

ARTICLE

The misplaced, if not erroneous, nature of many obligation attributions

Scott Scheall¹  and Parker Crutchfield²

¹Center for Economics, Politics, and History, University of Austin, Austin, TX, USA and ²Western Michigan University Homer Stryker MD School of Medicine, Kalamazoo, MI, USA

Corresponding author: Scheall Scott; Email: sscheall@uauustin.org

(Received 23 October 2024; revised 21 January 2025; accepted 3 March 2025)

Abstract

In the present paper, we consider the implications of our work on the *logical priority of the epistemic*, the thesis that persons' options are determined in the first instance by their relevant knowledge and ignorance, for the legitimacy of claims that some decision-maker bear an unconditional obligation to make a particular decision or perform a specific action (i.e., *categorical obligation attributions*). We argue that the logical priority of the epistemic implies that many such attributions are misplaced, if not erroneous. We defend *Naturalistic Moral Error Theory*, a novel theory which, though it bears many of the same consequences for moral discourse, is distinct in its metaphysical implications from J. L. Mackie's (1977) famous moral error theory.

Keywords: Logical priority of the epistemic; ignorance; epistemic burdens; moral error theory; normative judgments

1. The logical priority of the epistemic

For an actor *A* to be obligated to do *X* requires, at a minimum, that doing *X* is something that *A* can choose to do. In other words, doing *X* must be an *option* for *A* if *A* is to be properly obligated to do *X*.

Two important questions, then, are *Where do options come from?* and *What determines the rank ordering of options in a decision-maker's initial conscious preference ranking?* What determines that one course of action is an option for *A* and another course of action that might be attempted under given circumstances is not an option – and, thus, not a potential obligation – for *A*? Further, what determines that one option *X*₁ is ranked higher than another option *X*₂ when a decision-maker initially consciously reflects on their preferences in some decision context? Practitioners in the decision-theoretical disciplines tend to consider only those options encompassed in *given* preference rankings, as if options and their rank orderings were brute facts not determined by other considerations.

We have argued that, at least in the first instance, options are pre-consciously determined and ordered according to the actor's relevant *epistemic burdens*, i.e., the nature and extent of the actor's ignorance regarding courses of action that might be

attempted in the given decision context (Scheall and Crutchfield 2021a).¹ Some pre-conscious epistemic activity must occur for an actor to have an ordered incentive structure of various options from which to consciously choose.² At least in the first instance, relevant epistemic burdens determine options and, more to the present point, potential obligations.³

We have offered two philosophical arguments for the thesis of the logical priority of the epistemic. The first is a simple introspective argument. The second considers the meaning of the word “can” in generally accepted principles like *ought implies can* (Scheall and Crutchfield 2021a).

The introspective argument begins from the fact that the options that decision-makers consciously consider rarely, if ever, encompass every course of action that might be attempted in the given decision context, but only actions that the decision-maker understands themselves to know (or be able to learn) well enough to pursue with some effectiveness, which is to say, as a means of achieving the end(s) at which they aim.⁴

One obvious example of this phenomenon occurs when a person considers various means of traveling from one place to another. No sane person consciously considers *teleportation* or *flying like a bird without mechanical assistance* to be travel options. If we are right about the priority of the epistemic, then it is the nature and extent of the relevant ignorance, including incapacities – the relevant epistemic burdens – which decision-makers recognize, if only pre-consciously, they would have to overcome in order to effectively teleport or fly like a bird that explains why these courses of action do not appear in (human) decision-makers menus of travel options.

More generally, we argue that introspection reveals that the options that decision-makers consciously consider in any given decision context have been, at least roughly, pre-consciously sorted and ranked according to their relative epistemic burdens (Scheall and Crutchfield 2021a). Courses of action that decision-makers take themselves to be too ignorant to realize do not appear in their menus of options, while courses of action with respect to which they take themselves to be comparatively ignorant are typically ranked below options with respect to which they take themselves to be relatively knowledgeable.

Some pre-conscious assessment of one’s epistemic circumstances must occur before a decision-maker can consciously consider an array of options and make a choice. Moral, prudential, pecuniary, and any other non-epistemic normative considerations

¹We understand *knowledge* (and related concepts like *ignorance*) to encompass both *propositional knowledge-that* and *non-propositional knowledge-how*, and, moreover, to include both *explicit* and *implicit* (or *tacit*) knowledge (on tacit knowledge, see Polanyi 1958 and 1966). In other words, knowledge includes not merely knowledge of facts and theories, but also talents, capabilities, capacities, and abilities, and both knowledge that the epistemic subject can explicitly state and knowledge that the actor possesses without being able to recite.

²We use the phrases “incentive structure,” “preference ranking,” and “menu of options” interchangeably to mean the array of options that a decision-maker consciously considers and from which they ultimately choose in some decision context.

³We agree with any readers inclined to think that the logical priority of the epistemic is an obvious fact of life. However, its conspicuous nature does not mean it is either unoriginal or without interesting implications, such as those we try to explicate here. Indeed, we tend to think that the best philosophy often consists of elucidations of previously unexamined considerations eventually revealed upon examination to have always been “obvious.”

There is a burgeoning literature on the significance of ignorance for decision-making. Beyond our work, see, e.g., Karni and Viero 2013, 2015, 2017, Bradley 2017, Piermont 2017, Walker and Dietz 2018, Roussos 2020, Mahtani 2021, Steele and Stefánsson 2020, 2021a, 2021b.

⁴Including any relevant subsidiary ends to be adopted as intermediate means of achieving their ultimate goal.

enter the decision-making picture only after (logically and temporally) some pre-conscious epistemic work has occurred. Epistemic burdens are logically fundamental in decision-making.

Our second argument for the logical priority of the epistemic is that the presumed relation between *ought* and *can* – whatever its logical strength, i.e., whether it is *logical implication*, *presupposing*, *making plausible*, *conversationally implicating*, or some other relation that ultimately holds between *ought* and *can* – makes sense only if the epistemic is logically prior to other normative considerations.⁵ If knowledge encompasses both the propositional and non-propositional, then the word “can” is inherently epistemic. There is no non-epistemic meaning of “can.”⁶ What a person can do is necessarily a function of what they know *of* and know *how to do*. Thus, *ought implies* (or whatever) *knows enough to*.⁷ Anyone committed to the notion that some relation obtains between *ought* and *can* is also committed to the logical priority of the epistemic. Indeed, in the last analysis, *ought implies* (or whatever) *can* may simply mean that, in decision-making, knowledge is prior to normativity.

