Patient preference predictors and the problem of naked statistical evidence

Nathaniel Paul Sharadin

Correspondence to

Dr Nathaniel Paul Sharadin,
Department of Philosophy, The
College of New Jersey, Ewing,
NJ 08628, USA;
natesharadin@gmail.com

Received 2 August 2017

Revised 20 April 2018

Accepted 9 May 2018

Published Online First

12 June 2018

ABSTRACT

Patient preference predictors (PPPs) promise to provide medical professionals with a new solution to the problem of making treatment decisions on behalf of incapacitated patients. I show that the use of PPPs faces a version of a normative problem familiar from legal scholarship: the problem of naked statistical evidence. I sketch two sorts of possible reply, vindicating and debunking, and suggest that our reply to the problem in the one domain ought to mirror our reply in the other. The conclusion is thus conditional: if we think the problem of naked statistical evidence is a serious problem in the legal domain, then we should be concerned about the symmetrical problem for PPPs.

INTRODUCTION

Treatment decisions should be made by patients on the basis of their informed preferences. Many times, patients are incapable of making such decisions. How ought treatment decisions for incapacitated patients be made? Two familiar answers involve the use of advance directivesⁱ and of surrogates.ⁱⁱ Recently, a third answer has been proposed: using statistical modelling, we can attempt to predict patients' preferences with respect to treatment and then make treatment decisions on the basis of these predicted preferences.ⁱⁱⁱ In this article, I discuss the normative status of these *patient preference predictors* (PPPs).

The detailed plan is this: First, in 'Patient preference predictors' section, I provide more detail on PPPs: what their structure is supposed to be and what their use in practice might look like. Then in 'The problem of naked statistical evidence: law and PPPs' section, I sketch a problem, familiar from legal scholarship, in coming to legal verdicts only on the basis of statistical evidence—the so-called problem of naked statistical evidence. I show that, given the nature of PPPs, their use potentially faces a symmetrical problem. In the penultimate section 'Two diagnoses of the problem of naked statistical

evidence: autonomy and debunking', I consider two sorts of diagnoses of our problem and explain what would be required for each to be true. The 'Conclusion' section is conditional: *if* we are troubled by the problem in the legal domain, *then* we should be troubled by it here, in the medical domain *and* we ought to expect the same diagnosis of the problem in both cases.

PATIENT PREFERENCE PREDICTORS

The problem of making treatment decisions on behalf of incapacitated patients is widespread and widely discussed. Typically, medical professionals (and academics) have argued for solutions involving either advance directives, patient surrogates or some combination of the two. Such solutions face familiar difficulties, which can helpfully be thought of as comprising two sorts: epistemic and non-epistemic normative problems.

Very briefly: *Epistemic* problems are problems for the use of advance directives and surrogacy that are grounded in those methods' epistemic credentials, where a method's epistemic credentials can helpfully be thought of as how well that method does at reflecting incapacitated patients' actual preferences. These are normative problems, but their shape is essentially epistemic; hence, the label. Epistemic problems are familiar: perhaps most strikingly, surrogates are *much* worse at reflecting patients' preferences than one might have pretheoretically expected.^{iv} *Non-epistemic* problems are problems for the use of advance directives and surrogates that arise independently of those methods' epistemic credentials, that is, they are not grounded in facts about how well those methods do when it comes to accurately reflecting incapacitated patients' actual preferences. For instance: we sometimes resist substituting a surrogate's judgment for a patient's own judgment, even when the surrogate's choices (are likely to) accurately reflect the patient's preferences—perhaps because the surrogate is somehow unsuitable *qua* surrogate. Different non-epistemic problems, such as the normative relevance of past preferences to current decisions, face the use of advance directives. It does not matter here whether you think these issues with surrogates and advance directives can be resolved—I only mention them to set the stage.

Recently, a new solution to the problem of making treatment decisions on behalf of incapacitated patients has been proposed: so-called PPPs.^v

ⁱAdvance directives are ubiquitous both in practice and in theorising about patient care. Of course, advance directives are subject to criticism on a number of fronts. See Orentlicher¹; Levi and Green²; Miles *et al*³; Lo and Steinbrook⁴ for helpful overviews of proposals and criticisms.

