TAKING RESPONSIBILITY FOR UNCERTAINTY

RICHARD HOLTON & ZOË FRITZ

Our topic here is uncertainty, especially as this arises in medicine. We are concerned with the uncertainty of diagnosis and of prognosis, and about how this should be communicated and shared between a doctor and patient. We are concerned with the idea that the doctor might rightly take responsibility for some of that uncertainty, in the sense that they may manage it when the patient is unwilling or unable to do so. And we are concerned with the implications that this has for informed consent, for shared decision making, and for other normative dimensions of the doctor-patient relationship. Most centrally we argue that it is best understood, not in terms of informed consent, but in terms of informed trust.

INTRODUCTION

We start with the realities of contemporary medicine. When a patient consults a doctor about a set of symptoms, the doctor will typically form not a definitive diagnosis, but one that is, in the jargon, 'differential': that is, a list of possible diagnoses, with some rating of the likelihood and severity of each, and with some indication of what would need to be done to narrow down these possibilities.[1]

But the patient will not typically be given the full differential diagnosis. Instead they will typically be given the most likely possibility, perhaps two, or else a 'putative' or working diagnosis of the most likely cause . Why is this? Part of the reason stems from the complexity. Understanding a single diagnosis can be hard enough; understanding a whole list of possible diagnoses will be much harder. But in addition, the notion of uncertainty, and of the probabilities in which it is typically framed, brings a whole set of further difficulties. How seriously should one take a 2% probability? Or a 2% probability of something really bad? How should one weigh a test with a 10% chance of false positives? How should one factor in the relevant base rates? We know from many studies that even doctors are not good on many of these questions (especially the last).[2] For patients they risk becoming unmanageable. We know that people typically react in particular ways, over-weighting small probabilities, underweighting large, ignoring base rates and so on. Some of these difficulties can be minimised by a sensitive presentation of probabilistic information: natural frequencies ('Of 100 people who test positive, 7 will have the disease, 93 will not.') are easier to understand than probabilities and base rates ('The rate of the disease in the population is 1%; the proportion of false negatives in the test results is 0; the proportion of false positives in the test results is

[1] This is a common use of the term among practicing doctors. The next stage, the act of distinguishing one cause from the other possibilities, is colloquially referred to as ' working through the differential' or 'narrowing down the differential'.

[2] Gerd Gigerenzer *Rationality for Mortals* (New York: Oxford University Press, 2008) ch. 9.

10%.').[3] For others though there is no technical fix. There is no neutral answer to when we should worry about a possibility, and when we should ignore it as too small to worry about.

More broadly, we know that people do not simply find uncertainty intellectually challenging—it is emotionally challenging too, unsettling and aversive. People do not, in general, like being in a state of uncertainty, and consequently, of indecision, especially when this involves their own fate. A good illustration of this comes from the experiences of a doctor, Franz Ingelfinger. He was diagnosed with a carcinoma in the gastroesophageal junction, an illness about which he was an expert. Nevertheless, worrying about the best treatment, he became, in his own words

> increasingly confused and emotionally distraught. Finally, when the pangs of indecision had become nearly intolerable, one wise physician friend said: "What you need is a doctor."[4]

If even an expert like Ingelfinger can react to the uncertainty of what to do in this way, how much greater the risk for someone who finds it hard to understand what is going on. There has been some discussion of the 'diagnostic moment': the point at which a patient comes to know what they have, and can reconceptualise themselves and their future accordingly.[5] A differential diagnosis presents no such moment. Rather than leaving appropriately reassured, the patient may leave with heightened anxiety. Rather than deciding on a reasonable treatment, they may ask for further unwarranted and potentially dangerous investigations; or, conversely, they may be scared into refusing treatment altogether. We have heard anecdotally from quite a number of clinicians of outcomes in both of these directions. So it is not surprising that clinicians who are concerned for the well-being of their patients might want to restrict the information that they give them. In fact we are sure that almost all already do, if only because of the time constraints on the typical consultation. Possible diagnoses that are very unlikely are typically not reported to patients; nor are possible but very unlikely outcomes of treatment. "I rarely mention anything with a chance of less than 5%" one consultant told us.

But of course, this raises worries. Withholding information so as not to alarm a patient looks like paternalism at its most overt. Perhaps more worrying, it looks to be incompatible with two of the standard ethical requirements on modern medicine: (i) informed consent, the idea that the patient's consent must be given for any medical intervention, and that to properly give it the patient must fully understand what they are consenting to; and (ii) shared decision making, the idea clinicians and patients should share the best available evidence when deciding on a

---

[3] Gigerenzer *op.cit.* ch 12.