It is important for what follows to explicate the logic of principles like *ought implies can*. In particular, it is important to recognize that the contrapositive of *ought implies* (or whatever) *can* is *cannot implies* (or whatever) *not ought*. The relevant contrapositive is not, as an uncareful logician might think, *cannot implies ought not*. In other words, if the subject of an obligation attribution cannot perform the allegedly obligated action, then they are not obligated to perform that action; the purported obligation evaporates. It is not the case, however, that an obligation to *not* do the action (an “ought not”) emerges in its place. Decision-makers who are sufficiently ignorant or incapable with regard to *X* are not obligated to *X*, but this does not mean that they are obligated to *not-X*.

The thesis that epistemic considerations are logically prior to normative ones, generally, and to moral considerations, specifically, is also supported by the results of a recent series of psychological experiments in which research participants’ responses to various moral dilemmas under conditions of relevant ignorance were recorded (Crutchfield, et al. 2023). We observed in these experiments that the nature and extent of the ignorance associated with different options was the most significant factor affecting participants’ choices in a trolley problem context. Further experiments have also supported the logical priority of the epistemic. Most important for the arguments to follow, we have found in these subsequent experiments that research participants who are provided with information about ignorant decision-makers relevant epistemic burdens tend to refrain from attributing obligations to the latter, obligations that they readily assign to more relevantly knowledgeable actors (Scheall and Crutchfield Manuscript). These results provide empirical evidence that obligation attributions critically depend upon what would-be obligation attributors know about decision-makers’ relevant epistemic burdens. They show, in other words, that obligation attributions are regulated by the logical priority of the epistemic.

⁵The principle that *ought implies can* is widely accepted, but some moral philosophers accept only that a weaker logical relation than *implication* obtains between *ought* and *can*. See, e.g., Hampshire, 1951; Hare, 1951, 1963; Sinnott-Armstrong, 1984; Mizrahi, 2015; Littlejohn, 2009, Vallentyne, 1989; Vogelstein, 2012.

⁶In our earliest work on the logical priority of the epistemic (Scheall and Crutchfield 2021a), this second argument was slightly different. The earlier, superseded, argument is an implication of the later argument offered in the present text, so we have not rejected the first argument, as much as developed a more general argument for the logical priority of the epistemic.

⁷Naturally, if no relation obtains between *ought* and *can*, then this argument is moot, though the argument from introspection would remain.

2. Epistemic burdens and their implications for obligation attributions

Many normative judgments are systematically misplaced, if not entirely erroneous. More carefully, normative judgments that attribute obligations to decision-makers to either pursue ends or adopt means that a purportedly obligated decision-maker cannot pursue or adopt, because of their deficient epistemic circumstances, are systematically misplaced, if not erroneous, to the extent that such judgments are not qualified to account for these epistemic deficiencies.⁸

To be properly obligated to do *X* is (perhaps *inter alia*) to be normatively required to cause *X* to occur. If one is obligated to, say, vote for Candidate *C*, this means (again, perhaps among other things) that the decision-maker is normatively required to intervene in the relevant causal nexus in such a way that an additional vote – their own vote – for Candidate *C* is tallied. However, in order to cause their own vote for Candidate *C*, the decision-maker has to occupy an adequate epistemic position.

Epistemic burdens have several components. There are various respects in which decision-makers might be ignorant. Decision-makers might be ignorant of potential goals relevant in their decision context. For instance, a decision-maker may be ignorant of the election in which *C* is a candidate and, thus, not know that voting for Candidate *C* is a potential goal. Alternatively, decision-makers might know of particular ends, but be ignorant either that there are causal processes that lead to or how to manipulate these processes well enough to realize some end. A decision-maker might not know that there are causal processes that would produce their vote for Candidate *C* or maybe ignorant of how to intervene in these processes such that their vote for Candidate *C* is tallied.

A goal regarding which the relevant decision-maker is to some extent ignorant can be realized only if spontaneous forces that are no part of the decision-makers deliberations (e.g., luck, providence, serendipity, the “invisible hand”) intervene to mitigate the consequences of their ignorance. Decision-makers who confront an epistemic burden with respect to voting for Candidate *C* – i.e., who are ignorant to some degree of either the relevant election, pertinent causal processes, or how to manipulate these causal processes – can cause an additional vote for Candidate *C* to be tallied only if spontaneous forces beyond their ken and control compensate for the effects of their relevant ignorance. Thus, any unconditional obligation attribution to vote for Candidate *C* that fails to account for the target decision-makers limited knowledge regarding and control over relevant causal processes is inappropriate—at least misplaced, if not entirely erroneous. However, the conditional judgments “To the extent your knowledge permits you to vote for Candidate *C*, you ought to try to vote for Candidate *C*,” and “Provided that your epistemic condition satisfies the requirements of causing a vote for Candidate *C*, you ought to vote for Candidate *C*” are neither misplaced nor erroneous *per se*. In general, qualified obligation attributions such as, “If you can, you ought to vote for Candidate *C*,” are not necessarily misplaced or erroneous, at least, not for epistemic reasons.

The extent to which you can be *potentially* unconditionally obligated to do *X* and the extent to which an unconditional attribution of an obligation to you to do *X* is *well-placed*, is primarily a function of the relevant knowledge you possess about both *X* and causal processes that generate *X*, and, secondarily, a function of other, non-epistemic, circumstances. Inversely, the extent to which an unqualified attribution of an obligation to you to do *X* is *misplaced* is primarily a function of your ignorance concerning either *X* or causal processes that generate *X*. If your knowledge in this respect is adequate, then

⁸We are primarily concerned with second-person judgments that attribute obligations to decision-makers other than the attributor, though the kind of misplacement we point to can, at least in principle, arise with respect to self-attributed obligations.

doing *X* is a potential unconditional obligation of yours; an attribution to you of an unconditional obligation to do *X* is neither misplaced nor erroneous *per se*, at least, not due to your relevant ignorance, though it might be misplaced for non-epistemic reasons.

If your knowledge of either *X* or the causal processes that lead to *X* is limited, but not entirely lacking, then you can be obligated to do *X* only up to the extent of your relevant knowledge, i.e., only up to the extent that you do not need to rely on luck, fortune, divine intervention, or other spontaneous forces to make *X* occur. In this case, an attribution to you of a categorical obligation to do *X* is misplaced to the extent that you cannot do *X* on your own initiative, which is to say, to the extent that spontaneous forces are required to intervene to realize *X*. The misplacement here is a matter of attributing to you an unconditional obligation to cause something to happen that you can only, at best, conditionally make happen, that you can only make a partial and, by itself, inadequate contribution to causing to happen.

