Penultimate draft. Please cite final draft forthcoming in Episteme.

The Priority of the Epistemic

Scott Scheall and Parker Crutchfield

#### Introduction

Some people are impressively skilled. Skilled individuals can do various things that many of us cannot. For example, in the late 1980s and early 1990s, Bo Jackson was one of the premier athletes in the world, excelling as both an especially agile Major League Baseball centerfielder and the leading running back in the National Football League. Faced with a ten-foot wall blocking his path, Bo knew how to scale it. Confronted with a ten-foot wall, one of the options available to Bo was to climb over it and move on to the next athletic challenge or Nike commercial shoot. For most of us, however, climbing such a wall with only our hands, legs, and brute strength would not appear as an option. Climbing the wall would appear in Bo's menu of options, but not in either of the present co-authors' menus of options, for example. More generally, climbing the wall would be ranked in a different position in Bo's menu of options than in the menus of anyone who believed themselves insufficiently athletic to scale it. Differences in the rank placement of this option across distinct individuals, we contend, are due to the different *epistemic burdens* that various individuals take themselves to face in pursuing this option.

Scheall (2019) defines the epistemic burden of some objective as "simply everything that the actor must know (that and how), which the actor does not already know, in order to realize the objective deliberately as a result of related actions directed to its realization, i.e., not spontaneously, or otherwise in virtue of luck or fortune." Epistemic burdens thus encompass

both propositional and non-propositional knowledge (or, more exactly, ignorance). Epistemic burdens pre-consciously shape our incentive structures. Courses of action that seem to bear impossibly heavy epistemic burdens are typically not counted as options in the menu over which an actor deliberates and from which she eventually makes a conscious choice, while courses of action that seem to bear comparatively heavy epistemic burdens are systematically discounted in

<sup>1</sup> Epistemic burdens, Scheall (2019, 47 fn 5) further notes, are "relative to both knowers and their potential objectives. Epistemic burdens are always relative to some actor with respect to some goal."

[A]ctors can evaluate their epistemic burdens with regard to various objectives with some facility, but not necessarily infallibly. Thus, epistemic burdens encompass both "known unknowns," i.e., considerations that the actor knows he does not know, and "unknown unknowns" that the actor cannot evaluate in advance of action.

Epistemic burdens should not be conflated with *epistemic costs*. Rather, epistemic costs are incurred in the process of attempting to overcome epistemic burdens. Indeed, epistemic costs need not track epistemic burdens very closely. Just as one might climb a mountain of a given height in a more or less circuitous fashion and incur greater or lesser physical costs in the process, there can be more or less epistemically costly ways in which a given epistemic burden might be met.

<sup>2</sup> We use the phrases "incentive structure," "menu of options," and "preference ranking" interchangeably to mean the array of courses of action from which an actor ultimately makes a decision. We also use the terms "actor" and "decision-maker" interchangeably.

the actor's menu relative to options that appear less epistemically burdensome. Ignorance shapes the psychology of decision-making. Furthermore, that ignorance shapes the psychology of decision-making in the way that we posit means that epistemic considerations are logically prior to moral, prudential, and pecuniary considerations. In order to have moral, prudential, or pecuniary obligations, one must have options, and epistemic burdens – the nature and extent of one's ignorance (that and how) with respect to various courses of action – both bracket one's options and serve to determine where surviving options are ranked. In short, one cannot have obligations without doing some epistemic work, even if, as we posit, such work typically occurs pre-consciously.

In the first section, we further defend the claim that ignorance brackets and serves to determine the initial ranking of options. We argue that introspection indicates the reality of the phenomena. We also examine how the fact of pre-conscious filtering of options for the decision-maker's ignorance bears distinct consequences in cases of personal and surrogate decision-making. The logical priority of the epistemic is general across all varieties of decision-making; however, the unique epistemology of surrogate cases makes this priority more apparent.

