Abolishing Morality in Biomedical Ethics

Parker Crutchfield, Western Michigan University Homer Stryker M.D. School of Medicine

Scott Scheall, Arizona State University

In biomedical ethics, there is widespread acceptance of moral realism, the view that moral claims express a proposition and that at least some of these propositions are true. Biomedical ethics is also in the business of attributing moral obligations, such as "S should do X". The problem, as we argue, is that against the background of moral realism, most of these attributions are erroneous or inaccurate. The typical obligation attribution issued by a biomedical ethicist fails to truly capture the person's actual obligations. We offer a novel argument for rife error in obligation attribution. The argument starts with the idea of an epistemic burden. Epistemic burdens are all of those epistemic obstacles one must surmount in order to achieve some aim. Epistemic burdens shape decision-making such that given two otherwise equal options, a person will choose the option that has the lesser of epistemic burdens. Epistemic burdens determine one's potential obligations and, conversely, their non-obligations. The problem for biomedical ethics is that ethicists have little to no access to others' epistemic burdens. Given this lack of access and the fact that epistemic burdens determine potential obligations, biomedical ethicists often can only attribute accurate obligations out of luck. This suggests that the practice of attributing obligations in biomedical ethics is rife with error. To resolve this widespread error, we argue that this practice should be abolished from the discourse of biomedical ethics.

The primary aim of this paper is to support the idea that biomedical ethics, and biomedical ethicists, should abolish moral language and moral reasoning from the practice of biomedical ethics. The force of the 'should' in the previous sentence is practical rather than moral. The idea is that it is in the practical interests of biomedical ethicists to expunge moral language and reasoning from their practice. A further qualification is necessary, however. We don't argue that *all* moral language and moral reasoning should be abolished. Instead, biomedical ethicists should abolish *moral obligation attributions*. Since attributing moral obligations constitutes a significant component of the contemporary practice of biomedical ethics, abolishing moral obligation attributions constitutes a significant turnover of the practice.

The claim that biomedical ethicists (ourselves included) should refrain from attributing moral obligations rests on a type of moral error theory, though not the familiar moral error theory associated with Mackie.[1] According to Mackie's work, there simply are no properties of the sort that could make moral propositions true, and, even if there were such properties, we would lack cognitive access to them. Thus, Mackie's error theory supports a variety of moral anti-realism. Moral claims simply do not express propositions of a sort that can be true.

Our moral error theory, on the other hand, is neutral with regard to the debate between moral realists and anti-realists. Our moral error theory holds that, in order to be properly obligated to perform some action, a would-be actor must occupy a particular epistemic position with regard to the action. The actor must know *of* the action and its relevance in prevailing circumstances, and know *how to* perform the action such that its goal is achieved, its purpose satisfied. Courses of action about which an actor is relevantly ignorant or incapable of performing effectively cannot be obligations for the actor.

On our view, the error theory that supports the abolition of moral obligation attributions in biomedical ethics is that such attributions are typically inaccurate, due to the fact, for which we argue, that obligation attributors are often insufficiently knowledgeable of the epistemic conditions under which the subject of the attribution (hereafter, 'subject') makes and acts upon options. Biomedical ethicists (we) often attribute moral obligations to others. When we do so, it is typically from a position of ignorance of the subject's relevant epistemic circumstances. Given that biomedical ethicists typically aim to get ethics right, obligation attributions usually miss their mark, leaving our practical interests unsatisfied.

A moral obligation attribution is a claim of the sort, "S should do X." Such attributions constitute a significant proportion of the practice of biomedical ethics. Perhaps one of the more famous moral obligation attributions is that one should save the child drowning in a shallow pond.[2] A person writing in research ethics, or in advising

[1] Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin
[2] Singer, P. (1972). Famine, Affluence, and Morality. *Philosophy & Public Affairs*, *1*(3), 229–243.

policy, might claim, “One should not edit the human germline.” A clinical ethicist might advise that “The medical staff should refuse to provide the requested treatment.” Over the course of the pandemic, we heard numerous moral obligation attributions related to public health, and not just from biomedical ethicists, such as “You should get your vaccine,” or “States should have plans in place to mitigate suffering from public health emergency.” And one need not look very hard in journal articles to find an attribution to the public, such as the seemingly ubiquitous, “Society should debate X.”

Our claim is that bioethicists should neither utter such statements nor reason upon such premises, unless they know the adequacy of the epistemic circumstances of the targets of their attributions. Since, as we argue below, bioethicists are rarely, if ever, so knowledgeable, obligation attributions should be, for practical reasons, abolished from biomedical ethics. However, there is a significant difference between recommending that no one make moral obligation attributions and recommending merely that biomedical ethicists refrain from such language. The difference is that, at least in the long term, it is probably possible for biomedical ethicists to make such a radical change. Indeed, restricting the abolitionist recommendation to biomedical ethicists resolves one of the more pressing challenges to moral abolitionism.[3]

Abolitionism is a response to the “Now what?” problem for error theories. The more familiar error theory holds that all moral statements are neither true nor false—there is nothing that they refer to. There just are no moral facts, and so anytime a person issues a moral statement of any kind, she is saying something that has no referent. But if that’s the correct metaethical position, what then? One idea is that we should simply abolish all morality. This abolitionism, however, must contend with the challenge that it is not plausible that society can abolish morality from language and reasoning. This apparent impossibility undermines the plausibility of abolitionism, even if all of the other arguments in its favor are sound.

Restricting the abolitionist recommendation to biomedical ethics can help to meet this challenge. The community of biomedical ethicists does not include very many people. And of those in it, like many academic disciplines, an even smaller group of

[3] Although the paper is about biomedical ethicists, the argument can also work for applied ethics more generally.

people generate most of the ideas that pervade the discipline and practice. The small size and ease at which ideas become pervasive makes it plausible that a radical change such as abolishing moral language and moral reasoning *can* take hold. We argue that it *should*.