If your knowledge of either *X* or relevant causal processes is entirely lacking, then you cannot be obligated to do *X*; an attribution to you of an unqualified obligation to do *X* is entirely erroneous. The error here is due to attributing to you an obligation to cause something to happen that you can make no contribution to cause to happen.

Consider the paradigmatic example of a purportedly categorical obligation, i.e., that a person *ought never tell a lie*. In a world where everyone infallibly knows the truth values of every proposition and can express their beliefs, this is obviously a potential (and, if an ethical theory like Kant's is sound, an actual) categorical obligation. Everyone in such a world can avoid telling lies without the help of luck or any other spontaneous force. However, in a world more like our own, where propositions are rarely, if ever, known infallibly (even setting aside Cartesian concerns), some degree of spontaneous assistance is often required to avoid asserting untruths. A person who, in 1869, asserted that "Humans will someday walk on the moon," spoke a truth, but only to some extent fortuitously, not entirely deliberately. Had this same person asserted the negation of this proposition, they would have told a lie, but it is hard to see how they would have been infringing a well-placed obligation attribution to categorically avoid telling lies. Imagine a world like, but even more dystopian than, that described by George Orwell in *1984*, where the totalitarian government of Oceania has managed to implant in its citizens' voice boxes devices that transmute state-rejected truths into state-approved lies. In such a world, one cannot speak against the intentions of the government of Oceania without telling a lie. Attribution of a categorical obligation to not tell lies to a skeptic about the Oceanian state is clearly erroneous. However, what seems unproblematic in all of these cases is the conditional obligation attribution, "To the extent you can both distinguish truth from falsity and express your beliefs, you ought never tell a lie."

When moral spectators attribute obligations, they presumably aim to say things that are true. However, the truth of an obligation attribution depends (*inter alia*) on the relevant actor's epistemic circumstances. An attribution of an unconditional obligation systematically fails to be true when the relevant actor's epistemic condition fails to meet the requirements of the thing they are allegedly obligated to do, though it may be partially true up to the extent of the actor's relevant knowledge.

To adopt a helpful simile, think of obligation attributors as dart players throwing darts at a dartboard. Dart players aim to strike the bullseye, but, if not the bullseye, at least, the dartboard, and never, in any case, the wall, liquor bottles, or a passing person. Like a well-placed dart-throw, a well-placed obligation attribution strikes its target square in the middle, while a misplaced, but not entirely erroneous, obligation attribution, though it misses the bullseye, at least strikes the dartboard. An entirely erroneous attribution misses the dartboard altogether, hits the wall, the liquor bottles, the bartender, a cocktail server, or another bar patron. The best obligation attributors,

like the best dart-throwers, are those who regularly strike the (moral) bullseye, i.e., those who either make only qualified obligation attributions, or are adequately familiar with others' epistemic circumstances to attribute unconditional "ought" claims only to those whose relevant knowledge is adequate to the would-be obligation. Those who are ignorant of others' epistemic circumstances, but nevertheless attribute unqualified obligations to other persons, are a nuisance to the other persons at the bar.

3. Arguments from surrogate cases

The priority of the epistemic is most apparent in cases of surrogate decision-making, where one person decides on behalf of another person. These cases also display the misplaced nature of many widely asserted unconditional obligation attributions in stark relief. In many contexts of surrogate decision-making, e.g., in medical, legal, and political contexts, where one person, such as a physician or family member, lawyer, or politician, must decide on behalf of a patient, defendant, or constituent, it is frequently claimed that the surrogate has an unconditional obligation to pursue the surrogated person's best interests or, what might be the same thing, to decide as the surrogated person would decide, if they were in a position to do so.⁹ The priority of the epistemic, however, reveals the quixotic nature of many of these categorical obligation attributions (Crutchfield and Scheall 2019).

A physician placed in the unenviable position of deciding on behalf of an incapacitated person might *want* nothing more than to meet the commonly attributed categorical obligation to decide in the incapacitated patient's interests or as the patient would decide. Indeed, we can stipulate for the sake of the example that one might be the most ethical physician to ever live, a walking Hippocratic archetype. However, if this ethically perfect physician were ignorant either of the incapacitated patient's interests or of the decisions the patient would make, or, if the physician knew of the patient's interests or would-be decisions, but was ignorant of means to pursue these interests or to effect these decisions, it is apparent that – for purely epistemic reasons unrelated to their moral probity – the physician could not purposefully satisfy the obligation categorically attributed to them. Under such severely impoverished epistemic circumstances, courses of action related to the attribution would not even be options, much less legitimate obligations, for our Hippocratic epitome.

A judgment that a relevantly ignorant physician has an unconditional obligation to decide in an incapacitated patient's interests or as the patient would decide is, depending on the weight of the physician's epistemic burdens, at least misplaced, if not entirely erroneous. Whatever, if anything, a medical surrogate *is* obligated to do when epistemically deficient, it cannot be what standard medical-ethical traditions unqualifiedly claim it to be (Crutchfield and Scheall 2019, 2024).

Unconditional obligations to pursue their constituents' interests are frequently attributed to policymakers. However, in many relevant contexts, policymakers have, at

⁹An anonymous reviewer encouraged us to emphasize that surrogate decision-making need not involve an expert deciding on behalf of a non-expert. The fact mentioned in the text above that a family member might serve as a surrogate would seem to imply the possibility of non-expert surrogates. That said, in many surrogate contexts, e.g., medical contexts, an untrained family member may well know the patient's interests better than the attending physician. Who is the real "expert" in such circumstances? We prefer to remain silent on this question. In any case, the argument from surrogate decision-making does not hinge on the decision-maker being an expert in any sense of the word. As discussed in the closing section of the paper, even experts can fail or be otherwise limited in their expertise, a possibility of which those inclined to attribute obligations to experts should keep in mind, to the extent they can learn of the limits of expertise.

best, a vague hint of their constituents' interests or, what might be the same thing at least some of the time, of the policy goals their constituents most demand (DeCanio 2014). Even if they know of their constituents' interests or policy preferences, policymakers may be ignorant of means adequate for their realization (Scheall 2019, 2020). In any case, policymakers rarely, if ever, possess sufficient control of the causal levers of society to deliberately realize goals either in their constituents' interests or demanded by constituents. Policymakers often have to rely significantly on the intervention of spontaneous forces to promote their constituents' interests or policy preferences.