ⁱⁱLike advance directives, surrogate decision making is ubiquitous both in practice and in theorising about patient care. For a classic source, see Buchanan and Brock.⁵ And like advance directives, surrogate decision making is subject to a range of criticisms, most forcefully that the so-called 'substituted judgment standard' is either unworkable or delivers unacceptable results. For a nice overview of the issue, see Johansson and Broström.⁶

ⁱⁱⁱShalowitz, Garrett-Mayer and Wendler (2007); Rid and Wendler⁷; Rid and Wendler.⁸

^{iv}For instance, see Smucker *et al*⁹ and Shalowitz *et al*.¹⁰

^vTo be clear, the use of PPPs is *not* widely advocated, though it is advocated by some. For an early suggestion

► <http://dx.doi.org/10.1136/medethics-2018-104941>

To cite: Sharadin NP.
J Med Ethics
2018;**44**:857–862.

PPPs, at least in the first place, present as a novel solution to the *epistemic* problems faced by the use of advance directives and surrogates. In the remainder of this section, I will sketch the structure of PPPs; along the way, it will become clear why PPPs are meant to improve on advance directives and surrogates' epistemic credentials.

Recall the problem of making treatment decisions for an incapacitated patient. We often know a lot about an incapacitated patient—in addition to medical history and current condition: age, gender, race, socioeconomic status and so on. But, crucially, what we often do not know is *what the patient wants*. PPPs represent a novel approach to this problem. Simply put, PPPs take what we *do* know about a patient and attempt to solve for, or predict, what we do not know using statistical modelling. Understood in the broadest sense, then, PPPs are principles, or algorithms, that take us *from* known facts about a patient, for example, age, gender, race, socioeconomic status and so on *to* a particular set of unknown facts about that patient, viz. her preferences over particular treatment options, *via* statistical evidence. Since moving in this way from known facts about a patient to unknown preferences via statistical evidence is inherently a probabilistic task, PPPs include a CI or probability distribution over the assignment of preferences to treatment options. Hence, PPPs can helpfully be understood generically as functions the input of which are (1) a set of known facts about a patient and (2) statistical facts connecting these known facts to treatment preferences and the output of which is a set of preferences over particular treatment options together with a CI or probability distribution, over those preferences.

Let us take a toy example to illustrate the idea. Suppose we know nothing about an incapacitated patient P's preferences between two possible treatment options, O_1 and O_2 and that we must decide between them.^{vi} How confident should we be regarding P's preferences over O_1 and O_2 ? The answer is that we ought to split our credences cleanly between the three possibilities; effectively, for all we know, it is equal odds that P prefers O_1 to O_2 , that P prefers O_2 to O_1 and that P is indifferent between O_1 and O_2 .^{vii} If we find ourselves in this situation, we can do no better than *chance* at providing P the treatment option P actually prefers. But now suppose that we also know the following facts about P: he is a 34-year-old unmarried man with a college degree. Finally, suppose we know the following statistical fact: among 34-year-old unmarried men with a college degree, 90% prefer O_1 to O_2 , 8% prefer O_2 to O_1 and 2% are indifferent between them. The simplest version of a PPP would then take as input these facts about P, together with the statistical facts, and yield as output that P prefers O_1 to O_2 , with a 90% confidence.^{viii}

regarding the accuracy of PPPs as compared with surrogates, see Smucker *et al.*⁹ For more detailed discussion and advocacy, see Shalowitz, Garrett-Mayer and Wendler (2007); Rid and Wendler⁷; Rid and Wendler (2014b). For critical discussion, see Ditto and Clark¹¹; Kim¹²; Dresser¹³; John¹⁴; Lindemann and Nelson¹⁵; Brock.¹⁶

^{vi}If it helps in supposing that we must choose, suppose one option, O_2 say, is *doing nothing*, so that failing to choose O_1 amounts to choosing O_2 .

^{vii}The assumption that this is a case where the so-called 'principle of indifference' ought to govern our prior probability assignments is not necessary for the argument to go through. If you think this is a case in which the principle of indifference is for some reason the wrong principle, feel free to substitute one's preferred way of assigning probabilities under uncertainty.

^{viii}It will also yield as output a judgment regarding P's preference for O_2 over O_1 , viz. that this is 8% likely. Similar remarks go for indifference between O_1 and O_2 . We can safely ignore these complications at present.

This is a toy example. The statistical details get messy quickly. But for our purposes here, we can work at this level of abstraction and assume we have clean statistical data on how patients' memberships in demographic categories, for example, race, ethnicity, socioeconomic status, marital status and so on, are (strongly) correlated to preferences across treatment options.^{ix} Here, I am not interested in taking issue either with the feasibility of collecting such data or whether such data are statistically relevant. I will assume that we do, or at least could, have access to statistically relevant data suitable for playing the role PPPs assign to such data: that is, of taking us from demographic facts about an incapacitated patient to a probabilistic estimate of the patient's preferences across treatment options.

