Articles

Call of Duty at the Frontier of Research:

Normative Epistemology for High-Risk/High-Gain Studies of Deep Brain Stimulation

QA MERLIN BITTLINGER

Abstract: Research participants are entitled to many rights that may easily come into conflict. The most important ones are that researchers respect their autonomy as persons and act on the principles of beneficence, nonmaleficence, and justice. Since 2014, research subjects from numerous states in the United States of America also have a legal "right to try" that allows them, under certain circumstances, to receive experimental (i.e., preliminarily tested) interventions, including medical devices, before official approval from the United States Food and Drug Administration. In the context of experimental interventions, such as deep brain stimulation (DBS) for Alzheimer's disease, this article argues that research participants ought never to have a legal "right to try" without a corresponding "right to be sure." The latter refers to external epistemic justification construed in terms of *reliance on reliable evidence*. This article demonstrates that the mere complexity of intervention ensembles, as in the case of DBS for Alzheimer's disease which serves as a paradigm example, illustrate how unanswered and/or unasked open questions give rise to a "combinatorial explosion" of uncertainties that require epistemic responses that no single research team alone is likely able to provide. From this assessment, several epistemic asymmetrical relations between researchers and participants are developed. By elucidating these epistemic asymmetries, this article unravels the reasons why open science, transparent exhaustive data reporting, preregistration, and continued constant critical appraisal via pre- and postpublication peer review are not scientific virtues of moral excellence but rather ordinary obligations of the scientific work routine required to increase *reliability* and *strength of evidence*.

Keywords: neuroethics; deep brain stimulation; clinical translation; Alzheimer's disease; uncertainty; epistemic obligation; social epistemology; culpable ignorance

The translation of insight from the basic sciences into clinical practice is complex, tedious, and expensive. Alzheimer's disease is a paradigm case. Despite enormous research efforts, 99.6% of novel, potential therapeutic agents failed at some point during the clinical translation process between 1990 and 2014.[1] Key players such as Pfizer are already leaving the arena of pharmacological Alzheimer's disease research.[2] As such setbacks eventually concern the translation of findings from the brain sciences into clinical medicine, this raises particularly important questions for clinical neuroethics. The translational difficulties are tremendous and the complexity vexing. In consequence, professionals from all areas involved—be they scientists, funding agencies, science journalists, or research ethicists—may feel lost in the maze at some point.[3]

The author receives financial support from the Federal Ministry of Education and Research of Germany (01GP1621A) and works as research fellow for the international research consortium "Psychiatric Neurosurgery – Ethical, Legal, and Societal Issues" 2015; available at http://www.neuron-eranet.eu/en/630.php (last accessed 1 Apr 2018).

Cambridge Quarterly of Healthcare Ethics (2018), **0**, 1–13.

doi:10.1017/S0963180118000142

The argument put forward here makes the point that the epistemic duties of scientists play a crucial role in the translational process from bench to bedside. I attempt to show that doing exploratory research obliges researchers to take advantage of the open science framework. This requires fully transparent, systematic, and comprehensive publication of all information. Furthermore, I argue that beyond independent prepublication peer review, the scientific community has a duty to engage in meticulous postpublication review of any exploratory high-risk/high-gain research. I will do so by drawing from the example presented by the particular difficulties that arise in the context of "first-in-human" studies investigating deep brain stimulation (DBS) as novel, experimental intervention for new medical indications such as Alzheimer's disease. In this context, the scientific community as a whole has a *collective epistemic duty* to maximize the evidence, the strength of the evidence, and its reliability, and to use all means available to do so.

It is likely that this may strike some readers as trivial and odd at the same time. From a scientist's perspective, it may seem redundant to call for evidence as the basis of belief formation. Is not the endeavor of science defined as striving for knowledge? In addition, it may also appear to scientists strangely old fashioned to prescribe epistemic obligations at all, and worse, in the language of duties rather than of the virtues of scientific excellence. However, this may immediately change if we introduce some important conceptual distinctions.

In epistemology, the "ethics of (secular) belief" captures the idea that there are *epistemic* duties with regard to what one believes. According to the evidentialist stance, one ought to maintain and only to maintain beliefs for which one has sufficient reasons.[4] In this context, "sufficient reasons" means very roughly that, were the evidence that constitutes the reasons entirely true, then the belief would be beyond reasonable doubt when judged by the established standards and criteria for good, credible evidence. In contrast, the pragmatist stance maintains that one should hold those beliefs that maximize pragmatic values such as practical utility, and that we are sometimes positively obliged to form beliefs on insufficient evidence. Now, after having made this distinction, scientists may have very different inclinations to either side. Striving for practical utility seems as much a valid scientific value as striving for good evidence, and both pragmatist and evidentialist positions are abundant among scientists, research ethicists, and philosophers of science. Therefore, further distinctions are required to justify the evidentialist claims.