In the subsequent section, we offer a second argument for the priority of the epistemic. Many people, including most philosophers, believe that *ought implies can* or some similar, but logically weaker, principle like *ought supports can* or *ought makes plausible can*. Few would accept that there is no connection whatsoever between what one can do and what one ought to do. In any case, it is not our business to take a stance on whether ought and can are bound by *implication* or some weaker logical relation. We argue instead that, whatever version of such a principle one accepts, if it is to have any significance as a practical maxim – if it is to help us get around in the world without bumping into the furniture – the word "can" in such principles must mean *deliberately can*. Other candidate meanings of "can" make such principles trivial and

practically useless. However, we argue further that "deliberately can" ultimately means *knows* enough (that and how) to. Thus, whether one accepts ought implies can or a weaker version of a principle with similar import, one is ultimately committed to the notion that it is the nature and extent of one's ignorance, i.e., the epistemic burdens one confronts with respect to various courses of action, that determine one's potential moral, prudential, and pecuniary oughts. In other words, provided one accepts something like ought implies can, one is committed to the logical priority of the epistemic.

# **Ignorance and Menu Options**

The question that we want to answer in the present section is, what considerations determine whether a course of action counts as an option for an actor? Our answer is that epistemic considerations, in particular, the nature and extent of one's ignorance with respect to various courses of action, determine options. If it is true that one must have options to have obligations and that epistemic considerations determine options, then the epistemic is logically prior to obligations, be they moral, prudential, pecuniary, or otherwise normative in some (non-epistemic) sense.

Simple introspection shows that things about which one is seemingly completely ignorant – i.e., things that seem to bear impossibly heavy epistemic burdens – typically do not appear to the actor as options, and that the nature and extent of one's ignorance with regard to particular options serves to determine where the remaining options are initially ranked when first consciously considered. When a person's vehicle breaks down and the person is entirely ignorant of vehicle repair, do-it-yourself car repair either does not appear as an option at all or is

systematically discounted relative to the option of taking the car to the reliable and honest mechanic down the street.

In any particular decision context, the menu of options from which we choose when planning a course of action seems to arrive in consciousness pre-sorted for our ignorance. When considering, say, what to have for lunch, what film to see, or what mode of travel to adopt, it rarely, if ever, happens that options enter these menus which we think impossible on the basis of our existing knowledge, powers, abilities, skills (and our capacities for learning new knowledge, acquiring new powers, etc). We do not consider as options seemingly impossible lunches, apparently impossible films, or ostensibly impossible modes of travel. We consider as options only things that we believe (if only un- or sub-consciously) we either know or can learn enough to do. More generally, it rarely happens that options judged to bear a comparatively heavy epistemic burden are initially ranked higher in these menus than less epistemically burdensome courses of action.<sup>3</sup> Other things equal, we do not rank epistemically difficult lunches, films, or modes of travel above epistemically simple ones. If we do not believe that we know or can learn what we need to realize some option, we either exclude it from or deeply discount it in our menu of options. We either do not consciously consider or do not take seriously courses of action that, on an evaluation of the knowledge we possess and our capacity to learn, seem epistemically impossible, or especially burdensome.

Cases of surrogate decision-making, where one party decides on behalf and ostensibly in the interests of some other party, exhibit the significance of pre-conscious filtering for ignorance

<sup>-</sup>

<sup>&</sup>lt;sup>3</sup> Thus, technically, it is not our ignorance *per se*, but *how we evaluate our ignorance*, that determines options. For a discussion of how first-order and second-order epistemic considerations interact to affect the pre-conscious sorting of options, see Scheall (2019, 44-45).

more starkly than do cases of personal decision-making, where one decides for oneself. A sane person will never permit a seemingly impossible option to enter her menu and will discount epistemically burdensome options accordingly. However, no such guarantee exists where one person must decide on behalf and in the interests of some other. Simply put, there is no guarantee in such cases that the surrogate knows the options in the interests of the other party or how to achieve them. You will never leave yourself an option that you cannot know enough to achieve, but a person on whose behalf you must decide might assign such an option to you.<sup>4</sup>

