

Epistemic Burdens and the Incentives of Surrogate Decision-makers

Parker Crutchfield and Scott Scheall

We aim to establish the following claim: *other factors held constant, the relative weights of the epistemic burdens of competing treatment options serve to determine the options that patient surrogates pursue*. Simply put, surrogates confront an incentive, *ceteris paribus*, to pursue treatment options with respect to which their knowledge is most adequate to the requirements of the case. Regardless of what the patient would choose, options that require more knowledge than the surrogate possesses (or is likely to learn) will either be neglected altogether or deeply discounted in the surrogate's incentive structure. We establish this claim by arguing that the relation between epistemic burdens and incentives in decision-making is a general feature of surrogate decision-making. After establishing the claim, we draw out some of the implications for surrogate decision-making in medicine and offer philosophical and psychological explanations of the phenomenon.

INTRODUCTION

According to relevant medical-ethical standards, patient surrogates, those who make medical decisions for patients unable to make such decisions for themselves, ought to base their

decisions on what they think the patient would want, if she could express such a preference.¹ It is often the case, however, that this standard can be met only if the surrogate first overcomes her ignorance of various considerations. These include, first, that the surrogate may be ignorant of the patient's values, beliefs, attitudes, and dispositions. When we make our own medical decisions, we don't need to think about what kind of person we are. If she is to meet the relevant standard, however, a surrogate must possess intimate knowledge of the patient's values, beliefs, etc. This knowledge may be difficult to acquire.² Second, the surrogate may be ignorant of what the patient would want in the context under consideration, i.e., the surrogate may not know what the patient's values, beliefs, etc., imply about the decision she would want made under prevailing circumstances. When we make our own medical decisions, we make them in the context of our own lives and how we want to live them. For surrogates, however, this deliberation, must be more explicit in order to make a decision that meets the relevant standard. Third, the surrogate may be ignorant of some of the medical facts knowledge of which is required to meet the standard of deciding as the patient would decide. If there is a disconnect between what the surrogate and the patient know about pertinent medical facts, there is likely to be a disconnect between the surrogate's decision and the patient's unexpressed preference.

Scheall defines the *epistemic burden* of some objective as "simply everything that one must know...which one does not already know, in order to realize [the objective] deliberately as

¹ The substituted judgment standard is widely accepted in medical ethics as the primary standard of surrogate decision-making. Some may prefer instead the best interests standard, whereby the surrogate makes decisions according to what she thinks the patient's interests are. These standards have been thoroughly adjudicated throughout the medical ethics literature. But readers will notice that whether the appropriate standard is substituted judgment or best interest has little bearing on whether our argument is sound.

² We mean 'knowledge' to refer to both knowledge-that and knowledge-how. Lacking either of these types of knowledge implies that one is ignorant of that proposition or of that skill—ignorance is not limited to ignorance of propositions.

a result of related actions directed to its realization” (Scheall, forthcoming). In any particular case, the patient surrogate’s epistemic burden is everything that she needs to know, which she does not already know, in order to purposefully make the decision that the patient would want to be made, if she could express a preference. The range of things with respect to which a surrogate may be ignorant is potentially wide. A person may be ignorant of relevant facts or propositions, lack practical knowledge relating to required skills, be ignorant of what another person’s interests or preferences are, or lack knowledge of the potential consequences, however likely, of a particular action. All of these instances of ignorance could potentially figure into a person’s epistemic burden, as overcoming this ignorance may be necessary to realize her objective.

Much has been written on the reliability of surrogates and their ability to make medical decisions (Berger et al., 2008; Moorman and Carr, 2008; Shalowitz DI, Garrett-Mayer E and Wendler D, 2006; Suhl J et al., 1994). That surrogates face significant epistemic difficulties is not a novel point (Brudney, 2018; Buchanan and Brock, 1989; Pope, 2012). What has not yet been recognized in the literature is the crucial role that epistemic burdens play in shaping the incentives that surrogates confront to choose various treatment options and, thus, insofar as incentives determine decisions,³ the role of epistemic burdens in determining the decision that surrogates ultimately make.⁴ That epistemic burdens shape incentives is true for all decisions, personal and surrogate. But, given the greater likelihood of the decision-maker’s ignorance of the

³ Our use of ‘determine’ is compatible with any of the popular theories of causation. The truth of our claim does not hinge on any particular theory of causation. Also, the word ‘incentive’ is used in the paper in its well-understood sense, especially as commonly used in economics and the social sciences. That incentives play a causal role in decision-making is common currency across the social sciences.

⁴ Kibbe and Ford’s (2016) concluding qualification approaches our point, but nevertheless stops well short.

important considerations described above, epistemic burdens are especially pertinent wherever a surrogate decides on behalf, and ostensibly in the interests, of a patient.