It should be clear why PPPs, so understood, are a promising new solution to our problem. After all, assuming as we are here that we have both robust statistical data and the necessary models that are jointly capable of taking us from demographic facts about a patient that are relatively epistemically transparent *to* typically opaque facts about what that patient would prefer in a range of decision scenarios, the epistemic problem of making treatment decisions that actually reflect an incapacitated patient's preferences becomes *much* more tractable. In fact, at the limit—that is, at the limit of quality for our data and our models, the problem becomes trivial: we simply plug in the relevant data and *read off* a high-confidence prediction regarding what the incapacitated patient wants. Hence, PPPs represent both a novel and exciting solution to the problem of making treatment decisions for incapacitated patients. In the next section, I will argue that PPPs, so understood, face a hitherto unnoticed normative problem that is symmetrical to a problem familiar from legal scholarship: the problem of naked statistical evidence. Before that, let me mention a complication in order to set it firmly aside in what follows.

Because PPPs deliver probabilistic estimates regarding a patients' preferences over treatment options, it will be an open question what *degree* of probability it is reasonable to require in making use of a PPP. In other words: it will be an open question how confident one needs to be, given the verdict of a PPP, that some incapacitated patient prefers one treatment to another in order for it to be reasonable to choose the likely favoured treatment on behalf of the patient. There is one plausible limit possibility: we might set the threshold quite low and only require that our PPP deliver a result that is *better than chance*. On this view, given the choice between two treatment options that are otherwise (eg, medically) on a par and neither of which we have any antecedent reason to think the patient prefers, if our PPP says that there is a better than 50% chance the patient prefers one treatment to another, we ought to choose that treatment.^x Where to set the required confidence threshold is, as I have said, an open question and reasonable people might disagree over the details.^{xi}

^{ix}There are a range of ways such statistical data might be gathered. One way is that these data can be extracted from the treatment decisions of patients who are not incapacitated, either through explicit questioning or through the preferences revealed in their actual choices. Another is through the use of widely administered surveys intended to capture both the demographic details and the preference orderings of *potential* patients, that is, of the population at large. Rid and Wendler⁷ are explicit about this latter possibility.

^xNote that the alternative limit possibility, that we require 100% confidence, is *not* plausible, given that no PPP, using as it does the kind of statistical data it does, can cross that bar.

^{xi}See Rid and Wendler (2014, esp. section VII) for discussion of the issues involved in setting the relevant threshold.

Moreover—and this is the last I will say on the topic here—it is quite plausible to think that the correct view of the matter would hold that the confidence threshold required for the appropriate use of a PPP would vary with context and depend on a range of factors, for instance, on the severity of the patient's condition, on the range of possible medical outcomes provided by the various treatment options, on the availability of alternate methods for decision (eg, advance directives and surrogates) and so on. In any case, the issue I am interested in here does not depend on any particular view of where this threshold ought to be set, nor does it depend even on the view that the relevant threshold is in fact context-sensitive. I will turn to that issue now.

THE PROBLEM OF NAKED STATISTICAL EVIDENCE: LAW AND PPPS

In the last section we saw, in broad outline, what PPPs are and what their promise is: they promise to provide medical professionals with a new, epistemically powerful way to make treatment decisions on behalf of incapacitated patients. But PPPs, so understood, face a distinctive normative problem not faced by either advance directives or surrogates. The problem is symmetrical to a problem familiar from legal scholarship—the so-called ‘problem of naked statistical evidence’. Since I am going to argue that our response to the problem when it comes to PPPs ought to mirror our response to the problem in the legal domain, it will be helpful to look at an instance of the problem as it appears there first and then work our way back to the case of PPPs.

To that end, here is a common, stylized instance of the problem of naked statistical evidence in the legal domain, the so-called *Blue Bus Problem*:^{xii}

Jane's car is hit by a blue bus. The only issue being litigated is the owner of the bus. We know that 99% of the blue buses in Jane's city are operated by Blue Bus Company, that the remaining 1% of blue buses are operated by Red Bus Company, and that these are the only two companies that operate buses in her city. A somewhat reliable eyewitness agrees the bus was blue but testifies that it was a Red Bus Company bus.

Here is a question: Ought Jane to win a judgment against Blue Bus Company? Given that the civil standard of proof is a ‘preponderance of the evidence’, the answer appears to be a straightforward ‘yes’: the joint effect of the (positive) probabilistic evidence and the (negative) eyewitness testimony seems obviously to cross the threshold of preponderance. But that is not the judgment of the ‘overwhelming majority’ of American courts.^{xiii} Why not?