The thesis that I am going to defend is that scientists ought to prioritize the value of the *strength of evidence* for their beliefs over the value of the *utility of their beliefs* when performing high-risk/high gain clinical research involving invasive neurosurgical procedures. I defend this thesis on (normative) epistemic grounds that are largely independent from ethical concerns about potential hazards but, certainly, hazards would constitute additional reasons for concern.[5]

In stark contrast to scientists, research participants have no corresponding epistemic duty to prefer evidence in favor of utility with regard to their belief system, and this may be true for other stakeholders involved as well. Scientists play a particular role in translational research, and this role gives rise to the particular epistemic duty at stake. Against the pragmatist, I argue that scientists are never positively obliged to form beliefs about research on insufficient evidence. In the spirit of the evidentialist, I argue that researchers are often positively obliged to form beliefs on sufficient evidence. However, I drop the universal quantification

"always," as this makes the claim unnecessarily strict. Rather, there are specific conditions under which scientists are strictly obliged to seek further support to strengthen the available evidence.

Careful reflection about the social role of the scientist within clinical translation will sharpen scientists' understanding of their epistemic duties; that is, in technical philosophical terms, "practice-generated entitlements to expect."[6] The basic idea is that when scientists realize how much research participants depend epistemically on the integrity of critical appraisal through the science community, they *ought* to be motivated to accept the normative thesis. That is, to acknowledge the epistemic duty to apply all means and efforts to maximize the strength of evidence and to preempt the hazards and liabilities of *culpable ignorance*.

Case Study: Ms. Metis is Diagnosed with Alzheimer's Disease

For illustrative purposes, I will begin with casuistry to set out the general principles that rule the particular case of high-risk/high-gain research. I present the fictitious Ms. Metis, a 63-year-old woman recently diagnosed with probable Alzheimer's disease with mild cognitive deficits. Terrified by the diagnosis, Ms. Metis remembers her childhood and her experiences as a young woman when she learned from her grandmother what suffering from the relentless progression of neurodegeneration can mean. Ms. Metis is especially frightened because she does not want to burden her family. Her only son is a young father starting his own family and her husband is not in the best of health. Her physician explains that thus far no cure exists. There are only mildly effective pharmacological treatments addressing symptoms, and they are not disease modifying. However, there is a novel, innovative brain surgery that seems promising. If Ms. Metis accepts the risks associated, despite that this procedure is in a very early, experimental research phase, she may be eligible to participate in a "first-in-human" trial of DBS for Alzheimer's disease. Ms. Metis worked as an engineer before the diagnosis and strongly believes in the technological advances of science and medicine. Throughout her life she has taken pride in governing and disciplining herself through the use of reason. Her values are such that she despises idleness and inaction. She makes clear to her physician and her family that either she becomes part of that trial or she may choose suicide.[7]

The Ethical Dimension and Regulatory Oversight

Ms. Metis's situation is complex and uncertain because of multiple factors. There is not yet an available effective therapy. There is an expected disease burden of great magnitude and probability. There is an ambiguous signal evoked by the research context: caution because of the preliminary status. but perhaps also implicit signs of hope that research is conducted for a purpose, and a tacit assumption that even a very small chance is better than no chance at all. The epistemic challenge for Ms. Metis is to avoid severely underestimating the risks or overestimating the benefits or both (therapeutic misestimation).[8]

Avoiding therapeutic misestimation involves acknowledging that despite the lack of a medical remedy and the high disease burden, additional significant hazards are a possible outcome of "first-in-human" trials of interventions that have been only preliminarily tested. In truly innovative research, there is only scarce

and indirect prior knowledge to inform volunteers in early clinical trials about evidence-based information. Therefore, potential risks may be unknown and individual medical benefits may be uncertain. Because informed consent is a key cornerstone of human research, lack of evidence-based information raises ethical perils by conceptual necessity in "first-in-human" research. As a consequence, the question of regulatory oversight of such research has a conspicuous ethical dimension. Some ethical commentators hold that DBS research using new targets and novel indications should be regarded merely as clinical innovation,[9] whereas others criticize this proposal as "neurosurgical exceptionalism."[10] In consequence, the question of which ethical requirements need to be fulfilled for "first-in-human" DBS research is still open in clinical neuroethics. To cover a broader spectrum of these ethical dimensions, close attention to the context of DBS as a complex intervention ensemble is necessary. In particular, a close look at the specific hazards of "first-in-human" studies within the clinical translation process is needed.

