

[19] In the circumstances described, she applies Laplace's Rule of Succession and says that the probability of injury on the next procedure is equal to $(m + 1) / (n + 2)$, where m is the number of injuries observed out of n procedures. For more on the Rule of Succession, see [58, ch. 2].

[20] Traditionally, Bayesian credences have been interpreted as betting odds. If you are 70% sure that an event will occur, it means (roughly) you personally are willing to pay up to 70 cents for a bet that returns $1 if the event occurs and nothing otherwise. Although such an interpretation is illustrative, it is not without problems. In particular, it seems distasteful (at the least) to bet on whether a patient's operation will be successful. Nonetheless, we think the betting odds framework helps highlight the difference between the Frequentist and Bayesian approaches.

99% credible interval in her report in order to more accurately communicate the quantity and quality of her evidence. But despite the fact that Richard's doctor is an expert with extensive clinical experience, her figure of 70% was (likely) not arrived at through any careful or transparent means. And importantly, the basis of her initial credence is not open to inspection or criticism. Another physician who is equally qualified may very well report a different number. More generally, if the experts' credences are highly sensitive to their priors, then there will be a lot of possible variation between the credences of different experts. In such cases, experts have some duty to inform patients that other experts may have or likely *will* have different credences. After all, patients very likely would make different decisions depending on the doctor's reported credence, so patients should be alerted to situations where such credences are likely idiosyncratic or highly variable.[21] Again, the result is that even reporting an interpreted probability is not necessarily sufficient to secure informed consent.[22]

We have so far focused on issues of data *quantity*, but quality matters, too. A study having 100 participants selected at random from a population to which a target patient belongs is much better than a study having 100 participants selected systematically or selected from a population to which the target patient does not belong. Similarly, other things being equal, data collected in a controlled trial will be more informative than observational data, especially with respect to whether a proposed treatment is causally effective. In order to secure informed consent, physicians should often report about both the quantity and quality of their evidence.

Decisions, decisions

In the previous two subsections, we argued that reporting bare, uninterpreted probability values is frequently not sufficient to secure informed consent, and we argued that in addition to including an interpretation, patients often need to understand the quantity and quality of the evidence that grounds the probabilities that physicians report to give genuinely informed consent. Now, suppose a physician reports an interpreted probability that some procedure will be successful, and suppose she also adequately conveys the quantity and quality of the evidence for her probability claim. Will her reports be sufficient to secure informed consent? In this subsection, we argue that they will frequently not be. A report may be misleading because the probability value reported is not action-guiding.

Suppose that Jack develops kidney stones. He goes to see a local nephrologist, who describes two possible surgical procedures. One involves open surgery, and the other is minimally invasive. The nephrologist tells his patients that open surgery has a success rate of 72% and that

[21] A referee suggested that physicians could avoid the difficulty of expert disagreement by simply telling patients that other experts might have different views and then reminding patients that they can talk with other physicians before proceeding. In some cases, we think that would be sufficient, but there are many circumstances where it would not. For example, advising patients about the possibility of consulting with other physicians may not be adequate if one is a rural doctor seeing a patient who has very limited options or if one's patient is poor or has insurance that limits their ability to see other physicians or if the patient needs to act quickly. And we expect there are other similar kinds of cases. Even if the recommendation worked, however, it would concede the point we want to make here: that a bare numerical probability statement is not sufficient to secure informed consent.