We return to the challenges of abolitionism in subsequent sections. In the first section, we set the metaethical stage by pointing out that generally biomedical ethics, and biomedical ethicists, assume moral realism. We take moral realism to include any account according to which moral claims are propositions and that some of these propositions can be true. Accounts according to which moral claims can be true and the truthmakers are features of the mind-independent world are moral realist. But so are accounts according to which the truthmakers of moral claims are relative to some mind-dependent feature of the speaker's culture, society, etc. The presumption that most biomedical ethicists are moral realists (i.e., most are either objectivists or relativists), figures importantly in *our* (practical) recommendation. In biomedical ethics, moral obligation attributions seem to aim at truth. We want our attributions to get it right. This desire can be satisfied, as we argue, only if one attributes practical obligations, rather than moral obligations. The failure of moral obligation attributions to get it right is an epistemic failure, rather than, as Mackie or other error theorists would have it, a metaphysical one. Our account is neutral with regard to the metaphysical status of moral claims, so long as some of them express true propositions (regardless of whether it is the objective world that makes them true or the person's culture or society).

In the second section, we introduce the epistemic challenge to moral obligation attributions in biomedical ethics. We draw from our recent work in which we argue (and experimentally observe) that epistemic considerations are logically prior to moral considerations. This is to say that a person's epistemic circumstances fix that person's moral obligations. We argue for this priority in two ways. One way appeals to introspection of one's decision-making process, which is supposed to reveal that epistemic considerations constrain the structure of one's available choices. The second way exploits the plausibility of the principle *ought implies can*. If this principle holds, then it follows, we argue, that epistemic properties fix a person's moral obligations.

If we are right, then the truth of propositions expressing a moral obligation is fixed in part by epistemic properties. The problem, as we argue in sections three and four, is that when a biomedical ethicist makes moral obligation attributions, she typically lacks access to these epistemic properties. One needs to know, first and foremost, what the subject's epistemic circumstances are. But the attributor—the biomedical ethicist—is typically ignorant of the relevant circumstances. Thus, moral obligation attributions in biomedical ethics are typically inaccurate, or accurate only by accident. Thus, if we insist on maintaining the widespread use of moral obligation attributions in the practice of biomedical ethics, we may be better off flipping a coin and calling it in the air.

This rife error of moral obligation attributions creates its own "Now what?" problem. Our response is that that they should be abolished from biomedical ethics. But there are other options, such as the view that we should prop up the fiction that moral obligation attributions reflect moral truths. In section five, we discuss this fictionalism and other options. After dismissing these alternatives, in section six we turn to a more detailed defense of abolitionism in biomedical ethics as a response to the "Now what?" problem. In section seven, we address some of the lingering objections to the version of abolitionism we offer.

1. Metaethics of Biomedical Ethics

Biomedical ethics, and biomedical ethicists, often assume moral realism, even if individual ethicists fail to recognize that assumption or make it explicit. Moral realism includes all of those positions that allow for moral claims to be true, even if they are true only by virtue of features of the person's culture or society. That is, moral relativism is a moral realist position, by our lights. Metaethical considerations are rarely introduced in biomedical ethics. Instead, the moral obligation attributions and other moral claims that biomedical ethicists make indicate their metaethical commitments. For example, much of the moralizing in biomedical ethics proceeds from the notion that there is a morality common to various communities and cultures and that from this common morality four moral principles (i.e, autonomy, beneficence, non-maleficence, and justice) can be

derived (this is simply observational, not an endorsement of the view). The underlying assumption in this schema is that claims about these four principles can be true.

Also consider other examples. When an ethicist claims that factory farms should be closed down or that researchers should stop performing research on non-human primates, she does so on the basis that such claims can be true. Or when a clinical ethicist documents in a patient chart that "the ventilator should not be withdrawn, because to do so would violate the patient's autonomy," she is doing so on the basis that her claims can be true, either objectively or relatively. Biomedical ethicists presume that their invocation of moral values in reasoning is invocation of things that are real—autonomy is a real thing, and claims about it (a) express propositions and (b) those propositions can be true, even relatively so.[4]

We return to the metaethical commitments of biomedical ethics below, where we consider metaethical alternatives to abolitionism. But consider how different the field would be if biomedical ethics presumed Mackie's error theory or non-cognitivism. Biomedical ethicists don't seem to really mean, "Boo, human germline editing," when they claim that policymakers should prohibit germline editing. If they did mean this and revealed their true meaning to policymakers, they'd like elicit strange looks and puzzlement. And biomedical ethicists don't seem to be merely expressing a non-propositional attitude when they recommend that an institution impose a vaccine mandate. Similarly, it doesn't seem as though biomedical ethicists are covert fictionalists and really think they are saying something false when they make those same claims. In any case, as we claim below, pushing back on abolitionism by abandoning moral realism is not a strategy that will preserve the notion that biomedical ethicists should maintain the current practice of issuing moral obligation attributions.

Biomedical ethicists also have an interest in saying accurate, true, or correct things. This is true even if our claims aren't intended to express a true proposition about morality, but it is especially true when we do that. This is not a weak or light interest.

[4] Bourget, D., Chalmers, D. J., & Chalmers, D. (2023). Philosophers on philosophy: The 2020 philpapers survey. *Philosophers' Imprint*, *23*. In their paper, Bourget and Chalmers report on the prevalence of endorsement of moral realism among professional philosophers. Though it is not reported in that paper, the survey response on PhilPapers (https://survey2020.philpeople.org/survey/results/4866?aos=26) indicate that approximately two-thirds of 224 professional philosophers who claim specialization in biomedical ethicists are moral realists.

That biomedical ethicists aim to be reliable is one of the features that keeps us in the business of having influence over policies, patients, scientific progress, etc. In any case, this assumption is not likely to be challenged, because it seems incorrect to say that biomedical ethicists generally don't care about getting it right. There is a contingent of biomedical ethicists who think the goal for biomedical ethics, when trying to resolve value conflicts, is to achieve consensus.[5] But even this process aims at saying accurate things. It's just that what determines whether a claim is accurate also depends on consensus. Consensus-building ethicists don't sacrifice accuracy. Ethicists might disagree about the process by which to make ethical recommendations, but this disagreement doesn't cut so deep that there is disagreement about the commitment to accuracy (which entails a commitment to moral realism).