A bit more carefully, in trying to decide on behalf and in the interests of another person, the surrogate's epistemic burden potentially consists of two elements. First, the surrogate faces the burden of knowing the person's interests and the various options associated with these interests. In principle, a similar burden exists in cases of personal decision-making; however, introspection typically suffices to discover one's own interests, while pre-conscious filtering reveals a limited range of options associated with these interests seemingly deliberately realizable given the nature and extent of one's ignorance. The epistemic burden of discovering one's own interests, etc., would seem to be necessarily lighter than the comparatively heavy burden of discovering someone else's interests and the options relevant to them.<sup>5</sup> Second, *given* knowledge of the person's interests, the surrogate confronts another potential epistemic burden in acquiring adequate knowledge of *means* sufficient to realize relevant options. Even if the surrogate knows the interests of the person on whose behalf she decides, as well as the options

On

<sup>&</sup>lt;sup>4</sup> On the effects of ignorance on the surrogate decision-maker's incentive structure, see Crutchfield and Scheall (2019).

<sup>&</sup>lt;sup>5</sup> Crutchfield and Scheall (2019) provide both epistemological and psychological explanations of this phenomenon.

associated with these interests, she may not know of a way to effectively pursue these options such that the relevant interests are satisfied.

In short, pre-conscious filtering cannot prevent impossible or especially epistemically burdensome options in the case of surrogate decision-making: the option most in the interest of the person on whose behalf the surrogate decides may well be epistemically impossible or more epistemically burdensome for the surrogate than alternative courses of action less in the interests of the person on whose behalf a decision must be made. The point is that the surrogate's preference ranking will be determined by her own ignorance and not by the ignorance of the person on whose behalf she decides. The surrogate's menu of options need look nothing like the menu from which the person surrogated would choose, were the latter to choose for herself.

In order to undermine this argument from introspection, one would have to provide a counterexample that falsifies our analysis of the determination and ranking of options according to the actor's ignorance. A case in which a course of action about which the actor is entirely ignorant nevertheless appears to her as an option would be a counterexample to our argument. Instances of spontaneous inspiration might appear to be counterexamples to our analysis. For example, the person entirely ignorant of vehicle repair might spontaneously think, "I could fix it myself!" Without a specific case at hand, it is hard to know how one would respond to such an example. However, it seems open to us to claim of such a case that something must have happened to make the actor more knowledgeable about this course of action. Perhaps the actor witnessed a neighbor changing their own motor oil or saw a sign for a DIY vehicle-repair training program, triggering the thought that they could fix the vehicle themselves. No longer *completely* ignorant of DIY vehicle repair, the option appears in their menu (albeit, significantly discounted relative to less epistemically burdensome options).

Another kind of counterexample to our analysis would be a case in which an option that bears a heavier epistemic burden is nevertheless ranked above another option that bears a lighter burden, even though they are equivalent in other respects, i.e., it would be a case in which *ceteris* is *paribus*, but options are not pre-ranked in the way we posit.<sup>6</sup> It would not be enough to present a case in which a person chooses an option that bears a heavier epistemic burden over some other less burdensome option. Other considerations, such as the moral, prudential, or pecuniary value of competing options, figure in decision-making. Sometimes people pursue options with comparatively heavy epistemic burdens, because they take the moral, prudential, or pecuniary value of these options to outweigh their epistemic burdens. What we claim is that epistemic burdens are decisive for the pre-conscious ranking of options and, where these other considerations are equal across options, for the choice ultimately made.<sup>7</sup>

\_

<sup>&</sup>lt;sup>6</sup> An even more damaging counterexample would be a case in which one option, despite being both more burdensome *and* of lower value in other relevant respects is nevertheless ranked higher than another option both less burdensome and more valuable in other relevant respects. However, such a case would call into doubt not only our analysis but, seemingly, any attempt to think systematically about decision-making.