Complex Intervention Ensembles

Biology is becoming an increasingly complex information science[11] and the translational process itself is an information maximization process.[12] Information, in turn, is best understood as any "difference that makes a difference"[13] and can be defined as any pattern of data that resolve uncertainty. Therefore, the goal of clinical translation is to minimize uncertainty.

At the end of the translation process there is not only a new product to sell, there is also uncertainty removed as to whether the new product or intervention is sufficiently effective and safe.[14] Because high-quality clinical research requires expensive resources, as well as voluntary and informed human research subject participation, any scientific study that does not reduce uncertainty about either efficacy, safety, or both, is futile and questionable from an ethical perspective.[15] As such, it seems reasonable to assume that clinical translation is also a process of ambiguity aversion. If a novel research question is assessed in a study to resolve genuine uncertainty, then the study should ideally be designed to give a definite answer to that question, not a vague one. It seems prima facie better to have a chain of very definite answers to smaller but feasible questions than a sequence of rather vague answers to the next great "breakthrough." This view complements the very clear distinction between exploratory research and confirmatory research,[16] as what counts as a definite answer may vary in each of these different contexts. In addition, it also requires acknowledging the complexity of any effect brought about by DBS. Even if it is perhaps not necessary to fully understand the mechanism of action of an intervention in "first-in-human" research, there is always a complex interaction of different factors that needs to be taken into account.

For DBS, the "intervention ensemble"[17] consists at least of the brain target, the stimulation parameter, and the presumed aberrant brain circuits or brain function as affected by the disease pathology. Therefore, the relevant research question is necessarily quite intricate. In the case of Ms. Metis, the relevant question with regard to what is in her best interest as asked from a detached observer's position is arguably the following: What is the expected probability that performing DBS in some explicitly defined patient cohort C with the specific stimulation parameter settings S applied to the particular brain target T will effect some physiological

change, which is with known probability P_1 positively correlated with a clinically relevant outcome O and which is with known probability P_2 not positively correlated with any significant hazards, burdens, or harms that outweigh O?[18]

Ideally, the best parameter setting S, the best brain target T and both the probabilities P_1 for benefits and P_2 for hazards are well studied in some relevantly similar cohort C of research participants. Typically, in "first-in-human" studies, the latter cohort is some animal model of the respective disease that is investigated. Given this ideal scenario, the uncertainty reduces to the question of whether the effects that hold for cohort C also hold for the cohort of Ms. Metis and her fellow participants. This picture of transferring an effect from one cohort to another is at the heart of the metaphorical content conveyed by "clinical translation."

However, there are several practical complications to this picture. For example, review of ethically relevant questions shows that, in the case of DBS for Alzheimer's disease, there has been an absence of empirical evidence from preclinical studies prior to "first-in-human studies."[19,20] Certainly, the translational gap from rodents to humans is huge,[21] but there are important yet more modest questions that could be validly examined by high-quality translational animal studies.[22] An example is the question of how different sets of stimulation parameters interact with the disease pathology of various distinct animal models, and which stimulation parameters seem optimal to reach beneficial effects without potential harmful side effects (therapeutic window). Also, the choice of brain target could have been explored antecedently in animal models of Alzheimer's disease. This information would have been valuable to inform "first-in-human" studies, which were explored using three competing brain targets.[23,24,25] However, directly relevant studies in mice came only after the "first-in-human" studies.[26]

The lack of prior probabilities derived from animal models has some very serious methodological repercussions that are difficult to spot for scientists themselves, and are all the more perplexing for research participants in the situation of Ms. Metis.

Exploratory Research

Lack of accurate and reliable information about prior probabilities of relevant factors is a threat in exploratory research.[27] One reason is that it turns a conditional probability into a joint probability. Because of Bayes' theorem, the conditional probability would be the probability that the desired outcome O may also be observed in another cohort C of research participants; that is. participants of a different kind or species, *given the known probabilities* for all other relevant factors such as the certainty about having selected a potentially suitable brain target, stimulation parameter, or disease stage.

If these factors cannot be stipulated as *given*, the threat of a combinatorial explosion of uncertainties arises. The joint probability of independent events is multiplicative and Bayes' theorem does not apply.