[22] Professional or expert judgment about risk is an important issue in areas outside biomedical ethics, as well. For some discussion of expert judgment in engineering ethics, see [59]. In certain cases, physicians may well come close to consensus even if they have trouble articulating precisely what their evidential basis is. However, since physicians disagree relatively often, such cases are far from universal.

the minimally invasive surgery has a success rate of 83%. He tells his patients that these probabilities should be understood in terms of frequencies, and he tells his patients that his probability claims are based on a well-designed study involving more than 400 subjects. Everything he says is true, and as a result of his report, essentially all of his patients choose the minimally invasive option. Jack is not unusual in this regard, and he chooses the minimally invasive surgery as well. But unknown to Jack, the nephrologist is not an honest man. The nephrologist knows that this is a case of Simpson's Paradox and that among patients with small kidney stones (< 2cm), open surgery has a success rate of 92% compared to 87% for the minimally invasive option, and among patients with large kidney stones (≥ 2cm), open surgery has a success rate of 71% compared to 69% for the minimally invasive option.[23] However, he only ever reports aggregated probabilities because he makes more money performing minimally invasive surgeries, and he knows, let us say, that if he were to report the success rates separated out by size of stone, his patients would be much more likely to choose open surgery.[24]

We think it is clear that the nephrologist acted unethically and that Jack did not give genuinely informed consent for his surgery. The physician knows (and withholds) something that is relevant to Jack's decision. Specifically, it is plausible that conditioning on the size of the kidney stones discloses something important about the causal structure, and the physician knows (or suspects) that this is the case.[25] In the case of the dishonest doctor (and in many other cases where we need to choose an *action*), probabilities are relevant to decisions only insofar as they track the causal structure, and hence, the aggregated probabilities that the doctor reports are misleading.[26]

Our story about Jack and the unscrupulous nephrologist shows that even reporting a properly interpreted probability value together with an accurate statement about the quantity and quality of the evidence for that probability value might not be sufficient to secure informed consent. However, our tale is one of obvious abuse, and the problematic feature of the account is quite clear. Jack's physician knew something that was relevant to Jack's decision but intentionally withheld that information. Failing to provide information due to negligent ignorance would also block genuinely informed consent. But what if a physician neither withholds information intentionally nor fails to provide it due to negligence but simply does not know the relevant information? Physicians surely do not always need to know the causal structure of a given case in order to properly inform their patients. However, physicians should have (and should express that they have) no reason to think that the probabilities they report are not action-guiding. That is, they should think the probabilities they provide are the best information available for rational decision-making. Moreover, physicians should frequently explain how and why probability

[23] Simpson's Paradox occurs when an association between two variables disappears or reverses conditional on a third variable. The threat of Simpson's Paradox is non-trivial. Hanley and Thériault reported an example of Simpson's Paradox in meta-analyses of randomized controlled trials [60]. Nissen and Wolski found a significant increase in the risk of myocardial infarctions in groups taking Rosiglitazone over the control groups over a number of individual studies. However, when data from all the studies were pooled, the Rosiglitazone group had a slightly lower rate of MIs compared to the control [61].

[24] The numbers for our toy example are based on a famous real-life study by C.R. Charig et al. (1986) that compared different methods for treating kidney stones [62]. The unscrupulousness is novel to our story.

[25] For our purposes, a "causal structure" is a pattern of causal relations that hold with respect to some domain. Minimally, we take a causal relation to reveal how an outcome depends on changes in actions one could take. One would want to know, for example, what would happen to the success rate if one were to choose open surgery rather than a less invasive option. For introductions to some of the issues here, see [63-64].

[26] In the language of Pearl's 'do calculus,' a report of the value of $Pr(Y=y \mid X=x)$ is action-guiding only if it is a guide to the value of $Pr(Y=y \mid do(X=x))$ [65].

values sometimes are and sometimes are not decision-guiding. Patients may not be aware of Simpson's Paradox or related issues. If a patient knows only that success occurs with a specific probability value for some sample population, she may not know that the probability of success could be very different for an alternative subgroup to which she belongs and that subgroup membership tracks the decision-relevant causal facts.

A similar problem arises when either the sampled population is heterogeneous or idiosyncratic features of the patient are potentially highly relevant. Suppose Jack and Jill were deciding between no sterilization for either of them, a vasectomy for Jack, or a tubal sterilization for Jill. Suppose Jack would have another child if Jill wanted, but Jill is fairly sure she does not want one. Nonetheless, she worries that she will regret sterilization. Jack and Jill ask their doctor for advice. The doctor points to a study that found that after five years, 6.1% of women studied whose husband had a vasectomy said they regretted the procedure, whereas 7% of women expressed regret five years after tubal ligation [66].