<sup>&</sup>lt;sup>7</sup> It might be tempting to think that the fact that *ceteris* is rarely *paribus* in human decision-making trivializes that part of our thesis concerning the potential decisiveness of epistemic burdens. However, such "other things equal" clauses are essential in all domains of inquiry. The laws of Newtonian mechanics carry such conditions, yet *ceteris* is rarely *paribus* in the physical world outside vacuum chambers. Similarly, nothing in the economy remains constant very long, yet the Law of Demand operates regardless. As we have put this point elsewhere, "were the price of tea to increase ten-fold, a person who objects to the use of *ceteris paribus* conditions in

# The Practical Meaning of Ought Implies Can and Similar Principles

There is another route to the conclusion that the epistemic is logically prior to the moral, prudential, or pecuniary. Many people accept either that *ought implies can* or accept some similar, but logically weaker principle. For example, Stuart Hampshire (1951) and R. M. Hare (1951, 1963) argue that *ought presupposes can*, while Walter Sinnott-Armstrong (1984) argues that *ought conversationally implicates can*. Other, still weaker, candidate propositions are

scholarly inquiry would nevertheless be inclined to buy less tea" (Crutchfield and Scheall 2019, 6). Inquiry in the sciences and humanities cannot proceed without isolating causal factors and their specific effects on relevant phenomena.

Perhaps more to the point, the fact that other conditions may not be constant does not mean that the underlying principle fails to operate, though it can make its operation more difficult to observe. From the fact that moral, prudential, pecuniary (etc.) considerations can sometimes outweigh epistemic burdens and that relatively burdensome options can be chosen for these reasons it does not follow that ignorance fails to affect the ranking of options in such cases; it just makes its effects more difficult to observe. This is part of the reason why we discuss cases of surrogate decision-making in the current context, i.e., to display the role of ignorance in determining and sorting options in starker relief.

<sup>8</sup> Others who defend *ought presupposes can* include Cooper (1966), Martin (2009), Besch (2011), and Driver (2011). Mizrahi (2015) denies *ought presupposes can*, but describes it as "the best candidate for a relation between 'ought' and 'can'." Others who defend *ought* 

conceivable. Perhaps *ought supports can* or *ought makes plausible can*. It is not incumbent upon us to take a position on the correctness of any such principle; nothing in our argument for the logical priority of the epistemic hinges on *implication* (or any other particular relation) being the correct relation between ought and can. Our argument is simply that, if some such principle is both true and useful as a guide to practical action, then the epistemic is logically prior to other normative considerations. Anyone committed to such a principle, whether *ought implies can* or a weaker version, is also committed to the priority of the epistemic.

Consider what the word "can" must mean in the statements of such principles, if they are not to be utterly trivial, i.e., if they are not to make every course of action a potential obligation, but effectively bracket legitimate oughts according to the criteria of cans. If such principles have any substance at all as guides to successful decision-making, the relevant sense of "can" cannot be anything like *possibly can* or *can*, *with luck*, but must be something closer to *deliberately can*.

Ought implies possibly can means that every possibility is a potential obligation and that only impossibilities cannot be potential obligations. Of course, there are multiple senses of possibility – metaphysical, logical, physical, economic, legal, political, social, etc. – and if ought implies can means that ought implies possibly can, it is not obvious which sense of possibility is pertinent. More to the point, whatever sense of possibility might be relevant, it seems that the set of potential obligations cannot simply be the set of possibilities. If, say, the relevant sense is physical possibility, then, under appropriate circumstances, anything physically possible might be something one ought to do. But this is implausible. It is physically possible to voluntarily submit to a live flaying; but it is implausible that one should ever do this. It would seem that this

conversationally implicates can include Oppenheim (1987), Vallentyne (1989), Saka (2000), Littlejohn (2009), and Vogelstein (2012).

and many other physical possibilities cannot be things one ought to do, regardless of circumstances—they can never be, even merely potentially, obligations. On the other hand, if, say, the legal sense of possibility is relevant, then it seems that the set of potential obligations may well be larger than the set of possibilities. In other words, there may be legally impossible things (i.e., illegalities) which, nevertheless, one ought to do, under appropriate circumstances. It may be that a person ought to steal from a rich man to feed her starving family, though she cannot possibly do this legally. Given a particular sense of possibility, it seems that the set of potential obligations is not simply equivalent to the set of possibilities and, therefore, that possibly can cannot be the relevant sense of "can" in such principles.

