‘Why Do You Ask?’ Revisiting the Purpose of Eliciting the Public’s Moral Judgments About Emerging Technologies

Jared N. Smith, Anne Barnhill, Julian Savulescu, S. Matthew Liao, Matthew McCoy and Jennifer Blumenthal-Barby

ABSTRACT

It is increasingly common for bioethicists to consult with the public to solicit their judgments and attitudes about ethical questions and issues, especially ones that arise with new and emerging technologies. However, it is not always clear what the purpose of this engagement is or ought to be: do bioethicists seek the input of the public to help them arrive at a morally correct justified policy position, or do they seek this input to help them shape and frame their already-established moral position, or something else entirely? In this essay, we discuss four distinct possible functions of collecting moral judgments from the public: issue spotting, messaging for adherence and social stability, substantive moral guidance, and procedural fairness. For each function, we offer core examples from the literature before discussing the strengths and weaknesses attendant to each. We conclude with several preliminary questions bioethicists can ask themselves to clarify their own purpose for soliciting moral judgments from the public.

Introduction

It has become common practice in bioethics to consult with the public and other stakeholders to solicit their judgments and attitudes about practical ethical questions and issues. For example, as the Covid-19 pandemic proceeded, bioethicists sought to solicit and assess the attitudes of various stakeholders, especially community members, toward the ways scarce medical resources such as ICU beds, ventilators, and vaccines should be allocated (Bruno et al. 2021; Biddison et al. 2018; Schoch-Spana et al. 2020). Public engagement is not limited to crisis care applications; increasingly, bioethicists are soliciting the moral perspectives of the public about the development and use of emerging technologies. For example, the Moral Machine project used an interactive online platform to collect “a human perspective on decisions made by machine intelligence, such as self-driving cars,” (“Moral Machine” n.d.; Awad et al. 2018; see also Schneble and Shaw 2021, Takaguchi et al. 2022). In neuroethics, researchers often investigate or directly solicit the public’s perceptions about appropriate uses, and potential risks of, technologies like brain-computer interfaces (BCIs) and neuromodulation interventions (Cabrera et al. 2019; Sample et al. 2020; Cabrera, Courchesne, et al. 2021; Bluhm et al. 2023; Bluhm et al. 2024).

The ubiquity of this practice in the context of emerging technologies resurfaces a deeper discussion about the relationship between empirical bioethics research and normative theory: what is the *purpose* of soliciting moral judgments from the public in bioethics research? One thought is that it helps us get closer to a *moral truth* about how such moral issues should be handled. A very different purpose is *fairness*—the thought that soliciting people’s moral judgments helps bioethicists arrive at a fair view by consulting those who would be impacted by the relevant decision (Emanuel et al. 2020). On this view, the purpose is more procedural or

political. A third possibility is that knowing how the public thinks about certain moral issues will help us *better frame* our policies so that they are more acceptable to the general public. Still a fourth possibility involves ensuring bioethicists are identifying and addressing the *right* moral issues.

There is an extensive literature concerning the solicitation and use of key stakeholders' attitudes, particularly those of clinicians and patients. In this paper, we focus specifically on the purpose of soliciting *moral judgments* from *the public* regarding emerging technologies, offering a broad taxonomy of the functions of involving such judgments in moral/bioethical judgment and decision-making. We begin by distinguishing this taxonomy from previous attempts to organize the relationship between such empirical practices and normative prescriptions. To distinguish this taxonomy, we draw from our collective experience, including reviewing and conducting empirical bioethics research that involves surveying the public about their moral/normative judgments. We suggest four distinct possible functions of collecting moral judgments from the public: issue spotting, messaging for adherence and social stability, substantive moral guidance, and procedural fairness. We offer examples and discuss the strengths and weaknesses inherent in each approach. Finally, we conclude by suggesting several questions bioethicists can ask themselves as they engage in the consultation process with the public to better understand their rationale for engaging the public in the first place.

Background: the 'empirical turn' and moving from 'is' to 'ought'?

As a field, bioethics sprouted from theology and philosophy, transforming in the last several decades from a discipline concerned primarily with the application of moral theory and conceptual analysis to medical ethical dilemmas into one that is more evidence-driven and empirically perspicacious. Referred to as bioethics' 'empirical turn', this shift in methodology came at a time of tremendous interest in, and growth of, biomedical technologies (e.g., mechanical ventilation, dialysis, etc.). As Wangmo et al. point out, Borry, Schotsmans, and Dierickx characterize three forces at work in early bioethics which coalesced to form this turn: a growing dissonance between (1) discursive and conceptually-laden philosophical bioethics on the one hand and (2) the immediacy of the clinical context on the other hand, as well as (3) a growing receptivity to evidence-based medicine (Borry, Schotsmans, and Dierickx 2004, 2005; Wangmo et al. 2018). This empirical turn initiated theoretical and methodological debates in bioethics that continue to shape our research landscape, and prompts reflection on how bioethicists can formulate normative theory that is empirically-informed without also reducing the discipline to ethics by opinion poll (Salloch, Vollmann, and Schildmann 2014; Salloch and Ursin 2023). Some even suggest that bioethics is on the cusp of a similar 'digital turn' given the increasing use methods imported from data science (Schneider, Vayena, and Blasimme 2023; Salloch and Ursin 2023).

Bioethicists have tended to conceptualize the relationship between the empirical and normative in terms of the 'is/ought problem', or whether and how we can legitimately move from making empirical observations to making normative recommendations. The is/ought problem refers to a classical issue in normative

ethics, derived from the work of David Hume, that one cannot legitimately infer an ‘ought’ (a normative conclusion) from only ‘is’ premises (empirical facts). On at least one interpretation of Hume on this principle, “moral judgments cannot be deduced from statements of fact, not because the former are moral and the latter nonmoral, but because the former are evaluative and the latter are descriptive,” (De Vries and Gordijn 2009). In this way, the is/ought problem suggests that we cannot unproblematically form evaluative conclusions about what to *do* merely from descriptive facts about what *is*.1 What, then, is the purpose of seeking out such descriptive facts if our goal is to provide normative recommendations? Several bioethicists have attempted to address this question by articulating lists or hierarchies describing the relationship(s) between empirical research and normative recommendation. One prominent account, offered by Mildred Solomon, distinguishes ten distinct functions of empirical research in bioethics categorized under three broad headings: functions which *improve efficacy*, functions which *enrich moral justifications*, and those that *create or reveal new normative issues* (Solomon 2005). Solomon underscores that the various functions of empirical research primarily hold instrumental utility for bioethics, explicating how empirical research can be “deployed to serve bioethics’ [normative] goals,” (Solomon 2005).