Although these frequencies are potentially useful, they can also be highly misleading. Jack and Jill are, like any couple, idiosyncratic. They have features that make them importantly different from any member of the sampled population. For instance, suppose Jill is very decisive and rarely changes her mind but also has family who would be upset if they learned Jill decided to have any sterilization procedure done. She and Jack also have four children already — more than the average American family — and additional children would add a significant financial burden after Jack was recently laid off. Mere frequency data should be used with extreme caution in this case.

The upshot is that reporting an interpreted probability that some procedure will be successful might not be sufficient to secure informed consent even if the physician adequately conveys the quantity and quality of the evidence supporting her probability claim. Not all probabilities are action-guiding, and one can have better or worse reasons to think that a probability value is action-guiding. Hence, it is often important to convey information about the decision-relevance of probability values to patients in order to obtain genuinely informed consent.

Conclusion

The literature on informed consent in medicine recognizes that information about the probabilities of success and failure associated with possible treatments should be communicated to patients in order to obtain genuinely informed consent. In this paper, we have argued that to secure informed consent for a medical intervention, physicians often need to do more than report a bare, numerical probability value. When probabilities are given, securing informed consent generally requires communicating how probability expressions are to be interpreted and communicating something about the quality and quantity of the evidence for the probabilities reported. Patients may also require guidance on how probability claims may or may not be relevant to their decisions, and physicians should be ready to help patients understand these issues.

Up to this point, we have emphasized the negative, pointing out what seem to us to be barriers to securing informed consent. Indeed, one possibility that we noted at the outset is that patients may simply not understand the relevant concepts and issues, and thus it might not make sense to report probabilities, since doing so would not seem to accord with any of the leading

rationales for requiring informed consent in the first place. Accordingly, the informed consent requirement may need to be relaxed. But if we assume that patients can come to understand the relevant concepts and issues when properly advised, such that the concepts and issues we discuss are not a principled barrier to obtaining informed consent, how much information is really *required?* In closing, we want to make a tentative positive proposal about how the difficulties we have considered might be overcome.

Our suggestion is for physicians to engage in a dialogue that affords patients sufficient opportunity to learn about their risks. We assume here a reasonable patient standard, as opposed to a professional, physician-centered standard, for determining what constitutes sufficient opportunity overall and for determining what counts as a sufficient response at each stage in the suggested dialectic. Some people may prefer to be told *nothing* about their risks, trusting physicians to know what is best for them. And in practice, some patients do not want to know the details of especially bad diagnoses or the risks of dangerous procedures.[27] So, we suggest that discussion with any patient begin with the question whether the patient wants to be informed of the relevant risks. Assuming the patient wants to be informed of the relevant risks, the physician should first tell the patient of the most important and probable risks. After doing so, the physician should ask whether the patient would like to know the probability of those risks and how the physician determined which risks were counted as among the most important and probable. The point here is that patients should be informed that there is (potentially) more to know about the total collection of risks that they face and that there are reasons (perhaps vague but nonetheless present) for the presentation being made. Assuming the patient wants to be informed about the probability of the risks, the physician should provide (a) a probability estimate,[28] (b) an initial interpretation of the probability estimate (e.g., disclosing whether the probability represents a relative frequency or a credence), (c) a qualitative assessment of the evidence on which the estimate is based, and (d) either an assurance that the probability estimate is decision-relevant or a caveat that the probability estimate is not known to be action-guiding (but also not known *not* to be action-guiding). Then, the physician should check the patient's understanding and ask whether the patient would like further explanation about the probability of the risks. Physicians should answer questions and actively help their patients to understand the risks. Ideally, a physician will be sensitive to the demeanor of the patient and help the patient to formulate questions as well as providing answers. For example, the physician might help the patient to ask about confidence or credible intervals, about causal structure, or about experimental design in relevant medical research trials.