[4] Ingelfinger FJ (1980) Arrogance. *N Engl J Med* 303(26):1507–1511

[5] Jutel, A., 2014. When the penny drops: diagnosis and the transformative moment. In A.Jutel and K. Dew (eds.), *Social Issues in Diagnosis: an Introduction for Students and Clinicians*, Johns Hopkins University Press, Baltimore, MD, pp. 78–92; Heritage, John, and McArthur, Amanda: The Diagnostic Moment: A study in US Primary Care' *Social Science & Medicine* 228 (2019) 262–271

treatment.[6] For if information is withheld, the patient cannot be fully informed; and if they are not fully informed they cannot give informed consent, and decisions cannot be truly shared.

Of course informed consent and shared decision making do not require that the patient know everything about their treatment, down to the provider of the oxygen. A reasonable criterion of what they need to know might be: everything that, were they to know it, would reasonably affect their choice of whether to have the treatment. Something like that was articulated by the UK Supreme Court in *Montgomery*, specifically in relation to treatment decisions:

> The doctor is therefore under a duty to take reasonable care to ensure that the patient is aware of any material risks involved in any recommended treatment, and of any reasonable alternative or variant treatments. The test of materiality is whether, in the circumstances of the particular case, a reasonable person in the patient's position would be likely to attach significance to the risk, or the doctor is or should reasonably be aware that the particular patient would be likely to attach significance to it.[7]

But even with this reduced standard for what a patient needs to know, the problem remains. If the doctor's decision to withhold information is based on a concern about alarming the patient, then that is exactly to think that the information does meet this test of materiality: the patient would attach significance to the risk it were they informed of it.

We might interpret the judgment as proposing a straightforwardly objective test: it is not whether the patient would in fact be alarmed, but whether they would reasonably be. But as we have said, it is very hard to know what is a reasonable response to, say, a small chance of a bad outcome. And besides, it is unclear that that is how we should interpret the passage. While the the talk of the 'reasonable person in the patient's position' might encourage that interpretation, the requirement that the test places on the doctor—that they 'should reasonably be aware that the particular patient would be likely to attach significance to it'—makes no requirement of the reasonableness of the patient's response. It is clear from the passage cited that the judgment requires a doctor to inform a patient of a risk that they reasonably think the patient would be alarmed by, whether or not they in turn think that alarm reasonable (and they must in addition inform the patient if they think that they would have reason to be alarmed, even if they do not in fact think they would be).

Nevertheless two paragraphs earlier in the very same judgement we read:

> A person can of course decide that she does not wish to be informed of risks of injury (just as a person may choose to ignore the information leaflet enclosed

---

[6] Glyn Elwyn et al., 'Implementing shared decision making in the NHS', *BMJ*. 2010

[7] Montgomery v Lanarkshire Health Board [2015] UKSC 11 §87, per Kerr and Read

> with her medicine); and a doctor is not obliged to discuss the risks inherent in treatment with a person who makes it clear that she would prefer not to discuss the matter. Deciding whether a person is so disinclined may involve the doctor making a judgment; but it is not a judgment which is dependent on medical expertise.[8]

How do we fit these two ideas together? Sometimes there will indeed be an explicit request not to be told. But this will not remove the discretion from the doctor, for even patients who don't want to know certain things will rarely want absolute ignorance. Which things won't they want to know? At best they might be able to outline this in general terms—'I don't want to know the details'; 'Don't tell me about what might go wrong'—but the doctor will have to make decisions about how to understand such a request. Where the patient has not explicitly requested not to be told, things are even more complicated. Should the doctor ask them what they want to know? Questions themselves convey information, via their framing and presuppositions. To ask whether someone wants to know about other possibilities or risks is to communicate that there are other possibilities or risks. Even if the patient knows that there are going to be risks, to ask a question like this is to draw attention to them. If the patient had been wilfully ignoring them up till now, that stance is no longer available to them.

In the background of these concerns is what we will call *the irreversibility of knowledge*. Many treatments can be given on a trial basis: see if it works, and if it doesn't, discontinue. Knowledge is not like that. You cannot tell someone something, and then withdraw the knowledge if it distresses them, or if they don't want to know if for some other reason. You can perhaps tell someone not to worry about something that they have not been able to understand—tell them, that is, that they do not need to persist in trying to understand it—and that will sometimes be reassuring (though sometimes not). But that is a very different idea. You cannot simply tell someone to forget something that they have understood.