2. Epistemic Burdens and Obligation Attributions

In order to have an obligation to do *X*, doing *X* must at least be an option for the actor, *A*, i.e., *X* must be something that *A could* choose to do. The relevant question is *where do options come from*? What causes one course of action to be an option for *A* and another course of action to not be an option (and, thus, not a potential obligation) for *A*? Decision theorists, economists, ethicists, et al., tend to analyze preference rankings and incentive structures that encompass *given* options, as if these fall from the sky; the question of what determines whether a course of action is included in an actor's menu of options in the first place is rarely, if ever, asked.

We have argued in recent work that options are first pre-consciously determined and ranked according to the actor's relative *epistemic burdens*, i.e., the nature and extent of the actor's ignorance, with regard to various courses of.[6] In other words, some

[5] Brummett, A., & Salter, E. K. (2019). Taxonomizing Views of Clinical Ethics Expertise. *The American Journal of Bioethics*, *19*(11), 50–61.

[6] Scheall S. (2019). Ignorance and the Incentive Structure Confronting Policymakers *Cosmos + Taxis*, *7* (1–2): 39–51.; Crutchfield, P., & Scheall, S. (2019). Epistemic burdens and the incentives of surrogate decision-makers. *Medicine, Health Care and Philosophy*, *22*, 613-621.' Crutchfield, P., & Scheall, S. (2020). Epistemic Burdens, Moral Intimacy, and Surrogate Decision-Making. *The American Journal of Bioethics*, *20*(2), 59-61.; Scheall, S. (2020). *FA Hayek and the epistemology of politics: the curious task of economics*. Routledge.; Crutchfield, P., Scheall, S., Justin Rzeszutek, M., Dawn Brown, H., & Cardoso Sao Mateus, C. (2023). Ignorance and moral judgment: Testing the logical priority of the epistemic. *Consciousness and Cognition*, *108*, 103472

epistemic work must first occur, if only at a pre-conscious level, in order to have options from which to choose and, therefore, in order to have obligations to choose particular options. Our respective epistemic burdens serve to determine our options and, thus, our potential obligations. We have previously used this work to argue that surrogate decision-makers, especially those in medicine, are faced primarily with an epistemic challenge, rather than a moral one.[7]

We offer two main arguments for the thesis of the logical priority of the epistemic, one based on simple introspection and another grounded in reflection on the possible meanings of principles like *ought implies can*.

The argument from introspection proceeds from the observation that the options which persons consciously consider in any given decision context never encompass every course of action that might be pursued under the circumstances, but only those courses of action the person takes themselves to know (or to be capable of learning) enough to pursue more or less effectively, i.e., as a means to realize the relevant goal, whatever it might be.[8] More generally, we argue that introspection reveals that the options that a person consciously considers in any given decision context have been, at least roughly, pre-consciously sorted and ranked according to their relative epistemic burdens.[9] Courses of action that we take ourselves to be too ignorant to realize do not appear in our menus of options, while courses of action with respect to which we take ourselves to be relatively ignorant are typically ranked below options with respect to which we take ourselves to be relatively knowledgeable. When considering whether and how to travel overseas, for example, our initial menus of options, as pre-consciously determined, show that courses of action about which we take ourselves to be relatively more knowledgeable, say, airplane travel, typically rank above those with regard to which we take ourselves to be relatively ignorant, say, constructing and using a

---

[7] Op cit note 6, Crutchfield and Scheall, 2019, 2020

[8] See Scheall, S., & Crutchfield, P. (2021). The priority of the epistemic. *Episteme*, *18*(4), 726-737, p. 3-4. Including the subsidiary end of adopting some means toward some ultimate goal. That is, the realization of some ultimate end might require means the adoption of which is a subsidiary end to the ultimate end. For example., the ultimate end of, say, *driving to Los Angeles*, might require the subsidiary end of *acquiring an automobile* as a means to the ultimate end.

[9] Op cit note 8

transoceanic pontoon bridge. Moral, prudential, and pecuniary considerations only enter the decision-making picture after (logically and temporally) some pre-conscious epistemic work has occurred. In order to consider an array of options in the light of relevant moral, prudential, financial-pecuniary, etc., considerations, an actor first needs options, and it is the actor's relevant knowledge and ignorance in their given circumstances that determine their options. Some pre-conscious assessment of one's epistemic circumstances must occur before the person can consciously consider an array of options and make a choice.

Our second argument is that it is not possible to make sense of the presumed relation between *ought* and *can*, whatever its logical strength (i.e., regardless of whether it is *logical implication*, *presupposing*, *making plausible*, *conversationally implicating*, or some other relation that obtains between *ought* and *can*), unless the epistemic is logically prior to the normative.[10] This is because the word "can" is inherently epistemic.[11] There is no non-epistemic meaning of "can" – or, at least, none has been proffered in the literature – that does not make the things a person can do a function of the things they know *of* and know *how to do*. Thus, *ought implies* (or whatever) *knows enough to*.[12] In effect, even if there is disagreement about the relation that connects ought and can, everyone who accepts that some relation like those debated in the literature obtains between ought and can is already implicitly committed to the logical priority of the epistemic.