Notice that, in our suggested physician-patient dialogue, when it is required to report a probability, the reported probability value should usually be interpreted. As we have argued, providing a patient with a bare probability claim does not necessarily satisfy the justifications provided for informed consent in the literature. Consider the autonomy justification for informed consent: given the issues we have raised, if the physician provides an uninterpreted probability, she will not generally be providing information that is helpful to the patient in choosing a good course of action; in particular, she will not generally be providing information that will help in reasonably weighing alternatives or considering tradeoffs between costs and benefits. And this

[27] The case of *Arato v Avedon*, which deals explicitly with the relationship between informed consent and statistical information, turned in part on the perception among oncologists that many patients did not want to be told the hard truth about their condition [67].

[28] This is in line with empirical findings regarding effective communication of first-order uncertainty, for which consult the literature noted in fn. 2.

means that this type of information does not necessarily enhance patients' autonomy. Moreover, recall that proponents of alternative justifications of the requirement to obtain informed consent also had reason to endorse the claim that patients should generally be provided with information about probabilities of good and bad outcomes in order to secure informed consent. However, if the information that is provided to patients is unhelpful to their decision process, it is far from obvious how it might help patients assess whether *prima facie* trespasses on their bodies and *prima facie* violations of their self-ownership are justified or whether they should waive ethical and legal requirements related to medical procedures. Moreover, without any interpretation, probability claims do not generally serve to prevent abusive conduct or to build trust within the medical community. Without some interpretation, the best that can be said for probability claims is that they often provide patients with some *illusion* of understanding.

We recognize that these suggestions, if accepted, would impose a significant burden on physicians and on the medical community. Indeed, preparing physicians to do what we are proposing would require resources and planning. Physicians would need to be trained much more substantially in statistics and the philosophical foundations of statistics. Since even active researchers who use statistics daily often misunderstand basic statistical tools and techniques, such as confidence intervals and the like, clinicians cannot currently be expected to understand such concepts fully either. Moreover, physicians would need to be trained on the relationship between probability and statistics on the one hand and decision-making on the other. Additional courses in statistics, in causal reasoning, in decision theory, and in the philosophy of probability and statistics, along with courses or workshops on conveying statistical information to patients would have to be added to the curricula of medical schools (and other medical programs).[29] As an alternative, the medical community might consider employing third-party patient activist decision theorists to help patients understand the technical details relevant to their decisions. Such theorists could be members of broader patient activist groups and could help patients understand how the probabilities presented to them pertain to the decision at hand.

Acknowledgments We would like to thank the following individuals for invaluable feedback on an early draft of this paper: Richard Bradley, Kathleen Creel, Samuel Fletcher, Conor Mayo-Wilson, Colleen Murphy, and Jonah Schupbach. We would also like to thank audiences at workshops at the Romanell Center for Clinical Ethics and the Philosophy of Medicine and at the Rotman Institute of Philosophy for their excellent feedback. Finally, we are grateful to two anonymous referees for *Theoretical Medicine and Bioethics* for detailed comments that aided in substantially improving the paper.

[29] Paul Han et al. have developed and evaluated an experience-based clinical risk communication curriculum for medical students, which, although resource-intensive, was efficacious [68].

References

1. O'Neill Onora. 2002. *Autonomy and trust in bioethics*. Cambridge: Cambridge University Press.

2. Millum, Joseph, and Danielle Bromwich. 2018. Understanding, communication, and consent. *Ergo* 5: 45-68.

3. Manson, Neil C., and Onora O'Neill. 2007. *Rethinking informed consent in bioethics*. Cambridge: Cambridge University Press.

4. Beauchamp, Tom L., and James F. Childress. 2008. *Principles of biomedical ethics* (6th edition). Oxford: Oxford University Press.