Another influential example is the taxonomy articulated by Alexander Kon, who describes four successive categories of empirical research in bioethics that aim toward the ultimate goal of improving patient care (Kon 2009b) and thus with a focus on empirical work in the clinical context (Sugarman, Kass, and Faden 2009). Rather than four co-equal categories, Kon argues that these categories build upon one another in stepwise fashion, such that this hierarchy forms a definite process with distinct stages: the lay of the land, the ideal versus reality, improvement of care, and changing ethical norms. As Sugarman, Kass, and Faden note in their response to Kon’s proposed hierarchy, the focus is *positive* functions to the neglect of potential risks or pitfalls to incorporating empirical research into bioethics. They suggest there has been enough (or too much) ‘carving bioethics at the joints’ and not enough work determining when it is beneficial to incorporate empirical data into our normative recommendations, but especially when *not* to Sugarman, Kass, and Faden (2009). This contrasts with both Kon and Solomon’s approaches, who see it as a first step in a longer, ongoing discussion. Kon writes, “I will therefore end my response by quoting Professor Kirkwood, ‘[This] is but one good step, but more steps are needed before we can see where we are going,’ (Kirkwood 2009, 91). Let us hope that the conversation, debate, and inquiry continue,” (Kon 2009b quoting Kirkwood 2009). Here, we offer some next steps to continue this discussion.

Our goal in this article is neither to comprehensively summarize nor reconcile various taxonomies regarding the is/ought problem in bioethics with one another. Some of the functions we identify in this paper overlap with some of those identified by scholars like Solomon and Kon, while others we present do not appear on these hierarchies or lists. In truth, there is such a wealth of work in this area of bioethics that one could scarcely do the literature half the justice it deserves (if one is amenable to desert-based justifications). As Kon notes about his own endeavor:

> Such a discussion would likely take more space than the entire paper as written, therefore a complete and nuanced account [of foregoing work on empirical research and bioethics] would have engulfed the work [and] I would venture that my brief presentation of other scholars' work would have been met with criticisms that my simplistic renditions failed to adequately represent the sophisticated and nuanced schemed developed by these noted academicians. (Kon 2009a)

We are not the first bioethicists to consider the interplay between the empirical and the normative and, while we do discuss some previous taxonomies concerning the is/ought problem in bioethics where relevant, we envision ourselves stepping back from this discussion and engaging in a distinct but related inquiry: given the normative nature of the field, *why* do bioethicists look to the 'is' to inform our 'ought's' in the first place? What *purposes* do bioethicists have for engaging in this practice and how do these different purposes manifest? Importantly, we think this must include discussion of some of the potential pitfalls bioethicists may encounter as they consider their motivations for empirical research so that they can anticipate and address them, which is presently underrepresented aspect of the literature. Finally, we raise these questions in the context of bioethicists' increasing engagement of public opinion on moral issues in the context of new and emerging technologies such as artificial intelligence and neurotechnologies.

Much of empirical bioethics research relies on "public engagement" in the sense that it involves researchers soliciting a given public's views about some ethically relevant dimension of medicine, research, or public health. However, using the common term "public engagement" to refer to the methods employed across this body of work obscures the range of forms such research can take. Across the breadth of empirical bioethics, we find researchers engaging different publics for different reasons to answer different questions. We will limit our present inquiry in three critical ways. First, we restrict the scope of our analysis to research on emerging technologies. Second, we focus on research wherein bioethicists explicitly solicit the public's moral judgments or normative determinations as opposed to their knowledge or preferences. Third, we focus on research in which the engaged group comprises laypeople or members of the general public, as opposed to patients with lived experienced of a particular disease or stakeholders (e.g., clinicians working with a particular patient population, industry representatives) with a vested interest in some particular policy decision. Our operative definition of the public thus follows McCoy et al. who distinguish between patient and public involvement in research. We understand 'patients' to be persons involved in research or policy in virtue of their experience with a medical condition and/or treatments for it, whereas the 'public' are individuals involved in virtue of being a member of a relevant population (McCoy et al. 2019).

The context of public engagement regarding emerging technologies is an ideal one in which to revisit the relationship between the empirical and normative because as technological advances create novel ethical questions and possibilities, bioethicists have rushed to query the public about their views on every issue from what policies should guide self-driving car development and

implementation to the permissibility of using neurostimulation devices (such as deep brain stimulation or transcranial magnetic stimulation) for treatment and/or enhancement (Wagner et al. 2018; Zuk et al. 2018; Sample et al. 2020; Cabrera, Courchesne, et al. 2021; Bluhm et al. 2021; MacDuffie, Ransom, and Klein 2022; Sattler and Pietralla 2022). We argue that most instances of soliciting moral judgments from the public can be classified into one of four categories which are delineated chiefly by the purposes for querying the public. In what follows, we outline and discuss these categories in the context of bioethics of emerging technology, introducing the categories roughly in order of how comparatively controversial their purpose may be. While these are distinct categories, there will likely be some overlap. The first such category is what we term 'moral issue spotting.'

Moral issue spotting

Perhaps the least controversial function of soliciting public input about bioethical issues is 'issue spotting'. Broadly, the purpose of this manifestation is to direct the bioethics research process, setting the agenda for what moral questions and issues should take priority for investigation, what values are centralized in the research process, and so on. Sometimes it can manifest as bioethicists seeking out an 'on the ground' perspective about what research is not being done and what ethical questions are not being asked. On Kon's hierarchy mentioned earlier, moral issue spotting closely resembles what he calls 'the lay of the land' stage of research wherein bioethicists "seek to define current practices, opinions, beliefs, or other aspects that may be considered the status quo [...] This work can set the stage for further research and may help direct efforts at improving care," (Kon 2009b), a sort of fact-finding mission to collect information about what a stakeholder group might want, or believe, about a particular intervention or technology. Similarly, moral issue spotting broadly resembles most of the functions Solomon presents under the header of improving the efficacy of bioethics, especially documenting the gaps between theory and practice, highlighting the nature of moral reflection at the individual level, and accurately describing the contexts that most inform moral action. The function of informing interested parties with information about themselves or their constituents could likely fall within moral issue spotting, as could the function of identifying new or unforeseen moral issues (provided these are discovered *via* querying the public) (Solomon 2005).

In the context of emerging technologies, issue spotting is common in neuroethics, as one example. This is an outgrowth from a more general push for public engagement in neuroscience and neuroethics (Leshner 2005; Illes et al. 2005; Boyd and Sugarman 2022). There it has often taken the form of interviewing relevant stakeholders or the public to solicit their views about what the most pressing concerns are from their perspective. The goal of moral issue spotting is to find gaps in the literature that bioethicists and researchers have missed or which they lack exposure to. In the field of neuroethics, public engagement not only raises awareness of critical neuroscientific developments but also reveals "significant ethical and legal ramifications that require discussion not only among scientists but with all stakeholders," (Illes et al. 2005). One reason issue spotting is important is that sometimes the interests and

concerns of bioethicists differ markedly from those of the public who will be impacted such that the theoretical literature becomes skewed toward what bioethicists find most pressing or interesting and away from what the public identifies as missing from the conversation.

One such area where these differences occur is in in the context of research on different forms of *neuromodulation*, such as transcranial magnetic stimulation (TMS), electroconvulsive therapy (ECT), and deep brain stimulation (DBS). Neuromodulation consists in delivering electrical stimulation to targeted areas of the brain to induce functional changes as a way to treat the symptoms of certain disorders (such as Parkinson's disease, major depressive disorder, etc.). Neuromodulation has largely been found to be effective treatment pathway for some psychiatric conditions and most neuroethics studies have sought the attitudes of key stakeholders (patients with such conditions and their caregivers, clinicians who treat those conditions, etc.) over than the those of the public. Yet, as Cabrera et al. indicate in their 2021 study, "attitudes of the general public toward mental illness and novel psychiatric interventions can also have a powerful impact on the development and utilization of new treatments, and play an important role in influencing health policy," (Cabrera, Courchesne, et al. 2021). Issue spotting among the public concerning neuromodulation has the potential to inform an ongoing disagreement in neuroethics surrounding changes to the self from neuromodulation.