Alternatively, if *ought implies can, with luck*, then the principle is trivial and practically useless. If luck or fortune intervenes, virtually *anything* can happen. If this is the sense of "can" relevant to such principles, then anything that you can do, even if it requires the intervention of luck or of other spontaneous forces outside your purview, or beyond your control, is something that you potentially ought to do, under the right circumstances. In short, virtually anything is a potential ought, if *ought implies can, with luck*.

However, there is another reason to reject this reading of "can." Part of what it means for something to happen "fortunately" or "luckily" is that it happens for reasons, or in ways, that are not well understood or apparently under human control. "I just got (un)lucky" is how we often explain *ex post facto* personal circumstances that we do not understand and cannot explain otherwise. We are not very good at recognizing in advance where, when, or how luck or other spontaneous forces will intervene. If this is right, then *ought implies can, with luck* is worthless as a guide to successful decision-making. If you cannot distinguish the things that you cannot do, regardless of luck, from the things that you can do, with or without luck, then, if the relevant

meaning of "can" in such principles is *can*, *with luck*, then you cannot distinguish potential obligations from potential non-obligations, and any such principle is practically useless.

If this is right, then it seems that the relevant sense of "can" in such principles must be closer to *deliberately can*. However, to say that someone deliberately can do something is just to say that they possess or can learn enough knowledge (that and how) to realize the thing without need for the intervention of luck, fortune, or any other spontaneous forces. What we deliberately can do depends on our existing knowledge and our capacity to acquire more knowledge. Stated another way, in order to realize an objective with regard to which our knowledge and learning capacity is inadequate, luck, fortune, or other spontaneous forces must intervene.

It is important to emphasize that the conclusion of the logical priority of the epistemic, via this second dialectical route, hinges entirely on the notion that, in principles similar in relevant ways to those under investigation, "can" must mean *deliberately can*, which means *knows enough (that and how) to*, and not at all on an implicit assumption that implication is the appropriate relation between ought and can, in contravention of our explicit claim to assume nothing about the nature of this relation. One might be tempted to think that we need implication to be the appropriate relation, in order to draw the conclusion of the logical priority of the epistemic and that any weaker relation between ought and can licenses only a conclusion of a weaker kind of priority. However, this is an error. *Ought implies can* is sufficient, but it is not necessary, to conclude the logical priority of the epistemic.

What depends on the strength of the relation between ought and deliberately can is the determinacy of the set of potential obligations, not the logical order in which epistemic and other factors must be considered. Given some subject S, there is S's set of "deliberately can-dos," i.e., all of those options that S can realize on the basis of her knowledge and learning capacity, without need for the intervention of luck, fortune, or any other spontaneous forces. The stronger

the relation between ought and deliberately can, the more the set of potential obligations approximates the set of "deliberately can-dos." However, if the relation between ought and deliberately can is something weaker, say, *makes plausible*, then the set of potential obligations is less determinate. It becomes less clear whether any particular member of S's "deliberately candos" is also a member of the set of S's potential obligations. For example, if *makes plausible* is the right relation between ought and deliberately can, then whether some option O is a member of the set of S's potential obligations depends on how plausible it is that S deliberately can do O. O is more likely to be a member of the set of S's potential obligations the more plausible it is that S deliberately can do O. If it is definitely plausible that S deliberately can do O, then O is definitely a member of the set of S's potential obligations. However, if it is utterly implausible that S deliberately can do O, then O is not a member of the set of S's potential obligations.

Note that insofar as S can evaluate the extent to which she deliberately can do something, even an extremely weak principle like *ought makes plausible deliberately can* remains a practically useful maxim for S. So, the greater indeterminacy of potential obligations that follows if the appropriate relation between ought and deliberately can is relatively weak does not threaten that part of our argument that hinges on the assumption that some such principle is practically useful.

If this line of reasoning is sound, then epistemic considerations are logically prior to moral, prudential, and pecuniary considerations. In effect, whatever is the correct statement of the logical relation between ought and can, if it is to be a practically useful principle of practical action, then "can" must mean *deliberately can*, which means *knows enough (that and how) to*, and one's potential moral, prudential, and pecuniary oughts are determined by one's epistemic burdens. In order to make a moral, prudential, or economic decision, some epistemic work must be done. The epistemic is logically prior to other non-epistemic normative considerations.