We aim to establish one claim: *other factors held constant, the relative weights of the epistemic burdens of competing treatment options serve to determine the options that patient surrogates pursue.* Simply put, surrogates confront an incentive, *ceteris paribus*, to pursue treatment options with respect to which their knowledge (or their potential for learning) is most adequate to the requirements of the case. Regardless of what the patient, who confronts either a lighter epistemic burden or none at all, would choose, options that require more knowledge than the surrogate possesses (or is likely to learn) will either be neglected altogether or deeply discounted in her incentive structure. If, relative to other treatment options, the surrogate is less knowledgeable of the significance of the patient's preferred treatment for the patient's values, beliefs, attitudes, and dispositions, or of the significance of the treatment for the patient under prevailing circumstances, then the patient's preferred option will either be ignored entirely or ranked relatively lower in the surrogate's incentive structure. Other things equal, the surrogate will likely not choose the patient's preferred treatment option and the relevant medical-ethical standard will go unmet. Importantly, however, it will go unmet for *epistemic*, rather than *moral*, reasons.

Burdens of any kind, epistemic or otherwise, shape incentives. For example, suppose that before you are two paths, equally long and winding, that reach the same destination. However, while the path on the right is free of obstacles, the path on the left has across it a large wall that, were you to choose that path, would require you to get over it. This obstacle—a burden that must be overcome in order to travel effectively on this path—disincentivizes taking the left path, and, conversely, incentivizes taking the right path. Epistemic burdens influence incentives similarly. Our ignorance is a burden that must be overcome in order to realize certain objectives.

When there are competing options, their comparative epistemic burdens shape the incentives the actor faces to pursue different options. That is, the extent to which the epistemic burden of a particular option incents or disincentivizes the option is relative to the burdens of the other options, but the epistemic burdens of the various options themselves are independent of each other. You may find that learning Russian is more burdensome than learning French and this may inform your decision to learn French, but the burden that you confront in achieving French fluency is independent of the burden to be overcome in order to master Russian. More generally, given some subject *S* and options *O1* and *O2*, the epistemic burden that *S* confronts in attempting to realize option *O1* is independent of the epistemic burden that *S* confronts in attempting to realize option *O2*. This is true whether *S* is choosing between learning Russian rather than French, deciding which policy to adopt in order to most positively impact the lives of one's constituents, or choosing whether to continue or discontinue aggressive treatment for an incapacitated loved one. The ease or difficulty that *S* confronts in acquiring the knowledge (that and how) necessary to realize *O1* cannot be affected by the ease or difficulty that *S* confronts in acquiring the knowledge (that and how) necessary to realize *O2*. The burdens of each option just are what they are—they do not depend on each other in any way.

What *is* dependent on the relative weight of the epistemic burdens of competing options is the ranking of these options in *S*'s incentive structure. In effect, a subject's epistemic burden with respect to some option is the epistemic distance (if you will) separating the subject at that moment in time from realization of the option. Just as the physical distance separating a subject at any given moment from some physical location (say, Los Angeles) is not affected by the physical distance separating the same subject at the same moment from some other physical location (say, New York City), the epistemic distance that a subject needs to travel to realize option *O1* is not affected by the epistemic distance that the same subject needs to travel to realize some other option *O2*. Of course, the distance separating the subject from Los Angeles relative to the distance separating the subject from New York City might make

one or the other a more attractive travel destination, but the actual distances do not depend on each other. Epistemic burdens work the same way.⁵

Epistemic burdens play a role in determining the menu of options from which a person chooses and, thus, the option that she ultimately pursues. The scope of things that epistemic burdens determine may be wider than this, but we take no stance on this question. Nothing in our argument hinges on the possibility that epistemic burdens play a broader causal role in decision-making. In our view, following Scheall (forthcoming), the menu of options from which a person chooses in any given decision context has been pre-filtered on some (often, less-than-conscious) level for her ignorance. That is, options that are judged to bear an impossible epistemic burden are typically excluded from the menu and options judged to bear relatively heavy epistemic burdens are commensurately discounted in the menu. Epistemic burdens serve to determine where an option appears, if it appears at all, in a person's initial preference ranking (i.e., "incentive structure"). It is only once this pre-filtering for ignorance has occurred, we contend that a decision can ultimately be determined by considering the remaining menu options in light of other factors (moral, political, economic, what have you). Thus, our claim is that epistemic burdens serve to determine what a person is incented to do and not do, and, thus, insofar as incentives determine decisions, epistemic burdens serve to determine decisions as well.