To the extent that a particular policy goal is beyond the ken and control of policymakers, inasmuch as it cannot be realized through their own initiative, but requires the intervention of spontaneous forces, it is at least misplaced and, in the extreme case where policymakers are completely ignorant or incapable with respect to some potential policy objective, entirely erroneous, to attribute a categorical obligation to policymakers to pursue the relevant goal. A judgment that policymakers are unconditionally obligated to pursue policy objectives in their constituents' interests, which fails to account for policymakers' limited knowledge and practical capacities, is at least misplaced and, in the extreme case of policymakers' complete ignorance, entirely erroneous. Again, whatever policymakers are obligated to do when they are relevantly ignorant, it cannot be what the common conception of political ethics claims it to be (Scheall 2019, 2020).

However, surrogate cases, although they display the implications of the priority of the epistemic in particularly stark form, are not different in kind from cases of personal decision-making, where persons decide for themselves. The priority of the epistemic operates in the latter cases as in surrogate cases. One relevant difference, of course, is that the gap between the surrogated person's interests and the surrogate's knowledge of these interests is narrower, if not closed entirely, in cases of personal decision-making. In such cases, in effect, the surrogate and the surrogated person are identical. Assuming that a person has better, if perhaps still fallible, knowledge of their own interests, it follows that related epistemic burdens tend to be lighter in cases of personal decision-making. Another relevant difference is that, *given* knowledge of their own interests, a person often possesses better knowledge than any other party might of ends relevant to and of means appropriate for pursuing ends relevant to these interests. In virtue of an interest in, say, professional ice hockey, a person is likely to be in a better position than most potential surrogates who lack this interest to know of and realize objectives associated with it, e.g., discovering game scores, learning about transactions between teams, tracking current standings. However, even a person who knows their own interests perfectly well may be ignorant of either ends associated with or means adequate to realize goals associated with these interests. It is thus possible for self-ignorant persons to attribute misplaced or erroneous obligations to themselves.¹⁰

4. Impossibly heavy epistemic burdens

If a decision-maker, *A*, is utterly ignorant of either *X* or causal processes that produce *X*, then *A* cannot be obligated to do *X*. If some spectator, *S*, nevertheless judges that *A* has an unqualified obligation to do *X*, then *S*'s judgment is in error.

This error might be brought home to *S* by first asking them to place themselves in *A*'s completely impoverished (with respect to *X*) epistemic circumstances and then to consider whether, *under these utterly impoverished epistemic circumstances*, they would consider *themselves* categorically obligated to do *X* – something that will not even appear in their preference rankings as an option under the circumstances. As we put the point

¹⁰The analysis of these cases can be subsumed under the analysis of second-party cases.

elsewhere, “It is typically easier to be the person moralizing than the person obligated. We suspect that someone who claims of another person that they are obligated to do something that they cannot do would be more comfortable as the person moralizing than as the person obligated” (Scheall and Crutchfield 2021a).

Judgments of the sort that concern us silently assume that the epistemic circumstances of the person purportedly obligated are different than they are, in particular, they assume that the alleged obligor already possesses or is able to acquire the knowledge required to do the thing they are supposedly obligated to do. Where this assumption does not hold, where *A* cannot acquire the knowledge required to do *X*, doing *X* is not an option for *A*, and any judgment that *A* ought to do *X* is plainly in error.

Imagine an eight-year-old child grabbing a toy from her infant brother with the stern admonition, “You should share!” thereby leaving the infant in tears. The older child’s judgment is clearly erroneous. The infant’s epistemic burdens with respect to sharing are impossibly heavy for the infant to bear. He knows neither of sharing nor how to do it. A parent might reasonably say to the older child, “Give that toy back to your brother. He has no obligation to share with you. When he is older, we will teach him about sharing and then you can scold him when he does not share.” Attributing obligations to actors who confront impossibly heavy epistemic burdens is like assigning obligations to infants.

Or consider a rather more complicated and controversial example, the contemporary phenomenon of “cancellation,” wherein, for having committed some perceived social sin or *faux pas* – often only recently pronounced as such – a person is subjected to some combination of public rebuke, online humiliation, and open mockery, if not loss of friends, family, and career. We suspect that part of what many people find offensive about certain instances of the cancellation phenomenon is precisely that the targets of such attacks did not know – indeed, perhaps could not have known – that their actions might possibly offend anyone. In a world of rapidly and radically shifting moral standards, and of ever more sensitive moral sensibilities, the knowledge required to deliberately avoid cancellation is increasingly difficult to acquire; one might need considerable assistance from spontaneous forces to avoid cancellation.¹¹

5. Weighty (but not impossibly heavy) epistemic burdens

To the extent that *A* cannot do *X* or, what is the same thing, to the extent that doing *X* requires the intervention of luck, fortune, or other spontaneous forces beyond *A*’s ken and control, it is at least misplaced, if not entirely erroneous, to judge that *A* has an unconditional obligation to do *X*. Where *A* is to some extent ignorant of knowledge required to do *X*, *X* can be done only to the extent that spontaneous forces intervene to mitigate the consequences of *A*’s ignorance. Without these spontaneous forces beyond *A*’s ken and control, *A* cannot do *X*; to the extent that *A* cannot do *X*, to the extent that *A* must rely on spontaneous forces for the doing of *X*, *A* cannot be categorically obligated to do *X*. The limits of *A*’s relevant knowledge with respect to doing *X* are the limits of *A*’s causal responsibility and the limits of *A*’s potential obligation to do *X*. If *S* nevertheless judges that *A* has an unconditional obligation to do *X*, then *S*’s judgment, though it might strike the dartboard elsewhere, misses the bullseye.

Again, this misplacement might be brought home to *S* by asking them to place themselves in *A*’s relatively impoverished epistemic circumstances and then asking whether, *under these comparatively impoverished circumstances*, they would still consider themselves categorically obligated to do *X* or whether the sheer weight of the

¹¹We thank an anonymous reviewer of a previous version of this paper for encouraging us to consider the significance of our argument for the cancellation phenomenon.

epistemic burden associated with *X* might not necessitate a choice other than doing *X*. If *S* could imagine *A*'s relative epistemic burdens, in particular, the fact that other courses of action appear to require less spontaneous intervention than does *X*, the spectator would at least qualify, if not withdraw altogether, their judgment that *A* has an unqualified obligation to do *X*. Even though *X* might be an option for *A*, because *A* is not completely ignorant with respect to it, its epistemic burden compared to those of other courses of action may nevertheless be so weighty as to practically, if not logically or metaphysically, rule it out as a plausible choice for *A*.