There is a second complication concerning the 'test of materiality' that we outlined above: the test of what information is relevant, and hence what people should normally be told. A natural way of spelling that test out in more details makes use of a hypothetical retrospective test: if, at the end of the treatment, one were to ask the patient what, in retrospect, they would have wanted to have been told, and what not, then that which they say they would have wanted to know is material. The problem with this is that it is likely to be highly contingent on how things turn out. To see this, consider a parallel proposal for deciding on diagnostic procedures. Suppose there is a remote chance — 0.01% say — that a patient has something very nasty. There is a test for revealing whether or not they do, but it is itself intrusive and unpleasant. Suppose, therefore, that the standard procedure is not to test. The vast majority of patients, those for whom the test would have been negative, might, in retrospect, be pleased not to have been tested. But the unfortunate one in 10,000

[8] Ibid § 85. For a general discussion of how to apply Montgomerty to the issue of the communication of diagnosis, see Kathleen Liddell, Jeffrey M Skopek, Isabelle Le Gallez, Zoë Fritz, 'Differentiating Negligent Standards of Care in Diagnosis', *Medical Law Review*, 2022

who actually has the condition is likely, in retrospect, to be very regretful that they did not have the test. We can call this *the contingency of regret*.[9] It is a feature that can frequently be observed in debates about diagnostic testing. Very often patient groups are all in favour of such testing; but they, of course, are the ones who, suffering from the condition, would regret not having been tested. We tend not to hear from those who turn out not to have the condition, and would rather never have been tested. What is relevant to our concerns is that a similar consideration applies to knowledge: Would a patient want to have known about some remote possibility of a bad outcome? That might well depend on the outcome: if the possibility is not realized, then no, if it is, yes. This is important, because the cases that reach the courts are typically those—like Montgomery, cited above—in which the bad but unlikely outcome has come about. There is a risk that we will build our policy primarily to avoid such unexpected but unlikely bad outcomes. We need to ensure that in doing so we do not come up with a policy that is ill-suited to the vast majority of cases.

THE REALITIES OF THE CONSULTATION

Imagine a consultation in which a patient has presented with some symptoms, and the doctor has elicited their history and examined them. With this information, the doctor has determined a differential diagnosis, and they have reported (some of) it to the patient. The doctor has outlined the three most likely explanations, with different degrees of probability; and has told the patient that there are further tests that could be done, some of which carry risks, all of which would reduce the level of uncertainty. The patient listens carefully, nodding. Then they say:

> But doctor, what do you think I have?

Most clinicians will recognize this as a familiar response. What is the patient after? They have just heard the differential diagnosis. Should the doctor just repeat it, stressing their uncertainty? That might stop the patient persisting the question, but it doesn't seem that it will provide them with what they were after.

To throw some light on what the patient was really after, consider a slightly different case. Instead of further investigations to be undertaken, the patient is faced with three treatments which could be started, though still without a definitive diagnosis in place. Again the patient listens carefully to the three options, and then says:

> But doctor, what would you do?

---

[9] A related issue arises in discussions of *ex post* and *ex ante* construals of contractualism; see, for instance, Frick, Johann. 2015. Contractualism and social risk. *Philosophy and Public Affairs* 43: 176–223; Jussi Suikkanen, Ex Ante and Ex Post Contractualism: A Synthesis, *The Journal of Ethics* 23 (1):77-98 (2019). Our concern is related to some of the concerns about the *ex post* approach.

Again most clinicians will be very familiar with such questions. In this second case the request is more straightforward; the patient is after practical, not epistemic, advice. Our suggestion is that in fact the two cases are very closely related, and that even in the first it is something like practical advice that is wanted. The core thought is this: in the face of all the uncertainty, the patient is looking for a more secure path; they want to be on the other side of the diagnostic moment. This is so at the practical level, since the patient wants to decide on a treatment option. But it may also be so at the theoretical level. When the patient asks, in the face of differential diagnosis, what the doctor really thinks they have, they are asking for something like advice on what to believe. They want a diagnosis to use as a working hypothesis. For a patient who wants this epistemic simplification, the theoretical and the practical will typically go together, since the working hypothesis that they adopt will be that they have the condition on which the treatment is to be targeted.

Let us suppose then that the doctor says what they would do, and what, in that sense, they think the patient has; and that the patient is happy to adopt the treatment, and the corresponding working hypothesis, as the path forward. Does this mean that the other possibilities are ignored? In a sense that is so for the patient: the other possibilities are put out of mind. But a thoughtful patient will know that this is not how the doctor is thinking. More test results are pending; and there is of course no information yet on how the patient will respond to the treatment. The doctor needs to keep check on these, and the patient trusts them to do so.

So what we have is a kind of mental division of labour. The patient can focus on the treatment, on what they need to do for it to have a chance of working, and on what they need to keep an eye on and report ('Be sure to tell me if you have any new symptoms, especially if you feel …'). And they have a diagnosis that they can research, discuss with friends and family, and so on. They know that this diagnosis is not definitive: they know that there is a chance that what they have is something else, but that is left in the background. So they do not have to face the unsettling prospect of keeping more than one possibility live in their mind. The doctor, in contrast has a much more open epistemic field. The other elements of the differential diagnosis remain live, in that they are still being actively investigated. And if evidence comes in that starts to make one of them look somewhat more likely, the doctor will have to decide whether, and if so when, to tell the patient. This may be only if a change of treatment is called for, but it may be well before. That is a decision to be made in the light of the interests, needs, and capacities of the patient. We can say, in short, that the doctor takes responsibility for the uncertainty.