The answer from legal scholarship is complex; the guiding idea, in brief, is that, as one legal scholar puts it, the statistical ‘evidence *taken alone* can rarely, if ever, establish the crucial proposition with sufficient certitude to meet the applicable standard of proof’.^{xiv} But again: Why is this so? After all, our choice of bus-ownership ratios was arbitrary. Suppose we imagine instead that Blue Bus Company operates 999 out of 1000 blue buses in the city. Would this deliver ‘sufficient certitude’ in the crucial proposition? The consensus among legal scholars is that it still would not and that *something more* must be added to the evidentiary base in order to meet the burden of proof. In particular, what needs to be added, according to the consensus view, is

some kind of individualised or *non-statistical* evidence regarding the proposition in question. So, for instance, in the Blue Bus Problem, we might add to the evidentiary base a second eyewitness's testimony that the bus that struck Jane's car was a Blue Bus Company bus.^{xv} But once more it is fair to ask: *Why must we add such evidence?* I will return to this question in a moment. Before that, let me explain why the use of PPPs in making treatment decisions on behalf of incapacitated patients plausibly faces a problem symmetrical to this one. Then, I will argue that *whatever* answer to the problem we give to the problem in the legal domain (I will outline two sorts of possible answer) we ought to give that answer in the case of PPPs too. First, let me sketch the case for thinking there is a symmetrical problem for PPPs.

As a way in to the problem, let us imagine a use of PPPs that, to be absolutely clear, *no one advocates*.^{xvi}

You are in an accident and are incapacitated. Prior to becoming incapacitated, you completed a comprehensive advance directive covering a range of medical decisions in a range of circumstances, including one your doctors now face: whether to give you extensive palliative care. In your advance directive, you indicate that you do *not* want to receive such care in exactly the circumstances you are now in. The doctors consult the hospital's ethics board. On the basis of the board's recommendation the doctors provide you with palliative care since, according to their best information, people with your demographic characteristics have a 99.94% of preferring palliative care. In the ethics board's experience, this likelihood is high enough for them to think that—unbeknownst to you, and your advance directive's protestations to the contrary notwithstanding—you really do prefer to receive palliative care.

To repeat: no proponent of PPPs has argued in favour of this use of PPPs to override patients' advance directives. It might seem obvious *why* no one has advocated this use: you appear to have a serious complaint against being treated in this way. But what is the nature of your complaint?

Your complaint cannot be that you will not be getting what you want. At least, it cannot be that you will not be getting what you want *with respect to the choice between receiving and not receiving palliative care*. After all, like in the Blue Bus Problem, our choice of likelihoods was arbitrary. Since we should all agree that there is *some* non-negligible probability that you are mistaken about your own preferences in crafting your advance directive (introspection is not infallible), we should all also agree that, at least when it comes to the choice between receiving and not receiving palliative care, there is a version of this case where you are more likely to get what you want in virtue of the hospital's overriding your advance directive. Here is the recipe for generating such a case: simply modify the likelihoods delivered by the board's use of a PPP according to how fallible you think our introspective access to our preferences is.

Perhaps your complaint is instead that you are not getting what you want in a different respect: you are not getting what you want *with respect to having your advance directive be followed*. After all, your advance directive is clearly being disregarded in this scenario. This might be so; that is, it might be the case that you are not

^{xii} Adapted from Mass¹⁷. See Tribe (1971) for discussion.

^{xiii} Schauer.¹⁸

^{xiv} Tribe¹⁹ (1350, emphasis in original).

^{xv} Note that adding this second eyewitness testimony might rationally make us *less* confident in the proposition that Jane was hit by a Blue Bus Company bus (for instance, if the eyewitness is only somewhat reliable). Nevertheless, the consensus appears to be that adding this (rationally credence deflating) testimony to the evidentiary base might make that base suitable grounds for issuing a judgment against Blue Bus Company. See Buchak²⁰ (2014) for discussion.

^{xvi} Not even those who are otherwise quite optimistic about the use of PPPs (eg, Rid and Wendler (2014ab)⁷).

getting what you want with respect to having your advance directive followed. But this is so only if you are fetishistically interested in receiving such-and-such care *just because you said you wanted it and regardless of whether you do in fact want it*. That is certainly possible, but I hope we can agree it is a bit odd. It is like insisting on clearing your plate of disliked food at a buffet with endless choices. In any case, let us stipulate that your only aim in making an advance directive was to ensure, to the best of your ability, that your treatment while incapacitated reflected your preferences: in other words, all you are interested in is increasing the likelihood that your treatment mirrors your preferences and you have no fetishistic preference regarding your advance directive itself. If that is right, then your complaint cannot be that your advance directive is being ignored and for the same reasons as before: the hospital's use of the PPP is (by hypothesis) *more* likely to result in treatment that reflects your preferences.