For example, a person might have some knowledge of steps that they could take to secure world peace; nevertheless, under normal circumstances in the modern world, the weight of what the typical person does not know that they would need to know in order to deliberately secure world peace on their own initiative is so heavy that, although acting to secure world peace is always in some sense an option for the person – they could at least take the steps they believe conducive to world peace – they rarely, if ever, choose it, *given* the extent of the spontaneous forces that would have to intervene for their purposeful steps to make anything more than a trifling effect on world peace and their more extensive, and more relevant, knowledge with respect to other courses of action. The person might know well how to, say, make coffee, cook breakfast, go to the gym, or take a shower, so they do one or more of these things instead of acting to secure world peace, despite its obviously superior moral value. Such a person would seem justified in looking cross-eyed at anyone who attributed to them a full-throated obligation to act to secure world peace in the circumstances of their daily lives.¹² On the other hand, assuming appropriate non-epistemic circumstances, this person would have little reason to quarrel with the properly qualified judgments, “Provided that your epistemic condition satisfies the requirements of securing world peace, you ought to secure world peace,” or “To the extent your knowledge permits, you ought to try to secure world peace and hope that spontaneous forces intervene to secure world peace inasmuch as you cannot do so on your own initiative.”

In effect, obligation attributions of the relevant sort are disguised commentaries on the actor's epistemic circumstances. They assert *what might be potential obligations were the actor's epistemic circumstances different than they in fact are*. But, because the actor's epistemic condition is impoverished in some way and to some extent with respect to the end they are purportedly obligated to pursue or the means they allegedly ought to adopt, these judgments systematically miss their marks, in part or in whole. Properly conditioned obligation attributions, however, are neither erroneous nor misplaced. “If you know how to effect world peace, you ought to act to secure world peace” might be a potential obligation of a person under appropriate circumstances, say, a representative to the Paris Peace Conference in 1919. Indeed, “You should contribute what you can, on the basis of the relevant knowledge you possess, to secure world peace” might be a potential obligation for everyone under appropriate circumstances. Our obligations are proportionate to our relevant causal knowledge; to be well-placed, obligation attributions must be similarly proportioned. Singer's Principle (1972) that one should prevent something bad from happening so long as one does not have to sacrifice something of comparable moral significance should be amended: One should prevent something bad from happening *to the extent one knows enough to prevent the bad thing from happening* and if doing so does not require the sacrifice of something of comparable moral significance.

¹²Of course, none of this is to say that the person does not have an obligation to *not act to promote international conflict*, an obligation that they might satisfy by, say, making coffee, cooking breakfast, going to the gym, and taking a shower.

6. Weightless epistemic burdens

It is only when A’s epistemic burden with respect to doing X is weightless – it is only when A’s ken and control fully encompass doing X – that doing X is among A’s potential unconditional obligations. Whether X is among A’s *actual* (rather than merely potential) unconditional obligations depends on other, non-epistemic, circumstances. In order for A to be categorically obligated to do X, for example, A must occupy a context in which X has not already been done. A person who possesses all of the knowledge necessary to secure world peace is not actually obligated to secure world peace in a world where peace has already been secured. It is only in a world where some degree of conflict exists that such a person bearing a weightless epistemic burden with respect to securing world peace could be unconditionally obligated to do so.

However, even in such a context, matters are not so simple. Whether, in a world of strife, A is obligated *at any given moment* to act to secure world peace depends on a complex web of non-epistemic considerations. Other considerations (moral, prudential or pecuniary considerations of personal, familial, or societal scope) might swamp A’s potential categorical obligation to do X. It is least debatable, for example, that a person in a violent world, knowledgeable enough to end all conflict, who happens to, say, encounter a child dying of starvation who A knows enough to feed, is unconditionally obligated to act to secure world peace rather than to find food for the hungry child (at least to the extent these actions cannot both be pursued at the same time).

All that can be said about cases in which A’s epistemic burden with respect to doing X is weightless is that related unconditional “ought” claims are not *systematically* misplaced for failing to reflect the epistemically-conditioned nature of well-placed obligation attributions. They may, however, nevertheless miss the bullseye for other reasons, because they either are irrelevant in the actor’s particular context or fail to account for competing obligations. In general, we can say that unconditional obligation attributions are misplaced to the extent that they fail to reflect, in the first place, the alleged obligor’s epistemic circumstances, i.e., their relevant epistemic burdens, and, secondarily, the obligor’s non-epistemic circumstances.

| A’s epistemic burden with respect to doing X | Can A be unconditionally obligated to do X? | S’s judgment that A is unconditionally obligated to do X? |
|--|--|---|
| Impossibly heavy | Never | Erroneous |
| Not weightless, but not impossibly heavy | Only up to the extent of A’s relevant knowledge | Misplaced to the extent that spontaneous forces must intervene for X to occur |
| Weightless | Potentially, under appropriate non-epistemic circumstances | Neither systematically misplaced nor erroneous for epistemic reasons |

7. “Ought to know” claims do not undermine the logical priority of the epistemic

It might be argued that there are things that an actor categorically “ought to know,” that these obligations are logically prior to the things that an actor actually knows, and, thus, that such categorical obligations constitute counterexamples to the logical priority of the epistemic. This objection is a non-starter, however.

The relevant question is, *can* A deliberately learn, or must they rely on spontaneous considerations beyond their ken and control to learn, the things they are allegedly unqualifiedly obligated to know? If A cannot deliberately learn that which they are supposedly unconditionally obligated to know, if there is no course of learning that A might deliberately adopt to acquire the knowledge they are allegedly obligated to possess, then A cannot know these things (without the assistance of spontaneity) and the purported categorical obligation to know these things evaporates. “Ought to know” claims merely push the priority of the epistemic to another level, from the level of *what can be done* to the level of *what can be known*.

Consider again the example of “cancellation.” There are cases in which a person’s cancellation has been justified by the claim that they “should have known better” than to commit the purported social peccadillo. Perhaps there are indeed cases in which a soon-to-be-cancelled person should have recognized beforehand the problematic nature of some contemplated action. But, again, we suspect that part of what bothers many people about some cancellation events is just that the “could have known” requirement was not met in these cases, that the cancelled person could not have known of the problematic nature of the action before performing it, and that it was only in suffering the detrimental consequences of the action that they could have learned it was, or had recently become, censured by society (or, at least, by a “very online,” vocal, and intolerant part of it).