We contend that this priority of the epistemic is a feature of all decision-making, personal and surrogate. It is more obvious, however, in the case of surrogate decision-making. A sane person, whose possible courses of action in a particular context have been pre-filtered for her ignorance, will never find an epistemically impossible option in her menu. But there is no such

⁵ The argument is not materially affected even if the knowledge required to realize O1 is a subset of the knowledge required to realize O2, such that overcoming the epistemic burden of the latter requires overcoming the epistemic burden of the former. If our argument is sound, it remains the case that the relative epistemic burdens of the two options will affect how attractive they are to the subject and, thus, their placement in the subject's initial incentive structure.

guarantee in cases of surrogate decision-making. A patient surrogate tasked with doing what an incapacitated patient would want, were the patient capable of expressing a preference, may find it is epistemically impossible either to discover this preference or to realize it.

In the next section we draw from work that establishes an analogous thesis for policymakers, who are another type of surrogate decision-maker (Scheall, forthcoming). In the section that follows, we transpose this argument to patient surrogates. Whether the surrogates are family members, guardians, or physicians, they all confront heavy epistemic burdens that shape their incentives. Once we establish this phenomenon, we offer two explanations for it, one philosophical and one psychological.

EPISTEMIC BURDENS OF POLICYMAKERS

Policymakers are surrogate decision-makers—they decide on public policy on behalf, and ostensibly in the interests, of their constituents. Scheall considers the various factors that determine the extent to which policymakers pursue objectives that constituents demand. He makes no assumption about the manner, democratic or otherwise, in which policymakers achieve their positions in the imagined state of his model. Instead, he assumes that there is a widely-held belief among the general public, including among constituents, that policymakers ought to pursue their constituents' policy preferences. Thus, our present concern with the factors that determine the extent to which a surrogate acts according to the patient's wishes is analogous to Scheall's concern with the factors that determine the extent to which policymakers act on their constituent's policy demands. There is also a parallel in the assumption that policymakers ought to pursue their constituents' preferences and the medical-ethical standards relevant to patient surrogacy.

Scheall starts from the simplifying assumption that policymakers know the objectives that their constituents want them to pursue. By assumption, the policymaker faces an epistemic burden only in discovering the means to realize various policy objectives. This assumption is rather unique as simplifying assumptions go, because weakening it or dropping it altogether would seemingly bolster Scheall's ultimate conclusion: the heavier the epistemic burden involved in discovering their constituents' policy demands, *ceteris paribus*, the greater the incentive confronted by policymakers to pursue other objectives. That we make no such assumption about a surrogate's knowledge of the patient's wishes is a key difference in our analysis. Of course, the significance of dropping this assumption is that, *ceteris paribus*, the patient surrogate's epistemic situation is yet more burdensome than that of Scheall's idealized policymaker, who knows her constituents' policy demands automatically, without need for a potentially burdensome search process. The patient surrogate needs to know not only the means to realize the patient's wishes (i.e., the pertinent medical facts), but must also discover these wishes.

Scheall, Butos, and McQuade argue that the means to selfish ends are often (by no means always) more easily known than the means to altruistic goals (Scheall et al., forthcoming). The epistemic burden of self-interested policymaking tends to be lighter than that of constituent-minded policymaking. If this is right, they conclude, "We should expect to find more self-interested political behavior where (*ceteris paribus*) the epistemic burden of making effective public-minded policy is comparatively heavy" (Scheall et al., forthcoming). Generalizing this thesis, Scheall argues that "*Other factors held constant, the relative weights of the epistemic burdens of competing policy objectives serve to determine the objectives that policymakers pursue.*"

If this is right, then, *ceteris paribus*, the policymaker is incented to pursue policy objectives that bear the lowest epistemic burden, i.e., those goals she knows or has the best prospects of learning how to realize. Other factors held constant, the policy objectives that constituents demand will be pursued only if they impose a lighter epistemic burden on policymakers than alternative policy objectives. Otherwise, the policy objectives that constituents demand will be ignored in favor of other policy pursuits (perhaps accompanied by a public pretense of the pursuit of constituents' demands). In short, the nature and extent of her ignorance of the means required to realize her constituents' demands serve to determine the extent to which the policymaker pursues constituent-minded rather than other policy objectives. The relative epistemic burden of constituent-minded policymaking is a factor that contributes to determining how much of it we get.