So your complaint cannot be that (in either of these two ways) you will not be getting what you want. Nevertheless, I think some—and depending on my mood I count myself among them—will still think that you have a real complaint. It is hard to articulate the complaint, but we might find ourselves saying much the same things we saw legal scholars say in response to the Blue Bus Problem: that the statistical evidence taken alone can rarely, if ever, establish the patient's (*your*) preferences with 'sufficient certitude to meet the applicable standard of proof'. The similarities continue: your advance directive appears to be the analogue of contrary eyewitness testimony and, like in the Blue Bus Problem, the statistical evidence *points in the other direction*. The problem is thus exacerbated and we feel the need to add *something different*—that is, something non-statistical—to the evidentiary base in favour of palliative care in order to arrive at an actionable verdict regarding your preferences. If we had a surrogate ready to hand, perhaps that would do the trick.^{xvii} But in any case, like in the Blue Bus Problem, it is fair to ask: *Why must we add such evidence?*

So, we face an apparently symmetrical problem when it comes to the use of PPPs as we do in the case of merely statistical evidence in the law. On reflection, the existence of this symmetrical problem should not be too surprising. After all, what generated the problem in the legal case was reliance on naked statistical evidence in arriving at verdicts that were themselves normative or at least had some normative upshot. (The choice of whether to hold someone legally responsible is often, even if we can admit not always, morally fraught.) And in the case of PPPs, we find ourselves again relying on naked statistical evidence in arriving at verdicts that similarly have important normative upshots. (The choice of whether or what care to provide to an incapacitated patient is often, perhaps always, morally fraught.) What, then to do? In the next section, I will outline two possible strategies of reply to the problem based on two different diagnoses of what, precisely, the normative problem actually is. I will not argue in favour of one strategy over the other, but I will suggest that we ought to pursue the *same* strategy in replying to the problem in each domain, that is, both in the legal case and the case of PPPs.

TWO DIAGNOSES OF THE PROBLEM OF NAKED STATISTICAL EVIDENCE: AUTONOMY AND DEBUNKING

Reliance on naked statistical evidence in the course of arriving at normative judgments (in the legal or in the medical

domains) appears problematic, as we have just seen. What we now want to know is *why* this is problematic: we want a diagnosis of the problem. One diagnosis is vindicating; what I will call the *autonomy* diagnosis says that reliance on bare statistical evidence somehow constitutes a violation of agential autonomy. So, according to this strategy, our intuitions are at least partly revelatory of a genuine normative problem in need of a solution. A second diagnosis is debunking; what I will call the *debunking* diagnosis says that our intuitions mislead us in both cases and that, contrary to appearances, there *really is no problem of naked statistical evidence*.^{xviii} Hence, the debunking strategy offers an alternative diagnosis of our intuitions in the relevant cases and, on the basis of this diagnosis, concludes that, on reflection, there is nothing normatively troubling about reliance on naked statistical evidence. Let me explain these both in a bit more detail, starting with the autonomy diagnosis.

Here, roughly, is the idea behind the autonomy diagnosis. What seems problematic about making normative judgments on the basis of (merely) statistical evidence appears to be that such judgments treat individual or corporate agents—the Blue Bus Company or an incapacitated patient, say—as if their group membership in some way *determined* the normatively relevant information on which we ground the normative judgment—whether the bus belonged to Blue Bus Company or what the patient's preferences actually are. But this activity, again, appears normatively invidious. For it is not true, quite generally, that, for example, agents' preferences are *determined* in this way by their demographic membership: although there are strong correlations between demographic characteristics and preferences, we are all reluctant to say—not least of ourselves!—that the reason *why* we prefer (say) not to receive palliative care is *because* we are (say) straight, white, unmarried, college-educated 30-somethings. In other words, we view our preferences as merely *correlated with* and not *caused by* our group membership. Hence, it can appear to be a violation of our capacity to have, form and revise these preferences on the basis of reasons to treat our preferences as if they are determined by merely statistical information. And, moreover, that appears to be exactly what happens both in the legal case and in the use of PPPs. So, in short, such uses appear to be violations of our autonomy.