What we have sketched is a fairly explicit case of such taking responsibility. We suggest though that this is happening implicitly all the time. In a consultation a doctor will constantly be assessing how much of their own uncertainty they should communicate to the patient; and the patient will likewise be indicating how much they want to take on. At one extreme the doctor may convey very little of the uncertainty: they may give suggest a line of treatment, outline the diagnosis on

which this would be based, mention that there are other possibilities that remain under investigation, and ask the patient if they have any questions. At the other extreme they may share all of the uncertainty, and keep the patient fully apprised of every bit of evidence as it comes in. In between are many possibilities.

Patient and doctor need a shared understanding of who will take responsibility for the uncertainty. This is delicate because of the irreversibility of knowledge that we mentioned at the outset. So it must be done with care. Given that they have more experience in these areas, much of the onus will fall on the doctor. Asking

> Do you have any questions?

provides an opening for questions, but scarcely encourages them; much more encouragement is provided by an approach that presupposes that there will be questions:

> So which questions would you like to ask me?[10]

There is even more encouragement if the patient is given time to prepare, and the expectation that they will:

> When you get home I'd like you to get some paper and jot down any questions you might have, and we can go through them when I see you tomorrow.

Similarly

> There are some other possibilities that the tests might show.

invites less inquiry than:

> There are some other possibilities that the tests might show. Would you like me to go over them now, or shall we wait and see if the tests throw anything up?

Kerr and Read, in the Supreme Court ruling cited above, claimed that judgments about the issues involved here are 'not dependent on medical expertise'. We disagree. There are skills involved in these negotiations, and they are a central part of effective medical treatment. They are hard won. They are rarely taught explicitly

---

[10] For a wide-ranging discussion of the way that questions can invite certain sorts of responses see Elizabeth Boyd and John Heritage, 'Taking the History: Questioning During Comprehensive History-Taking' in John Heritage and Douglas Maynard (eds.) *Communication in Medical Care* (Cambridge: Cambridge University Press, 2006) 151–84, esp. pp. 177–8. Very often open-ended questions of the kind we are considering will elicit quite new concerns, so-called 'doorknob concerns'. For discussion see J.D. Robinson, et al., 'Agenda-setting revisited: When and how do primary-care physicians solicit patients' additional concerns?', *Patient Education and Counseling* (2015).

(though perhaps they could be); most doctors learn them as part of their medical experience. Some doctors are undoubtedly better at making them than others.[11]

TAKING RESPONSIBILITY

We have spoken of *taking* responsibility for uncertainty; here we want to say a little about what we mean by this. The expression is crucially ambiguous. One way to take responsibility is to let oneself be held responsible, and correspondingly to hold oneself responsible. A second way is to be given charge, or to take charge. We are primarily concerned with the second sense here.

In contrast, the familiar notion of moral responsibility, which has come to hold a very central place in moral philosophy[12], makes use of the first notion—that of being held responsible. This central understanding of moral responsibility has come to be framed in terms of accountability—indeed much of the recent debate has been over whether we can get by with just this, or whether we need in addition what are, by our lights, the fairly closely related notions of attributability or answerability.[13] This understanding is sometimes buttressed with a bit of etymology: to be responsible, on this understanding, is to be able, or liable, to respond; the responsible individual needs to be ready to give an account of themselves. Responsibility is response-ability, as John Gardner once put it.[14]

In spelling out the alternative conception that we have in mind, we'll start with that etymology, since there is something rather curious about it (to be fair to Gardner, it's not clear how seriously he meant it, but we think it's worth pursuing). The English suffixes 'ible' and 'able' are typically passive: the drinkable thing is something that can be drunk, not something that can drink, and so on.[15] In its earliest usages, 'responsible' followed the same path: a responsible person in Anglo-Norman law was not someone who had to respond; on the contrary, they were the person to whom a response had to be given. The responsible were those who issued writs;

---

[11] For discussion of the claim that these should be thought of as medical skills even though they are not explicitly taught, see Liddel et al., op. cit.

[12] The most widely used introductory philosophical textbook is entitled Reason and Responsibility; desires for alliteration notwithstanding, 'Responsibility' is clearly seen as so central that it can here stand for the whole of ethics. This centrality is really a 20th c. phenomenon; Hume and Mill scarcely discussed it; Sidgwick talks about it only a little when talking about free will.