Neuroethicists disagree concerning whether and to what degree there are 'changes-to-self ' following DBS treatment, where a pacemaker-like device provides electrical stimulation to targeted areas of the brain with surgically implanted leads. DBS is currently being researched as a potential form of treatment for certain neurological and psychiatric disorders such as Parkinson's disease, dystonia, and OCD. In the 2000s, several studies identified adjustment difficulties in patients with Parkinson's Disease receiving DBS (Schüpbach et al. 2006; Agid et al. 2006; Houeto et al. 2002). These difficulties became cast by some bioethicists as concerns about 'changes to the self ' in various ways, from changes in personality or behavior to changes to authentic agency and identity (Pugh 2020). Complicating this discussion in the ethics literature was a lack of standardization about the meaning and use of key terms and how they relate to one another.

In their analysis of the literature on 'changes to the self ' and identity following DBS, Gilbert, *via*ña and Ineichen point out a striking disparity between the focus of theoretical bioethics papers about DBS as compared to primary study or clinical papers (Gilbert, Viaña, and Ineichen 2021; Pugh et al. 2021). They point to a 'bubble' of bioethics theory papers about the possibility or nature of these changes and suggest it is insufficiently grounded in primary studies working directly with stakeholders to solicit their actual and lived ethical concerns. More worrisome is that these original papers have been cited to support bioethics conclusions which the authors of those clinical papers themselves are wary of endorsing: e.g. Schüpbach et al. and Agid et al. suggest the changes patients experienced were the result of difficulty adjusting to improvements in health and capacity stemming from treatment (Agid et al. 2006; Schüpbach et al. 2006). While some of these adjustments do relate to self-conception, others such as

marital or relationship difficulties as well as issues in their social and professional lives reflect a complex concern not necessarily reducible to a change to the self.

In a 2018 piece, Cabrera et al. compare and analyze themes raised by researchers in academic neuroethics literature with those raised by members of the public commenting on relevant news articles (e.g., about neuromodulation or brain computer interfaces). They find that researchers emphasized philosophical concerns, such as impacts to autonomy or identity, which were not common in public commentary (noting that most neurosurgery articles found concerned deep brain stimulation specifically). In the neuroethics literature, concerns about authenticity were raised in relation to both neuromodulation and pharmacological treatment, and which were almost nonexistent among public comments. Instead, the public raised issues relating to social and personal responsibility for treatment (Cabrera et al. 2018). One possibility they consider is that the myriad terms bioethicists and philosophers use to reference 'changes to self ' are not widely known or used by members of the public and so we do not see parallel concerns. Alternately, we might think that some moral philosophers and bioethicists became focused on the issue of changes to self/identity, topics which were less of a dominant concern for the public, those considering or experiencing DBS, caring for someone who is, or a clinician who is likely to interact with those considering it— they had other ethical concerns that they would have liked bioethicists to analyze and provide guidance on.

While this example may portray issue spotting as always beneficial, there is a tension between the positives of issue spotting and its potential drawbacks as a method of public engagement. On the one hand, bioethicists don't want to become divorced from real-world issues. Given the way research and development of medical treatment closely attend one another, bioethicists have a strong influence as compared to the public and this warrants perspicacity. At the same time, members of the public (even those with direct stake in the given issue) are not trained to conceptualize questions or pick apart lines of thought and so their guidance in what the relevant ethical issues are has limits. Furthermore, one could argue that just because the public does not see the issue on the ground does not mean it is not a valid or pressing ethical issue.

Messaging for adherence and social stability

A second distinct function of consulting the public about bioethical issues is that soliciting their moral views enables bioethicists to frame policies and recommendations to address the moral concerns and views of the public. Bioethicists figure out the most philosophically defensible position and then frame it in ways the public will likely buy into (given what we know about their moral views). Broader appeal might not only make it more likely that the public will abide by a policy (adherence), but that widespread adherence will result in increased social stability (in the form of reliable and continued uptake of, or respect for, established policies). One example of this approach to consulting the public can be seen in the potential use of 'nudges' as well as public perception about the permissibility of nudges in decision making. A 'nudge', according to Thaler and Sunstein who coined the term, is a decision about how choices are presented which reliably influences people's behavior without either limiting the choices or changing their relative incentives (Thaler and Sunstein 2008).

In “Do People Like Nudges? Empirical Findings” from *The Ethics of Influence,* Sunstein discusses the results of a national study that tested the ethical permissibility of 34 nudges with 563 Americans. He found that when people believe that the end goal of the nudge is legitimate and important, they favor it, morally. But people reject nudges when the ends involve political or religious favoritism, when they are inconsistent with the interests or values of most people, and when people’s decisional biases are used against them (Sunstein 2016, 116–58). Suppose that bioethicists used these findings to build certain nudges into policies. For example, electroconvulsive therapy (ECT) has garnered a negative public perception since its initial forms and association with serious and debilitating side-effects, with many likening it to a lobotomy or lesion therapy (Kring, Bergholt, and Midtgaard 2018; Griffiths and O’Neill-Kerr 2019; Cabrera, Courchesne, et al. 2021). Suppose that bioethicists helping craft state policy concerning treatment over objection have already decided that ECT cannot be ethically administered to an incapacitated patient. They might draw on empirical research demonstrating the public’s general distrust of the efficacy or safety of neuromodulation to help frame their policy which makes it against hospital policy to administer neuromodulation therapy to a patient who lacks decisional capacity.

Evidence of this dynamic of public engagement is identified by Stilgoe and Cohen (2021) concerning the dominant model of public engagement concerning driverless vehicles, drawn from the Driverless Futures project and other stakeholder engagement endeavors. They describe a framework dominating driverless vehicle debates “in which the challenge is seen as first public *awareness*, then public *acceptance*, and finally public *adoption* of new technology,” (Stilgoe and Cohen 2021, 852). This framework treats the public as future users of the technology who must be informed about driverless vehicles so as to prepare them to adopt such vehicles and accept their use in their community. Stilgoe and Cohen, who challenge this framing, demonstrate how pervasive this attitude is among various private and public actors, amounting to conceiving of the public as a stumbling block for the inevitable implementation of driverless vehicles.

There is another manifestation of the messaging function which we might call coalition building. Sometimes consultation with the public about their moral views is not necessarily for the sake of understanding how to frame or present a *particular* policy in the most successful way, but for the sake of making sure that the public is consulted, which itself may make certain policies or decisions more likely to be accepted. For example, a new U.S. AI Bill of Rights outlines a blueprint for a handful of principles or “rights” that humans should have in relation to AI (e.g., when used as part of their health care). The white paper begins by noting these principles were formed after consulting the public. Now, we know little about *how* exactly the public’s moral views informed the rights that the Blueprint articulates, but the very fact that the public *was* consulted might make the Bill more likely to be accepted (“Blueprint for an AI Bill of Rights - OSTP,” n.d.). While the rights articulated in the Bill may have been developed largely by policymakers and ethicists, they may also be looking to recruit and build support so that the work can be presented under the banner of public consensus with some valuable status (such as those who will need to abide by laws concerning human-AI interaction). This can lead to more public uptake of

philosophically defensible positions. This potential for increased uptake signals an important consideration when engaging in this approach: adherence and social stability. As bioethicists, we must often account for both the content of arguments and policies and the need for public adherence to the policy with an eye toward long-term stability (in terms of reliable uptake) and acceptance. Ideally, we seek to put forward policies that people will comply with but that are neither burdensome nor demand coercive mechanisms for uptake or application.

