

What Evidence? Whose Medicine? And on What Basis?

Rani Lill Anjum

When we get sick, we want to receive the best treatment that medicine can offer. Ensuring that this happens is the main purpose of evidence-based medicine (EBM). The basic idea is that the choice of treatment should be informed by scientific evidence. But what is the alternative? That doctors should base their medical decisions on evidence seems a no-brainer. How could anyone disagree?

The evidence-based medicine movement started in the early 1990s and was primarily intended as a methodological revolution. Its proponents suggested the best way to establish the effectiveness of treatment and new criteria to choose between available treatments without bias. Philosophically, however, these changes were not so innocent, ontologically speaking. In fact, in bringing itself closer to science, medicine has ultimately become far less suitable for dealing with complex illnesses, individual variations and, I will argue, with causation in general. Let me explain.

In many ways, evidence-based medicine is the perfect empiricist, positivist system. From that philosophical perspective, therefore, it is the best scientific approach. Instead of developing theory, which typically relies on some form of metaphysical speculation, one starts from purely observable data. And rather than considering individual propensities of the patient, probabilities are calculated using statistical frequencies of group averages. But perhaps most importantly, causation is treated purely empirically. Causal conclusions are derived from correlation data generated by statistical and comparative methods.

What is new for evidence-based medicine is thus not that decisions should be based on scientific evidence. In some sense they always were. The new part is emphasis on data over theory, frequencies over propensities and the adoption of David Hume's regularity theory of causation, the idea that causes are invariably followed by their effects.

Hence my questions: *What evidence? And on what basis?*

A normative motivation of evidence-based medicine is to guarantee fairness and equality concerning treatment. One way to do this is by avoiding implicit biases and prejudices in medical decision-making. It should not matter what particular experiences or preferences individual doctors have about what treatment works best. All patients should get the best available treatment, independently of who they are or what doctor they can afford. Ideally, then, if the best treatment for a medical condition is shown to be x, then everyone should be offered x.

Another commitment of evidence-based medicine is to empiricism: only what can be observed directly represents true knowledge. Scientifically, this means that data carry more epistemic weight than theories. Evidence-based medicine is explicitly downgrading knowledge about causal mechanisms, which is not strange since these are theoretically motivated and not dictated by data. The main causal question that evidence-based methods should help us answer is *what works* and *how often*; not *why* or *how* it works. This, it is assumed, can be found by looking at correlations and difference-makers, which means that statistical and comparative methods are the best tools for establishing causation.

By looking only at the observable data, we might not find out anything about the underlying causal processes of how an intervention is doing its causal work, but we can check whether the intervention makes a difference to the outcome. The gold standard for this type of knowledge is randomized controlled trials (RCTs). For an RCT, people are randomly distributed into two groups: the test group who receives the intervention (usually a drug), and the control group who receives nothing or a placebo. If there is a greater rate of improvement in the test group than in the control group, it seems fair to conclude that the treatment was causally effective.

The motivation for giving evidential priority to such studies over causal theories of mechanisms is that if our scientific claims cannot be backed up by data, we have little reason to trust them. Doctors might recommend for their patients certain treatments because they think they understand how and why they work, but these theories could be mistaken. One problem with theories is that they are not themselves observable. Another problem is that for any set of data, there could be many possible theories that explain them equally well; the problem of underdetermination.

By sticking to observable data, therefore, one avoids what Hume described as 'metaphysical speculation'. As a strict empiricist, Hume was skeptical of scientific theories, since they are attempts to conclude and predict beyond the evidence. If we observe some effect in a sample, small or large, we cannot automatically conclude that the effect will also occur outside the sample. Nancy Cartwright calls this the inferential gap from 'it worked here' to 'it will work there'. We can also call it external validity: how can we know that what we have seen in the sample applies to the general population?

