

### **Risk and Radical Uncertainty in HIV Research**

A central question in HIV research ethics is this: ‘Antiretroviral therapy for HIV works *pretty well*. HIV positive patients who interrupt it for the sake of participating in our research studies are often taking a risk. How much of a risk should we, the researchers, allow them to take?’

One reply: ‘we should allow them to take whatever risks they informed-ly consent to. They are autonomous adults. So long as we have told them all that we know about the possible costs and benefits associated with participating in the study, we should respect their decisions to participate.’

Another reply: ‘Fully informing research subjects will sometimes involve presenting to them a complex picture that they are unqualified to assess. In these situations, particularly when the possibility of a ‘cure’ is salient to them, some subjects may want to take risks that we, the researchers, judge to be *desperate*. It is not okay to be complicit in desperate risk-taking. It is doubly not okay to do that when we stand to benefit from the desperate risk-taking. There is an upper bound to how much risk we should allow our subjects to take.’

If we take the second reply seriously then we have the job of specifying just what the upper bound is, and explaining how we tell whether it is being exceeded. Roughly speaking, this will be a matter of comparing the patients’ chances if they participate in the study with their chances if they don’t. Less roughly speaking, it will be a matter of comparing the prospect-for-the-patient associated with participating in the study with the prospect-for-the-patient associated with continuing with antiretroviral therapy. (Generally, the *prospect-for-you* associated with an option is the set of all things that might happen to you if you take it,

weighted by the chances they will happen if you take it. So the prospect-for-the-patient associated with continuing with antiretroviral therapy might be a 10% chance of living a further 50-60 immuno-compromised years, a 20% chance of living a further 40-50 immuno-compromised years... etc.) If the *expected value* of the prospect associated with participating (think of the expected value of the prospect as the sum of the values-for-the patient of each of the things that might happen, weighted by the probability that it will happen) is *sufficiently close* (what counts as sufficiently close? – that will be depend on just how desperate or self-sacrificial we want to allow our research subjects to be) to the expected value of the prospect associated with not participating, then we should allow participation. Otherwise we should not.

But there's a problem. Antiretroviral therapy for HIV is thirty years old. Millions of people have used it. We have an enormous amount of excellent statistical data on its effectiveness. On the basis of that data, we can often assign very precise numbers to our patients' chances if they continue with the therapy. But typically we do not have data like this to go on, when assessing our patients' chances if they participate in the study. How likely is it that a patient will develop AIDS within a year if he or she participates? We can look to the results of animal trials, the hunches of a group of prominent scientists, our understanding of viral mechanisms..., and all that may be helpful, but it doesn't suggest a precise number.

Generally, in situations in which we have fragmentary, conflicting, or sparse evidence for a proposition, philosophers like to say that it is rational to have *imprecise levels of confidence* in the proposition. Contrast, for example, the proposition that the next two coins I toss will both come up heads, with the proposition that fully autonomous vehicles will be legal throughout the city of Boston by the year 2036. How confident am I in the former? I am precisely 25% confident. How confident am in the latter? I have some grounds for

confidence (Silicon Valley is marching ahead with the technology, the economic incentives for fully autonomous vehicles are considerable...etc.), and some grounds for skepticism (these technologies often take longer than one expects to mature, Boston is a challenging urban environment... etc.) Because my evidence is fragmentary, conflicting and sparse, it does not collectively license an attitude of 53.4688109353...% confidence, or any attitude so precise.

But to say that a level of confidence is imprecise is not to say that it cannot be represented by numbers. It is just to say that a single number won't do. I am more than 5% confident that autonomous vehicles will be legal throughout Boston by 2016. I am less than 95% confident of that. My confidence can be represented by rough-boundaried interval – from around 25% to around 45%.

The problem is how to assess whether or not participation is too risky, when our levels of confidence in the prospects associated with participation are imprecise. Three basic approaches spring to mind:

*The Optimistic Approach*

Take the most optimistic precisification of your imprecise levels of confidence in the prospects associated with participation. If participation is too risky, under that precisification, then do not allow it.

*The Moderate Approach*

Take the precisification at the mid-point of your imprecise levels of confidence in the prospects associated with participation. If participation is too risky, under that precisification, then do not allow it.

*The Pessimistic Approach*

Take the most pessimistic precisification of your imprecise levels of confidence in the prospects associated with participation. If participation is too risky, under that the precisification, then do not allow it.

Suppose my confidence that a patient will develop AIDS within a year if she participates in my research study is best represented by a rough-boundaried interval from around 5% to around 30%. The Optimistic Approach would have me imagine my being given a large volume of statistical evidence that licenses precise 5% confidence in her developing AIDS within a year. If and only if it would be okay for me then, with precise 5% confidence, to allow participation, then it is okay for me now, with imprecise 5-30% confidence, to allow participation. The Pessimistic Approach would have me imagine my being given a large volume of statistical evidence that licenses precise 30% confidence in her developing AIDS within a year. If and only if it would be okay for me then, with precise 30% confidence, to allow participation, then it is okay for me now, with imprecise 5-30% confidence, to allow participation. The Moderate Approach would split the difference.

Which approach is right? Philosophers who write about rationality tend to adopt a more optimistic approach to the rational permissibility of action. Their idea is that with imprecision in levels of confidence comes an expansion in the range of things that it is rationally permissible to do. (See, for example, Levi,[1] White,[2] Elga,[3] and Hare[4]) If you have 40% confidence in a proposition and I offer you an even bet on it (if it is true I pay you a dollar, if it is false you pay me a dollar) then it is rationally permissible for you to refuse the bet, rationally impermissible for you to take it. But if you have imprecise 20-60% confidence in the same proposition then it becomes both rationally permissible for you to refuse it (because it would be rationally permissible for you to refuse it if you were precisely 20%

confident) and rationally permissible for you to take it (because it would be rationally permissible for you to take it if you were precisely 60% confident).