For illustration, we can simplify the scenario and limit the scope of relevant factors to the number 10 with binary outcomes (yes/no). Then, if DBS experts would be able to estimate the prior probabilities that for each of the 10 relevant factors, a suitable choice can be made with 80% certainty, then the joint probability is calculated as 0.8 to the power of 10; that is multiplying 0.8 10 times. This results in an estimated chance to observe the desired outcome of only 11%, and, respectively,

an 89% chance to fail.[28] However, if we raise the estimated probability via sound preclinical research to 0.95, then the expected 89% chance to fail turns into a 60% chance of success (0.95 to the power of 10). Although these numbers are fictitious, this clearly calls for rigorous confirmation of exploratory preclinical research using 95% or higher confidence intervals to exclude "false positives."[29] It is an important step in risk mitigation to reduce uncertainty about the estimated probabilities of relevant factors and, by doing so, to reduce the risk of "uncertainty blindness."

Risk and "Uncertainty Blindness"

According to the United States Food and Drug Administration (FDA) "risk is defined as the combination of the probability of occurrence of harm and the severity of that harm."[30] This definition is useful to describe the situation of Ms. Metis, but it is also defective in an important way. It only applies to known risks. That is, to risks that can be assessed on the basis of already available empirical frequencies. Such empirical frequencies can require sufficiently large numbers of relevant observations to reach statistical reliable estimates for the probability of *occurrences*. Therefore, it seems that risks according to the FDA, can, by definition, not be estimated for "first-in-human" research, in which no prior observations exist.

In the case of Ms. Metis, it is obviously not in her best interest to have an exact estimate of potential harms ex post facto. She probably wants to know the relevant risk *prior* to undergoing DBS for a novel medical indication such as Alzheimer's disease in a "first-in-human" research context. In addition, Ms. Metis may also want to take the degree of uncertainty of the estimation into account. According to this, risk can then be defined as the combination of three factors: (1) the estimated magnitude of potential harmful consequences, (2) the estimated probability of the occurrence of any of these potential harmful consequences, and (3) the strength of evidence that justifies arriving at these estimates in a scientific, valid, and reliable way.

Decision theory is an established scientific field that studies risk taking under uncertainty and informs complex and difficult questions ranging from financial investments to policymaking and disaster management, but is surprisingly seldom used in discussion of "first-in-human" studies of medical interventions.[31,32] Instead the discussion is focused primarily on the low probabilities of benefits, which treats the probabilities as fixed (e.g., a random variable, where the mathematically expected value is the measure of potential benefits or harm [payoff] and the standard deviation is the measure of uncertainty of that benefit or harm to occur). In this simplistic picture, there are many different strategies for making as rational a decision as possible, depending on risk preference such as risk aversion or risk taking. One option is the optimistic strategy of maximizing the likelihood that the *best possible* outcome is as good as possible (risk taker). Another option is to maximize the likelihood that the *worst possible* outcome is as good as possible (risk avoider). It may seem appealing to assume that both researcher and participants who decide to play an active part in "first-in-human" research tend to be risk takers, who want to stay optimistic despite any unfavorable circumstances (see Table 1). Accordingly, risk-averse researchers and participants would seem to avoid "first-in-human" trials right away. Whereas the latter is probably true, the former assumption about risk takers is elusive.

Table 1. Risk Matrix: A Non-Probability-Based Payoff Table Depicting Therapeutic Misestimation Caused by "Uncertainty Blindness"

		Research participation	
		Yes	No
Scenario	Best possible	Some benefit compared with TAU	TAU
	Most likely	Research burden + effect of TAU	TAU
	Worst possible	Some harm compared with TAU	TAU

If the estimated payoffs were accurate and reliable, risk takers would maximize the best possible scenario that is expected from research participation; that is, some benefit compared with treatment as usual (TAU), whereas risk-averse candidates would minimize the worst outcome (TAU and nonparticipation). As alluring as this picture is, it neglects the uncertainty that is associated with estimating the probability of best and worst possible outcomes or the most likely outcome. It entirely omits any estimate of the strength of evidence.

The whole account discussed so far has an important limitation. It entirely ignores the *unknown* probabilities of any relevant events that are typical for "first-in-human" research, and neglects the *uncertainty* of any estimation of relevant probabilities.[33] To remedy this, the strength of evidence for all relevant factors that influence the probability of worst or best possible scenarios needs to be taken into account. A further implication is that factors that drive the uncertainty of the probability of some potential risk factor need to be explicitly modeled in order to determine an adequate estimate for the risk. Factors that drive uncertainty are the lack of reliable empirical evidence and the plausible existence of unknown relevant factors that influence the probability of the desired study outcome. An example of the lack of reliable evidence is sparse preclinical data prior to "first-in-human" trials.[34] The existence of yet-unknown relevant factors is the more plausible the less is known about the mechanistic explanations of the indication; for example, Alzheimer's disease, and the less is known about the mechanism of action of the intervention. To be clear, a mechanistic understanding or a lack thereof does not per se affect the likelihood of a favorable or unfavorable outcome. But the lack of a mechanistic understanding may increase uncertainty as to whether there are any relevant factors that have been overlooked while planning a "first-in-human" study.

