

CASE AND COMMENTARY

Should Aggregate Patient Preference Data Be Used to Make Decisions on Behalf of Unrepresented Patients?

Nathaniel Sharadin, PhD, MA

Abstract

Patient preference predictors aim to solve the moral problem of making treatment decisions on behalf of incapacitated patients. This commentary on a case of an unrepresented patient at the end of life considers 3 related problems of such predictors: the problem of restricting the scope of inputs to the models (the “scope” problem), the problem of weighing inputs against one another (the “weight” problem), and the problem of multiple reasonable solutions to the scope and weight problems (the “multiple reasonable models” problem). Each of these problems poses challenges to reliably implementing patient preference predictors in important, high-stakes health care decision making. This commentary also suggests a way forward.

Case

Mr T is an 88-year-old black man with squamous cell carcinoma of the throat metastatic to the brain, complicated by recurrent seizures. The patient is admitted from his medical care facility because he is no longer able to swallow due to tumor progression, and there is concern for pending airway occlusion. Mr T’s wife passed away many years ago and he has no children, other family, or friends. Mr T has had altered mental status, likely due to the brain metastases and recurrent seizures. He is unable to speak and intermittently makes uninterpretable vocalizations. The oncology team is not offering to continue Mr T’s treatment and has predicted that he will likely die within weeks without interventions, such as a percutaneous endoscopic gastrostomy (PEG) tube and tracheostomy, and in less than 6 months in any case.

In a multidisciplinary meeting, Dr J references an article featuring a predictive model that shows that being black is consistently associated with preferring more end-of-life treatment in medical settings.^{1,2} Dr J asks, “Should we conclude from this article that Mr T would want a PEG tube and tracheostomy?” Dr O, an oncologist, voices doubt about applying population-level data to decisions involving particular patients. Dr J says, “Yes, I wondered about that, too.” He then mentions a study showing that 78.9% of patients would prefer to have aggregate data incorporated into processes of making clinical decisions on their own behalf.³ They further discuss and wonder what to do.

Commentary

Care for incapacitated patients generates familiar moral problems. Foremost among these is the problem of making treatment decisions on their behalf. The cause of the problem is obvious. Typically, medical professionals defer to a patient's (informed) decisions, but an incapacitated patient clearly can't make the relevant treatment decisions. What to do? The obvious thing to do is for medical professionals to defer to some other source with the moral standing to speak on behalf of the incapacitated patient.

Traditionally, the sources thought to have the moral standing to direct treatment on behalf of incapacitated patients are of 2 broad sorts: advance directives and surrogates. In an advance directive, a patient expresses her preferences over a range of treatment decisions in advance of her incapacitation. It's obvious why advance directives plausibly have the relevant sort of moral standing: deferring to a patient's advance directive is a way of deferring to that patient's own decisions at one remove, as it were.⁴⁻⁷ In surrogate decision making, third parties make treatment decisions on behalf of the incapacitated patient. These surrogates are supposed to have the relevant moral standing either because of a formal relationship—say, medical power of attorney—or because of a more informal relationship that intuitively justifies such moral standing—say, in the absence of medical power of attorney, the surrogate's being a spouse or an adult child. (Of course, the two often go together, as relatives make natural candidates for power of attorney.)^{8,9}

Recently, a third source of moral standing for making treatment decisions on behalf of incapacitated patients has been proposed: so-called patient preference predictors (henceforth, PPPs).^{1,10,11} Very briefly, PPPs are statistical models that predict the treatment preferences of a patient described by a combination of known demographic variables—such as age, educational level, religion, and so on—based on aggregate data about the treatment preferences of persons with demographic profiles similar to the patient.¹²⁻¹⁴ Before turning to the moral standing of PPPs, let me mention one issue in order to set it firmly aside in what follows. One immediate reaction to the suggestion that we predict an incapacitated patient's preferences using patient demographic descriptors that are likely to be known—especially, race, gender, and age—is that it's radically unclear how robust such predictions might be. The first thing to say in response is that there are in fact some broad, relatively strong correlations between such factors and, for instance, preferences for palliative care at the end of life.¹⁵ Now, it's true that we lack the kind of broad-based statistical data over a range of treatment options that would be necessary to make PPPs widely useful in a clinical context. But it's trivial to imagine how we might go about acquiring such data. For instance, Rid and Wendler propose a national PPP survey of competent adults that could be used to correlate the most likely relevant demographic descriptors with patients' preferences over a wide range of treatment options in a wide range of scenarios.¹ There are [other possibilities](#),

too. In any case, this practicability question isn't one that I'll be concerned with in what follows.