The thought, then, is that we diagnose what is wrong with issuing a verdict against the Blue Bus Company and what is wrong with issuing a verdict contrary to a patient's advance directive *in the same way*: both represent the agent's activities as caused by something they are manifestly *not* caused by, assuming the patient has (or has had, in the case of incapacity) the capacity to cause *herself*. In the Blue Bus Company case, it is wrong to issue a verdict against the company because doing so would involve treating the negligent driving of the Blue Bus Company's bus (if it was in fact one of their buses) as if it were caused by their operating a certain percentage of blue buses in the city. But of course, that is both likely to be a spurious causal claim and certain to represent an assault on Blue Bus Company's autonomy, insofar as it actually treats the actions of its drivers as determined by statistical data rather than by (say) their training, the Blue Bus Company policies and so on.^{xix} Things are much the same in the

^{xvii}Or perhaps not, given the well-known problems with surrogate reliability.

^{xviii}I am grateful to an anonymous referee for suggesting this way of thinking about the two strategies of reply. I am indebted to comments from that referee in what follows.

^{xix}Do not let the issue of corporate agency and therefore corporate autonomy distract from the point; if it helps, we can imagine Blue Bus Company is a sole-proprietorship with the owner as the sole driver.

case of PPPs. It is wrong to decide on behalf of the patient via the PPP because doing so would involve treating the patient's preferences as if they were caused only by her membership in various demographic categories. But of course, and again, that is both likely to be a spurious causal claim (it is not your *being a certain age* or *being a certain sex* that makes you prefer palliative care) and certain to represent an assault on the incapacitated patient's autonomy, insofar as it actually treats her preferences as determined by statistical data rather than by (say) her autonomous choices or values.^{xx}

So, that is the autonomy diagnosis in rough outline. Of course, spelling it out requires much more detail. For instance, we will want to know what conception of autonomy, precisely, stands behind the diagnosis. And we will no doubt want to hear more about the way in which the legal and the medical cases differ, despite enjoying the same diagnosis. But let all of that pass for the moment. Here, I lack the space to explore the strategy in any more detail; and in any case, I am tempted by the thought that, despite appearances, *there really is no problem to be solved*. Hence, I think something like the *debunking* diagnosis might be correct. I will turn to an explanation of the debunking diagnosis now.^{xxi} Then, I will say a bit more about where this leaves the debate over which diagnosis is correct.

To see the idea behind the debunking diagnosis, recall the central idea behind the *autonomy* diagnosis: according to that view, acting on or forming a normative judgment on the basis of *merely correlative* statistical information (between demographic information and preferences, say) represents an assault on agential autonomy. According to the debunking diagnosis, this can sometimes *appear* to be so, but as a matter of fact, there is nothing normatively problematic per se about moving from (merely) demographic information to normative conclusions. And, moreover, according to the debunking diagnosis, we can explain why it appears that this is sometimes normatively problematic. So, the debunking explanation has two parts: the first aims at explaining why it is not normatively problematic per se to draw normative conclusions, such as conclusions about agents' preferences, from merely statistical data; the second aims at explaining why it appears normatively problematic to do this. Let us take these both in order.

The first part: Consider the following related use of (merely) statistical information that has the same structure as the use of PPPs. Suppose I discover a strong correlation between listening to National Public Radio (NPR) and preferring wine to ale and so bring my dinner host (who I know listens to NPR) a bottle of wine rather than a six-pack. This might be an odd way to behave—one does not normally decide on host gifts on the basis of listening preferences—but in so acting it is not true that I am necessarily assuming that my host prefers wine to ale *because* she listens to NPR. Now, it does seem true that in order for my bringing wine rather than ale to be reasonable, I must think there is *some* causal relationship between listening to NPR and preferring wine. But the thought that there is some such causal relationship between the two does not preclude the thought that an

individual's preference for wine over ale is autonomous. At least, it need not, and will not, unless we endorse a *remarkably* strong account of autonomy according to which *any* causal influence by *any* factor in one's choices is sufficient to make that choice non-autonomous. Of course, by that measure, plausibly almost none of our choices count as autonomous—and at the very least, none of the choices that we think of as most important count as autonomous. That is clearly too strict a measure. So, by the same lights, acting as if the person prefers wine to ale on the grounds that she listens to NPR need not involve treating that person as lacking any (relevant sort of) autonomous control over her preferences.

The idea, then, is that we have been led into confusion by thinking that *any* causal influence on our behaviour undermines our autonomy. But, on reflection, we do not think this is true, since a commitment to this would commit us to an implausibly impoverished view about which acts—ours and others'—count as autonomous. On the plausible assumption that *all* of our actions are influenced to at least some degree by causal factors outside our control, we would arrive at the implausible result that *none* of our actions count as in the relevant sense autonomous. If this is right, then, just as acting on the basis of a predicted preference for wine over ale on NPR-listening grounds might not represent an assault on my host's capacity to autonomously decide her alcohol preferences for herself, so choosing for an incapacitated patient on the basis of a predicted preference (say) for palliative care might not represent an assault on the patient's autonomy. Hence, the diagnosis is *debunking*: there is a genuine *prima facie* problem, but on reflection the problem disappears.