Some epidemiologists are reluctant to make causal claims on the basis of data and prefer instead to speak of raised incidence of an outcome or increased relative risk. Such skepticism towards causal claims is compatible with Hume's empiricism, from which he also analyzed the notion of causation. All we can know about causation, according to him, is what can be observed. Hume's perfect instance of causation was on the billiard table. We see a ball rolling, touching another ball, which then starts rolling. We think that the first ball caused the second ball to move, but this only comes down to habit of expectations, based on prior observation. Hume urges that we think of causation as nothing over and beyond two types of events, one regularly following the other. On this view, it is perfectly rational to search for causation via correlation data. The more robust the correlation, the stronger the evidence of causation. For Hume, there is nothing more to be found of causation than this. A is repeatedly followed by B in time and space; constant conjunction, temporal priority and contiguity. What we cannot observe is that B happens *because of* A, or that the next time A happens, B will follow.

Even for the most hard-core empiricist, causation remains a challenge for science, and medicine is no exception. If we follow Hume, then causal conclusions can be inferred from data, but they lack predictive power. All we can know is what we already observed. To predict anything beyond that would be an inductive inference. Only a perfect data set, which includes past, present and future instances, could satisfy the strict empiricist requirement for causal predictions. But then, of course, one wouldn't need to make predictions, since the future was already available to us. In reality, however, predictions are crucial. This is especially so when the aim of the scientific study is to make decisions about medical treatments.

In the absence of a time machine, our best method for making predictions is to gather as much data as humanly possible. The more data, the better the causal evidence. One individual study might not count for much, but if we can look at all the relevant studies performed, in a meta-analysis or systematic review, we have a better grounding for our causal conclusions.

Now here is a problem for the statistical approach to causation. Even if we grant that causal conclusions can be made using statistical and comparative methods, it's a question of to whom the results apply. Thus my question: *Whose medicine?*

Hume took it as a fact of causation that same cause should always give the same effect. In reality, however, we know that there are all sorts of individual variations according to particular circumstances where a cause C, that usually causes E, fails to do so in some cases. General causal claims are therefore often spelled out with a *ceteris paribus* clause, or as conditional upon some very specific or ideal conditions: C causes E, *ceteris paribus*, or all else being equal.

For medicine, this is a difficult starting-point. If everything else remains the same, then the same treatment should have the same effect in all patients. The problem for doctors, however, is that every patient is different. There are no 'ideal' or even 'normal' conditions. Still, when decisions need to be made about which treatment should be used, the doctor must assume that the patient in front of her is some sort of normal or average, at least of his relevant patient group. A relevant patient group can be thought of as a group consisting of your 'duplicates': same gender, age, medical condition, socio-economic status, lifestyle, and so on. This way, each sub-group has some degree of homogeneity, for which 'same cause, same effect' can be assumed. Information about the patient group is then used to inform decisions about an individual patient. How does this work?

Suppose a treatment cured 30 percent in a study. In each individual patient, one might then infer that it was a 0.3 probability for being cured by the treatment. But is that so? To make inferences in this way, from group level to individuals, is sometimes referred to as an ecological fallacy. If we follow Hume and the empiricist account of probability, however, then individual propensities are exactly what is given by statistical frequencies. How else could one find out about a person's chance of recovery, if not by looking at the rate of recovery in other people? The same goes for causation. If we want to know whether causation happened in this particular case, we first need to see what happened in similar cases elsewhere. Assuming this philosophical perspective, causal knowledge goes from the general to the singular, from the group average to the individual.

This approach does not seem to work for all medical conditions, however. In the Western industrialized world, about 30-50 percent of all symptoms presented to doctors are medically unexplained, which means that they do not have a known cause. These include conditions such as chronic fatigue syndrome, irritable bowel syndrome, low back pain, post-traumatic stress disorder, multiple chemical sensitivity, tension type headache and fibromyalgia. Three features in particular make these conditions a challenge for medicine: vast individual variations, complexity of symptoms and medical uniqueness. As a result, researchers have not been able to identify causes for these conditions. This claim comes with a qualification: causal explanations might be used in individual cases, but nothing that can be backed up statistically when looking at group level.