It is important to understand the logic of principles like *ought implies can*. The contrapositive of *ought implies* (or whatever) *can* is *cannot implies* (or whatever) *not*

[10] It is often asserted that *ought implies can*, though other moral philosophers have argued that a weaker relation than *logical implication* obtains between ought and can. Perhaps ought merely *presupposes* or *makes plausible*, or is a *conversational implication of*, can. Hampshire, Stuart (1951). Symposium: Freedom of the will. *Aristotelian Society Supplementary Volume* 25: 161-178.; Hare, R.M. (1951). Symposium: Freedom of the will. *Aristotelian Society Supplementary Volume* 25: 201-216; Hare, R.M. (1963). *Freedom and Reason*. Oxford: Clarendon Press; Sinnott-Armstrong, Walter (1984). 'Ought' conversationally implies 'can'. *The Philosophical Review* 93: 249-261.; Mizrahi, Moti (2015). Ought, can, and presupposition: An experimental study. *Methode* 4 (6):232-243; Littlejohn, Clayton (2009). 'Ought', 'can', and practical reasons. *American Philosophical Quarterly* 46 (4): 363-373; Vallentyne, Peter (1989). Two types of moral dilemmas. *Erkenntnis* 30: 301-318;
Vogelstein, Eric (2012). Subjective reasons. *Ethical Theory and Moral Practice* 15: 239-257.

[11] Recall that, for us, knowledge encompasses both propositional knowledge-that and non-propositional knowledge-how

[12] Of course, if no relation obtains between ought and can, then this argument loses its force. However, the argument from introspection would still stand, even if this were the case.

ought. It is not *cannot implies ought not*. In other words, if the subject of an obligation attribution cannot, because of their epistemic deficiencies, perform the allegedly obligated action, they cannot be obligated to perform that action (according to the principle); the purported obligation simply evaporates. It is not the case, however, that an obligation to *not* do the action (an "ought not") emerges in its place. Actors who are relevantly ignorant or incapable with regard to *X* are not obligated to *X*, but this does not mean they are obligated to *not-X*.

The thesis that the epistemic is logically prior to the normative, generally, and the moral, specifically, has recently been confirmed by a series of experiments observing research participants' responses to moral dilemmas in which they, or the actor in the dilemma, are ignorant. [13]In these studies, we found that the most significant factor in a participants' choice in the trolley problem was the degree of ignorance related to the various options. Subsequent studies further confirm the logical priority of the epistemic. Most recently, and highly relevant to argument that follows, we have observed that research participants who are fully knowledgeable of a subject's epistemic burdens are less likely to attribute moral obligations to ignorant subjects than to knowledgeable subjects.[14] This observation provides good additional empirical evidence that moral obligation attributions critically depend upon the subject's epistemic burdens and what the attributor knows about those epistemic burdens.

Epistemic burdens have two components. That is, in any given decision context, there are two things with regard to which an actor might be ignorant. An actor may be ignorant of various potential goals relevant in the context in which they find themselves or they might know of relevant ends, but be ignorant of means adequate to realize them. In other words, an actor might be ignorant *that* there is a relevant goal or an actor might be either ignorant *that* there are causal processes that lead to, or ignorant *how* to manipulate relevant causal processes in order to realize, the goal. Such ignorance can inhibit or altogether prevent an actor from pursuing various goals. Obviously, courses of action related to goals with regard to which an actor is entirely ignorant do not appear to

[13] Op cit note 6, Crutchfield et al., 2023
[14] Scheall, S., Crutchfield, P. The Logical Priority of Epistemic Considerations in Third-Person Normative Judgments**. (**under review)

them as options. Similarly, an actor who knows of a goal, but is completely ignorant of means to realize it, finds no options related to the goal in their preference ranking.

More to the present point, a goal with regard to which the actor is to some extent ignorant can be realized only if spontaneous forces (e.g., luck, providence, the "invisible hand") intervene to mitigate the effects of this ignorance. To the extent that the actor is ignorant of or how to realize a particular goal, it is these spontaneous forces, not anything due to the actor or their actions, that, if the goal is realized despite the actor's ignorance, account for the realization of the goal. If your epistemic burden with respect to *X* is weightless, which is to say, if your knowledge of the causal processes that lead to *X* is adequate – you know both that there are some relevant causal processes and how to intervene in them so as to bring about *X* – then *X* can be something you are potentially obligated to do; an attribution to you of an obligation to do *X* is more likely to be accurate. It seems uncommon for subjects of biomedical ethicists' obligation attributions to be in this situation. If they were, there would be little need for the ethicists.

If your epistemic burden with respect to *X* is neither weightless nor impossibly heavy, i.e., if you have some knowledge of relevant causal processes, but not enough (or not of the right kind) to bring about *X* entirely on your own initiative without help from spontaneous forces, then you can be obligated to *X* only up to the extent of your relevant knowledge. An attribution to you of an obligation to do *X* is likely to be inaccurate. This is the situation most subjects of biomedical ethicists' obligation attributions find themselves in.

But sometimes subjects find themselves in even more epistemically impoverished situations: If your epistemic burden with respect to *X* is impossibly heavy for you to lift, if you are entirely ignorant of any causal processes that might lead to *X* or how to manipulate them effectively, then you cannot be obligated to *X* and an attribution to you of an obligation to do *X* is guaranteed to be inaccurate.

3. Never Weightless

Biomedical ethicists frequently attribute moral obligations, and do so for subjects occupying a wide range of roles. In a clinical setting, they attribute obligations to

patients and families and health care providers. In organizational settings they attribute obligations to hospital administrators. Often ethicists attribute obligations to policymakers. Ethicists frequently also attribute obligations to researchers, either in their role as ethics advisor for the research team, or as a member of an institutional review board (IRB) overseeing research. Finally, biomedical ethicists may also attribute moral obligations to the public. Thus, biomedical ethicists attribute obligations to a wide range of subjects. Each of these subjects confronts epistemic burdens in satisfying the obligation that ethicists attribute to them. None of these epistemic burdens are weightless—subjects always have epistemic work to do—if they didn't there would be no need to call on the ethicist.