The second part: If the debunking diagnosis is correct, we still face a question. Why does the *prima facie* problem appear in the first place? That is, what explains why it *appears* normatively problematic to rely on statistical information to issue verdicts either in the legal or in the medical case? The proponent of the debunking diagnosis might answer this question in a range of ways; here I will just mention one thing I think it is plausible for the debunker to say before moving on, in the conclusion, to a consideration of where this leaves the debate over the use of PPPs. The plausible thing to say on behalf of the debunker that I have in mind is this: The use of PPPs appears normatively problematic because although the use of statistical information in this way does not itself represent an assault on agential autonomy, people tend to *hear* the use of statistical information in this way as an assault on their autonomy. In other words, the use of statistical information conversationally implies a sort of disrespect that (as we have just seen) it need not—indeed usually does not—actually involve.^{xxii} Let me explain.

Returning to our dinner host, suppose the host asks me *why* I brought wine (rather than ale) and I respond by citing her listening patterns: 'Well, I know you listen to NPR!' The host might thereby feel badly treated by me: it sounds an awful lot like I am assuming (what as we just saw above I need not actually be assuming) that it is *only because* she listens to NPR that she prefers wine to ale. After all, she asked after an explanation of my choice, presumably she was after a rationalising explanation and I cited her listening to NPR. Assuming something like the Gricean maxim of quantity, what I have just conveyed to her is that I think her preference for wine over ale exists only because

^{xx}In his 2014 study, John argues that *certain* uses of PPPs might be objectionable for this reason. But, importantly, he does not argue there that *all* uses of PPPs will face this problem. Presumably, that is because he thinks that the correct account of autonomy will *not* be one that precludes in *all* cases a reliance on statistical information. In that sense, John appears to be a *debunker* despite his insistence on the importance of agential autonomy. See the discussion of the debunking strategy below. Thanks to an anonymous referee for suggesting clarity on this point.

^{xxi}I am grateful to an anonymous referee for suggesting this strategy. The example below is adapted from one suggested by that referee.

^{xxii}Thanks to an anonymous referee for suggesting a version of this explanation.

of her listening habits (rather than as the studied, cultivated preference she considers it to be and it plausibly might in fact be). That is why reliance on merely statistical information can seem so normatively noxious. But of course, if the debunking diagnosis is correct and such reliance does not actually commit one to this noxious conclusion, but only appears to do so via implicature, then we should be able to *cancel* the implicature. And it appears we can do this. For, suppose I follow-up: 'Well, I know you listen to NPR! I wanted to ensure that I brought you something you would like and I recently learnt that NPR listeners tend to prefer wine to ale'. It seems to me this addendum is sufficient to cancel any implicature that I am assuming my host's preference is non-autonomously formed as a result of her listening habits. It makes explicit my reliance on the (merely) correlative information about her preferences provided by the (bare) statistics and explicitly denounces any corresponding causal claim about the (normative) source of those preferences. And while in certain contexts it can be reasonable for her to hear the first sentence of my utterance as an assault on her autonomy, it is not at all reasonable in normal contexts for the entire utterance to be taken as such. For, what I am explicitly articulating is a desire to get her what she (independently) wants and indicating my reliance on the statistical information as a (mere) tool for achieving that—obviously normatively non-problematic—end. If my host were to continue to take offence at this point, I think we can agree it would be unreasonable.

Hence, according to the debunking diagnosis, there is nothing per se wrong with reliance on bare statistical information in arriving at the sorts of normative judgments—about preferences or behaviour—we see in the legal and medical cases. Why, then, does such reliance appear so normatively noxious? The answer is that, typically, without further elaboration, such reliance seems to involve a normatively objectionable treatment of an agent as having preferences that are non-autonomously formed.

CONCLUDING REMARKS: THE WAY FORWARD AND PARITY

Where does this leave things? Here, I have not argued that we should prefer one of our two possible diagnoses to the other. (Nor have I argued that they are the only two possibilities.) But I hope to have put some pressure on anyone who wants to argue that there *really is* something normatively problematic in using PPPs to predict and act on predictions of incapacitated patients' preferences. At the very least, the argumentative burden is now on the defender of the problem to fully articulate an account of autonomy according to which the use of PPPs really does (always?) represent an assault on agential autonomy. I am not optimistic about the prospects for doing so; as we saw above, since *any* of our actions will to some degree be influenced by causal factors and since it is just this inevitable influence on which verdicts issued by PPPs rely, the account of autonomy required must be unconscionably strong in order to deliver the result. In any case, we shall want to hear more from the defender of the autonomy diagnosis. By contrast, the debunker seems to be in relatively good shape, though of course there is more that could be said there too. In particular, it would be worth exploring the ways in which the pragmatic signal sometimes conveyed by the use of PPPs—that the agent's preferences are non-autonomous—can be mitigated. This is especially pressing if we think PPPs have an important role to play in future medical practice.