### **Implications of the Priority of the Epistemic**

If either of our arguments is sound, then epistemic considerations are logically prior to moral, prudential, pecuniary, and other (non-epistemic) normative considerations. The set of a person's potential obligations is determined by the nature and extent of their ignorance. If this is right, then it seems either false or simply empty to insist that an actor ought to do something that, because she pre-consciously takes herself to be incapable of realizing it on the basis of her existing knowledge and learning capacity, will never appear to her as an option, or appear only as deeply discounted. At least, a person would seem to be tilting at windmills who attributes an obligation to someone whose cognitive architecture and experiential history is such that the purportedly obligated course of action will either not appear among her options or appear only as severely discounted relative to other options.

Bo Jackson may know (or may have known in his younger, more athletic, days) how to climb a ten-foot wall to steal a home run or how to trample opposing linebackers, but if he does not know how to save a child from drowning in a shallow pond, then it is either false or meaningless to assert that he ought to save the drowning child. Bo's potential obligations are determined by his relative epistemic burdens and if (*ex hypothesi*) his relative epistemic burdens are such that trying to save the drowning child does not appear to him as an option, it makes little sense to assert that he has a moral obligation to try or that he is immoral for not trying to save the child. Singer's (1972) case is compelling in large part because it strains credulity to think that a typical adult, much less Bo Jackson, would be ignorant of how to perform this act. But, place the child in deeper water and take away Bo's swimming know-how, and it becomes more difficult to see the sense in the assertion that he is morally obligated to try to save the child.

Or consider the CEO whose shareholders demand profits, but who runs a company in a dying industry (say, print news media). The CEO is ignorant of how to achieve the demanded profits, because they are impossible to achieve under prevailing market conditions. Again, the standard notion that CEOs are obligated to try to make profits for their shareholders is, under such circumstances, either false or meaningless. Something similar can be said about policymakers, who, like CEOs, are a kind of surrogate decision-maker (Scheall 2019). Policymakers' constituents may (indeed, perhaps often) demand of them impossible tasks, or tasks the realization of which require a well of knowledge deeper, or learning capacities more expansive, than the policymaker possesses. The policymaker whose constituents demand things that the policymaker cannot deliberately achieve is under no obligation to attempt to meet these demands.

Many failings that seem at first glance to be of a moral, prudential, or economic nature, are primarily epistemic. The imaginary Bo Jackson who is so ignorant that he cannot deliberately save a child drowning in a shallow pond fails for entirely epistemic reasons, and not at all because he is morally, prudentially, or economically deficient.

One might be inclined to argue of such a case that Bo "should have known" enough to deliberately save the child. That is, Bo should have taken the steps necessary to possess the required knowledge. Moreover, such an argument might be thought to undermine the logical priority of the epistemic, i.e., if there are things that one should know, such considerations must be prior to questions of what the actor does know. However, if our position is sound, such "should have known" claims merely push the priority of the epistemic to a higher order. On our way of thinking, whether Bo should have taken the steps necessary to possess the knowledge required to save the drowning child hinges crucially on whether he deliberately can have taken the necessary steps. If Bo's epistemic burdens were such that taking these steps did not appear to

him as an option – perhaps because he was ignorant of the practice known as swimming or of the possibility that human beings can drown in water – then it is either false or simply meaningless to assert that Bo "should have known" enough to deliberately save the child. Potential epistemic obligations are not determined by non-epistemic considerations but by higher-order epistemic considerations, i.e., not by the limits of what the actor deliberately can do, but by the limits of what the actor deliberately can know. It is potentially true that an actor "should have known" only those things that she deliberately can have known.

Consider the following example from Sanford Goldberg (2017, p. 2889):

SOCIAL WORKER: S is a very overworked social worker, whose case load is more than 150 (typical for the underresourced Department of Family Services in her state). It comes to pass that one of the foster children in S's caseload suffers a tragic—and, it is deemed, preventable—death. People clamor that S should have known that the child was at serious risk! But when the facts come to light, it becomes clear that there was no reasonable way S could have known this, given the way that DFS organized itself on its underresourced budget. And it becomes clear as well that there was no reasonable restructuring of DFS itself, consistent with its level of funding and the totality of its responsibilities, that would have made it so that S would have been in a position to know.