There is a danger that this process could turn into tokenism, where the consultation is less than sincere. Fairlee warns of this danger, describing "…the process of seeking opinion on options before decisions are reached. Those opinions may not affect the ultimate outcome and in such instances can feel like tokenism and result in disengagement," (Fairlie 2015, 267). This practice, at its worst, could give the appearance that the work (and content) has emerged from interaction with the community such that it is reflective of their moral views and voices when this is not the case. This problem can also arise when there is a mismatch between the expectations of bioethicists engaging with the public and the expectations of members of the public themselves. If the public thinks they are being approached to help in the moral decision-making process, but researchers are approaching them primarily to maximize public buy-in on what has already been decided, people are more likely to feel slighted or that their involvement amounts to tokenization.

Substantive moral guidance

Another possible function of soliciting moral judgements from the public is to seek substantive moral guidance. In this approach, public moral judgments are used to help *shape* bioethical views (to some extent) because they help us get closer the most morally justified view (some might say the moral truth). As an example of consultation for substantive moral guidance, consider Savulescu et al.'s "Collective Reflective Equilibrium in Practice (CREP) and controversial novel technologies." Seeking substantive moral guidance from the public *via* a method like reflective equilibrium has been a component of bioethics since the introduction of Principlism. For example, in a recent piece reflecting on the 40th anniversary of *Principles of Biomedical Ethics*, Beauchamp and Childress write:

> We have argued in addition that these principles can and should be regarded as deriving historically and philosophically from the common morality, that is, the universal morality to which all morally committed persons subscribe […] In our common-morality account we justify our moral claims by calling on both the common morality and the method of reflective equilibrium. (Beauchamp and Childress 2019, 11; Beauchamp and Childress 2001, 1–24)

Since Rawls' influential work on Justice, reflective equilibrium has become a well-known methodology for moral decision-making. It involves back and forth negotiation and testing between our intuitive judgments about individual cases or issues on the one hand and moral principles on the other. Along the way, we refine both our moral principles as well as our intuitive judgments—the two inform each other. In bioethics, this is typically done by an individual bioethicist or clinical ethicist. But recently, this methodology has expanded to include the moral judgments of groups of stakeholders and the public. This is referred to as

collective reflective equilibrium (CREP): "…CREP provides a principled way of using some public preferences as an input for forming concrete policy that is practically implementable, politically legitimate, and ethically defensible, while justifiably disregarding other public preferences," (Savulescu, Gyngell, and Kahane 2021, 653).

In the Moral Machine Experiment mentioned earlier, Awad et al. apply CREP to the question of what moral preferences or judgments to assign to autonomous vehicles (AV) (Awad et al. 2018). Available data from the Moral Machines project shows shared moral preferences for the governance of the behavior of autonomous vehicles (AV) across cultures. Awad et al. think bioethicists should prioritize the public's three strongest moral preferences. Ranked by strength, some of the top preferences expressed are: (1) prioritizing human lives over animal, (2) prioritizing the many over the few, (3) prioritizing the young over the elderly, and (4) prioritizing those following the law versus those breaking it.

One looming concern about the practice of soliciting the public's moral preferences is that some preferences appear reasonable to guide our theorizing while others appear wrong to include on principle; for AVs this might look like prioritizing the lives of the physically fit over the disabled or obese, or prioritizing law-abiding citizens over criminals, and so on. Savulescu et al. recognize this when they indicate that "To make ethical decisions a matter of [mere] referendum is to eschew ethical expertise and professional responsibility," (Savulescu, Gyngell, and Kahane 2021, 655). If we rely on the aggregate preferences of polled respondents to derive principles for how AVs should be programmed, we must ask what role ethicists play (if any) in moderating these preferences against existing, broad ethical principles. At the same time, we also cannot ignore public preferences when making morally relevant policy, especially when they constitute one group that will be impacted by that policy.

How can CREP and similar approaches mitigate this concern over ethical moderation of these preferences while still respecting the input of the public? Savulescu et al. argue that the public's moral intuitions about particular cases can *justify moral views* if the intuitions or preferences meet particular requirements. First, responses (moral judgments) from the public must be screened to remove unreliable data such as that produced by framing bias. Second, soliciting judgments should prioritize informed respondents **(**those who take their time and consider the moral questions carefully) such that these are their *considered* judgments rather than snap choices. Third, we must take care that the body of solicited respondents is representative of the whole affected population rather than focusing on just one group or region.2 As a result of these requirements, the public's - filtered – moral intuitions can be balanced and negotiated with ethical theories and principles as in traditional reflective equilibrium:

> … in CREP, competent judges [like bioethicists or policy makers] are not using their own intuitions – they are applying the reflective equilibrium to the 'will or preferences of the people'. It is their responsibility to apply ethical theory, principles, concepts and justification to the preferences of people affected by their policy. (Savulescu, Gyngell, and Kahane 2021, 658)

In this way, CREP represents more than simply polling the public to see what the 'right' view is, but rather an incorporation of *what matters to them* into our moral deliberation and decision-making. The judgements of the public are one input (raw data) into our reflective equilibrium but there is also clear role for bioethicists trained in normative analysis and moral theory (as competent judges).

One potential concern about deriving substantive moral guidance or seeking moral truth from large-scale public polling about moral dilemmas (even assuming a process of reflective equilibrium) is that it may not represent the *best* (or most urgent) framing of the issue. Formulating ethical problems in certain ways is convenient for philosophers and bioethicists but may fail to serve the public if it (1) misdirects the public's attention, (2) obscures or deprioritizes more pressing ethical problems, or (3) artificially simplifies the ethical issue in ways that have substantive moral implications. In the case of AVs, some might worry that modeling the concern along the lines of a "trolley problem" as above demonstrates these framing issues. By posing the ethical issue that requiring public judgments as a dilemma about what principle(s) AVs should use in the event of unavoidable collisions, we presume the issue of 'who-to-save' is the most ethically pressing issue surrounding AVs. Instead, perhaps the most urgent issue is designing AVs such that they are more attractive for use, i.e., they cause fewer accidents than human drivers. Addressing *this* formulation of the issue may reveal that there are options that exist outside of the traditional framing of the trolley problem.