I will close with a brief remark on parity. As we saw, the normative problem with PPPs—if indeed it is a problem—mirrors a

normative problem familiar from legal scholarship. The structure of the problems is the same: in each case, we are asked to draw a normative conclusion (a legal verdict or a decision about preferences over medical treatment) on the basis of merely ('naked') statistical evidence. The problem as it appears in legal scholarship is decades-old and versions of the *autonomy* and *debunking* diagnoses are available there too.^{xxiii} This suggests that the solution to the problem in both cases, whatever it is, will be the same. Of course, if we are *debunkers* in the legal domain, we will not be too concerned about the use of PPPs! Hence, my conclusion is conditional: *if* we are worried about the problem as it arises in the legal domain, *then* we should be worried about it here too. More generally, we can hope for a solution in each domain that mirrors the solution in the other.

Funding The authors have not declared a specific grant for this research from any funding agency in the public, commercial or not-for-profit sectors.

Competing interests None declared.

Patient consent Not required.

Provenance and peer review Not commissioned; externally peer reviewed.

© Article author(s) (or their employer(s) unless otherwise stated in the text of the article) 2018. All rights reserved. No commercial use is permitted unless otherwise expressly granted.

REFERENCES

- Orentlicher D. Advance Medical Directives. *JAMA: The Journal of the American Medical Association* 1990;263:2365–7.
- Levi BH, Green MJ. Too soon to give up: re-examining the value of advance directives. *Am J Bioeth* 2010;10:3–22.
- Miles SH. Advance End-of-Life Treatment Planning. *Arch Intern Med* 1996;156:1062–8.
- Lo B, Steinbrook R. Resuscitating advance directives. *Arch Intern Med* 2004;164:1501–6.
- Buchanan AE, Brock DW. *Deciding for Others: The Ethics of Surrogate Decision Making*. Cambridge, UK: Cambridge University Press, 1990.
- Johansson M, Broström L. Counterfactual reasoning in surrogate decision making -- another look. *Bioethics* 2011;25:244–9.
- Rid A, Wendler D. Treatment decision making for incapacitated patients: is development and use of a patient preference predictor feasible? *J Med Philos* 2014;39:130–52.
- Rid A, Wendler D. Use of a patient preference predictor to help make medical decisions for incapacitated patients. *J Med Philos* 2014;39:104–29.
- Smucker WD, Houts RM, Danks JH, et al. Modal preferences predict elderly patients' life-sustaining treatment choices as well as patients' chosen surrogates do. *Med Decis Making* 2000;20:271–80.
- Shalowitz DI, Garrett-Mayer E, Wendler D. The accuracy of surrogate decision makers: a systematic review. *Arch Intern Med* 2006;166:493–7.
- Ditto PH, Clark CJ. Predicting end-of-life treatment preferences: perils and practicalities. *J Med Philos* 2014;39:196–204.
- Kim SY. Improving medical decisions for incapacitated persons: does focusing on "accurate predictions" lead to an inaccurate picture? *J Med Philos* 2014;39:187–95.
- Dresser R. Law, ethics, and the patient preference predictor. *J Med Philos* 2014;39:178–86.
- John S. Patient preference predictors, apt categorization, and respect for autonomy. *J Med Philos* 2014;39:169–77.
- Lindemann H, Nelson JL. "The Surrogate's Authority". *Journal of Medicine and Philosophy* 2014;39:161–8.
- Brock DW. Reflections on the patient preference predictor proposal. *J Med Philos* 2014;39:153–60.
- Mass. "Smith V. Rapid Transit. Inc". 1945. 317 Mass. 469, 58 N.E. 2nd 754.
- Schauer F. *Profiles, Probabilities, Stereotypes*. Belknap Press, 2006.
- Tribe LH. Trial by Mathematics: Precision and Ritual in the Legal Process. *Harv Law Rev* 1971;84:1329–93.
- Buchak L. "Belief, Credence, Norms.". *Philosophical Studies* 2014;169:285–311.

^{xxiii} In the legal literature, the autonomy diagnoses is sometimes framed as a 'due process' diagnosis, but the justification of due process in turn (usually) runs through agential autonomy.