In the research project *CauseHealth – Causation, Complexity and Evidence in Health Sciences* – we start instead from an assumption that the problem of medically unexplained symptoms is not a purely epistemic problem, to be solved by doing more of the same. Instead, we see it as a symptom of a deeper philosophical problem of scientific methodology, namely: *how to deal with causation in cases of complexity and heterogeneity?* Most available research methods start from the assumption that what we are looking for are single interventions suitable for all members of a homogeneous group.

Could there be another way to approach causation scientifically? There is. But for this, we need a different philosophical starting point than Hume and empiricism. Our suggestion is that the best

philosophical basis for medicine is an ontology of dispositions, together with a dispositional notion of causation. On this view, causation is an intrinsic, singular and tendential matter. In contrast, Hume's regularity theory treats causation as extrinsic, universal and perfectly regular. Does it matter which one we choose? Absolutely.

Understood as dispositional, a cause is something that has a disposition to bring about its effect, in virtue of its intrinsic properties. Paracetamol has a disposition to alleviate headaches because of its properties and the way in which it interacts with the human body; not because there is an observed regularity between people taking paracetamol and their headaches going away. This is only a symptom of causation. The real causal work happens within each person's body, in their mutual interaction with the pill. In this sense, we might say that the drug is only doing a part of the causal work. Outside the context of the body, the drug is causally impotent. We therefore have to acknowledge the patient's own causal contribution to the outcome. The effect is produced by the intervention together with the local context to which it is added. Causation is thus a complex matter, but also sensitive to context.

We saw that evidence-based medicine is motivated by the norms of fairness and equality. The treatment that has been shown to work for most should be given to all. This is what it means that decisions are guided by evidence. But if each patient is different, then how do we know that what works best for most, on average, is what will work best for a particular individual? Certainly, we would not advise someone with peanut allergy to eat peanuts as a source of protein, even if this works well for most people. Similarly, paracetamol is safe and effective for most, but in certain patients, it can provoke kidney damage and consequently cardiovascular problems. It must therefore be used with caution in patients with kidney problems, or patients who have been dehydrated or malnourished for a long time.

Two lessons can be learned from these examples about decision-making in medicine. One is that it is important to understand how a medical intervention interacts with the human body; what we call causal mechanisms. Recall that evidence-based medicine downgrades this type of evidence, since all one needs to know is whether and how often an intervention works. The second lesson is that once we know how a treatment works, we also need to know more about the patient who will receive the treatment. Is there causally relevant information that could affect which treatment is best in this case? For instance, we know that many drugs interact with each other, sometimes with harmful outcomes. This is why the doctor needs to know what other medications a patient is taking before prescribing a treatment.

From the dispositionalist perspective presented here, we must change the way in which we approach causal evidence methodologically. To spell out how exactly this should be done, and offer full justification for this, is the project of a forthcoming book by Stephen Mumford and myself, called *Causation in Science: On the Methods of Scientific Discovery*. In the meantime, however, here is a quick list of some main ideas, as they apply to the case of medicine:

1. Evidence about causal mechanisms – that is, how the intervention does its causal work – must be given epistemic priority over how often it works in a population.
2. Qualitative methods are better suited than quantitative methods for understanding such mechanisms, since they allow us to study complex interactions.
3. Case-studies are important for picking up on unexpected effects, and to develop our causal knowledge about a treatment, especially if these are too rare to show up in population trials.

4. If we understand how the treatment works, we might also be able to assess risk better in the individual case, and for long-term effects.
5. Individual propensities can be found by considering causal mechanisms and intrinsic properties, not derived from statistical frequencies.
6. Treatment ought to be tailored to the individual rather than advised based on statistical evidence alone.
7. The same intervention in different contexts should not be expected to give the same effect.
8. Knowledge about the individual patient is essential before making decisions about treatment.
9. That a condition is medically unexplained does not mean that no cause is found. Rather, it means that a cause cannot be backed up by statistical evidence.

I am not questioning the idea that medical decisions ought to be based on evidence. All doctors should do this. But from a dispositionalist perspective, the evidence on which medical decisions should be based must have a different priority than that which is currently the case in evidence-based medicine. The nine points suggest in which direction I think medicine needs to go to improve the quality of causal evidencing.