5. Sreenivasan, Gopal. 2003. Does informed consent to research require comprehension? *Lancet* 362: 2016-2018.

6. Miller, Franklin G., and Alan Wertheimer. 2011. The fair transaction model of informed consent: An alternative to autonomous authorization. *Kennedy Institute of Ethics Journal* 21: 201-218.

7. Walker, Tom. 2012. Informed consent and the requirement to ensure understanding. *Journal of Applied Philosophy* 29: 50-62.

8. Christopher, M. M., C. S. Hotz, S. M. Shelly, and P. D. Pion. 2010. Interpretation by clinicians of probability expressions in cytology reports and effect on clinical decision-making. *Journal of veterinary internal medicine* 24: 496-503.

9. Collins, Peter J., and Ulrike Hahn. 2018. Communicating and reasoning with verbal probability expressions. *Psychology of Learning and Motivation* 69: 67-105.

10. Hamann, J., R. Mendel, M. Bühner, S. Leucht, and W. Kissling. 2011. Drowning in numbers–what psychiatrists mean when talking to patients about probabilities of risks and benefits of medication. *European Psychiatry* 26: 130-131.

11. Hanauer, David A., Yang Liu, Qiaozhu Mei, Frank J. Manion, Ulysses J. Balis, and Kai Zheng. 2012. *Hedging their mets*: The use of uncertainty terms in clinical documents and its potential implications when sharing the documents with patients. *AMIA Annual Symposium Proceedings* 2012: 321-330.

12. Jenkins, Sarah C., Adam J. L. Harris, and R. M. Lark. 2018a. Understanding 'unlikely (20% likelihood)'or '20% likelihood (unlikely)'outcomes: The robustness of the extremity effect. *Journal of Behavioral Decision Making* 31: 572-586.

13. Karelitz, Tzur M., and David. V. Budescu. 2004. You say "probable" and I say "likely": improving interpersonal communication with verbal probability phrases. *Journal of Experimental Psychology: Applied* 10: 25-41.

14. McDowell, Michelle, and Perke Jacobs. 2017. Meta-analysis of the effect of natural frequencies on Bayesian reasoning. *Psychological Bulletin* 143: 1273-1312.

15. Spiegelhalter, David. 2017. Risk and uncertainty communication. *Annual Review of Statistics and Its Application* 4: 31-60.

16. Alan R. Tait, Terri Voepel-Lewis, Brian J. Zikmund-Fisher, and Angela Fagerlin. 2010. Presenting research risks and benefits to parents: Does format matter? *Anesthesia and Analgesia* 111: 718-723.

17. Vahabi, Mandana. 2010. Verbal versus numerical probabilities: Does format presentation of probabilistic information regarding breast cancer screening affect women's comprehension? *Health Education Journal* 69: 150–163.

18. Dieckmann, Nathan F., Robert Mauro, and Paul Slovic. 2010. The effects of presenting imprecise probabilities in intelligence forecasts. *Risk Analysis: An International Journal* 30: 987-1001.

19. Jenkins, Sarah C., Adam J. L. Harris, and R. Murray Lark. 2018b. When unlikely outcomes occur: the role of communication format in maintaining communicator credibility. *Journal of Risk Research* 22: 537-554.

20. Longman, Thea, Robin. M. Turner, Madeleine King, and Kirsten J. McCaffery. 2012. The effects of communicating uncertainty in quantitative health risk estimates. *Patient Education and Counseling* 89: 252-259.

21. Sladakovic, Jovana, Jesse Jansen, Jolyn Hersch, Robin Turner, & Kirsten McCaffery. 2016. The differential effects of presenting uncertainty around benefits and harms on treatment decision making. *Patient Education and Counseling* 99: 974–980.

22. Lipkus, Isaac M. 2007. Numeric, verbal, and visual formats of conveying health risks: suggested best practices and future recommendations. *Medical Decision Making* 27: 696-713.

23. Spiegelhalter, David, Mike Pearson, and Ian Short. 2011. Visualizing uncertainty about the future. *Science* 333: 1393-1400.