One might complain that the use of a *ceteris paribus* condition trivializes our claim, as in politics and medicine *ceteris* is never *paribus*. But such a complaint would be mistaken. First, *ceteris paribus* conditions figure in all branches of inquiry across the sciences and humanities. Newton's laws of motion have *ceteris paribus* conditions attached to them, as do all of the laws of the social sciences. Of course, it is no less the case in, say, economics, that other factors are *not* held constant. Yet, the Law of Demand remains in operation: presumably, were the price of tea to increase ten-fold, a person who objects to the use of *ceteris paribus* conditions in scholarly inquiry would nevertheless be inclined to buy less tea. Science cannot proceed by considering phenomena in their entirety. Rather, science proceeds by considering the effects of one factor or of one set of factors, on the phenomenon of interest, while controlling other factors as far as possible. Second, we invoke the *ceteris paribus* assumption in order to focus on the effects of ignorance on decision-making. It is a simplifying assumption. It is important to note, however, that the assumption simplifies not only our epistemic circumstances, but those of the surrogate

decision-maker in our model. Dropping the assumption that all other potentially relevant factors are being artificially held constant would not necessarily improve the surrogate's epistemic condition and, thus, need not diminish her epistemic burdens or their significance for her decisions.

Among the further implications that Scheall draws from this thesis is the priority in the political realm of epistemic to normative considerations, i.e., the priority of questions concerning what policymakers *know enough* to do over questions of what they *should* do. "The policymaker's knowledge and ignorance serve to determine what policy objectives [she] is incited to pursue or to not pursue, but the opposite is not true: that a policymaker is incited to some objective is irrelevant to whether [she] possesses or can acquire the knowledge required of its realization...Epistemic burdens are factors in the determination of incentives, but incentives are irrelevant to epistemic burdens." Importantly for our purposes, this priority would seem to be fully general in decision-making, albeit largely obscured in personal decision-making for the reasons mentioned in the previous section.

To illustrate the point, consider the significance for decision-making of the principle that *ought implies can*.⁶ We typically apply this principle such that, when we consider what we ought to do, our options are pre-sorted for ignorance of relevant factors. We include among our options only things that we think we know (or can learn) enough to do. If we do not think that we can acquire adequate knowledge to achieve some end, we exclude it from, or deeply discount it in, our incentive structure. We either do not consider or do not take seriously options that we think

⁶ The main thesis does not ultimately hinge on acceptance of this principle, which is not without controversy. One might think that a relation weaker than logical implication holds between *ought* and *can*, and still accept that ignorance shapes our incentive structures. The ought-implies-can principle merely illustrates the point in an especially stark way.

impossible on the basis of the knowledge we possess or are likely to acquire in the process of their pursuit.

Such deliberations often occur sub-consciously, but occur they surely do. It is not likely to occur to someone ignorant of baking that they ought to make a loaf of bread when they find flour, water, salt, and yeast in front of them. But put those four ingredients in front of a baker, and baking a loaf will appear as an option, albeit perhaps only at a less-than-fully conscious level. The difference between the options available to two such people is determined by differences in their ignorance. We also typically exclude or discount options that require for their realization the substantial intervention of *luck* or other spontaneous forces. That is, the practical meaning of “can” in “ought implies can” is “can *deliberately*.” Thus, the practical significance of the principle, Scheall concludes, is that “*ought to X implies can X, which implies knows (or can learn) enough to X.*”⁷ It is empty to insist that an actor ought to pursue some end which, because she takes herself to be ignorant of pertinent knowledge, will either not appear or appear only as deeply discounted in her incentive structure. For instance, suppose that a person is swimming in deep water surrounded by circling sharks and that you are standing on a nearby ship next to an inflatable life boat, with no other life-saving devices in reach. Unfortunately, the instructions for inflation of the life boat are written in Sanskrit. You cannot save the person. There is no course of action available to you that will result in the preservation of the swimmer’s life. It is either false or simply meaningless to assert that you *ought* to save the person, because you *cannot* save the person. The epistemic burden that must be overcome in order to save the swimmer is too heavy for any person to lift in the time available. As the soon-to-be shark food calls for help,

⁷ That the significance of “can” in “ought implies can” must be something like “deliberately can” follows also from the fact that any other meaning seemingly trivializes the claim. That is, if the meaning of “can” in “ought implies can” is something like “possibly can” or “can, with luck,” then the principle prohibits no possible (in some undefined sense) course of action. We assume that the principle is not trivial.

spontaneously learning Sanskrit is not an option and will not occur to you as an option. Your ignorance gets you off the moral hook, but only because the epistemic burdens shape the options available to you and, in turn, what you can do.

All of this is to demonstrate that ignorance shapes both *options* and the *incentives associated with those options*. Ignorance and epistemic burdens serve to determine the menu of options and, thus, the option one ultimately chooses. If ought to X ultimately implies that the actor knows enough to X, then there can be no normative obligation for the actor to pursue ends the realization of which require more, or different, knowledge than she can acquire. If the policymaker cannot know enough to realize her constituents' policy demands, there can be no obligation for her to pursue them. If the principle that ought implies can (or some weaker, but relevantly similar, principle) is applicable to the ethics of patient surrogacy, then epistemic considerations are also prior to normative ones in these contexts. The surrogate must meet her epistemic burdens before she can reasonably be expected to satisfy any moral obligation associated with her surrogate decision.