[13] SeeWatson, Gary. 2004. *Agency and Answerability*. Oxford: Oxford University Press; Smith, Angela. 2012. "Attributability, Answerability, and Accountability: In Defense of a Unified Account." *Ethics* 122: 575–89.; David Shoemaker Shoemaker, David. 2011. "Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility." *Ethics* 121: 602–32.

[14] 'The Mark of Responsibility' in his *Offences and Defences* (Oxford: Oxford University Press, 2007), at p. 182.

[15] We see the force of this clearly in new coinages: a writeable disc, for instance, is a disc upon which one can write, not a disc that can write; someone is Googleable if they can be found using Google, not if they are able to use Google; and so on.

others then responded to them.[16] The responsible person was thus passive when it came to responding, but of course they were active in a deeper sense, in that of demanding the answer. It was the responsible person who was calling the shots.

This deeply active notion of responsibility has continued on in the rather different notion of responsibility we are after. If someone is given responsibility for the drinks at the party, it is true that you might go to them to complain if it turns out that the drinks ran out. But to focus there is to focus in the wrong place. The core of the issue is that they were in the position of managing the drinks, and that would be so even if no one were ever in a position to hold them to account—if no one were in a position to *hold* them responsible. One way of putting the distinction we have in mind is thus to contrast (i) being *given responsibility* for something; with (ii) being *held responsible* for something.

Being given responsibility though is only one way to get it: one might assume responsibility oneself. So to give a canonical usage that is broad enough to cover the cases, from now on we will contrast

being responsible for something

with

being held responsible for something.

Drawing the distinction this way, though somewhat stipulative, is close enough to normal usage. The first is the notion we are after; the second is the notion that dominates the philosophical discussions of moral responsibility.

There have been a few philosophical discussions of this notion of being responsible. Hart terms it 'role responsibility', but we find this somewhat misleading since it suggests that it comes automatically with a given role, whereas it is much more plastic than that—plastic in what is required, and plastic in the situations in which it can arise.[17] It has though received far less attention that the idea of being held responsible.

---

[16] See, for instance *Year Books of the Reign of Edward III, Trinity Term of the 16th Year* (1342), trans. and ed. Luke Owen Pike (Cambridge: Cambridge University Press 1900) No 36, pp. 158–9. Pike translates 'responable' as 'entitled to an answer'.

[17] H L A Hart *Punishment and Responsibility* (Oxford: Oxford University Press, 1968) pp. 212–13. Hart did want to generalize the idea to a task 'assigned by agreement or otherwise',although most subsequent discussion has kept the idea of an institutional role; see, for instance, P. Cane, 'Role Responsibility *The Journal of Ethics* , (2016), Vol. 20, pp. 279–298. The other recent discussion has been in the works of hard determinists, who want to use it in lieu of the more traditional accounts of responsibility; see e.g. Bruce Waller, *Against Moral Responsibility*, (Cambridge, MA: MIT Press, 2011) where it is termed 'take-charge responsibility'. We want to retain both notions.

We do not deny that these two states can go together, and typically they do: if one has responsibility for something, one is typically held responsible if it goes wrong, and conversely. But they can come apart.

One can be held responsible, or take responsibility, for an outcome, even if one was not responsible for the events: one could take responsibility for events that happened before one was born, or for the doings of one's subordinates etc. The classic British example of the latter is of a government minister taking responsibility for things done in his or her department, even without, quite legitimately, having had any knowledge of them.[18]

Conversely, one can have responsibility without being held responsible, for instance if one was too young, too inexperienced, out of one's depth etc. For instance, to take another British example, the initial responsibility for handling the disastrous Grenfell Tower fire in West Kensington in 2017 fell to a relatively junior watch officer, who had no experience of high-rise fires and had been poorly informed about the particular circumstances of that building. As became clear from reactions to his heartfelt testimony in the Public Enquiry that followed, it is very plausible that though he certainly had responsibility for the fire-brigade's initial response, he shouldn't be held responsible for it, perhaps because he shouldn't have been given that responsibility in the first place

Even when someone both has responsibility and is properly being held responsible, there is no guarantee that the two be proportional. Both come in degrees. If I have given responsibility for you to you, the two of us together may be held responsible for the outcome. The link between the two notions is not an analytic one, but a substantial moral one: it a substantial moral claim that someone should be held responsible to the degree that they are responsible. While we accept that there is much that is attractive in that claim, the examples that we have given should show that it is far from obvious that it is true.