So, why think PPPs have any moral standing to determine treatment decisions on behalf of patients? The intuitive argument is straightforward. Recall that the best-case scenario is one in which a patient's treatment reflects her own informed decisions. Advance directives and surrogates are ways of trying to epistemically access—in more or less direct ways—what we think those informed decisions would have been. Moreover, that they are directed toward the patient's own preferences is what explains why they have the kind of moral standing they do in the process of making treatment decisions on behalf of incapacitated patients. Likewise, the use of PPPs is a way of trying to epistemically access an incapacitated patient's own preferences with respect to her treatment. Intuitively, then, if advance directives and surrogates enjoy the relevant sort of standing, then PPPs ought to enjoy that same standing.

In any case, I'm going to assume we're sometimes justified in using PPPs to make treatment decisions on behalf of incapacitated patients. This assumption is compatible with restrictions on their permissible use—perhaps we think they can only be used to supplement surrogate decision making or that they can only be used in the absence of both an advance directive and a suitable surrogate.¹² These restrictions won't concern me here. Instead, my interest is more narrowly focused on conceptual challenges to reliably implementing patient preference predictors in such high-stakes health care decision making.

Suppose we are justified in sometimes using PPPs; there are 2 interesting problems that arise in the context of using PPPs that don't arise for advance directives or surrogates. (There are, of course, other problems. For instance, Rid and Wendler¹—2 proponents of the use of PPPs—point to the **danger of stereotyping** inherent in this sort of statistical model using readily apparent demographic descriptors. In other work,¹² I've also argued that PPPs face a closely related problem—analogue to a familiar problem in legal scholarship—having to do with moving from bare statistical evidence to normative conclusions. These problems, while important, won't concern me here.) I'll argue that the presence of these 2 interesting problems generates a third, quite serious difficulty for the use of PPPs in a clinical setting. I'll then close by briefly explaining one strategy for solving this third difficulty as a way of laying out a fruitful direction for future research. But first, I discuss the 2 initial problems.

Two Initial Problems With PPPs

The scope problem. Recall that a PPP can usefully be thought of as a function that takes as input known characteristics of the patient and produces a probabilistic prediction regarding the patient's preference for (or against) a particular medical intervention based on aggregate data from people who share similar characteristics with the patient. Not all

known information about a patient ought to be used in producing these probabilistic predictions. In particular, we ought to exclude uncontroversially false (but suitably important) normative beliefs that might be correlated with patients' preferences, and we ought to exclude some (but not all) normatively irrelevant characteristics of patients that might be correlated with patients' preferences. Hence the *scope problem*. The scope problem is the problem of restricting, in a principled way, what sort of information can be used as inputs to a PPP. Let me comment briefly on false normative beliefs and irrelevant characteristics as a way of making intuitive the case that these inputs ought to be excluded from PPPs.

Take false normative beliefs first. Suppose we knew that victims of long-term domestic abuse were much less likely to prefer palliative care in particular circumstances. Suppose further that we had a normative explanation of this fact, viz, that victims of domestic abuse have a diminished sense of their own worth as compared to the worth of others—they (incorrectly) believe that they somehow “deserve” suffering or that their suffering somehow counts for less. I take it for granted that we would want to exclude such a false belief from our PPP when using it to deliver a verdict regarding a patient's preferences in treatment. The argument for this claim is straightforward; spelling it out in detail would require too much space in the present context. Roughly, you can't derive a correct normative verdict regarding how you ought to treat a patient, ie, what medical treatment to provide to her, from that patient's preference grounded in a manifestly false normative belief, ie, that her suffering is less morally important than others'.