In a clinical context, the subject of a biomedical ethicist's obligation attribution may be a patient or a patient's family member. For example, when a patient is incapacitated and potentially nearing end of life, an ethicist (perhaps not a very good one) might attribute an obligation to the family who is trying to decide whether to withdraw life-sustaining technologies: "You should do what the patient would want done." Or to a patient who is struggling with whether to refuse physically burdensome treatments, the ethicist might recommend, "You should talk to your family and think about what is important to you." To a physician wondering whether to side with the patient's validly executed but old advanced directive or the patient's adult children, the ethicist might recommend that the physician should explore the family's reasons for disputing the advanced directive. Or they may recommend to the social worker that they should not call child protective services (CPS) on an infant's parents.[15]

None of these obligations (i.e., do what the patient would want; talk to family, etc.) carries a weightless epistemic burden. And in some cases, such as doing what the patient wants, the epistemic burden is impossibly heavy. Even ordinary obligations that are comparatively easy to satisfy carry significant epistemic burdens. The social worker

[15] The claim that ethicists attribute moral obligations intersects with the literature related to whether ethicists *should* make morally substantive recommendations. Some think that ethicists should refrain from doing so, and should issues recommendations only related to process. Notwithstanding the fact that making recommendations regarding process are still moral recommendations, this view is not incompatible with the more descriptive claim that we rely on here, namely that ethicists *do* attribute moral obligations. The view that biomedical ethics should be a process-only practice is itself revisionist. See note 5 (Brummett and Salter, 2019) for a helpful taxonomy of views and related literature.

needs to know how to contact CPS specifically and work efficiently with overburdened government workers. The physician needs to know how to explore the family's reasons (a skill not effectively included in medical school curriculum and often not residency). The patient needs to know what the likely consequences are for withdrawal and how those consequences will affect their own life.

When ethicists attribute obligations to researchers, the researchers similarly confront (sometimes impossibly) heavy epistemic burdens in satisfying the obligation. For example, the ethicist on the IRB might attribute the obligation that the researchers should get local approval for research proposed to be conducted in a developing nation. To satisfy the obligation, researchers are likely to confront the epistemic burden of knowing how to navigate the local society to solicit approval. Or an ethicist might recommend to the principal investigator on their research team that they should stop the study. Satisfying these obligations requires overcoming various epistemic burdens.

Policymakers are in the same position. When an ethicist recommends that a drug in short supply should be prioritized for children, or for racial and ethnic minorities, policymakers confront extremely heavy epistemic burdens in satisfying the obligation attributed to them. For example, they would need to know how to proportionally allocate the drug and get it to the intended facilities. If the policymaker is elected, they would additionally need to know how to justify and explain it to their constituents, or else how to conceal their role in the policy. Policymakers are notoriously ignorant.[16] Attributing obligations to them requires that they repair this ignorance.

The public is similarly notorious for lacking knowledge, or at least for allowing misinformation to be taken as reliable and action-guiding. When an ethicist recommends that the public not eat factory-farmed meat, the public must know how to sustain themselves using more expensive meat (which may require knowing how to find another job) or how to prepare plant-based foods.

In these cases, the subjects of the obligation attribution must confront significant epistemic burdens in satisfying the obligations. Not only must they confront significant epistemic burdens, because epistemic properties are logically prior to normative properties—the 'shoulds' depend on the 'knows'—for the subjects to do what the ethicist

[16] Op cit note 6, Scheall 2020

says they should, it must also be the case that the other options available to the subject are associated with, other things being equal—lighter epistemic burdens. That is, for subjects to satisfy the obligation the ethicist attributes to them, satisfying the obligation must rise to the top of their preference ranking. Otherwise, they'll do something with a lighter epistemic burden.

More generally, ethicists are typically consulted, or issue recommendations in the form of obligation attributions, only when the person soliciting the ethicist's advice is significantly ignorant of what they should do. This is to say that anytime an ethicist gets involved in a recommendation, it is extremely likely that the potential subjects of an obligation attribution confront a good deal of ignorance. When an ethicist is involved, it is usually in the context of others' ignorance—they lack the know-how or the know-that. Thus, when an ethicist attributes an obligation to a subject, the obligation the ethicist attributes is typically associated with weighty epistemic burdens.

4. Darts in the Dark

Given that (i) the epistemic is prior to the moral and (ii) that subjects typically confront weighty epistemic burdens associated with an ethicist's attributed obligation, a lot has to happen for the ethicist's attribution to be accurate. To say that an attribution is accurate is to say that (a) the obligation to do X is attributed to S and (b) S is in fact obligated to do X. Whether a subject of an obligation attribution is in fact so obligated is largely an epistemic matter, as we have been arguing. There is, of course, much on the metaphysical end of the obligation that must go right: the obligation has to be grounded in the true moral theory, whatever that happens to be. But lots must also go right epistemically for the ethicist to *deliberately* accurately attribute an obligation. That lots has to go right epistemically for obligation attributions to be deliberately accurate is true even if moral relativism is the correct metaethical theory.

What, exactly, must go right epistemically for an ethicist to issue a deliberately accurate obligation attribution?

First, any obligation attributed (by a biomedical ethicist or anyone else) to an actor must, at a minimum, be among the actor's potential obligations. It cannot be among

those actions that, for epistemic reasons, because the actor is sufficiently relevantly ignorant or incapable, is not even an option for the actor and is, therefore, among the actor's non-obligations. As we previously argued, it is rare that targets of a biomedical ethicist's obligation attribution possess the required knowledge. This undermines the likelihood that attributions will be accurate. Whatever obligation the biomedical ethicist attributes, there is a good chance that the subject is not so obligated, because there is a good chance that satisfying the obligation is beyond the subject's epistemic capacities. By way of *ought implies can*, cannot implies not ought. If the subject cannot, because of their epistemic deficiencies, do the thing the obligation requires, they are just not obligated to that thing. This means that an obligation attribution is likely to be inaccurate—it is likely to miss its target.

Second, the ethicist must possess adequate knowledge of the decision-maker's epistemic circumstances. The ethicist must know that the action they ascribe to the actor as an obligation is in fact among the latter's potential obligations, that it is not a member of the set of the actor's non-obligations. Given the difficulties we all confront in peering into each other's minds, it is always likely that the biomedical ethicist lacks this knowledge.