The philosophers have reasons for going this way. They see just three alternatives:

*The Lax Theory of Rational Permissibility*

It is rationally permissible for you to take an option iff it would be rationally permissible for you to take it under *some* precisification of your imprecise levels of confidence.

*The Moderate Theory of Rational Permissibility*

It is rationally permissible for you to take an option iff it would be rationally permissible for you to take it under *the precisification at the mid-point* of your imprecise levels of confidence.

*The Strict Theory of Rational Permissibility*

It is rationally permissible for you to take an option iff it would be rationally permissible for you to take it under *every* precisification of your imprecise levels of confidence.

They find the Strict Theory unsatisfactory because it says that very often there is nothing that it is rationally permissible to do. So, for example, when you have 20-60% confidence in a proposition the Strict Theory says it is neither rationally permissible to accept my even bet on it nor rationally permissible to refuse my even bet on it. They find the Moderate Theory unsatisfactory because it says there is no practical difference between having imprecise 20-60% confidence in a proposition and having precise 40% confidence in a proposition. Surely there is a difference. So they are left with the Lax Theory.

These general considerations about rational permissibility might seem to tell in favor of the Optimistic Approach to enrolling research subjects. After all, enrolling research

subjects is just a special case of rational action. And there are obvious further policy reasons to favor the Optimistic Approach. The bar for participation will be much lower. We will get more research subjects.

But here I want to suggest one less obvious policy reason to favor the Pessimistic Approach to enrolling research subjects. A peculiar feature of research studies is that, typically, they generate good statistical data on the prospects of their own subjects. Because of this, we can often reasonably expect that our attitudes, once the study is concluded, will be different and more precise than our attitudes now. The intervals representing our imprecise levels of confidence will in some way *shift* and *contract*. Just how will they shift and contract? Where our present confidence that the patient will develop AIDS within a year on this medication is 5-30%, will it later be 2-4%? 4-10%? 20-25%? 25-40%? 80-90%? Typically we do not know. Indeed typically we do not have precise levels of confidence about our future levels of confidence. But there is something we can say. Here are three propositions about our future levels of confidence in bad outcomes for our research subjects:

*Rising Lower Bound*

The lower bound of our future levels of confidence will exceed the lower bound of our present levels of confidence.

*Rising Mid-Point*

The mid-point of our future levels of confidence will exceed the mid-point of our present levels of confidence.

*Rising Upper Bound*

The upper bound of our future level of confidence will exceed the upper bound of our present levels of confidence.

Because we expect the intervals representing our imprecise levels of confidence to contract, we should be more confident in *Rising Lower Bound* than in *Rising Mid-Point*, and we should be more confident in *Rising Mid-Point* than in *Rising Upper Bound*. (Why? Note that, if the interval contracts, any shift that raises the mid-point also raises the lower bound, but some shifts that raise the lower bound do not raise the mid-point, and any shift that raises the upper bound raises the mid-point, but some shifts that raise the mid-point do not raise the upper bound.)

This means that, if we adopt the Optimistic or Moderate Approach, we should be more confident that we are enrolling subjects in a study that will generate evidence in light of which it will be impermissible, by our own standards, to enroll them, than if we adopt the Pessimistic Approach. But it is bad, other things being equal, to enroll subjects in a study that generates evidence in light of which it is impermissible, by our own standards, to enroll them. What should we do, as this evidence begins to accumulate? Our options are effectively to ignore the accumulating evidence until the study is complete, or to refuse to allow the now-too-risky subjects to participate further in the study. Neither of these options is good. Just as it is not okay to allow subjects to take risks that we consider desperate, so it is not okay to allow subjects to take risks that we know we would consider desperate, if we were not now willfully ignoring evidence. And releasing research subjects mid-study is a waste of time and resources – better never to have enrolled them at all.

So we have a reason to be pessimistic in our assessment of risks to our research subjects. That will allow us to expect to keep them in our studies without exposing them to risks that are by our own lights unacceptable.

*Funding*

This work was supported by the NIAID grants 1 R01 AI114617-01A1 and 1 R56 AI114617-01. The sponsor had no role in the writing of this paper and in the decision to submit this paper for publication.

*References*

- 1 Levi I. *The Enterprise of Knowledge*. Cambridge MA: MIT Press 1980.
- 2 White R. Evidential Symmetry and Mushy Credence. In Gendler T, Hawthorne J, eds. *Oxford Studies in Epistemology 2009*, Oxford University Press 2009:161-186
- 3 Elga A. Subjective Probabilities Should be Sharp. *Philosophers' Imprint* 2010;5.
- 4 Hare C. Take the Sugar. *Analysis* 2010;2:237-247

*Abstract*

How much risk can we expose our research subjects to? There is a special challenge answering this question when the evidence on which we base our assessments of risk is fragmentary, conflicting or sparse. Such evidence does not support *precise* assignments of risk (e.g. there is a 24.8% chance that this patient will develop AIDS in the next year if she participates in my study). At best it supports *imprecise* assignments of risk (e.g. there is between a 5% and 35% chance that this patient will develop AIDS in the next year if she participates in my study). Here I discuss three approaches to evaluating risk when probability assignments are imprecise – an optimistic approach, a moderate approach and a pessimistic approach. I offer a practical reason to favor the pessimistic approach.