So when we talk of responsibility for uncertainty in medicine we mean it primarily in terms of having responsibility, rather than of being held responsible for it. To take responsibility for the uncertainty is to manage it. This involves both an epistemic dimension, and a practical one. The person who manages the uncertainty needs to keep abreast of the information they have, and to handle it properly—no ignoring of base rates. And they need to act on this information, and the inferences they draw from it in the right way. And in addition there is also an emotional load. As we saw earlier, uncertainty can be emotionally draining. Doctors are typically

---

[18] Even in the UK, this seems to be little honoured nowadays: the standard example concerns Austen Chamberlain's resignation after the disastrous siege of Kut in WWI; somewhat more recent (though less clear cut) is the resignation of Lord Carrington, the British Foreign Secretary, after the invasion of the Falkland Islands. More recently UK ministers have tended to follow the practice of assuming that if they can show they did not know what was going on, this absolves them of being held responsible—that is, in our terminology, they think that if they did not have responsibility, they cannot be held responsible. In this they follow most other jurisdictions.

better placed than their patients to cope with this: better informed and less involved. But we need to be aware of the costs.

Should the doctor also be held responsible for these decisions? We remain uncommitted, since there are difficult issues here. Some concern the issue of who is to hold them responsible—the patient? the family? the authorities? Others concern what we earlier called the contingency of regret. Suppose that a doctor does take responsibility for uncertainty, decides not to follow up a remote possibility, and things go horribly wrong because of it. It is the stuff of medical nightmare. No wonder there is an emotional cost to taking responsibility. Should the doctor be held responsible for the outcome? They can of course justify their decision in terms of the expected outcomes, and that is certainly what they would be likely to do if there is any malpractice issue. Suppose though there was no malpractice: a reasonable decision was made, and the patient was unlucky.

In retrospect, of course both doctor and patient will wish that the remote possibility had been followed up. The doctor will doubtless feel agent regret—we would be concerned if they did not.[19] Designing policy to avoid such cases leads to defensive medicine with all its costs, which might suggest that the doctor simply not be held responsible. But that in turn brings concerns about accountability, and the incentive structure that such a system would produce. We have, in short, a difficult tangle of both moral and policy issues that we do not seek to untangle here. How to respond in such circumstances—who, for instance, should pay to ameliorate the damage—and whether this should be understood following a broadly 'civil' or 'criminal' model, is a question for a very different paper.[20]

JUSTIFICATION

We have tried to describe what it is to take responsibility for uncertainty; but we have said little about the circumstances under which it might be defensible. The question here is tied to the broader one of when a doctor can be justified in not telling a patient something, a topic that has generated a fair bit of discussion.[21] We start with some cases, keeping the discussion focussed as far as possible on the specific issues around uncertainty.

---

[19] Bernard Williams 'Moral Luck'

[20] For a good discussion of the distinction see Peter Cane, *Responsibility in Law and Morality* (Oxford: Hart Publishing, 2002)

[21] See especially the discussion around nocebos, for instance, Wells, R. E., And T. J. Kaptchuk. 2012. To Tell the Truth, the Whole Truth, May Do Patients Harm: the Problem of the Nocebo Effect for Informed Consent. *American Journal of Bioethics* 12 (3):22–29; Cohen, S. 2014. The Nocebo Effect of Informed Consent. *Bioethics* 28 (3):147–54.

Case 1
A woman in her late 30s presents to the hospital with a cough and some chest pain. She is worried that she might have a clot on her lung—she is in the early weeks of a pregnancy. She is fit and active, and has walked up the stairs to the clinic without difficulty. On further discussion, it transpires that she has had some general aches and pains and a slight fever; her husband has had similar symptoms. Her blood tests are consistent with a viral infection (and she is covid negative). The consultant simply tells her that she does not need a scan, since it is very unlikely that she has a clot, and the scan carries a risk for both her and the foetus. The consultant reassures her that if she herself were in this position, she would not be having a scan. She tells her to come back if it gets worse, and gives her the number of the clinic to call.

Note what is happening here. The consultant does not just *offer* a course of action to the patient, or make a *suggestion* as to what to do. Rather she makes what Stivers et al. call a *pronouncement*: she tells the patient that she does not need a scan.[22] There is a reason for this: the patient is anxious, the consultant seeks to reassure, and by taking charge in this way she does so. A suggestion or an offer—'What if we send you home without a scan?'; 'We could skip the scan; what do you think?'—would be less effective as assurance. It remains open to the patient to insist that she wants a more thorough investigation, but that would require an active intervention, taking back the responsibility. The consultant is responding to clinical considerations, but, if she is doing her job well, she is responding to interactional ones too. She has judged that the patient needs the kind of guidance and reassurance that her pronouncement provides, and that it would be welcome, not an unwanted imposition. Such judgments are subtle, and it is all too easy to get them wrong, but they are the stuff of day to day medicine.

Case 2
A 90 year-old patient enters a hospice on their GP's referral. They have a cancer that they have chosen not to have treated. They expect this to be their last few weeks, and make their plans and set their mind accordingly. The consultant, who is less convinced than the GP that their cancer is so advanced, tells the patient that, unless their condition deteriorates rapidly, they will be sent home after a couple of weeks of observation. The patient spends the week in a state of anxiety —fearful that they are a fraud, and fearful that they will be sent home—before it becomes clear that the cancer is well advanced, and there is no question of them being sent home.