In the absence of this knowledge, attributing obligations is like throwing darts in the dark. The ethicist, like the dart thrower, might hit the target. But if they do so it will be out of luck. In a darkened room, a dart thrower might hit triple twenty, but they would only do so out of luck. Similarly, an ethicist blind to the subject's epistemic burdens associated with the target option might accurately attribute an obligation, but they could only do so out of luck. This is to say that the ethicist, in the absence of access to the subject's epistemic burdens, cannot hit the target deliberately (presumably requiring luck rules out deliberate accuracy; lotto winners might intend to pick the winners, but winning doesn't mean they did so deliberately). They would need to be lucky enough that the subject's epistemic burdens are such that they could themselves deliberately pursue that option. For any particular moral obligation attribution, the epistemic burdens associated with the attribution will necessarily vary, and will do so based on what the subject and attributing ethicist know in those particular circumstance. But the logical

priority of the epistemic implies that in all of these circumstances the ethicist hitting the target is unlikely.

Lacking access to the subject's epistemic burdens means that one lacks access to what the subject knows and to what the subject is ignorant of. If the ethicist additionally knows that one lacks access to the subject's epistemic burdens, then she knows she doesn't know—it's a known unknown whether the attribution will hit its mark. This known ignorance defeats the ethicist's reasons for issuing the attribution in the first place. A dart thrower in a dark room has little reason to throw the dart at all. Indeed, to throw in spite of knowing that one is ignorant of the target is imprudent and potentially dangerous.

In order for the ethicist to deliberately accurately attribute a moral obligation, the subject's epistemic burdens must be illuminated for her. One might thus think that it is sufficient for the subject of the attribution to inform the ethicist of their epistemic burdens. But this isn't good enough, because merely illuminating one's epistemic burdens does still not entail that the subject has the required know-how and know-that to pursue the option the ethicist wants to obligate. Whatever those epistemic burdens are, the subject needs to be able to overcome them. Granted, the ethicist may be able to do this, to some degree. But in the absence the ethicist's successful attempt at eliminating or overcoming the subject's epistemic burdens associated with the option the ethicist wants to obligate, the ethicist's accuracy is a matter of luck, and they therefore cannot achieve accuracy deliberately.

To extend the analogy to moral metaphysics, our claim is that biomedical ethicists are throwing darts in the dark. This is true regardless of whether the dart board is fixed in place (i.e., objectivism) or it's moving around according to the speaker or subject's culture (i.e., relativism). Though if the dart board is moving around—if relativism is true—making deliberately accurate obligation attributions is even more difficult.

5. Now What?

Given that the ethicist will typically not be able to deliberately accurately attribute a moral obligation to a subject, what *should* the ethicist do?[17] This 'should' is not necessarily moral, however. Ethicists, for the most part, aim to say accurate things. When those things are moral obligation attributions, there are good reasons to be dubious of that accuracy. Attributing moral obligations undermines the ethicist's practical goal of saying accurate things. One option is to attempt to secure the practical aim of saying accurate things by manipulating the conditions under which obligation attributions are inaccurate. Darts in the dark are inaccurate because the lights are off. Flip on the lights and dart throws become more accurate. For ethicists attributing moral obligations, this means improving the subject's epistemic position relative to the thing the ethicist wants to obligate and in improving their own epistemic position relative to the subject's epistemic burdens. It means helping the subject overcome the epistemic burdens associated with the target option.

Doing so, however, obviates the moral obligation attribution. By helping the subject overcome the epistemic burdens associated with action the ethicist wants to obligate, the ethicist is indirectly moving that option up in the subject's preference ranking. In helping them confront and overcome the epistemic burdens, the ethicist is making the subject's selection of that option far more likely. That's what it is for the epistemic to be prior to other normative considerations. Suppose the ethicist thinks that S should X, where the 'should' is moral. In attributing the moral obligation to X, the ethicist is throwing a dart in the dark. So, instead the ethicist helps S confront and overcome all epistemic burdens associated with doing X. But in doing so, the ethicist has helped S move doing X all the way to the top of their preference ranking. To then attribute an obligation to X, the ethicist adds nothing. Indeed, the attribution will be accurate, but it will also be trivial. By attributing the obligation after helping S overcome the associated epistemic burdens, the ethicist is, in effect, saying that S should do what they have most reason to do. In helping the subject overcome the epistemic burdens associated with doing X, the ethicist has made the attribution of the obligation to X meaningless. While

[17] Given that the 'Now What?' problem here is slightly different from that which follows Mackie's traditional error theory, the solutions are also slightly different. However, often they run in parallel. For a review of responses to the problem, see Lutz, M. (2014). The 'Now What' Problem for error theory. *Philosophical Studies*, *171*(2), 351–371.

ethicists presumably have practical interests in saying accurate things, we presumably also have practical interests in not saying meaningless things. Saying meaningful, accurate things is what justifies our continued employment and the expansion of professionalized biomedical ethics.

This is not to say that the ethicist should refrain from helping subjects overcome the epistemic burdens associated with the option that the ethicist wants to obligate. In fact, this might be the primary function of the ethicist and how their practical interests are best satisfied. But helping subjects overcome epistemic burdens doesn't require attributing moral obligations to them also. Instead, doing so conflicts with the ethicist's practical interests in not saying superfluous, trivial things. Responding to the likely inaccuracy of obligation attribution by asserting that the ethicist should "turn on the lights" for the subject therefore fails to provide reasons to think that the ethicist should attribute moral obligations.

A second option is to adopt fictionalism, as others have in reaction to error theory.[18] Fictionalism is a way of maintaining current practices, which in the case of biomedical ethics is offering moral obligation attributions, but recognizing that these practices are not grounded in moral truth. There are supposed to be some practical benefits, the most significant of which is that it requires no revision of a common practice.