We have observed in a recent series of experiments that subjects tend to revoke “ought to know” claims about relevantly ignorant decision-makers (Scheall and Crutchfield Manuscript). We provided subjects with scenarios in which decision-makers were either capable or incapable of acquiring knowledge of things that one might be inclined to think everyone ought to know – e.g., either a basic fact (“George Washington was the first President of the United States”) or an uncontroversial moral rule (“A person ought not fart at the dinner table”) – and found that experimental subjects “ought to know” judgments depended upon decision-makers’ ability to learn the relevant fact or rule. Subjects tended to refrain from “ought to know” judgments about decision-makers who could not acquire knowledge of a particular proposition.

8. Moral error theories, old and new, metaphysical and naturalistic

Other philosophers, J. L. Mackie (1977) most famously, have argued that many, if not most, categorical normative judgments are erroneous. Mackie’s (1977: 38–42) well-known *argument from queerness* for this conclusion hinges on the allegedly *sui generis* – and, therefore, implausible – nature of objective moral properties and, consequently, of our possible knowledge of such properties, were they to exist.

If there were objective values, then they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe. Correspondingly, if we were aware of them, it would have to be by some special faculty of moral perception or intuition, utterly different from our ordinary ways of knowing everything else (Mackie 1977: 38).

It is true that, like Mackie’s error theory, the priority of the epistemic implies a deflated realm of legitimate normative judgments, but it does not imply, as Mackie’s (1977: 29) view does, that all categorical imperatives are erroneous. The priority of the epistemic implies that the set of categorical obligations is likely of lower cardinality than may be commonly thought, but not that it is empty. Provided that the relevant epistemic

pre-requisites are satisfied by a particular decision-maker, provided that they possess adequate knowledge of the thing they are purportedly obligated to do and sufficient knowledge of, and control over, relevant causal processes, a decision-maker can be unconditionally obligated.

The priority of the epistemic does not imply moral anti-realism, as does Mackie's argument from queerness. The twin facts that epistemic burdens constrain our obligations and, concomitantly, that unconditional "ought" judgments which fail to adequately account for these constraints, are at least misplaced, are entirely consistent with a realist ontology that includes objective moral properties.¹³ What the logical priority of the epistemic implies about these moral properties is not *per se* that they are metaphysically and, thus, epistemologically queer (though they may be) but that, inasmuch as they exist, they can be, and at least occasionally are, misused, misattributed, and infelicitously invoked.

If Mackie's error theory is built, at least in part, on metaphysical premises, the error theory defended in the present paper is built on the logic of everyday discourse about knowledge and normativity, on what this discourse suggests about the implausibility of impossible, cannot-be-done, obligations, and, relatedly, on experimental evidence concerning widespread judgments about this discourse. If Mackie's arguments lead to a metaphysical moral error theory, our arguments inform what might be called *Naturalistic Moral Error Theory* (NMET).

9. Now what? The obligations of obligation attributors

Moral error theories hold that moral discourse – either all claims that invoke objective moral properties, as according to Mackie's theory, or some subset of moral claims, namely, categorical obligation attributions, as according to NMET – is rife with error. If some moral error theory is correct, what is to be done about moral discourse? How, if at all, should current moral discursive practices change, given their tendency to rampant error? In particular, if NMET is sound, how should the practice of attributing categorical obligations change?

One possibility, frequently invoked in discussions of potential responses to the possible truth of Mackie's error theory, is to simply proceed *as if* moral claims were true, even though (*ex hypothesi*) they are not. *Fictionalism* (see Joyce 2001, 2005; Kalderon 2005) maintains at a surface level our current practices of moral discourse, while acknowledging at a deeper level their lack of grounding in moral truth. The benefit of fictionalism is that it requires no drastic revision of, does no damage to, our current practices. For this reason, fictionalism would seem appealing, if Mackie's error theory were true, but this same reason makes it decidedly less attractive if only NMET (but not Mackie's error theory) is true.

Those who attribute categorical obligations to decision-makers who cannot satisfy their purported obligations are, quite frankly, obnoxious. Metaphorically speaking, they are like people who wantonly toss darts around crowded barrooms without regard for the actual placement of the bullseye. Just as we want to hinder such indiscriminate dart-throwers, we also want to limit epistemically ill-informed obligation attributions. Adopting fictionalism in response to NMET, however, would prevent this. Indeed, adopting the fiction that unconditional obligation attributions can be true regardless of the epistemic circumstances of the decision-makers to whom they are attributed ultimately means indulging, instead of inhibiting, those prone to incautious categorical

¹³The logical priority of the epistemic is also neutral regarding competing versions of realist moral ontology (e.g., deontology *versus* consequentialism) and moral epistemology (intuitionism, etc.).

obligation attributions. Fictionalism is not a solution to the problem of unconditional obligation attributions implied by the logical priority of the epistemic, but a gutless acquiescence in the face of it.

Given that the problem with categorical obligation attributions arises from decision-makers ignorance about and incapacity to do many things that might be assigned to them as categorical obligations, one possible solution would be to simply make decision-makers more knowledgeable and capable. If ignorance means that many categorical obligation attributions are misplaced, if not erroneous, then making decision-makers more knowledgeable and capable would move relevant obligation attributions nearer to the bullseye.

Naturally, there are many reasons why we might want to make decision-makers less ignorant. The logical priority of the epistemic provides little, if any, independent warrant for this desire on the margin. If we knew how to make decision-makers more knowledgeable and capable, we would have done so long ago, for reasons having nothing to do with NMET. Unfortunately, we obviously do not know how to deliberately realize this ambition. Of course, given the arguments of the present paper, the fact that we do not know how to effect this would-be solution implies that there can be no categorical obligation to pursue it. We might be conditionally obligated to put what relevant knowledge we possess to work, and desire the enduring and, hopefully, increasing assistance of spontaneous forces, in the ultimate realization of this goal, but we cannot be obligated to pursue it beyond the limits of our limited relevant knowledge.

Rather than making decision-makers more knowledgeable, a more plausible solution to the problem for categorical obligation attributions raised by the logical priority of the epistemic would require moral spectators to be more careful in the categorical obligations they attribute.

This solution is somewhat related, but not identical, to the *abolitionist* response to Mackie's moral error theory (see the essays collected in Joyce and Garner 2019). As a reaction to Mackie's error theory, moral abolitionism would mean doing away with all discourse in which objective moral properties figure. As a response to NMET, however, moral abolitionism would be rather less drastic; it would mean doing away merely with categorical obligation attributions, in particular, epistemically ill-informed categorical obligation attributions.