EPISTEMIC BURDENS OF PATIENT SURROGATES

In this section, we aim to establish that patient surrogates occupy an epistemic position analogous to that of policymakers. If it is true both that, *other factors held constant, the relative weights of the epistemic burdens of competing policy objectives serve to determine the objectives that policymakers pursue*, and that the epistemic position of policymakers is analogous to the epistemic position of patient surrogates, then it is also true that, *other factors held constant, the relative weights of the epistemic burdens of competing treatment options serve to determine the options that patient surrogates pursue*.

There are several relevant points of similarity between the epistemic positions of policymakers and patient surrogates. First, both policymakers and patient surrogates ostensibly make decisions on behalf of their constituents / patients. Their epistemic ends—possession of the knowledge necessary to realize the preferences of the persons for whom they decide—are the same; yet, both kinds of surrogate decision-maker seem similarly misplaced to possess knowledge of these preferences. At least, there is no reason to think in either case that decision-makers always (perhaps ever) automatically know the content of these preferences without first engaging in some sort of learning process.

It is possible for a surrogate's own preferences to align with those of her constituents / patient, but such alignment should not be assumed. A policymaker's own personal preferences might be most easily satisfied by pursuing her constituents' policy demands or by merely making a public pretense of their pursuit. Similarly, a physician's preference for providing meaningful teaching experiences for her residents may also realize her patient's preference. A family member may benefit from the inheritance that results from her decision to withdraw life-sustaining treatment, but this might also be the decision that best accords with the patient's preferences. In any case, even such alignment of preferences does not ensure knowledge of aligned preferences—that a surrogate's preferences in fact align with those of the person for whom they decide is not sufficient for knowledge of this fact. The surrogate always has some work to do in order to realize her epistemic ends.

The problem of knowledge of relevant preferences is quite complex in contexts of medical-surrogate decision-making. A physician may have insight into a patient's *medical* interests, but these are only a subset of the patient's total interests (including the patient's interest in having her preferences satisfied). Physicians are typically ignorant of those interests that outstrip the merely medical, such as a patient's preference to not be intubated. This is why

physicians rely on family and friends of incapacitated patients. These are the people who know the patient intimately and can most reliably convey her preferences. But even the *most* reliable sources of a patient's preferences may be nevertheless unreliable. There are many ways for a loved one's belief about a patient's preferences to be confounded. There may be no consensus about her preferences among a patient's friends and family. But even supposing that loved ones reliably express a patient's preferences, they may not understand the relation between these preferences and available treatment options, or how to express relevant preferences to physicians in such a way that the latter proffer fruitful advice.

Second, beyond ignorance of relevant preferences, both policymakers and patient surrogates are often ignorant to some extent of how to achieve relevant goals implied by these preferences. We noted above that Scheall assumes policymakers to be fully knowledgeable of their constituents' policy demands. This assumption is made merely for the sake of simplifying the analysis. In fact, policymakers are often ignorant of what their constituents want from them. They have access to polls that provide limited evidence of these demands, but that are, in any case, infamously defeasible. They have anecdotal evidence by way of contact with individual voters, extrapolations from which are difficult, indeterminate, and regularly falsified. And they may have evidence from historical trends that often change in ways that can only be recognized retrospectively.

While physicians may be able to inform patient surrogates of the relevant treatment options, they typically do not act as surrogates themselves. Family and friends may know the patient's preferences, but not how to achieve these preferences, even after the physician has informed them of the treatment options available. For example, surrogates may justifiably believe that the patient's preferences are best served by attempting to preserve life at all costs, but they might be ignorant of the violence and invasiveness required to serve these preferences.

This ignorance is often barely mitigated by the physician informing them of this violence and invasiveness, which is more persuasively felt than merely spoken of. Or surrogates might think that the patient's preferences are best served by bringing her home to die (relatively) peacefully under hospice care, but they may be ignorant of some of what is required to respect these preferences, such as the steps necessary to ensure medical stability while the patient is transported. So, although patient surrogates may be in the best position to express the patient's preferences, they may be ignorant of how to achieve them. Conversely, physicians may be knowledgeable of how to achieve patient preferences insofar as they know them, but not know them very well. Effective communication between physician and surrogate can partially repair both parties' ignorance, but never so much that the surrogate's epistemic burden is reduced to the level of a capacitated patient able to make her own medical decisions. The surrogate's epistemic burdens will always be greater.