Goldberg thinks this is a case in which S should have known that the child was in danger. On our analysis, however, S should have known only if S deliberately can have known that the child was in danger, and it does not seem as if this was an option for S. One might be tempted to say that S could have made such a discovery and we would agree that, surely, S could have *accidentally* or *fortunately*, or *luckily*, discovered that the child was in danger. But, given the description of the

case, it does not seem as though S *deliberately can* have discovered that the child was in grave danger. Given S's caseload and the structure of the DFS organization, she cannot have deliberately gone through each case and assessed each child's risk. There was simply too much work. The only way that S can have made such a discovery was if she had luckily guessed that the child was in danger or randomly picked the child's file for review. S could only have made such a discovery with luck or by accident. Given this, it does not seem that the child's susceptibility was something that S should have known. Goldberg's claim that S should have known is in error. If anyone failed to satisfy their obligations, it was not S, but the government for which she works, which – assuming its officials deliberately can have organized the DFS in a way that would have facilitated S's knowledge of the child's risk and, thus, presumably prevented the child's death – had a potential obligation to organize the DFS in the way required to prevent the child's death and, presumably the deaths of still other children.

Or consider another of Goldberg's examples (p. 2887):

BASEBALL Jones has just begun to play baseball. Given the lack of players in his high school, he made the team as a back-up. After the seasoned third- baseman gets injured just before the start of the high school season, Jones is asked to play third base in the season's first game. Consequently, he has little time to practice and to get to know the game. The result is that he doesn't learn some of the game's less familiar rules. For example, he isn't taught that a hit ball is still in play if it hits the third base bag. As a result of a play in which this happens at a crucial moment in the first game of the season, he lets the winning run score. Fans are furious, and are not mollified by his pleading ignorance: how could he have failed to know this? He's a ballplayer after all, he should have known! When the fans find out that he is new to the game, however, they no longer

blame him, but instead re-direct their anger at the manager, Ralphie Schneck: why hadn't Ralphie made Jones aware of this?!?

Like the case of the social worker, unless Jones deliberately can have known about the rule, it is difficult to see why he should have known. His only possible source of knowledge of the rule was his coach, Ralphie. Jones cannot deliberately have known about the rule, so, on our analysis, it is not true that he should have known. However, our analysis also implies that, given the way the example is set up, Ralphie had no obligation to inform Jones of the relevant rule. Surely, given enough time before Jones's first big game, Ralphie deliberately can have informed Jones of the rule and this would have thus been one of Ralphie's potential obligations. With enough time to spare, Ralphie deliberately can have provided Jones with a comprehensive set of the rules of baseball and walked him through various game situations, including the relatively uncommon one covered by the rule. However, the example makes it seem like Ralphie simply did not have time for such deliberate training. He could have informed Jones only luckily or accidentally, i.e., only if, as in the social worker case above, he had simply guessed that the rule would be relevant to the game or randomly selected it for Jones's pre-game edification.

The logical priority of the epistemic implies that much opprobrium of an ostensibly moral, prudential, or economic nature is misplaced, and that what such misconceived expressions of disapproval actually express is a wish that the epistemic circumstances of the person criticized were different than they in fact are.

#### **Further Potential Objections and Counterexamples**

We conceive of decision-making as proceeding in the following (logical) order: first, the actor's ignorance pre-consciously sets brackets around options that the actor knows enough to deliberately achieve and ranks these options in accordance with their relative epistemic burdens for the actor. Then – second – the actor considers and chooses from the remaining options in accordance with relevant moral, prudential, pecuniary, and any other pertinent considerations. The epistemic is logically prior to these other non-epistemic normative considerations. Epistemic burdens – the nature and extent of one's ignorance with respect to courses of action – determine what counts as a potential obligation.