This case brings more complications. Clearly the patient had built up an understanding with the GP, on the basis of which they had moved into the hospice. The consultant has sought to construct a rather different relationship, one based on a more explicit presentation of the possibilities. We do not argue that they were

[22] Tanya Stivers , John Heritage, Rebecca K. Barnes , Rose McCabe, Laura Thompson & Merran Toerien (2017): 'Treatment Recommendations as Actions', *Health Communication*.

obviously wrong to do so; there might have been something important to be gained from the greater transparency. But given the predictable distress that ensued from it, we do suggest that it would have been at least permissible to have maintained something like the GP's attitude, at least until the course of the disease was clearer. That would have been to take on the responsibility for uncertainty from the GP, rather than giving it back to the patient.

> Case 3
> A 27 year-old man is admitted to hospital after a motorcycle accident . He has suffered a catastrophic brain injury; while he is not 'brain dead' there is very little brain activity that can be picked up on the many scans that have been taken. He is on a ventilator, which is keeping him breathing. The neurosurgeon looking after him cannot reliably put a number on his chances of meaningful recovery. He doubts they would be as good as 1/100; he would hazard a guess at 1/1000. From past experience he knows that if he were to give any such figure to the family they would be more likely to want to keep him on the ventilator than if he were simply to say that he thinks it is almost certain that he won't make a meaningful recovery; and he knows what existence is like for the 999 who do not. So he has a blanket policy of not giving probabilities when the chances are small. He simply tells the family that he is very sorry, but he doesn't think the patient will recover.

Here again the neurosurgeon has taken responsibility for the uncertainty. Here though we are less convinced of the rightness of the approach, and others have shared our concern. The issue is the blanket policy. It is one thing to negotiate an understanding with a patient or with their family, and on the basis of this to take on the responsibility for the uncertainty. It is quite another to have a unilateral policy of doing so, even if this is motivated by legitimate beliefs about the best outcome.

We tentatively propose then that the key requirement for taking responsibility for uncertainty stems from a relationship with the patient, or, if they are in no position to form it, with those close to them. We think it is best characterised as a relationship of trust. To start to spell it out, let us return to the issue of informed consent.

INFORMED CONSENT, INFORMED TRUST

At this point, there are two obvious approaches that we might take. One is to say that, in the some of the cases described above, a doctor can be morally justified in taking responsibility for uncertainty; this is incompatible with informed consent, since by hypothesis the patient isn't properly informed; and hence informed consent must give way as a necessary condition. The other is to say that taking responsibility for uncertainty is really quite compatible with informed consent, once the latter is properly understood: it is just that the patient needs to be informed at some more abstract level, and then to consent to that. In fact our approach runs

somewhere between the two. We do think that there is an understanding of what it is to be informed that justifies the claim that there is informed consent in these cases. But rather than being the touchstone of good medical practice, the corresponding notion of informed consent must find its place amidst a welter of different requirements, that need to be suitably balanced. Most centrally, we think that we need a notion of *informed trust*.

We start with Manson and O'Neill's influential discussion in *Rethinking Informed Consent*. One of their targets is the 'conduit/container model' of information flow: the idea that, in a consultation, information held by the doctor flows to patient. The patient needs to understand the information for the flow to work; but once it has worked, the patient is able to provide informed consent, and this is the basis for legitimate treatment.[23]

Manson and O'Neill think that the conduit model is hopeless: it cannot be universally applied, and even where it can it doesn't give the right results. In its place they propose an 'agency model'; their stress is on the need to recognize not just what is said (the speech content) but also what is done (the speech act). We want to explore this idea further. A consultation is indeed not just a transfer of information. Like any conversation it involves a set of illocutionary acts (Manson and O'Neill's speech acts), and a set of perlocutionary consequences. Assertion is one such illocutionary act; a change in the hearer's belief—a coming to be informed — is a typical perlocutionary consequence. Consenting is also an illocutionary act; the creation of an agreement, or a contract is a typical perlocutionary consequence. Focussing just on those gives a single model for a successful consultation: the patient informs the doctor; the doctor informs the patient; this continues until all the necessary information has been transferred. The doctor then proposes a treatment. The patient is then in a position to give, or to withhold, informed consent.