Policymakers and patient surrogates occupy similar epistemic positions: they are frequently ignorant of either relevant constituent / patient preferences or how to achieve them. Realizing their ostensible obligations requires overcoming their epistemic burdens. If it is true of policymakers that these epistemic burdens shape their incentives, then it is also true of patient surrogates. *Other factors held constant, the relative weights of the epistemic burdens of competing treatment options serve to determine the options that patient surrogates pursue.*

IMPLICATIONS

Epistemic burdens are likely to differ depending on whether the patient surrogate is a loved one or a physician. As surrogates, family members and friends are likely to confront heavier epistemic burdens in understanding the relevant medical circumstances and translating

their knowledge, such as it is, of the patient's preferences into effective treatment. These burdens might be partially met in cases where the causal mechanism underlying the patient's condition is relatively simple and, thus, comparatively easy to understand, or by a physician filling in some of the missing knowledge. But, assuming that patient surrogates without medical expertise confront their greatest epistemic burdens in overcoming ignorance of relevant medical facts and in translating these facts into a treatment conducive to the patient's preferences, the path of least epistemic resistance is likely to flow around, rather than through, this ignorance. Assuming other relevant factors are equal across treatment options, friends and family members acting as patient surrogates will pursue treatment options that are epistemically easier, i.e., those that are already understood or are relatively easy to understand.

Of course, treatment options might vary along any number of different, non-epistemic, dimensions. These dimensions may include the likelihood of realizing the patient's preferences, financial cost, severity of the patient's pain or the length of time the patient must spend in pain, or the physician's time and effort. It is rarely the case that other relevant factors are truly equal across treatment options. But this does not mean that epistemic burdens do not shape a surrogate's incentives.

Suppose that a surrogate justifiably believes that the patient would not want to be permanently dependent on life-sustaining technology, but is ignorant whether the patient would prefer short-term dependence with a small chance of recovery to withdrawal of treatment and likely death. Further suppose that the recommended treatment leading to the greatest likelihood of recovery requires a tracheostomy and a percutaneous endoscopic gastrostomy, followed by transfer to a long-term acute-care facility for multiple months. The alternative is withdrawal of life-sustaining treatment and likely death in the near-term. There are significant epistemic burdens in this case that shape the surrogate's incentives. The first option requires that the

surrogate understand not only the medical facts associated with the treatment, but also how the treatment integrates with their existing beliefs about the patient's preferences. The second option requires less understanding of the relevant medical facts—relative to the consequences of the treatment option, death is a comparatively stark and easily understood result—but still requires that the surrogate understand these consequences in light of her knowledge of the patient's preferences. The epistemic burdens of the two options may be different, but so too are the values with which they intersect. The first option requires the patient to bear greater physical pain, as recovery from the treatment involves significant physical and emotional suffering. It will also be more expensive. But its greatest potential cost may be a possible violation of the patient's dignity, if treatment is contrary to her preferences. Of course, given that the second option leads to likely death, the potential costs associated with it might be even greater.

Our point is simply that the relative epistemic burdens of the two options will be a significant factor in determining how the surrogate evaluates them. That the first requires epistemic work (e.g., knowing what is involved in being dependent upon tracheostomy and ventilator, or extracorporeal membrane oxygenation, or long-term dialysis, etc.) that the second does not is relevant to where the options will appear in the surrogate's incentive structure. Whether the surrogate explicitly considers them or not, the epistemic burdens she must confront in pursuing various options will guide her decision-making. When other factors are *not* equal across options, the relative epistemic burdens of competing options may or may not be the deciding factor, but when *ceteris is paribus*, the relative weights of competing epistemic burdens will determine the option that the surrogate selects.

Qua surrogate, a physician is likely to confront a heavier epistemic burden in discovering the patient's preferences than in knowing the medical facts or in determining an effective

treatment *given* knowledge of relevant preferences and facts.⁸ When a physician acts as a surrogate, it typically means that there is no one else able and willing to fill the role. In such cases, the attending physician is likely to be especially ignorant of the patient's preferences, though perhaps knowledgeable of how to satisfy a range of interests once these are determined. When other factors are constant across treatment options, the physician will pursue the least epistemically burdensome option.