Members of the public with personal experience related to a particular bioethical issue have particular moral knowledge to contribute (Nelson et al. 2023; Pratt 2021; Hubbard, Kidd, and Donaghy 2008). Consider the following:

> Of particular relevance to bioethicists is moral knowledge, which some have suggested is uniquely tied to personal experience. McGrath (2019) argues that experience can help us grasp moral concepts, confirm or disconfirm our moral views, and trigger moral knowledge [...] This suggests that, in a given domain, some people may make more reliable moral judgments than others—not because they are better reasoners, but because they have the relevant experience from which to reason, which others lack (DePaul 1993). (Nelson et al. 2023, 5)

Some research domains are already engaged in such inclusion of those with lived experience. It is increasingly common in philosophy of psychiatry to involve service-users/survivors (those with personal experience of psychiatric conditions and treatment) as epistemic and moral authorities (Rose and Kalathil 2019; Russo 2012; Sweeney and Beresford 2020; Taylor et al. 2022). Further still, some calls for inclusion articulate the need for *researchers* with lived experience to be involved in the process and trained in the discipline (Jones et al. 2021; Kimhy, Jones, and Dixon 2022).

But perhaps the most significant deficit of the substantive moral guidance usage of the public's views is that it is unclear how much weight ought to be given to the public's input versus that of the experts in practice and in theory. The lurking concern here is of reinforcing a technocracy whereby those privileged with

expertise and access have a disproportional influence on research and policy. For example, in “The end of Personhood” Blumenthal-Barby argues that if the public finds the concepts or policy bioethicists generate “incomprehensible, off-putting, or in direct conflict with their own lived and moral experiences” then bioethicists should reconsider employing such concepts (Blumenthal-Barby 2024). The concern is in developing or reinforcing a technocracy within bioethics research and policy, with Evans concluding that “it is fine for bioethicists to be out of tune with the citizens in cultural bioethics. To avoid a static culture, someone needs to be promoting unpopular ideas. However, when it comes to public policy bioethics, when the writer has more influence than an average citizen does, the values of the public should be the driver,” (Evans 2012, 2024).

Procedural fairness or political proceduralism

A fourth possible function of soliciting moral judgments from the public is to satisfy a procedural fairness requirement to get input or participation into a decision or policy process. Unlike other approaches which seek out substantive moral guidance or a roadmap for marketing established bioethical policies, procedural requirements are about ensuring that the public has the proper kind of voice in decisions. One way of understanding this purpose for soliciting moral views from the public is to fulfill a process requirement (rather than an outcome requirement). That is, if a fair process is followed then a fair outcome must result. Norman Daniels suggests that when we cannot identify substantive criteria for fair deliberation, a fair process may be the best that we can hope for (Daniels 2000; Daniels and Sabin 2008).

These procedural fairness requirements may vary – that is, the kind of input required may vary – depending upon the issue at hand, as well as the actor in question. In particular, governments may face stricter requirements to get public input than private institutions. Political theorists have argued that public input and participation in pandemic policymaking, for example, advanced or was required by various moral and political values. People’s participation in political life and in decisions that affect them is claimed to be a democratic ideal, a democratic right, and a recognized human right (Norheim et al. 2021; Smith et al. 2021). Norheim et al. (2021) argue that fairness and procedural justice require that decisions affecting people’s interests are made “in a way that gives equal consideration to everyone’s interests and takes account of their perspectives,” (Norheim et al. 2021, 11). Other theorists argue that it’s the fact of disagreement – including moral disagreement in addition to disagreement about political decisions – that creates a normative need for public input; giving each person a say is how political and moral disagreements can be respectfully and fairly resolved (“How Should Governments Make COVID-19 Policy? - The Philosophers’ Magazine,” n.d.; Thompson, Cochrane, and Hopma 2020). Political theorists also advocate for different forms of input from the public. For example, some argue that members of the public should engage in formal deliberation about specific policy issues before policymaking (for example, through citizens’ assemblies or similar bodies) (“Oregon Citizen Assembly on COVID19 Recovery - Report from Oregon’s Kitchen Table Survey” 2020), whereas others advocate for various

forms of “informal deliberation,” such as conversations at community events and in online forums, opinion articles and organized demonstrations (Scheinerman and McCoy 2021).

As this short discussion has hopefully shown, political theorists have argued for different kinds of procedural fairness requirements on government decision-making and have provided distinct justifications or normative foundations for these requirements. It should not be assumed that the same requirements apply to private actors and policies that will only impact private institutions. Governments may have obligations to citizens which have no clear corollary within private institutions.

One example of political proceduralism in bioethics is offered by Benjamin Gregg, who contends that bioethics, as an endeavor, is political in nature rather than substantively moral: “Likely no answer to a bioethical question can claim universal validity inasmuch as bioethical questions are matters of normative preference,” (Gregg 2022, 2). Gregg’s view is that bioethical regulations or policies cannot be morally true given widespread disagreement about bioethical questions and the contingent, perspectival nature of moral preferences. Instead, what we can strive for is procedural legitimization. Thus, vindicating an argument or policy is not about locating the objectively correct answer but rather securing a legitimate policy by following proper, democratic procedures. This proceduralism is a process-based justification. That is, assuming the process is followed correctly then whatever outcome obtains must be correct (politically speaking, not morally speaking). It should be noted that while Gregg denies that there are objectively correct answers to bioethics questions, we needn’t embrace that view in order to motivate a procedural requirement to obtain public input.

One benefit to this kind of proceduralist approach is that the dominant perspective of one set of values is not able to control the policymaking discussion without engaging in rational debate and disagreement with appeals to reason. More importantly, this reason and debate is such that it must engage with the views of a wide range of people, not all of whom will have background knowledge relevant to the policy question at hand. This means the arguments and evidence provided must be constructed by researchers and bioethicists such that laypeople understand the ethical questions involved as well as the impact of the proposed policies. This can be a two-edged sword, however: experts may struggle to persuade citizens if there are too many options available, and if citizens lack the foundational knowledge that allows experts to make informed judgments and policies. To combat this, Gregg proposes combining two different types of proceduralism: expert committees paired with deliberative democracy such that the committees provide detailed background considerations and arguments which the public then weighs to develop their own viewpoints.

Another mark in favor of reaching bioethical conclusions through procedures that incorporate members of the public – such as Gregg’s proceduralism as well as other procedures – may be that this approach gains legitimacy through a connection to democratic ideals and procedures. If we are comfortable with using various kinds of democratic processes for making decisions about regulation and public policy outside the sphere of bioethics then, the proceduralist may argue, we should be comfortable uses processes that incorporate public opinion in order to answer bioethical questions about the ethics of regulation and public policy.

However, this may be a contentious claim, since democratic governance by a government is a different kind of activity than bioethicists reaching ethical conclusions about regulation and public policy. Do arguments that democratic processes are a legitimate way for a government to make policy decisions tell us anything about the legitimacy of incorporating the public's moral judgments into bioethicists' activity of reaching ethical conclusions about policy?

Plausibly there is something of moral import about giving the public a voice in that it respects them and provides recognition of their standing. At the same time, it matters quite a bit how the voice of the public is used. We should distinguish between public engagement aimed at improving perceived legitimacy of a bioethics judgment or conclusion, and public engagement aimed at improving the normative legitimacy of a judgment or conclusion. The former borders on using the public for 'cover' or 'clout' while the latter is aimed at infusing our conclusions with moral legitimacy.

Implications and discussion

While we have identified four possible distinct functions of soliciting moral judgments from the public in bioethics debates, these are not hard and fast boundaries. In practice, there are often areas of significant overlap. As we have noted, in neuroethics, for example, some bioethicists engage in issue spotting through qualitative interviews and surveys but then frame the findings as substantive moral guidance (a misstep). Another potential dynamic is that once research is published, it can be used by others for purposes the authors might not have considered. This is one reason why it is important for bioethicists and policymakers to have a clear understanding of their own method and motivation when engaging with the public.