We have recently defended abolitionism in a particular, limited, domain in which categorical obligation attributions commonly figure, namely, that of medical ethics (Crutchfield and Scheall 2024). However, the several conditions that make abolitionism a plausible solution in bioethical contexts make it rather less plausible as a general solution. Bioethics is a relatively small area compared to that of the entire realm of moral discourse. What might be relatively easy – the abolition of uncared-for obligation attributions – in a small subset would be much more difficult to realize across the entire domain of moral discussion. Moreover, practitioners in the field of medical ethics are both relatively well-informed about metaethical arguments relevant to their practices and uniquely concerned with avoiding moral error. It would be comparatively simple, in bioethics alone, to both raise an alarm about categorical obligation attributions and catalyze the abolitionist response; it would be more difficult to either stoke such worries as a general concern or ensure the abolition of incautious obligation attributions as a universal reaction to the problem implied by the priority of the epistemic.

On these grounds we defend a general solution – or, more exactly, a general mitigation strategy – to the problem that would seem to be, in any case, a necessary precursor to full-fledged abolitionism. This mitigation strategy – which, if we must describe it as an “ism,” we might call moral *cautionism* (though moral *humility* or *thoughtfulness* come to mind as less neologistic alternatives) – would leverage our

arguments concerning the common human desire to avoid either moral hypocrisy or its appearance. As we suggest above, it may be possible to discourage persons from attributing categorical obligations to relevantly ignorant decision-makers by pointing out that they themselves would reject such obligations attributed to them when in similarly deficient epistemic circumstances. As a practical matter, moral cautionism would advise moral spectators to always remember their limited knowledge about other persons' epistemic circumstances and to proceed cautiously in attributing categorical obligations, to opt instead for properly limited conditional obligation attributions.

The question arises, however, whether would-be unconditional obligation attributors can themselves be obligated to refrain from attributing unconditional obligations to actors whose epistemic circumstances they are ignorant of to some degree. If our arguments are sound, the answer to this question depends on whether such a person knows that they are ignorant of the relevant epistemic burdens of the decision-makers to whom they might assign an unconditional obligation. If the obligation attributor knows that they are ignorant of a decision-maker's relevant epistemic burdens, then they ought to refrain, at least up to the limit set by their recognized ignorance. However, if the obligation attributor is ignorant of their own ignorance of the decision-maker's relevant epistemic burdens, but can deliberately repair this ignorance, then they ought to refrain from attributing unqualified obligations to the decision-maker, at least up to the extent of their deliberately remediable ignorance. In either of these cases, would-be obligation attributors have conditional obligations to refrain from attributing unconditional obligations to actors whose epistemic circumstances they know that they are to some degree ignorant. However, if they are ignorant of their own ignorance and cannot deliberately repair any of this ignorance, then they are not obligated to refrain from attributing unconditional obligations to actors whose epistemic circumstances they are ignorant of to some degree.

If this seems a paradoxical result, notice that someone who falsely believes their knowledge of a decision-maker's epistemic circumstances to be adequate to a categorical obligation attribution is like the infant who cannot be obligated to refrain from stealing his sister's toys. Such pretenders to knowledge positively believe they know enough about the relevant decision-maker's specific epistemic position to attribute a categorical obligation. Indeed, not only do they not know better than to misattribute a categorical obligation to a relevantly epistemically-deficient decision-maker, they might even think themselves obligated to attribute such an obligation.

However, it might be argued instead that everyone knows on some level, even if they rarely acknowledge it to themselves, that they are epistemically limited in many respects, perhaps especially, with respect to the contents of other minds, and, therefore, that humility in – cautionism about – attributing unconditional obligations is always an unconditional obligation of all would-be obligation attributors.

10. Epistemic asymmetries between expert decision-makers and non-expert obligation attributors

One might question whether in cases involving expert decision-makers it is reasonable for would-be obligation attributors to base their attributions on the presumed adequacy of the relevant expertise, or whether in such cases would-be obligation attributors have a potential obligation, to the extent they are able, to investigate the limits of this expertise, and to base their attributions on what they discover about these limits.

There seem to be at least two issues to consider in cases of expert/non-expert epistemic asymmetry, namely, (1) the potential obligations of non-experts to not overstep their relevant knowledge in attributing unconditional obligations to experts,

and (2) the potential obligations of the experts themselves to be honest (with both non-experts and other experts, including, perhaps most importantly, themselves) about the limits of their relevant knowledge and the implications of this limited knowledge for the success of any actions they might recommend.

Regarding the latter, it seems likely that experts themselves are typically, perhaps always, in a better position than non-experts to evaluate both the limits of their relevant knowledge and the implications of these limits for related actions. If this is right, then, according to our account, experts confront a potential obligation to explicate their epistemic deficiencies and the implications of these deficiencies to the non-expert. If the expert knows that their expertise is limited and that, because of this, their advice is unlikely, without the help of spontaneous forces, to suffice for the success of related actions, then they have a potential obligation to tell this to the non-expert. Assuming appropriate non-epistemic circumstances, experts in such an epistemic position – aware of their limited relevant knowledge – should be honestly humble about their relevant knowledge and ignorance.

Regarding the potential obligations of non-experts to not go beyond their relevant knowledge in attributing unconditional obligations to experts, one might argue that it is excessively demanding for would-be obligation attributors to know the epistemic burdens that experts confront within their narrow range of expertise. It may be a bridge too far to expect non-experts to evaluate what, say, a public health expert does not know about public health. Nevertheless, we would push back on the notion that, in such circumstances, non-experts should just take experts at their word and base obligation attributions on the assumed adequacy of experts' relevant knowledge within their purported domain of expertise. Assuming that non-experts are in a position to learn about the limited value of expertise – and we would suggest that evidence of this limited value is everywhere one looks these days – then non-experts confront a potential obligation to learn about the limits of expertise; and some degree of skepticism about expertise, and accordingly, some caution in attributing unconditional obligations to experts, is in order. In other words, it is not necessary for non-experts to evaluate how much or how little some public health expert really knows about public health. To be potentially obligated to moderate the unconditional obligations they attribute to public health (or any other) experts, all that is required of non-experts is that they know or be able to learn of the manifest evidence of prior expert failures.

Related to both of these points, we have argued in another context that, at least with regard to the very complex phenomena of society, economy, and public health, there ultimately are no experts.

Scientists learn their trades through specialized study of specific phenomena, which they investigate in isolation under carefully controlled conditions, holding other factors constant. The best scientists in their fields are experts about a very limited range of phenomena that are assumed not to interact or integrate with any other kinds of phenomena. Outside of the graduate seminar and theoretical model, beyond the laboratory and controlled trial, however, other things are not constant.