Now take normatively irrelevant facts about patients. We know that religion is strongly correlated with preferences regarding end-of-life care.¹⁶ And religious identity is precisely the kind of thing we intuitively want to include as input to a PPP. So far, so good. But now notice that there's nothing that rules out the possibility that (say) whether one prefers the NFL to the NBA correlates quite strongly with one's preferences regarding end-of-life care. One response would be to simply include that sports preference among the demographic descriptors that, if known, could be used by a PPP to predict a patient's preferences over the relevant treatment options. But this is strongly counterintuitive. While there doesn't appear to be any—or, at least, much—intuitive resistance to taking religion as a relevant input to a model designed to predict what an incapacitated patient would want, there is a strong intuitive case against taking sportsball preference into account in the same way. Again, spelling out the details of this argument would require too much space in the present context, but here, roughly, is the idea. Some facts about one's self are more or less central to one's identity. And when it comes to making life-altering—and potentially life-ending—decisions, it's perfectly natural to want those decisions to be made on the basis of those facts about one's self that are central—for most people, their religious identity—rather than on the basis of those facts that float around the periphery—again, for most people, their preferring the NBA to the NFL, or vice versa.

Hence, the scope problem is this: How should we restrict, in a principled way, the sort of information that can be appropriately used as the input to a PPP? (If you're not particularly bothered by this problem because you think it's obvious that all information ought to be allowed as input to a PPP, the scope problem still generates the problem that below I call the multiple models problem.)

The weight problem. The second problem is what I'll call the *weight problem*. The weight problem is the problem of explaining, in a principled way, how to correctly weight the information that serves as input to a PPP. Even if we have a solution to the scope problem in hand, this solution does not tell us how to weight the various inputs to a PPP in coming to a final verdict. For instance, while we might agree that (say) someone's ethnicity ought to serve as input to a PPP, you might think that predictions based on ethnicity ought to be weighted less as compared to (say) predictions based on religion, and I might think the opposite. This disagreement over the correct weight assignment for these 2 factors could obviously lead to disagreement over the correct treatment decision. (And assuming there's some fact of the matter about how we ought in fact to treat the patient, it can lead predictably to patient harm.)

One natural response to this problem is to posit that there is some objective principle for weighting the inputs to the PPP that assigns some specific weight to each—0.3 to ethnicity, 0.6 to religion, say. But it's hard to see what the principled way of determining the values of such an objective weighting might be. (This is what I meant, above, in saying that the problem was one of explaining in a principled way how to weight the various inputs to a PPP.) The difficulty arises because it's hard to imagine what sort of data we could acquire—either by sampling existing data or by gathering new data—that would lead us to reasonably conclude that always, everywhere, for patients who are (say) both Asian and Catholic, their ethnicity ought to be weighted (say) $\frac{1}{3}$ as much as their religion. This is not to say that we might not gain some information that would lead us to a range of likely values for these relative weights. But then we face the multiple models problem, which I'll turn to now.

Multiple reasonable models problem. In order for the use of PPPs in a clinical context to be morally permissible, the scope and the weight problems require solutions; as yet, they have none. But things are even more complicated than this. For, together with some commonsense observations about the difficulty of the issues involved, these problems generate what I'll call the *multiple reasonable models problem*. The multiple reasonable models problem is this: there will be reasonable disagreement over how to solve the scope and the weight problems.

A very quick word about *reasonable* disagreement. To say that these disagreements are reasonable is to say that well-informed persons reasoning together in good faith over what information should be included in our PPPs (the scope problem) and how to weight

that information (the weight problem) could in principle continue to disagree and that neither need be making a normative or nonnormative mistake correctable by further analysis of the values at stake or the data available.¹⁵ Such reasonable disagreement is presumably due to what John Rawls calls, in a related context, the very difficult “burdens of judgment” when it comes to normative questions.¹⁷ The sort of reasonable disagreement I have in mind is analogous to (indeed, might be a species of) Rawls’s idea of reasonable disagreement in a liberal political society over what comprehensive doctrine is correct. In the present context, the claim is not that the presence of reasonable disagreement informs the correct political arrangement; instead, it is simply that in virtue of reasonable people’s reasonable disagreement over the correct “comprehensive doctrine,” there will be downstream reasonable disagreement over how to solve the scope and weight problems, ie, disagreement over what factors ought to be used in predicting agents’ preferences and how to weight those factors. This is not to say that all such disagreement over how to solve these problems would be reasonable. This is just to say that some such disagreement would be.