We discussed above what would constitute a counterexample to our first argument. It remains merely to consider what would qualify as a counterexample to our second argument, according to which the logical priority of the epistemic follows from the fact that some principle like *ought implies can* is true and useful, and the argument that, in order to be useful, the word "can" in such principles must mean *deliberately can*, which means *knows enough (that and how) to*.

One possible counter to this second argument is to reject any and all principles like *ought implies can*, and insist that there is no relation whatsoever between ought and can. Someone who insists that there is no relation between ought and can will eventually find themselves asserting plainly impossible obligations. As above, it is difficult to know how one might respond to a merely imaginary counterexample, but given the general contours as just stated of such a counterexample, it is interesting to consider how one might respond to the assertion that they are

\_

<sup>&</sup>lt;sup>9</sup> We remind the reader again that, other than accepting that some relation obtains between ought and can, we take no stance on the exact nature of this relation. Thus, arguments against *ought implies can* or against any other specific version of this principle are impotent against us.

obligated do something that, for whatever reason, they simply cannot do, deliberately or even with a healthy heaping of spontaneity. It is even more interesting to consider how one who denies any connection between ought and can would react to the assertion of themselves that they are obligated to do something that they cannot do. 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. It is simple to claim that there is no connection whatsoever between ought and can. It is much harder to live one's life in accordance with this notion.

Alternatively, one might argue that our analysis of the meaning of the word "can" in principles like *ought implies can* is faulty and that "can" need not mean *deliberately can*, in order for such principles to be true and practically useful. Unless one is prepared, as above, to simply deny the truth and usefulness of such principles, this would seem to require establishing that another candidate meaning of "can" renders some such principle practically useful, and, moreover, arguing that this meaning is more likely to be true than a similar principle that invokes *deliberately can* as the meaning of "can."

Or, one might accept our argument that the meaning of "can" in such principles is deliberately can, but reject the notion that "deliberately can" means knows enough to. However, this would seem to require arguing that there is no connection between a person's knowledge and the things that she deliberately can do. If there is any such connection – if, say, "deliberately can" does not mean knows enough to, but rather, it is the case that being able to deliberately do something merely makes it likely that the actor knows enough to realize the option in question, our argument for the logical priority of the epistemic still stands. Under such circumstances, it

remains the case that having moral, prudential, or pecuniary obligations requires some epistemic work and that one's epistemic burdens are a necessary determinant of these obligations.

## **Bibliography**

- Besch, Thomas A. (2011). Factualism, normativism and the bounds of normativity. *Dialogue* 50 (2): 347-365.
- Cooper, Neil (1966). Some presuppositions of moral judgments. Mind 75 (297): 45-57.
- Crutchfield, Parker and Scott Scheall. (2019). Epistemic Burdens and the Incentives of Surrogate Decision-makers. *Medicine, Health Care, and Philosophy*. https://doi.org/10.1007/s11019-019-09899-2
- Driver, Julia (2011). Promising too much. In H. Scheinman (Ed.), *Promises and agreements*.

  Oxford: Oxford University Press, 183-197.
- Goldberg, Sanford. (2017). Should have known. Synthese 194: 2863-2894.
- 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.
- Littlejohn, Clayton (2009). 'Ought', 'can', and practical reasons. *American Philosophical Quarterly* 46 (4): 363-373.
- Martin, Wayne (2009). Ought but cannot. *Proceedings of the Aristotelian Society* 109 (2): 103-128.

- Mizrahi, Moti (2015). Ought, can, and presupposition: An experimental study. *Methode* 4 (6):232-243.
- Oppenheim, Felix E. (1987) National interest, rationality, and morality. *Political Theory* 15 (3): 369-389.
- Saka, Paul (2000). Ought does not imply can. American Philosophical Quarterly 37 (2): 93-105.
- Scheall, Scott (2019). Ignorance and the Incentive Structure confronting Policymakers. *Cosmos* + *Taxis* 7 (1-2): 39-51.
- Singer, P. (1972). Famine, Affluence, and Morality. Philosophy & Public Affairs, 1(3), 229–243.
- Sinnott-Armstrong, Walter (1984). 'Ought' conversationally implies 'can'. *The Philosophical Review* 93: 249-261.
- 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.