24. Zipkin, Daniella A., Craig A. Umscheid, Nancy L. Keating, Elizabeth Allen, KoKo Aung, Rebecca Beyth, Scott Kaatz, Devin M. Mann, Jeremy B. Sussman, Deborah Korenstein, Connie Schardt, Avishek Nagi, Richard Sloane, and David A. Feldstein. 2014. Evidence-based risk communication: a systematic review. *Annals of internal medicine* 161: 270-280.

25. Nelson, Wendy, Valerie F. Reyna, Angela Fagerlin, Issac Lipkus, and Ellen Peters. 2008. Clinical implications of numeracy: Theory and practice. *Annals of Behavioral Medicine* 35: 261–274.

26. Brock, Dan W. 1987. Informed consent. In *Health care ethics: An introduction*. ed. Donald VanDeVeer and Tom Regan, 98-126. Philadelphia: Temple University Press.

27. Katz, Jay. 1984. *The silent world of doctor and patient*. New York: The Free Press.

28. *Natanson v. Kline*. 1960. 186 Kan. 393, 350 P.2d 1093.

29. Young, Robert. 2009. Informed consent and patient autonomy. In *A companion to bioethics* (2nd edition). ed. Helga Kuhse and Peter Singer, 530-540. Malden, MA & Oxford: Wiley-Blackwell.

30. Berg, Jessica. W., Paul S. Appelbaum, Charles W. Lidz, and Lisa S. Parker. 2001. *Informed consent: Legal theory and clinical practice* (2nd edition). New York: Oxford University Press.

31. Faden, Ruth. R., and Tom L. Beauchamp. 1986. *A history and theory of informed consent*. New York: Oxford University Press.

32. Beauchamp, Tom L. 2010. Autonomy and consent. In *The ethics of consent: Theory and practice.* ed. Franklin. G. Miller and Alan Wertheimer, 55–78. New York: Oxford University Press.

33. Sawicki, Nadia N. 2016. Modernized informed consent: Expanding the boundaries of materiality. *University of Illinois Law Review* 2016: 821-871.

34. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. 1979. The Belmont Report. *Federal Register* 44: 23192–23197.

35. President's Commission for the Study of Ethical Problems in Medicine and Biomedical and Behavioral Research. 1982. *Making health care decisions: The ethical and legal implications of informed consent in the patient-practitioner relationship*. Washington, D.C.: U.S. Government Printing Office.

36. Dworkin, Gerald. 1988. *The theory and practice of autonomy*. Cambridge: Cambridge University Press.

37. Archard, David. 2008. Informed consent: Autonomy and self-ownership. *Journal of Applied Philosophy* 25: 19-34.

38. Bromwich, Danielle, and Joseph Millum. 2015. Disclosure and consent to medical research participation. *Journal of Moral Philosophy* 12: 195-219.

39. Tännsjö, Torbjörn. 1999. *Coercive care: the ethics of choice in health and medicine*. London & New York: Routledge.

40. Jackson Jennifer. 2001. *Truth, trust and medicine*. London & New York: Routledge.

41. Eyal, Nir. 2014. Using informed consent to save trust. J*ournal of Medical Ethics* 40: 437-444.

42. Jamieson, Denise J., Caroline Costello, James Trussell, Susan D. Hillis, Polly A. Marchbanks, and Herbert B. Peterson. 2004. The risk of pregnancy after vasectomy. *Obstetrics & Gynecology* 103: 848-850.

43. Cook, T. M., D. Counsell, and J. A. W. Wildsmith. 2009. Major complications of central neuraxial block: Report on the Third National Audit Project of the Royal College of Anaesthetists. B*ritish Journal of Anaesthesia* 102: 179-190.

44. Han, Paul K. J. 2013. Conceptual, methodological, and ethical problems in communicating uncertainty in clinical evidence. *Medical Care Research and Review* 70: 14S-36S.