Since bioethicists elicit moral judgments from the public in a variety of contexts for a variety of reasons, no single approach or method is necessarily superior to another. Instead, some approaches may be ideal for certain contexts but inappropriate for others. For example, it may be that moral issue spotting and agenda-setting is called for in areas of research where there are particularly strong social stigmas for those with relevant research, as we see in the contexts of mental health and psychiatry, or in new bioethical research. Conversely, there may be some domains of bioethics research which are so ethically and/or medically complex that they demand primary guidance from knowledgeable researchers such that seeking substantive moral guidance from laypeople might be less appropriate. As a thoughtful reviewer notes, throughout we have referenced examples of public engagement that differ in their method (e.g., public survey or polling, analyzing public commentary, qualitative interviews), raising the further question of whether such methods all constitute engagement in the same or similar sense. There is a further dimension to public engagement, perhaps relating to depth or the closeness of engagement, which could mark out approaches like community-engaged participatory research from the more remote analysis of public comments, for example. We are inclined to view this dimension contextually as well, although 'deeper' engagement (in the sense gestured at) is likely preferable to avoid some of the deficits we outline, such as tokenism.

Finally, we suggest the beginnings of several questions bioethicists can put to themselves to assess the function public engagement is likely serving in how they conceive of their research.

- *Which* public attitudes are you interested in understanding?
- *Why* are you interested in understanding these particular attitudes?
    - For more targeted consultation, why are you interested primarily in the attitudes from stakeholder group X in particular?
- Does your interest in public engagement relate to strategic communication with the public (or with specific stakeholder groups)?
- Do you intend to reach conclusions (or make decisions/recommendations) that align with the values, interests, or preferences of the public?
    - If so, how do you envision public attitudes influencing these conclusions (or decisions/ recommendations)?
- How will public attitudes be evaluated or otherwise analyzed via professional expertise?

These questions are certainly not exhaustive, but they provide examples of ways bioethicists can begin to interrogate their own purposes for public engagement.

Through surveying the landscape of the purposes bioethicists and health policy researchers have for soliciting public judgments and discussing examples of these approaches, we hope to prompt reflection on the part of those doing empirical bioethics to be thoughtful about *what* we are doing and *why* we are doing it when we solicit public moral views and judgments. Furthermore, such reflection can reveal existing gaps between how we as bioethicists conceive of our practice of soliciting public views vs. the desires of the public regarding the *use* and *purpose* of our soliciting their views.

Notes

1. We bracket here a more nuanced discussion of the relationship of the is-ought problem to the naturalistic fallacy and the fact-value distinction, a nuance beyond the scope of the present investigation.

2. Of course, polling only an affected population can result in bias due to local variation or unique experience.

References

Agid, Y., M. Schüpbach, M. Gargiulo, L. Mallet, J. L. Houeto, C. Behar, D. Maltête, V. Mesnage, M. L. Welter, M. B. H. Youdim, et al. 2006. Neurosurgery in Parkinson’s Disease: The doctor is happy, the patient less so? In *Parkinson’s Disease and Related Disorders*, edited by P. Riederer, H. Reichmann 409–14. Vienna: Springer Vienna. doi:10.1007/978-3-211-45295-0_61.

Awad, E., S. Dsouza, R. Kim, J. Schulz, J. Henrich, A. Shariff, J.-F. Bonnefon, and I. Rahwan. 2018. The moral machine experiment. *Nature* 563 (7729):59–64. doi:10.1038/s41586-018- 0637-6.

Beauchamp, T., and J. Childress. 2019. Principles of biomedical ethics: Marking its fortieth anniversary. *The American Journal of Bioethics* 19 (11):9–12. doi:10.1080/15265161.2 019.1665402.

Beauchamp, T. L., and J. F. Childress. 2001. *Principles of Biomedical Ethics*. Oxford, USA: Oxford University Press, Incorporated. http://ebookcentral.proquest.com/lib/ tmclibrary/detail.action?docID=5763592.

Biddison, E. L. D., H. S. Gwon, M. Schoch-Spana, A. C. Regenberg, C. Juliano, R. R. Faden, and E. S. Toner. 2018. Scarce resource allocation during disasters. *Chest* 153 (1):187–95. doi:10.1016/j.chest.2017.08.001.

“Blueprint for an AI Bill of Rights – OSTP.” n.d. The White House. Accessed October 17, 2022. https://www.whitehouse. gov/ostp/ai-bill-of-rights/.

Bluhm, R., E. D. Sipahi, E. D. Achtyes, A. M. McCright, and L. Y. Cabrera. 2024. Stakeholders’ ethical concerns regarding psychiatric electroceutical interventions: Results from a US Nationwide Survey. *AJOB Empirical Bioethics* 15 (1):11–21. doi:10.1080/23294515.2023.2224592.

Bluhm, R., E. Castillo, E. D. Achtyes, A. M. McCright, and L. Y. Cabrera. 2021. They affect the person, but for better or worse? Perceptions of electroceutical interventions for depression among psychiatrists, patients, and the public. *Qualitative Health Research* 31 (13):2542–53. doi:10.1177/10497323211037642.

Bluhm, R., M. Cortright, E. D. Achtyes, and L. Y. Cabrera. 2023. ‘They are invasive in different ways’: Stakeholders’ perceptions of the invasiveness of psychiatric electroceutical interventions. *AJOB Neuroscience* 14, (1):1–12. doi: 10.1080/21507740.2021.1958098.

Blumenthal-Barby, J. 2024. The end of personhood. *The American Journal of Bioethics* 24 (1):3–12. doi:10.1080/15265161. 2022.2160515.

Borry, P., P. Schotsmans, and K. Dierickx. 2004. What is the role of empirical research in bioethical reflection and decision-making? An ethical analysis. *Medicine, Health Care, and Philosophy* 7 (1):41–53. doi:10.1023/B:MHEP. 0000021844.57115.9d.

Borry, P., P. Schotsmans, and K. Dierickx. 2005. The birth of the empirical turn in bioethics. *Bioethics* 19 (1):49–71. doi:10.1111/j.1467-8519.2005.00424.x.

Boyd, J. L., and J. Sugarman. 2022. Toward responsible public engagement in neuroethics. *AJOB Neuroscience* 13 (2):103–6. doi:10.1080/21507740.2022.2048736.

Bruno, B., H. M. Hurwitz, M. Mercer, H. Mabel, L. Sankary, G. Morley, P. J. Ford, C. C. Horsburgh, and S. L. Rose. 2021. Incorporating stakeholder perspectives on scarce resource allocation: Lessons learned from policymaking in a time of crisis. *Cambridge Quarterly of Healthcare Ethics* 30 (2):390–402. doi:10.1017/S0963180120000924.

Cabrera, L. Y., C. Courchesne, M. Bittlinger, S. Müller, R. Martinez, E. Racine, and J. Illes. 2021. Authentic self and last resort: International perceptions of psychiatric neurosurgery. *Culture, Medicine and Psychiatry* 45 (1):141–61. doi:10.1007/s11013-020-09679-1.