If attributing moral obligations in biomedical ethics is groundless, yet we maintain the fiction that those attributions are deliberately accurate, we can maintain the practical benefits that come with the practice. One such benefit of pretending that our obligation attributions are deliberately accurate is that we can continue to justify the presence of ethicists in hospitals, on research teams, in corporations, and the continued expansion of the profession. It's much harder to justify the United States' Joint Commission requirement that there be a mechanism and personnel to resolve ethical conflicts in health care facilities if the people doing that are merely pretending that their recommendations (i.e., moral obligation attributions) are accurate.

[18] Joyce, R. (2001). The myth of morality. Cambridge: Cambridge University Press; Joyce, R. (2005). Moral fictionalism. In M. E. Kalderon (Ed.), Fictionalism in metaphysics. Oxford: Oxford University Press. Joyce, Kalderon, M. E. (2005). Moral fictionalism. Oxford: Oxford University Press.

However, pretending that our moral obligation attributions are deliberately accurate gives up on the goal getting things right, of saying true things. It would require the ethicist to attribute moral obligations to others knowing that they have little justification to do so. Perhaps maintaining the fiction and the associated practical benefits is worth it. Is it ethical to tell someone they should do something knowing that telling them that is nothing more than a fictional practice meant to keep oneself employed? We can't answer that question without violating the main claim of this paper—we can't say on ethical grounds that ethicists should refrain from attributing moral obligations without violating our main claim. But to say that the practice fails to secure ethicists' practical interest in telling the truth is not off-limits. Nor is it off-limits to point out the fact that it would appear rather ironic for ethicists to maintain a fiction for the purposes of maintaining their influence and salaries.

There may be additional costs of fictionalism, namely the violation of moral and epistemic norms. W.K. Clifford is well known for claiming that believing anything—and by extension testifying—without sufficient evidence is not only a violation of an epistemic norm, but also a moral one.[19] His example is of a shipowner who sends a ship out without proper evidence for its seaworthiness. It sinks; everyone dies. The shipowner's epistemic norm violation was also a moral one. But people don't have to die for the epistemic norm violation to give rise to the moral norm violation. Inconsequential claims held without sufficient evidence are also morally wrong. If, without evidence, I believe I have a penny in my pocket and testify to that, I hold myself out as reliable. But I am not, so to pretend like I'm reliable is a lie. The same may be true of ethicists attributing moral obligations: we hold ourselves out as reliable in doing so, but actually we perform the practice without evidence sufficient for its justification. If Clifford is right, this is a moral violation as well as an epistemic one.

Also among the epistemic norms violated by continuing to attribute moral obligations is the norm of assertion. It is a controversial matter, but some authors claim that knowledge is the norm of assertion, which means, at least, that one ought to only assert

[19] Clifford, W. K. (1877). The Ethics of Belief. In *The Ethics of Belief and Other Essays*. Amherst, New York: Prometheus Books.

propositions that one knows. If knowledge is the norm of assertion, and epistemic properties are logically prior to moral properties, then it is straightforward that ethicists' moral obligation attributions violate the norm of assertion. If in the hospital an ethicist says, "The attending physician should change the patient's code status so that CPR is withheld," (where the 'should' is moral), they have violated the norm of assertion, because the ethicist doesn't know the physician's epistemic burdens related to withholding CPR in that patient, and thus doesn't know that that is what the attending physician should do. To make such an assertion—an attribution of a moral obligation—thus violates the norm of assertion, if knowledge is the norm of assertion.

Fictionalism is likely the easiest thing to do in response to the "Now what?" question. It is far less intrusive to the current practice and preserves our professional roles and the associated benefits. It's simply the continuation of what we've been doing all along. But it comes with a lot of moral and epistemic baggage. And it undermines the pursuit of the ethicist's practical interest in being accurate.

Manipulating a subject's epistemic burdens so that a subsequent moral obligation attribution hits the target is trivial. Maintaining the fiction of accuracy comes at a high moral, epistemic, and prudential cost. If manipulation of the circumstances affecting accuracy isn't an option, and maintaining the current practice isn't an option, then significant revision of the practice of attributing moral obligations is in order. This revision can occur in two ways. One way is to abandon the commitment to accuracy. We assume that abandoning this commitment is unacceptable for biomedical ethicists individually and the profession generally. The only other option, then, is to abandon the practice of attributing moral obligations. If biomedical ethicists' commitment is to accuracy, then it is in our practical interests to abolish attributions of moral obligations. The argument for abolitionism thus occurs by process of elimination.

6. Abolitionism

Abolitionism carries numerous benefits.[20] First, it allows ethicists to maintain a commitment to accuracy in their recommendations, by eliminating the sources of inaccuracy. Second, it allows ethicists to maintain a commitment to the claim that moral claims have a truth-value and that some of those claims are true. Third, and perhaps most importantly, it forces the ethicist to focus on what is perhaps their most important function and professional role: to identify reasons, evidence, and help others weigh these alongside a wide range of values. While it may be inaccurate to attribute to a physician the obligation to withhold CPR, there is nothing inaccurate about bringing to their attention the patient's advance directive, or in facilitating a discussion with the patient's family about what the patient's preferences are. It may be inaccurate to claim that in a resource shortage a hospital should prioritize "frontline" health care workers; but it is not inaccurate to point out that the continued functioning of the hospital requires a sufficiently large labor pool or that reciprocity for their sacrifice has moral value. Any use of 'should' or 'ought' ought to be merely prudential or epistemic. The second 'ought' in the previous sentence is practical: satisfying that prescription maintains the commitments to accuracy while preserving the value of well trained and knowledgeable ethicist.

Our claim has no implications for what the metaphysical status is of moral claims, or for which moral theory ends up being true. Instead, our claim is that so long as biomedical ethicists continue to (a) presume that moral claims can be true (again, regardless of whether they are true only relatively) and (b) aim to say true things themselves, they should be epistemically humble in their attributing obligations to others. They should be so epistemically humble that they stop doing so.