45. Hájek, Alan. 1996. "Mises Redux"—Redux: Fifteen arguments against finite frequentism. *Erkenntnis* 45: 209-227.

46. Hájek, Alan. 2009. Fifteen arguments against hypothetical frequentism. *Erkenntnis* 70: 211-235.

47. Good, I. J. 1971. 46656 varieties of Bayesians. *The American Statistician* 25: 62-63.

48. O'Neill, Onora. 2003. Some limits of informed consent. *Journal of Medical Ethics* 29: 4-7.

49. Clarke, Christopher E., Brooke Weberling McKeever, Avery Holton, and Graham N. Dixon. 2015. The influence of weight-of-evidence messages on (vaccine) attitudes: A sequential mediation model. *Journal of Health Communication* 20: 1302-1309.

50. Peirce, Charles S. 1878. The probability of induction. *Popular Science Monthly* 12: 705-718.

51. Keynes, John M. 1921. *A treatise on probability*. London: Macmillan and Company.

52. Good, I. J. 1985. Weight of evidence: A brief survey. *Bayesian Statistics* 2: 249-269.

53. Bradley, Richard. 2016. *Decision theory with a human face*. Cambridge: Cambridge University Press.

54. Weed, Douglas L. 2005. Weight of evidence: A review of concept and methods. *Risk Analysis* 25: 1545-1557.

55. Hoekstra, Rink, Richard D. Morey, Jeffrey N. Rouder, and Eric-Jan Wagenmakers. 2014. Robust misinterpretation of confidence intervals. *Psychonomic Bulletin & Review* 21: 1157-1164.

56. Gigerenzer, Gerd. 2004. Mindless statistics. *The Journal of Socio-Economics* 33: 587-606.

57. Steel, Daniel. 2003. A Bayesian way to make stopping rules matter. *Erkenntnis* 58: 213-227.

58. Zabell, S. L. 2005. *Symmetry and its discontents: essays on the history of inductive probability*. Cambridge: Cambridge University Press.

59. Murphy, Collen, and Paolo Gardoni. 2008. The acceptability and the tolerability of societal risks: A capabilities-based approach. *Science and Engineering Ethics* 14: 77-92.

60. Hanley, James A., and Gilles Thériault. 2000. Simpson's paradox in meta-analysis. *Epidemiology* 11: 613-614.

61. Nissen, Steven E., and Kathy Wolski. 2007. Effect of rosiglitazone on the risk of myocardial infarction and death from cardiovascular causes. *New England Journal of Medicine* 356: 2457-2471.

62. Charig, C. R., D. R. Webb, S. R. Payne, and J. E. A. Wickham. 1986. Comparison of treatment of renal calculi by open surgery, percutaneous nephrolithotomy, and extracorporeal shockwave lithotripsy. *British Medical Journal (Clinical Research Edition)* 292: 879-882.

63. Cartwright, Nancy. 1979. Causal laws and effective strategies. *Nous* 13: 419-437.

64. Woodward, James. 2003. *Making things happen*. Oxford: Oxford University Press.

65. Pearl, Judea. 2009. *Causality: Models, reasoning, and inference* (2nd edition). Cambridge University Press.

66. Jamieson, Denise J., Steven C. Kaufman, Caroline Costello, Susan D. Hillis, Polly A. Marchbanks, Herbert B. Peterson, and U.S. Collaborative Review of Sterilization Working Group. 2002. A comparison of women's regret after vasectomy versus tubal sterilization. *Obstetrics & Gynecology* 99: 1073-1079.

67. *Arato v. Avedon*. 1993. 5 Cal. 4th 1172, 23 Cal. Rptr. 2d 131, 858 P.2d 598.

68. Han, Paul K. J., Katherine Joekes, Glyn Elwyn, Kathleen M. Mazor, Richard Thomson, Philip Sedgwick, Judith Ibison, and John B. Wong. 2014. Development and evaluation of a risk communication curriculum for medical students. *Patient Education and Counseling* 94: 43-49.