But, then, given that some such disagreement is reasonable, in clinical contexts medical professionals will be forced to decide between equally reasonable PPPs that deliver equally reasonable but incompatible verdicts regarding patients’ preferences, ie, each PPP is an equally reasonable function that takes us from known information about a patient to (competing) verdicts regarding her preferences for medical treatments. Hence, the multiple reasonable models problem arises.

We can see that the scope problem allows for reasonable disagreement by focusing on either of the 2 illustrations of it I offered above: the case of false normative beliefs and the case of normatively irrelevant facts about patients. I offered what I took to be uncontroversial instances of the sorts of things we want to exclude from our PPP: on the one hand, the correlation between patients’ normatively false belief that they don’t merit care (because of being victims of domestic abuse) and their care preferences, and, on the other hand, the correlation between patients’ normatively irrelevant preference for the NBA over the NFL and their care preferences. But it should be obvious that there can be reasonable disagreement of the sort just described over both these types of cases. For instance, you and I might reasonably disagree over whether some patient’s normative belief is actually false or not; or we might disagree over whether this or that characteristic is suitably central to patient identity to merit inclusion in the PPP.

To see that the weight problem allows for reasonable disagreement, notice that even if we agree on whether to allow both ethnic and religious identity to inform our PPP, we might reasonably disagree over how to weight those factors. Returning to the discussion above, you might think we ought to assign ethnicity a weight $\frac{1}{2}$ that of religion in the present context; I might disagree. This disagreement would not be unreasonable;

instead, it would presumably be due to what is itself reasonable disagreement on how to value religious as compared to ethnic identity.

The multiple models problem is hence where the scope and the weight problems truly earn their keep: what these observations show is that, even assuming we've provided *some solution or other* to the scope and weight problems, this solution will just be one among several reasonable alternative solutions. Hence medical professionals will still face reasonable disagreement over *what PPP to use*; in a clinical context, this is a real barrier to the use of PPPs in making treatment decisions on behalf of incapacitated patients.

Prospects for Solving the Multiple Reasonable Models Problem

Let me close with a couple of quick remarks on how I think we should move forward. We can begin by noticing that there appears to be an analog of the multiple reasonable models problem when it comes to the use of surrogates. This is clearest in the absence of a patient's formal designation of a surrogate: in that case, it's possible for reasonable people to reasonably disagree over who should serve as the patient's surrogate. For instance, we might reasonably disagree over whether it should be (say) the patient's relatively new spouse or her adult child. When this happens, we appear to have a way to resolve the disagreement or at least to move forward. Assuming the patient resides in a state that does not specify an order in which relatives should be identified as surrogates, we can split the decision making between the surrogates and encourage them to decide together.

What might the analog of this solution look like in the case of PPPs? My suggestion, which for reasons of space I can only gesture to here, is straightforward: we should give equal weight to competing PPPs that are reasonable in the way identified above. Two PPPs will count as competitors in any given case when they deliver differing probabilistic judgments regarding a particular patient's preferences over the available treatment options.^{18,19} Two comments are in order. First, what this means in practice is that we shall need to weigh the verdicts of competing PPPs by updating our prior credences about a patient's preferences in the ordinary way by, for instance, conditionalizing on those competing verdicts as independent pieces of evidence. Second, the ethical benefit to patients of this approach should be clear. For notice that the trouble here is caused by the fact that there's reasonable disagreement over which PPP—which way of solving the scope and weight problems—might be correct. But if that disagreement really is reasonable, then our practice ought to reflect the underlying normative uncertainty, and patients deserve treatment that does so. After all, patients might themselves reasonably have one or the other view of the matter. So an approach that instead simply plumped for one or the other PPP that delivered incompatible verdicts regarding a particular patient's preferences on a particular occasion appears viciously arbitrary—not just from the theoretical point of view, but also from the point of view of patients themselves.