Cabrera, L. Y., M. Brandt, R. McKenzie, and R. Bluhm. 2019. Online comments about psychiatric neurosurgery and psychopharmacological interventions: Public perceptions and concerns. *Social Science & Medicine* 220:184– 92. doi:10.1016/j.socscimed.2018.11.021.

Cabrera, L. Y., M. M. C. Gilbert, A. M. McCright, E. D. Achtyes, and R. Bluhm. 2021. Beyond the Cuckoo's Nest: Patient and public attitudes about psychiatric electroceutical interventions. *The Psychiatric Quarterly* 92 (4):1425– 38. doi:10.1007/s11126-021-09916-9.

Cabrera, L. Y., M. Brandt, R. McKenzie, and R. Bluhm. 2018. Comparison of philosophical concerns between professionals and the public regarding two psychiatric treatments. *AJOB Empirical Bioethics* 9 (4):252–66. doi:1 0.1080/23294515.2018.1512534.

Daniels, N. 2000. Accountability for reasonableness: Establishing a fair process for priority setting is easier than agreeing on principles. *BMJ* 321 (7272):1300–1. doi:10.1136/bmj.321.7272.1300.

Daniels, N., and J. E. Sabin. 2008. Accountability for reasonableness: An update. *BMJ* 337 (1):a1850. doi:10.1136/bmj.a1850.

De Vries, R., and B. Gordijn. 2009. Empirical ethics and its alleged meta-ethical fallacies. *Bioethics* 23 (4):193–201. doi:10.1111/j.1467-8519.2009.01710.x.

Emanuel, E. J., G. Persad, R. Upshur, B. Thome, M. Parker, A. Glickman, C. Zhang, C. Boyle, M. Smith, and J. P. Phillips. 2020. Fair allocation of scarce medical resources in the time of Covid-19. *The New England Journal of Medicine* 382 (21):2049–55. doi:10.1056/NEJMsb2005114.

Evans, J. H. 2012. *The history and future of bioethics: A sociological view*. OUP USA.

Evans, J. H. 2024. Personhood and the Public's Definitions of a Human. *The American Journal of Bioethics* 24 (1):49– 51. doi:10.1080/15265161.2023.2278543.

Fairlie, S. 2015. Coproduction: A personal journey. *Mental Health Review Journal* 20 (4):267–78. Edited by Michael Clark. doi:10.1108/MHRJ-06-2014-0021.

Gilbert, F., J. N. M. Viaña, and C. Ineichen. 2021. Deflating the 'DBS Causes Personality Changes' Bubble. *Neuroethics* 14 (S1):1–17. doi:10.1007/s12152-018-9373-8.

Gregg, B. 2022. Political bioethics. *The Journal of Medicine and Philosophy* 47 (4):516–29. doi:10.1093/jmp/jhac008.

Griffiths, C., and A. O'Neill-Kerr. 2019. Patients', carers', and the public's perspectives on electroconvulsive therapy. *Frontiers in Psychiatry* 10(May). :304. doi:10.3389/fpsyt. 2019.00304.

Houeto, J. L., V. Mesnage, L. Mallet, B. Pillon, M. Gargiulo, S. T. Du Moncel, A. M. Bonnet, B. Pidoux, D. Dormont, P. Cornu, et al. 2002. Behavioural disorders, Parkinson's disease and subthalamic stimulation. *Journal of Neurology, Neurosurgery, and Psychiatry* 72 (6):701–7. doi:10.1136/ jnnp.72.6.701.

"How Should Governments Make COVID-19 Policy? – The Philosophers' Magazine." n.d. Accessed February 22, 2023. https://www.philosophersmag.com/ essays/262-how-sho uld-governments-make-covid-19-policy.

Hubbard, G., L. Kidd, and E. Donaghy. 2008. Involving people affected by cancer in research: A review of literature. *European Journal of Cancer Care* 17 (3):233–44. doi:10.1111/j.1365-2354.2007.00842.x.

Illes, J., C. Blakemore, M. G. Hansson, T. K. Hensch, A. Leshner, G. Maestre, P. Magistretti, R. Quirion, and P. Strata. 2005. International perspectives on engaging the public in neuroethics. *Nature Reviews. Neuroscience* 6 (12):977–82. doi:10.1038/nrn1808.

Jones, N., K. Atterbury, L. Byrne, M. Carras, M. Brown, and P. Phalen. 2021. Lived experience, research leadership, and the transformation of mental health services: Building a researcher pipeline. *Psychiatric Services* 72 (5):591–3. doi:10.1176/appi.ps.202000468.

Kimhy, D., N. Jones, and L. Dixon. 2022. Investing in a research workforce with personal experience of serious mental illness. *JAMA Psychiatry* 79 (9):841–2. doi:10.1001/ jamapsychiatry.2022.2026.

Kirkwood, K. 2009. Empiricism, values, and bioethics. *The American Journal of Bioethics* 9 (6-7):91–2. doi:10.1080/ 15265160902874478.

Kon, A. A. 2009a. It is settled: The 'is' can (and should) inform the 'ought'!. *The American Journal of Bioethics* 9 (6–7):W4–6. doi:10.1080/15265160902966902.

Kon, A. A. 2009b. The role of empirical research in bioethics. *The American Journal of Bioethics: AJOB* 9 (6-7):59–65. doi:10.1080/15265160902874320.

Kring, I. S., M. D. Bergholt, and J. Midtgaard. 2018. The perspectives of former recipients and experts on stigmatization related to electroconvulsive therapy in Denmark: A focus group study. *Journal of Psychiatric and Mental Health Nursing* 25 (5-6):358–67. doi:10.1111/jpm.12470.

Leshner, A. I. 2005. It's time to go public with neuroethics. *The American Journal of Bioethics* 5 (2):1–2. doi:10.1080/ 15265160590970419.

MacDuffie, K. E., S. Ransom, and E. Klein. 2022. Neuroethics inside and out: A comparative survey of neural device industry representatives and the general public on ethical issues and principles in neurotechnology. *AJOB Neuroscience* 13 (1):44–54. doi:10.1080/21507740.2021.18 96596.

McCoy, M. S., J. Warsh, L. Rand, M. Parker, and M. Sheehan. 2019. Patient and public involvement: Two sides of the same coin or different coins altogether? *Bioethics* 33 (6):708–15. doi:10.1111/bioe.12584.

"Moral Machine.". Moral Machine. n.d. Accessed December 1, 2023. http://moralmachine.mit.edu.

Nelson, R. H., B. Moore, H. F. Lynch, M. R. Waggoner, and J. Blumenthal-Barby. 2023. Bioethics and the moral authority of experience. *The American Journal of Bioethics* 23 (1):12–24. doi:10.1080/15265161.2022.2127968.

Norheim, O. F., J. M. Abi-Rached, L. K. Bright, K. Bærøe, O. L. M. Ferraz, S. Gloppen, and A. Voorhoeve. 2021. Difficult trade-offs in response to COVID-19: The case for open and inclusive decision making. *Nature Medicine* 27 (1):10–3. doi:10.1038/s41591-020-01204-6.