Epistemic Asymmetries

Given the outlined epistemically complex situation of Ms. Metis, there are important epistemic asymmetries between her situation as a potential research participant and the epistemic situation of the scientists involved in the study. As a first approximation, it seems intuitive to conceptualize these asymmetries using the concept of "trust." Accordingly, participants enter research with an attitude of trust that research is based on sound principles and that it is more or less safe and reliable. There is some empirical support showing that participants are greatly influenced by the information they receive on medical decisionmaking irrespective of whether that information is evidence based or not.[35] This can be interpreted as indicative of a certain basic reliance or trust.

To make a well-informed decision, potential research subjects do not only need to know, among other things, their subjective interests and risk preferences. A well-informed decision requires also some cognitive access by research subjects such as Ms. Metis to what would be in her best interest from a more objective point of view. That is, independent of how Ms. Metis *evaluates* the facts, she must have some understanding of *the facts*. However, to know what is in one's own best interest in an epistemically complex situation as outlined may go well beyond one's cognitive agency. This is hampered by external influence such as time pressure, urgency, or affective states such as fear and anxiety or even mild cognitive impairments. Moreover, the degree of desperation constitutes a vulnerability of people diagnosed with a severe medical condition for which no effective therapy is yet available. This is "vulnerability" in a nonstigmatizing sense, because it does not devalue the rights of these persons to make their own decisions, given legal capacity. Rather it entitles them to the additional right that healthcare providers and physicians abide to even higher degrees to the principles of beneficence and nonmaleficence; for example, by setting a "lowered standard of acceptable risk" as some argue.[36] But if the abovedescribed account is correct, this should rather include higher standards for the *strength of evidence* that support the probability of "acceptable risk."

Epistemic Dependence

The simpler a theory is, the better, but sometimes the simplest adequate theory is complex and all simple theories are false. As scientific theories tend to be complex, scientists and clinical researchers themselves acquire specialized epistemic skills and take advantage of a fine-grained professional division of epistemic labor. As such, scientific progress is a collective enterprise. Social practices such as peer review and the critical appraisal of the work of others is essential for the quality assurance and improves the reliability of scientific evidence. In consequence, reliable evidence is the process of a collective, interdependent, and social process. Similarly, Ms. Metis depends on the cognitive capacities of others; that is, she depends on insights and evidence provided not only by individual researchers working on a relevant research question but also on the scientific community at large that provides the epistemic "infrastructure" and quality control. Ideally, patients can trust that the information provided by researchers is honest, accurate, and reliable and that potential open questions and uncertainties are systematically addressed with due diligence. If this can be granted, it would allow potential candidates to go beyond what they can know or vindicate by their own cognitive resources independently and self-sufficiently (thesis of epistemic dependence).

Epistemic Duties

A key task of scientists is to retrieve relevant information on an investigational intervention and to minimize uncertainties for hypothetical benefits and potential hazards. Recent insights from social epistemology in analytic philosophy are valuable for analyzing the duties that arise in such situations. According to reliabilism, a (true) belief is justified if and only if the belief has the right kind of causal history; that is, if the belief formation is reliable to distinguish between truth and falsity.[37]

For Ms. Metis, this account has several implications. She can be justified in her true beliefs about trial participation, if only the information provided during informed consent process is reliable and is communicated in an intelligible and reliably truth-preserving way; that is, without manipulation, nudging, pressure, deception or any other form of "truth-modulation." Researchers have developed different tools to increase reliability of belief-formation. The most uncontroversial techniques include the scientific method itself but also cultural, institutional, and organizational structures that increase the reliability of evidence. Beyond that, the most important measure for quality assurance is independent peer review[38] and critical appraisal of published information in systematic reviews and meta-analyses on the basis of clear reporting criteria.[39] The open science movement is a further breakthrough in improving reliability of evidence unraveled by research practices.[40] In addition, metaresearch is forcefully transforming the scientific enterprise by scrutinizing research methods and "testing empirically their effectiveness at producing the most reliable evidence."[41]

Because of the high standards of transparency, open science may have some practical inconveniences. The discussion about open science is often framed about individual benefits that override these repercussions[42] and about which external incentives are best suited to complement intrinsic motivation and support scientific virtues. This focus on external incentives and individual benefits is pragmatic but theoretically misdirected. For example, in explorative research such as "first-in-human" DBS research, openness, reproducibility, and transparency are not only virtues of epistemic excellence of outstanding individuals, but rather ordinary epistemic duties that can be directly derived from the practical utility of increasing the reliability of evidence. Given the complexity and interdisciplinarity of the whole clinical translation process from bench to bedside, no individual scientist or team can rationally take full responsibility for minimizing all uncertainties about relevant open question. Nor does it seem feasible to systematically oversee at a given moment in time the potential relevance of published results for future research.