Of course, there is much more to say about the merits of (and problems with) this approach. And we shall still need some account of which ways of solving the scope and the weight problems count as reasonable as opposed to unreasonable (and so are taken into account in the way just suggested). That is a profitable direction for future theoretical research on the use of PPPs in clinical contexts.

References

1. Rid A, Wendler D. Use of a patient preference predictor to help make medical decisions for incapacitated patients. *J Med Philos*. 2014;39(2):104-129.
2. Barnato A, Anthony DL, Skinner J, Gallagher PM, Fisher E. Racial and ethnic differences in preferences for end-of-life treatment. *J Gen Intern Med*. 2009;24(6):695-701.
3. Wendler D, Wesley B, Pavlick M, Rid A. A new method for making treatment decisions for incapacitated patients: what do patients think about the use of a patient preference predictor? *J Med Ethics*. 2016;42(4):235-241.
4. Orentlicher D. Advance medical directives. *JAMA*. 1990;263(17):2365-2367.
5. Levi BH, Green MJ. Too soon to give up: re-examining the value of advance directives. *Am J Bioeth*. 2010;10(4):3-22.
6. Miles SH, Koeppe R, Weber EP. Advance end-of-life treatment planning. A research review. *Arch Intern Med*. 1996;156(10):1062-1068.
7. Lo B, Steinbrook R. Resuscitating advance directives. *Arch Intern Med*. 2004;164(14):1501-1506.
8. Buchanan AE, Brock DW. *Deciding for Others: The Ethics of Surrogate Decision Making*. Cambridge, UK: Cambridge University Press; 1990.
9. Johansson M, Broström L. Counterfactual reasoning in surrogate decision making—another look. *Bioethics*. 2011;25(5):244-249.
10. Shalowitz DI, Garrett-Mayer E, Wendler D. How should treatment decisions be made for incapacitated patients, and why? *PLoS Med*. 2007;4(3):e35.
11. Rid A, Wendler D. Treatment decision making for incapacitated patients: is development and use of a patient preference predictor feasible? *J Med Philos*. 2014;39(2):130-152.
12. Sharadin NP. Patient preference predictors and the problem of naked statistical evidence. *J Med Ethics*. 2018;44(12):857-862.
13. John S. Patient preference predictors, apt categorization, and respect for autonomy. *J Med Philos*. 2014;39(2):169-177.
14. John SD. Messy autonomy: commentary on Patient Preference Predictors and the Problem of Naked Statistical Evidence. *J Med Ethics*. 2018;44(12):864.
15. Miesfeldt S, Murray K, Lucas L, Chang CH, Goodman D, Morden NE. Association of age, gender, and race with intensity of end-of-life care for Medicare beneficiaries with cancer. *J Palliat Med*. 2012;15(5):548-554.
16. Burdette AM, Hill TD, Moulton BE. Religion and attitudes toward physician-assisted suicide and terminal palliative care. *J Sci Study Relig*. 2005;44(1):79-93.

17. Rawls J. *Political Liberalism*. New York, NY: Columbia University Press; 1996.
18. Elga A. Reflection and disagreement. *Noûs*. 2007;41(3):187-217.
19. Kelly T. Peer disagreement and higher order evidence. In: Goldman A, Whitcomb W, eds. *Disagreement*. Oxford, UK: Oxford University Press; 2010:111-174.

Nathaniel Sharadin, PhD, MA is an assistant professor of philosophy at the College of New Jersey in Ewing Township. His work is primarily on normative issues in ethics and epistemology.

Editor's Note

The case to which this commentary is a response was developed by the editorial staff.

Citation

AMA J Ethics. 2019;21(7):E566-574.

DOI

10.1001/amajethics.2019.566.

Conflict of Interest Disclosure

The author(s) had no conflicts of interest to disclose.

The people and events in this case are fictional. Resemblance to real events or to names of people, living or dead, is entirely coincidental. The viewpoints expressed in this article are those of the author(s) and do not necessarily reflect the views and policies of the AMA.