Suppose an incapacitated patient is unrepresented by a surrogate, so that the physician must make all medical decisions. Suppose also that the patient is terminally ill. One option is to continue aggressive treatment with the aim of prolonging life. The other option is to focus on comfort and palliation at the end of life. The physician knows how to achieve these goals. What the physician doesn't know is which option best satisfies the patient's preferences, because this information is not available—the patient is unrepresented. Whether the physician chooses aggressive treatment or palliation is partly dependent on the relative weights of their respective epistemic burdens. The default assumption in such cases is often that a patient wants to continue to live. To overcome this assumption—to know that the patient prefers the second option—the physician requires evidence of the patient's preference for comfort over continued life. Thus, under these circumstances, the epistemic burden of palliative treatment is weightier and, if other relevant factors are equal, the treatment option that the physician will avoid. With more information, however, the physician's epistemic burdens may shift in such a way that palliation becomes the more attractive option. Suppose that a friend of the patient—someone with limited knowledge of the patient's interests, but who is otherwise unable to make decisions for the patient—reveals that the patient valued her personal independence above all else. This new

⁸ If one prefers the language of know-how and know-that, the physician is more likely to lack knowledge that the patient's preferences are X, rather than to lack knowledge of how to achieve a particular medical goal.

information repairs some of the physician's ignorance of the patient's preferences, lowering the epistemic burden associated with choosing palliation (i.e., the treatment associated with less dependence on life-sustaining technology).

As a final example, consider advanced directives, or other types of documentation that a patient might compose to indicate their preferences for medical treatment in the event they are incapacitated. These documents derive their moral authority from the notion that treating a patient according to their known capable preferences is to respect their autonomy. But when they are needed, such as when a person is unconscious and on a ventilator, surrogates often dispute their accuracy, testifying to medical staff that the patient changed her mind or that she lacked the capacity to make rational decisions when composing the directive. Such claims may sometimes be correct. But advanced directives also function to reduce the epistemic burdens of surrogate decision-makers. In the case of disputed directives, there are two options: pursue the veracity of the claim that the directive is unreliable or do what the directive says. There is little or no epistemic burden in doing what the directive says. However, investigating the veracity of the claim that the directive is unreliable is relatively epistemically burdensome, perhaps significantly so. To pursue this option, the surrogate must do other things to repair her ignorance of the patient's preferences. The relative burdens of the two options shape the incentives of these options and, thus, serve to determine the option that the surrogate pursues.

EXPLANATION

It is widely recognized that patient surrogates face epistemic burdens. We claim something stronger: these burdens structure and shape the patient surrogate's incentives in a way that significantly affects the treatment options she pursues. Making a surrogate medical decision is

not simply a matter of substituting one person's judgment for another's, it is also a matter of the knowledge which informs a surrogate's judgment. That epistemic burdens shape incentives and, thus, decisions, is a feature of all decision-making, personal and surrogate.⁹ The unique epistemology of surrogate decision-making only makes this feature more obvious. In this section, we offer two complementary and mutually reinforcing explanations for the fact that the significance of epistemic burdens is more apparent in surrogate cases. That the phenomenon can be explained is further evidence that it is a real feature of surrogate decision-making.

First, the philosophical explanation. We all have preferences and beliefs about these preferences. In most cases, one's own belief that one has a preference for X is a confident belief. One's own credence that one wants a cup of coffee is typically very high. This higher-order belief—a belief about another mental state—has a high credence. In determining whether one wants a cup of coffee, the epistemic burden is very low. One needs only to introspect to overcome this burden. However, suppose that one is wondering about whether a friend wants a cup of coffee. The credence in the belief that one's friend wants a cup of coffee will be significantly lower. That is, one's confidence in the belief that one's friend wants a cup of coffee will be less than one's confidence that one wants a cup of coffee. To equalize the credences of the two higher-order beliefs, more is required of the belief about one's friend's preference for a cup of coffee. There is additional epistemic work to do. One must do things like ask whether the friend wants a cup of coffee or otherwise collect further evidence for or against the proposition. This is more burdensome than the mere introspection required to discover one's own preferences. This relation between higher-order beliefs about lower-order preferences generalizes

⁹ Decisions follow the path of least epistemic resistance in the case of personal decisions, also. Suppose that a physician gives you two treatment options, but you are uncertain which option best serves your own interests. Our claim is that, other things being equal, the option you pursue will be the one that imposes the lowest epistemic burden—the option about which your ignorance is most easily repaired.

to any surrogate decision. Relative to a lower-order preference for X, the epistemic burden of believing that another person has a preference for X will be greater than believing that one has a preference for X, *because* the credence of the belief in another person's preference for X is lower than the credence of the belief in one's own preference for X. This is why surrogate decisions are especially epistemically burdensome—more is required to bring the credence of the higher-order belief up to a degree that is sufficiently high to guide action.