Peer-scientists therefore seem to have an epistemic right—if perhaps not a legal one— to comprehensive and unrestricted access to relevant information constituting scientific evidence. This can be seen as a corollary from realizing that pre- and postpublication peer review is partly constitutive of the reliability of the very evidence that builds the basis of scientific belief formation

It is always conceivable that there is some relevant evidence for a particular research question that investigators should look for as further support or negation of their hypothesis, even if they do not possess any evidence at the moment that there is such evidence.[43] The evidence that one expert holds given that particular person's contingent epistemic situation may not be exhaustive. The remedy is repeated critical thinking by many experts. Ideally, every expert who publishes an article on a hypothesis should benefit from the critical appraisal by its readers, who should be encouraged to raise open questions as part of a continued postpublication review. If performed for scientific ends (and not as power game among scientists), competing constant critical appraisal increases the reliability of evidence, and may efficiently spot yet-unaddressed uncertainties. An article that holds up to this constant public and critical appraisal may still be empirically disproven, but the original evidence provided by the article will stand strongly on theoretical grounds. In this way, the *reliability of the evidence provided* can be

increased, although it may not increase the evidence provided by the scientific article per se.

Problem Summary

High-risk/high-gain research at the frontier of science involves *by definition* unknown probabilities about clinically relevant factors. These relevant factors may sometimes bear on questions about life and death such as brain hemorrhage or serious brain infections. Exploratory research is an approach that severely impedes evidence-based decisionmaking, and involves ineliminable uncertainties about potential harms, potential absences of benefits, and potential futile trial participation. Therefore, *exploratory research* involving a high magnitude of burden such as brain surgery is hardly justifiable if uncertainty about relevant risks is high. The complexity of estimatng the joint probabilities of accumulating uncertain factors exacerbates evidence-decisionmaking about trial participation, and may lead to misestimation of risks.

In the perspective of a risk taker, it is often the case that potential participants can decide only once whether or not to participate in a high-risk/high gain trial. They have one attempt to hit a home run. In contrast, researchers examining novel investigational interventions; for example, for Alzheimer's disease, may have several opportunities to engage in high-risk/high-gain research. Trials are expected to frequently fail, and therapeutic interventions are urgently needed. In addition, it is sometimes sufficient to already detect some signal of efficacy only once ("proof of concept"), so that each participant examined is another opportunity to still be successful. It is noteworthy that participants take considerable risks in "first-in-human" studies involving neurosurgery, but—by the definition of "first-in-human"—cannot rationally expect medical benefits from the intervention based on any prior probability that is supported by strong directly relevant empirical evidence. Although the pragmatist stance; that is, to believe in the intervention nonetheless (optimistic bias), is viable for research subjects, researchers ought only to believe and communicate what is *strictly* supported by the evidence to avoid the "uncertainty blindness" of potential research subjects.

Strengthening the reliability of evidence or identifying yet-unknown uncertainties is a difficult scientific task that is a social endeavor of the scientific community as a whole and that is best reached by constant critical appraisal and reuse of comprehensively, transparently, and openly published data.

Strength and Limitations

Evidently, a valid argument is only as strong as its premises, and to resist one of its premises easily allows resistance to the whole argument. This conceded, I am confident that resisting the premises is not an easy task. Even opponents of the view that scientists are epistemically obliged to transparency, openness, sincerity, and honesty at least agree that scientists should not engage in "wishful speaking."[44] It is even harder to challenge the view that potential participants of high-risk/high-gain studies are not epistemically dependent on researchers. But if so, the argument is in good standing. The epistemic dependence gives rise to the epistemic obligation of the researcher "to communicate only those claims which are well established"[45] and to strictly stick with the evidentialist agenda to avoid epistemic

perils such as "uncertainty blindness." However, being an evidentialist entails being a researcher who strives to live up to not only the ideals and virtues of sincerity and honesty but also to those of transparency and openness.

A clear limitation of the presented account is its scope. I have only defended the existential statement that there are epistemic duties and started to preliminarily characterize the reasons for such potential duties. However, little has been said about the moral force or the enforceability of such epistemic duties.