Oregon's Kitchen Table. 2020, October 20. "Oregon Citizen Assembly on COVID19 Recovery - Report from Oregon's Kitchen Table Survey." 2020. https://www.oregonskitchentable. org/results/oregon-citizen-assembly-covid19-recovery-repo rt-oregons-kitchen-table-survey.

Pratt, B. 2021. Achieving inclusive research priority-setting: what do people with lived experience and the public think is essential? *BMC Medical Ethics* 22 (1):117. doi:10.1186/s12910-021-00685-5.

Pugh, J. 2020. Clarifying the normative significance of 'personality changes' following deep brain stimulation. *Science and Engineering Ethics* 26 (3):1655–80. doi:10.1007/s11948- 020-00207-3.

Pugh, J., L. Pycroft, H. Maslen, T. Aziz, and J. Savulescu. 2021. Evidence-based neuroethics, deep brain stimulation and personality – Deflating, but not bursting, the bubble. *Neuroethics* 14 (Suppl 1):27–38. doi:10.1007/s12152- 018-9392-5.

Rose, D., and J. Kalathil. 2019. Power, privilege and knowledge: The untenable promise of co-production in mental 'health'. *Frontiers in Sociology* 4:57. doi:10.3389/fsoc. 2019.00057.

Russo, J. 2012. Survivor-controlled research: A new foundation for thinking about psychiatry and mental health. January 30.

Salloch, S., and F. Ursin. 2023. The birth of the 'digital turn' in bioethics? *Bioethics* 37 (3):285–91. doi:10.1111/ bioe.13099.

Salloch, S., J. Vollmann, and J. Schildmann. 2014. Ethics by opinion poll? The functions of attitudes research for normative deliberations in medical ethics: Table 1. *Journal of Medical Ethics* 40 (9):597–602. doi:10.1136/medethics- 2012-101253.

Sample, M., S. Sattler, S. Blain-Moraes, D. Rodríguez-Arias, and E. Racine. 2020. Do publics share experts' concerns about brain–computer interfaces? A trinational survey on the ethics of neural technology. *Science, Technology, & Human Values* 45 (6):1242–70. doi:10.1177/0162243919 879220.

Sattler, S., and D. Pietralla. 2022. Public attitudes towards neurotechnology: Findings from two experiments concerning Brain Stimulation Devices (BSDs) and Brain-Computer Interfaces (BCIs). *PLoS One* 17 (11):e0275454. doi:10.1371/ journal.pone.0275454.

Savulescu, J., C. Gyngell, and G. Kahane. 2021. Collective Reflective Equilibrium in Practice (CREP) and controversial novel technologies. *Bioethics* 35 (7):652–63. doi:10.1111/bioe.12869.

Scheinerman, N., and M. McCoy. 2021. What does it mean to engage the public in the response to Covid-19? *BMJ* 373:n1207. doi:10.1136/bmj.n1207.

Schneble, C. O., and D. M. Shaw. 2021. Driver's views on driverless vehicles: Public perspectives on defining and using autonomous cars. *Transportation Research Interdisciplinary Perspectives* 11:100446. doi:10.1016/j.trip.2021.100446.

Schneider, M., E. Vayena, and A. Blasimme. 2023. Digital bioethics: Introducing new methods for the study of bioethical issues. *Journal of Medical Ethics* 49 (11):783–90. doi:10.1136/medethics-2021-107387.

Schoch-Spana, M., E. K. Brunson, H. Gwon, A. Regenberg, E. S. Toner, and E. L. Daugherty-Biddison. 2020. Influence of community and culture in the ethical allocation of scarce medical resources in a pandemic situation: Deliberative democracy study. *Journal of Participatory Medicine* 12 (1):e18272. doi:10.2196/18272.

Schüpbach, M., M. Gargiulo, M. L. Welter, L. Mallet, C. Béhar, J. L. Houeto, D. Maltête, V. Mesnage, and Y. Agid. 2006. Neurosurgery in Parkinson Disease: A distressed mind in a repaired body? *Neurology* 66 (12):1811–6. doi:10.1212/01.wnl.0000234880.51322.16.

Smith, G., T. Hughes, L. Adams, and C. Obijiaku. 2021. *Democracy in a Pandemic*. University of Westminster Press. University of Westminster Press. doi:10.16997/book57.

Solomon, M. Z. 2005. Realizing bioethics' goals in practice: Ten ways 'is' can help 'ought'. *Hastings Center Report* 35 (4):40–7. doi:10.1353/hcr.2005.0048.

Stilgoe, J., and T. Cohen. 2021. Rejecting acceptance: Learning from public dialogue on self-driving vehicles. *Science and Public Policy* 48 (6):849–59. doi:10.1093/ scipol/scab060.

Sugarman, J., N. Kass, and R. Faden. 2009. Categorizing empirical research in bioethics: Why count the ways? *The American Journal of Bioethics* 9 (6-7):66–7. doi:10.1080/15265160902874379.

Sunstein, C. R. 2016. *The ethics of influence: Government in the age of behavioral science*. Cambridge: Cambridge University Press.

Sweeney, A., and P. Beresford. 2020. Who gets to study whom: Survivor research and peer review processes. *Disability & Society* 35 (7):1189–94. doi:10.1080/09687599. 2019.1664062.

Takaguchi, K., A. Kappes, J. M. Yearsley, T. Sawai, D. J. C. Wilkinson, and J. Savulescu. 2022. Personal ethical settings for driverless cars and the utility paradox: An ethical analysis of public attitudes in UK and Japan. Edited by Xiaoqiang 'Jack' Kong. *PLoS One* 17 (11):e0275812. doi:10.1371/journal.pone.0275812.

Taylor, A., C. McMellon, T. French, A. MacLachlan, R. Evans, R. Lewis, M. McCann, L. Moore, S. Murphy, S. Simpson, et al. 2022. Defining research priorities for youth public mental health: reflections on a coproduction approach to

transdisciplinary working. *Health Research Policy and Systems* 20 (1):72. doi:10.1186/s12961-022-00871-w.

Thaler, R. H., and C. R. Sunstein. 2008. *Nudge: Improving decisions about health, wealth, and happiness*. Yale: Yale University Press.

Thompson, M. S., A. Cochrane, and J. Hopma. 2020. Democratising food: The case for a deliberative approach. *Review of International Studies* 46 (4):435–55. doi:10.1017/ S0260210520000017.

Wagner, K., H. Maslen, J. Oakley, and J. Savulescu. 2018. Would you be willing to zap your child's brain? Public perspectives on parental responsibilities and the ethics of enhancing children with transcranial direct current stimulation. *AJOB Empirical Bioethics* 9 (1):29–38. doi:1 0.1080/23294515.2018.1424268.

Wangmo, T., S. Hauri, E. Gennet, E. Anane-Sarpong, V. Provoost, and B. S. Elger. 2018. An update on the 'empirical turn' in bioethics: Analysis of empirical research in nine bioethics journals. *BMC Medical Ethics* 19 (1):6. doi:10.1186/s12910-018-0246-9.

Zuk, P., L. Torgerson, D. Sierra-Mercado, and G. Lázaro-Muñoz. 2018. Neuroethics of neuromodulation: An update. *Current Opinion in Biomedical Engineering* 8 (December):45–50. doi:10.1016/j.cobme.2018.10.003.