One might recognize that our argument doesn't uniquely apply to biomedical ethics: perhaps it's true *no one* should be making obligation attributions. But here is where the argument does apply, especially to biomedical ethicists. First, the field is small enough,

[20] Hinckfuss, I. (1987). The moral society: Its structure and effects (Vol. 16). Canberra: Australian National University, Departments of Philosophy; Garner, R (2007). Abolishing Morality. *Ethical Theory and Moral Practice*, *10*, 499–513; Ingram, S. (2015). After moral error theory, after moral realism. *Southern Journal of Philosophy*, *53*(2), 227–248; Dockstader, J. (2019). Nonassertive Moral Abolitionism. *Metaphilosophy*, *50*(4), 481–502; Schwarz, L. (2020). Error Theory and Abolitionist Ethics. *Southern Journal of Philosophy*, *58*(3), 431–455.

and influenced by few enough people, that abolitionism *could* take hold. It could not take hold in society, generally, and this likely impossibility blocks the argument for abolitionism among everyone. Second, one of the primary functions of biomedical ethicists is to make obligation attributions, which makes our argument especially relevant.

Abolitionism is an answer to the "Now what?" that arises from the recognition that moral obligation attributions are likely to be inaccurate. But one might ask the question again. Now that moral obligation attributions ought to be abolished from biomedical ethics, what happens? Our claim is that although the practice of attributing moral obligations is abolished, the proper role for ethicists is to bring to light reasons and evidence and help others weigh and balance these. Rather than throwing darts in the dark, the ethicist should turn on the lights and stop there.

7. Objections

One might object that this is nothing new. We've known for a rather long time that we have imperfect access to others' mental states and we've never taken that imperfect access to undermine our own ability to attribute obligations to others.

This objection is tantamount to the claim that the logical priority of the epistemic adds nothing to the argument. But the logical priority of the epistemic is key. The reason it's important is that it entails that epistemic properties fix moral properties. Specifically, a person's ignorance and the epistemic burdens to overcome it determine what they are not obligated to do, and in turn determine what they are obligated to do. As this ignorance and these epistemic burdens shift, so do one's moral obligations. Their obligations are a moving target, shifting with the person's knowledge and ignorance. You can't get that consequence out of garden-variety problems related to other minds.

Another objection is that epistemic properties are not logically prior to moral properties—the epistemic is *not* logically prior to the moral.

We provided above a short summary of the logical priority of the epistemic. In other journals, we have defended the logical priority of the epistemic, both from the armchair as well as the lab. The theoretical support for the logical priority of the epistemic is

extremely strong, as is the empirical research. This is not the venue to re-litigate those discussions. Suffice it to say that the available evidence for the logical priority of the epistemic is strong. But one reminder might be helpful: denying the logical priority of the epistemic requires denying that ought implies can, or whatever derivation ends up being true. Denying the logical priority of the epistemic requires denying that one's obligations are largely a matter of one's abilities.

Abolitionism in biomedical ethics is revisionist, but the recommendation to turn on the lights and stop short of attributing moral obligations doesn't change the practice so much, and requires very little actual revision in behavior. The revision is primarily in virtue of how we all think about and justify our behavior. Some might thus think that the recommendation amounts to a triviality, a distinction without a difference. While we disagree that it's a trivial recommendation to abolish moral obligation attributions from biomedical ethics, if it's not revisionist enough for some readers, we invite them to consider this high standard for being sufficiently revisionist and reflect upon the fact that almost every single biomedical ethics article ever published will fail to meet it.

Some objections attack the substance of the argument, such as by asserting that we do know others' epistemic burdens quite well. We know well others' minds, including what it is they know.

In response, first, that's plainly false. We don't know what knowledge others possess. This is perhaps especially true in biomedical ethics, where the subject is, relative to the ethicist and ethics, in an epistemically impoverished position. But, second, and more important for the present argument, we don't know what others are ignorant of. It's the knowledge one doesn't possess, one's ignorance, that constitutes epistemic burdens and shapes one's options and their relative weights. We don't know what others don't know. We can hardly identify our own gaps in knowledge, let alone others' gaps.

Claiming that we do have sufficient knowledge of others' epistemic burdens doesn't defeat the present argument. Neither does the claim that many of our obligation attributions are true, or accurate. It doesn't work because, first, to offer such an assertion one would need to know what the subject's moral obligations are, and to do this one must also know what the subject is ignorant of. Second, if they are accurate,

they are accurate by luck. If this is what ethicists are relying on, when deciding which obligation to attribute, it might be better to flip a coin.

A potential complication of the view we have offered here is that abolitionism is compatible, and may even endorse, a view that many ethicists hold: the aim of biomedical ethics is to broker consensus among interested parties, not to attribute moral obligations.[21]

However, this isn't an unwelcome relation, though we are highly skeptical of the value of consensus regarding ethical claims. One person is likely to be wrong about what a person morally ought to do. Getting a bunch of people who are just as likely to be wrong to come to a consensus does not seem to us a good way of increasing accuracy. Rather, it almost guarantees inaccuracy. But perhaps consensus is not about accuracy. If that's true, then the present argument and the view that ethics is about consensus are orthogonal to each other, because they don't share in the foundational assumption that accuracy in moral claims is a goal.

8. Summary

Moral obligation attributions should be abolished from the practice of biomedical ethics. Supporting this claim are the observations that the practice of biomedical ethics maintains a commitment to moral realism and to saying accurate things. Combined with the logical priority of the epistemic, biomedical ethicists are unlikely to deliberately make accurate obligation attributions. Typically, such attributions will only be accurate by luck. Thus, the practice of attributing moral obligations to subjects whose ignorance we are ignorant of should be abolished. This leaves the ethicist to only illuminate the relevant reasons and evidence and help subjects weigh and balance these. It is impractical for the ethicist to attribute moral obligations. On this basis they should not do so.

[21] See note 5 (Brummett and Salter, 2019) for a taxonomy of such views.