Conclusion

Research participants are entitled to be informed about the *strength of evidence* available to estimate the probabilities of relevant factors potentially influencing the safety or efficacy of an intervention. The mere complexity of intervention ensembles such as DBS for Alzheimer's disease illustrates how yet-unidentified open questions give rise to a combinatorial explosion of uncertainties. To mitigate this risk demands collective social epistemic practices that no single research team is likely able to provide alone.

Open science, transparent and exhaustive data reporting, preregistration, and continued constant critical appraisal via pre- and postpublication peer review seem, therefore, not to be scientific virtues of moral excellence but rather ordinary obligations of the scientific work routine to increase *reliability* and *strength of evidence*.

Notes

1. Cummings JL, Morstorf T, Zhong K. Alzheimer's disease drug-development pipeline: few candidates, frequent failures. *Alzheimer's Research & Therapy* 2014;6(4):37.
2. Johnson CY. Why coming up with a drug for Alzheimer's is so devilishly hard. *Washington Post*, 2018; available at https://www.washingtonpost.com/news/wonk/wp/2018/01/12/why-coming-up-with-a-drug-for-alzheimers-is-so-devilishly-hard/?utm_term=.ff3af498c17b, (last accessed 13 Mar 2018).
3. The preparation of this present manuscript was performed under the influence of attending the conference "Lost in the Maze? Navigating evidence and ethics in translational neuroscience" in Hannover, Germany (February 14–16, 2018).
4. Clifford WK. The ethics of belief. In: Burger AJ, ed. *The Ethics of Belief*, Revised ed. Createspace Independent Publishing Platform; 1879/2008:9–40.
5. Emanuel EJ, Wendler D, Grady C. What makes clinical research ethical? *JAMA* 2000;283(20): 2701–11.
6. Goldberg SC. Should have known. *Synthese* 2017;194(8):2863–94.
7. Although fictitious, the case is not far-fetched. The first trials with patients with probable early-onset Alzheimer's disease have already been conducted and some people's motivation to participate was superimposed by the frightening thought that otherwise only suicide might remain as the last resort. See "What should I have done, then? Either I take myself to the water [German saying for suicide] or what do I know?" [Translation], September 27, 2011; available at https://www.ndr.de/fernsehen/sendungen/visite/schwerpunkte/Hirnstimulation-gegen-Alzheimer,visite6207.html (last accessed 3 Mar 2018).
8. Horng S, Grady C. Misunderstanding in clinical research: Distinguishing therapeutic misconception, therapeutic misestimation, & therapeutic optimism. *IRB: Ethics & Human Research* 2003;25(1):11–6.
9. Bell E, Leger P, Sankar T, Racine E. Deep brain stimulation as clinical innovation: An ethical and organizational framework to sustain deliberations about psychiatric deep brain stimulation. *Neurosurgery* 2016;79(1):3–10.
10. Fins JJ. Commentary: Deep brain stimulation as clinical innovation: an ethical and organizational framework to sustain deliberations about psychiatric deep brain stimulation. *Neurosurgery* 2016;79(1):11–3.