But now consider some of the huge range of other illocutionary actions, and other conversational interactions, that occur in a typical consultation, and the perlocutionary states that result. These have been well documented in sociolinguistics. They include:

Questioning
Advising
Reassuring
Warning
Offering
Bargaining

A cooperative conversation involves the attempt by all parties to understand each other, and from this to build an agreed understanding: a common ground. If things

[23] Neil Manson and Onora O'Neill, *Rethinking Informed Consent in Bioethics* (Cambridge: Cambridge University Press, 2007). See esp. Ch. 4.

go well, from all of this comes a growth in trust: each of the parties comes to trust the other. Inviting trust is not typically a direct illocutionary act: it is possible to say 'Trust me', but this will often be either redundant or ineffectual. Instead trust grows through the variety of other interactions. If a patient and doctor have worked together for years, then it can be largely grounded in a history of effective treatment. But if no treatment has happened yet, the trust must be grounded in the conversations around diagnosis: in the kinds of speech interactions that we have been discussing up till now.

Our basic contention is this: for the trust to be legitimate, both patient and doctor need to be relevantly informed. That is, they need to be informed to the point that can legitimate that trust. Call the resultant trust 'informed trust'. Clearly you cannot build legitimate trust on self-serving lies. Nor can you build it by withholding information on self-serving grounds. But the whole point of trust is that it grants a degree of discretion to the one you trust, and that discretion will mean that very often there will be important things that they know and you don't. This happens all the time, in perfectly routine cases. You trust a professional to fix some electrical thing in your house whose workings you do not understand. You trust your partner to organise a surprise event for your birthday; telling you would spoil the surprise. Or you trust them to take the children rock-climbing without telling you quite what they are doing, because the whole idea of the children exposed on the mountain terrifies you.

Do you have informed consent in these cases? Not really. You consent, but your consent is not fully informed. Because you trust the person, you consent to them doing things about which you are not informed. But the trust itself has to be informed: you need to know that your electrician knows what they are doing; that your partner will organise a party that, even if it doesn't delight you, will not appal you when it happens; and that they will be responsible on the mountains, even if they do things with the children that you would never dream of doing.[24] Note though that there are differences between the cases: when you trust the electrician, you are primarily trusting them to be technically proficient. When you trust your partner you are trusting them to make decisions that are much broader than that.

We think that, in good medical practice, things will be similar with the relationship between patient and doctor; here again we have been much influenced by Manson and O'Neill's discussion.[25] The doctor trusts the patient: trusts them to be honest in their reporting of symptoms, trusts them to keep to the treatment, to make sensible decision on when to call if things start to go wrong. Equally, the patient trusts the doctor. This has been much more discussed in the medical literature, but too often the trust has been talked about as though it is just trust in technical

[24] Of course, if things go wrong it might turn out that your views on what was responsible will come apart—but this, we think, is another instance where we shouldn't understand things in terms of what happens in case of failure. It may well be that the failure will change your attitude to the whole thing.

[25] Op cit, Ch. 7.

proficiency: the patient trusts the surgeon to operate on the right organ, complete the procedure, not leave their instruments inside them, and so on. But as we hope our discussion has shown, much of the trust is more like that shown between partners. The patient will often accord much discretion to the doctor: to decide on treatment, to handle the uncertainty, perhaps to act as a custodian of information, selecting and sharing that which is is in their interests, and withholding information which they think will not be in their interests to know. Here there should be consent, but again it will not always be fully informed consent. What needs to be informed is the trust.

As with relations between partners, there is no uniform level of discretion that will suit all cases. Some patients will want more control than others, as indeed will some doctors; one of the functions of the verbal interactions we have been talking about is to arrive at the right level here.

SHARED DECISION MAKING, SHARED UNDERSTANDING

We turn finally and briefly to shared decision making. There has been a great deal of enthusiasm to involve patients in the decision making process.[26]As that approach has been developed, the idea has been truly share a decision, both parties must be equally well informed. Sometimes that is achievable, and desirable, and many of the proponents of shared decision making have convincingly argued that we often give up on it too early: patients are put off by a lack of self-confidence, or a culture of deference.

But in many of the cases we have been talking about it is neither possible nor desirable for all of the parties to be equally well-informed; that has been our theme. So in so far as full shared decision making requires full information, it will not be possible. What is achievable, and desirable, is a shared understanding of how the labour will be divided: of who will manage the uncertainty, who will make the decisions, and so on. At that point there should be shared decision making—a shared decision on the division of labour. But that may equally be a decision about where decisions will not be shared.[27]

CONCLUSION

We have stressed here the role of informed trust in legitimating a doctor's taking responsibility for uncertainty. We could equally have stressed the role of informed

---

[26] See for instance Glyn Elwyn et al.'Shared Decision Making: A Model for Clinical Practice' J Gen Intern Med 27(2012):1361–7; Glyn Elwyn, Adrian Edwards, and Rachel Thompson, *Shared Decision Making in Health Care: Achieving evidence-based patient choice* (Oxford, Oxford University Press 2016)

[27] We are grateful to Richard Lehman for discussions here.

trust in many other medical contexts; we do not mean to suggest that it only has application here. But this will do for now.