When the phenomena of multiple scientific fields interact, such as when it is necessary to trade off the health costs of a virus against the economic and other costs of a lockdown, policymakers can turn to experts about isolated phenomena. But there are no experts about the interaction of different kinds of phenomena or about the proper weighting of some against others. Policymakers can ask epidemiologists to weigh in on epidemiology, infectious disease specialists to weigh

in on infectious disease, and economists to weigh in on economics. But there are no experts about how these subjects interact or how to balance them (Scheall and Crutchfield 2021b).

For both expert and non-expert, humility about expertise and caution in attributing unconditional obligations to experts is the order of the day.

References

- Bradley R.** (2017). *Decision Theory with a Human Face*. Cambridge: Cambridge University Press.
- Crutchfield P. and Scheall S.** (2019). 'Epistemic Burdens and the Incentives of Surrogate Decision-makers.' *Medicine, Health Care and Philosophy* **22**, 613–21.
- Crutchfield P. and Scheall S.** (2024). 'Abolishing Morality in Biomedical Ethics.' *Bioethics* **38**(4), 316–25.
- Crutchfield P., Scheall S., Rzeszutek M.J., Brown H.D., and Cardoso Sao Mateus C.** (2023). 'Ignorance and Moral judgment: Testing the Logical Priority of the Epistemic.' *Consciousness and Cognition* **108**, 103472.
- DeCanio S.** (2014). 'Democracy, the Market, and the Logic of Social Choice.' *American Journal of Political Science* **58**(3), 637–52.
- Hampshire S.** (1951). 'Symposium: Freedom of the Will.' *Aristotelian Society Supplementary Volume* **25**, 161–78.
- Hare R. M.** (1951). 'Symposium: Freedom of the Will.' *Aristotelian Society Supplementary Volume* **25**, 201–16.
- Hare R. M.** (1963). *Freedom and Reason*. Oxford: Clarendon Press.
- Joyce R.** (2001). *The Myth of Morality*. Cambridge: Cambridge University Press.
- Joyce R.** (2005). 'Moral Fictionalism.' In M.E. Kalderon (ed), *Fictionalism in Metaphysics*, pp. 287–313. Oxford: Oxford University Press.
- Joyce R. and Garner R.** (2019). *The End of Morality: Taking Moral Abolitionism Seriously*. New York and London: Routledge.
- Kalderon M. E.** (2005). *Moral Fictionalism*. Oxford: Oxford University Press.
- Karni E. and Vierø M-L.** (2013). 'Reverse Bayesianism: A Choice-Based Theory of Growing Awareness.' *American Economic Review* **103**, 2790–810.
- Karni E. and Vierø M-L.** (2015). 'Probabilistic Sophistication and Reverse Bayesianism.' *Journal of Risk and Uncertainty* **50**, 189–208.
- Karni, E. & Vierø M-L.** (2017). 'Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance.' *Journal of Economic Theory* **168**, 301–25.
- Littlejohn C.** (2009). "'Ought," "Can," and Practical Reasons.' *American Philosophical Quarterly* **46**(4), 363–73.
- Mackie J. L.** (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Mahtani A.** (2021). 'Awareness Growth and Dispositional Attitudes.' *Synthese* **198**, 8981–97.
- Mizrachi M.** (2015). 'Ought, Can, and Presupposition: An Experimental Study.' *Methodos* **4**(6), 232–43.
- Piermont E.** (2017). 'Introspective Unawareness and Observable Choice.' *Games and Economic Behavior* **106**, 134–52.
- Polanyi M.** (1958). *Personal Knowledge: Towards a Post-Critical Philosophy*. Chicago: University of Chicago Press.
- Polanyi M.** (1966). *The Tacit Dimension*. Chicago: University of Chicago Press.
- Roussos J.** (2020). *Policymaking Under Scientific Uncertainty*, Ph.D. thesis, London School of Economics and Political Science.
- Scheall S.** (2019). Ignorance and the Incentive Structure Confronting Policymakers. *Cosmos + Taxis: Studies in Emergent Order and Organization* **7**(1–2), 39–51.
- Scheall S.** (2020). *F. A. Hayek and the Epistemology of Politics: The Curious Task of Economics*. London: Routledge.
- Scheall S. and Crutchfield P.** (2021a). 'The Priority of the Epistemic.' *Episteme* **18**(4), 726–37.
- Scheall S. and Crutchfield P.** (2021b). 'There Are No Experts on That for Which We Really Need Experts.' *The Daily Economy*, February 4, 2021. Accessed March 12, 2025. <https://thedailyeconomy.org/article/the-re-are-no-experts-on-that-for-which-we-really-need-experts/>

- Scheall S. and Crutchfield P.** Manuscript. The Logical Priority of Epistemic Considerations in Third-Person Normative Judgments.
- Singer P.** (1972). 'Famine, Affluence, and Morality.' *Philosophy & Public Affairs* **1**(3), 229–43.
- Sinnott-Armstrong W.** (1984). "Ought" Conversationally Implies "Can".' *The Philosophical Review* **93**, 249–61.
- Steele K. and Stefánsson H. O.** (2020). 'Decision Theory.' In E.N. Zalta (ed), *The Stanford Encyclopedia of Philosophy (Winter 2020 Edition)*. <https://plato.stanford.edu/archives/win2020/entries/decision-theory/>
- Steele K. and Stefánsson H. O.** (2021a). 'Belief Revision for Growing Awareness.' *Mind* **130**(520), 1207–32.
- Steele K. and Stefánsson H. O.** (2021b). *Beyond Uncertainty: Reasoning with Unknown Possibilities*. Cambridge: Cambridge University Press.
- Vallentyne P.** (1989). 'Two Types of Moral Dilemmas.' *Erkenntnis* **30**, 301–18.
- Vogelstein E.** (2012). 'Subjective Reasons.' *Ethical Theory and Moral Practice* **15**, 239–57.
- Walker O. and Dietz S.** (2018). 'A Representation Result for Choice under Conscious Unawareness.' Grantham Research Institute on Climate Change and the Environment, Working Paper No. 59.

Scott Scheall, PhD is Associate Professor of Philosophy & Economics in the Center for Economics, Politics & History. His work has appeared in journals in philosophy, political economy, history of economics, bioethics, and cognitive psychology. Scott's research considers the significance of human ignorance for decision-making, particularly in the political realm.

Parker Crutchfield, PhD is Professor of Medical Ethics, Humanities, and Law at Western Michigan University Homer Stryker M.D. School of Medicine. His research is in philosophy, bioethics, and psychology. In his clinical role, he commonly makes recommendations on what physician, patients, and families should do.