11. Insel TR. Join the disruptors of health science. *Nature* 2017;551(7678):23–6.
12. Kimmelman J, London AJ. The structure of clinical translation: Efficiency, information, and ethics. *Hastings Center Report* 2015;45(2):27–39.
13. Bateson G. *Steps to an Ecology of Mind: Collected Essays iAnthropology, Psychiatry, Evolution, and Epistemology*. Northvale, NJ, London: Jason Aronson Inc.; 1972.
14. Kimmelman J. A theoretical framework for early human studies: Uncertainty, intervention ensembles, and boundaries. *Trials* 2012;13:173.
15. See note 5, Emanuel 2000.
16. Kimmelman J, Mogil JS, Dirnagl U. Distinguishing between exploratory and confirmatory preclinical research will improve translation. *PLoS Biology* 2014;12(5):e1001863.
17. See note 14, Kimmelmann 2012.
18. Bittlinger M, Müller S. Opening the debate on deep brain stimulation for Alzheimer disease—A critical evaluation of rationale, shortcomings, and ethical justification. *BMC Medical Ethics* Accepted for publication. https://dx.doi.org/0.1186/s12910-018-0275-4
19. Viaña JNM, Vickers JC, Cook MJ, Gilbert F. Currents of memory: Recent progress, translational challenges, and ethical considerations in fornix deep brain stimulation trials for Alzheimer's disease. *Neurobiology of Aging* 2017;56:202–10.
20. Bittlinger M, Müller S. P 137 an ethical perspective on deep brain stimulation as an investigational treatment for Alzheimer's disease. *Clinical Neurophysiology* 2017;128(10):e395–6.
21. van der Worp HB, Howells DW, Sena ES, Porritt MJ, Rewell S, O'Collins V, et al. Can animal models of disease reliably inform human studies? *PLoS Medicine* 2010;7(3):e1000245.
22. Kimmelman J, Henderson V. Assessing risk/benefit for trials using preclinical evidence: a proposal. *Journal of Medical Ethics* 2016;42:50–3.
23. Scharre DW, Weichart E, Nielson D, Zhang J, Agrawal P, Sederberg PB, et al. Deep brain stimulation of frontal lobe networks to treat Alzheimer's disease. *Journal of Alzheimer's Disease* 2018;62:621–33.
24. Kuhn J, Hardenacke K, Lenartz D, Gruendler T, Ullsperger M, Bartsch C, et al. Deep brain stimulation of the nucleus basalis of Meynert in Alzheimer's dementia. *Molecular Psychiatry* 2015;20(3):353–60.
25. Lozano AM, Fosdick L, Chakravarty MM, Leoutsakos JM, Munro C, Oh E, et al. A Phase II study of fornix deep brain stimulation in mild Alzheimer's disease. *Journal of Alzheimer's disease* 2016;54(2):777–87.
26. See note 18, Viaña et al. 2017.
27. See note 16, Kimmelmann et al. 2014.
28. Talk Lozano A. Two failed trials of DBS for depression; what went wrong? Presented at the *17th Quadrennial Meeting of the World Society for Stereotactic and Functional Neurosurgery. Pre-Meeting Workshop. Surgery for Psychiatric Disorders*. Berlin, 2017.
29. See note 16, Kimmelmann, et al. 2014.
30. United States Food and Drug Administration. Guidance for industry: Q9 Quality risk management. In: *U.S. Department of Health and Human Services Center for Drug Evaluation and Research (CDER) Center for Biologics Evaluation and Research (CBER)* Washington, DC: United States Food and Drug Administration; 2006.
31. Eyal N. How to keep high-risk studies ethical: Classifying candidate solutions. *Journal of Medical Ethics* 2017;43(2):74–7.
32. Hug K, Johansson M. Challenges to informed consent in first-in-human trials involving novel treatments: A case study of Parkinson's disease. *Journal of Parkinson's Disease* 2017;7(4):695–702.
33. Goerlandt F, Reniers G. On the assessment of uncertainty in risk diagrams. *Safety Science* 2016;84:67–77.
34. See note 18, Viaña et al. 2017; note 20 van der Worp et al. 2010.
35. Wegwarth O, Wagner GG, Gigerenzer G. Can facts trump unconditional trust? Evidence-based information halves the influence of physicians' non-evidence-based cancer screening recommendations. *PLoS ONE* 2017;12(8):e0183024.
36. Beauchamp TL, Childress JF. *Principles of Biomedical Ethics*, Seventh ed. New York: Oxford University Press; 2013.
37. Goldman AI. What is justified belief? In: Pappas G, ed. *Justification and Knowledge*. Dordrecht: D. Reidel Publishing Company; 1979:1–23.
38. Wicherts JM. Peer review quality and transparency of the peer-review process in open access and subscription journals. *PLoS ONE* 2016;11(1):e0147913.

39. Moher D, Simera I, Schulz KF, Hoey J, Altman DG. Helping editors, peer reviewers and authors improve the clarity, completeness and transparency of reporting health research. *BMC Medicine* 2008;6(1):13.
40. Nosek BA, Alter G, Banks GC, Borsboom D, Bowman SD, Breckler SJ, et al. Promoting an open research culture. *Science* 2015;348(6242):1422–5.
41. Ioannidis JPA, Fanelli D, Dunne DD, Goodman SN. Meta-research: Evaluation and improvement of research methods and practices. *PLoS Biology* 2015;13(10):e1002264.
42. McKiernan EC, Bourne PE, Brown CT, Buck S, Kenall A, Lin J, et al. How open science helps researchers succeed. *eLife* 2016;5:e16800.
43. See note 6, Goldberg 2017.
44. John S. Epistemic trust and the ethics of science communication: Against transparency, openness, sincerity and honesty. *Social Epistemology* 2018;32(2):75–87.
45. See note 43, John 2018.

AUTHOR QUERIES

QA	The distinction between surnames can be ambiguous, therefore to ensure accurate tagging for indexing purposes online (eg for PubMed entries), please check that the highlighted surnames have been correctly identified, that all names are in the correct order and spelt correctly.
AQ1	Per your request, I am leaving this reference as accepted for publication. Please update in proof.