Second, the psychological explanation. Epistemic burdens influence incentives in all cases, including in cases of personal decision-making where one decides for oneself. However, with respect to personal decision-making, much of this influence is experienced only sub-consciously. The options from which we consciously choose have already been pre-processed for epistemic considerations on a less-than-fully conscious level. This is easily recognized in the fact that, when deciding for oneself, no one includes options that they consider impossible among the array of possibilities from which they ultimately make some choice (e.g., non-bakers do not consider baking bread an option given some flour, water, salt, and yeast). However, in cases of surrogate decision-making, the options are not automatically pre-processed for the decision-maker's ignorance. A sane person will never leave himself an epistemically impossible option (e.g., spontaneously learn Sanskrit) and will pre-discount epistemically difficult options accordingly, but there is no such guarantee in surrogate cases.¹⁰ You are never going to give yourself a goal that you cannot know enough to achieve, but somebody else might give you such a goal. You would never set for yourself the goal of spontaneously learning Sanskrit, but your friend surrounded by sharks might impose it upon you. If our argument here is sound, you will ignore this request.

¹⁰ An option judged to be impossible is *ipso facto* judged to be epistemically impossible, i.e., to bear an infinitely heavy epistemic burden.

Consider these explanations in the context of surrogate decision-making. For example, it often happens in acute care hospitals that even though a patient is incapable of making her own medical decisions, her family will agree that she has expressed a preference not to be intubated. They might further agree that the patient would want to be resuscitated in the event of cardiac arrest. However, respecting the first interest often requires violating the second and *vice versa*. Assuming the relevant attitudes accurately reflect the patient's preferences, she has left her surrogates an epistemically impossible task. The surrogate's credences about her beliefs concerning the patient's preferences are likely to be quite low. If the decision concerned her own medical treatment, such a conflict could be easily resolved via introspection and rational deliberation. But this method is not available when the surrogate has inherited the conflict from the patient. To decide on behalf of the patient, the surrogate must resolve the conflict by raising the credence of either her belief that the patient's preference for avoiding intubation outweighs the patient's preference for being resuscitated following cardiac arrest or the contrary belief. The surrogate confronts a unique epistemic burden merely in virtue of being a surrogate. Our claim is that, if other factors are not equal across options, the relative weight of the epistemic burdens of competing options will partly determine, and, if other things are equal, fully determine, which option the surrogate ultimately pursues.

CONCLUSION

It can be difficult to discover what another person wants and how to achieve these preferences. The weight of this challenge serves to determine the option that the surrogate pursues. This is not a prescription or a guide to surrogate decision-making—it is a glimpse of the structure of these decisions. That surrogate decisions are structured this way is a feature of our

own psychology, and of how we evaluate preferences and beliefs about preferences. That surrogate decisions flow around the path of least epistemic resistance is not something we can change or avoid. But recognizing this fact helps explain how and why surrogates make the decisions that they do.

REFERENCES

Berger JT, DeRenzo EG and Schwartz J (2008) Surrogate Decision Making: Reconciling Ethical Theory and Clinical Practice. *Annals of Internal Medicine* 149(1). DOI: 10.7326/0003-4819-149-1-200807010-00010.

Brudney D (2018) The Different Moral Bases of Patient and Surrogate Decision-Making. *Hastings Center Report* 48(1): 37–41. DOI: 10.1002/hast.809.

Buchanan AE, and Brock DW (1989) *Deciding for Others: The Ethics of Surrogate Decision Making*. Studies in Philosophy and Health Policy. Cambridge University Press.

Kibbe B and Ford PJ (2016) What's Knowledge Got to Do with It? Ethics, Epistemology, and Intractable Conflicts in the Medical Setting. *The Journal of clinical ethics* 27(4): 352–358.

Moorman SM and Carr D (2008) Spouses' Effectiveness as End-of-Life Health Care Surrogates: Accuracy, Uncertainty, and Errors of Overtreatment or Undertreatment. *The Gerontologist* 48(6): 811–819. DOI: 10.1093/geront/48.6.811.

Pope TM (2012) Legal Fundamentals of Surrogate Decision Making. *Chest* 141(4): 1074–1081.

DOI: 10.1378/chest.11-2336

Scheall, S. forthcoming. Ignorance and the Incentive Structure confronting Policymakers.

Cosmos + Taxic. Issue TBD Draft available at:

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3240070

Scheall, S, Butos, W, and McQuade T. Forthcoming. “Social and Scientific Disorder as Epistemic Phenomena, or The Consequences of Government Dietary Guidelines.” *Journal of Institutional Economics*, Issue TBD

Shalowitz DI, Garrett-Mayer E and Wendler D (2006) The Accuracy of Surrogate Decision Makers. *Archives of Internal Medicine* 166(5): 493. DOI: 10.1001/archinte.166.5.493.

Suhl J, Simons P, Reedy T, et al. (1994) Myth of Substituted Judgment. *Archives of Internal Medicine* 154(1). DOI: 10.1001/archinte.1994.00420010122014.;