Chapter 3

Probability for the Clinical Encounter

Elena Rocca

3.1 Uncertainty and Probability in the Single Case

‘Healthcare professional’ are words we use to describe a group of professionals with a variety of different specialisations, approaches and backgrounds. And yet, one can draw out some commonalities from this diversity. There are some aspects that everyone working through consultation with suffering individuals will recognise, in different degrees, as a daily part of their job. One of these aspects is uncertainty.

No matter their specific field of expertise, every healthcare professional must cope with the fact that all patients are different. Not only is every biological setup unique; the multiplicity of contextual conditions, lives, habits and stories make every patient a special case. Because of this, the practice of inferring information from one patient (or group of patients or experimental model) to another patient always entails some unknown margin of error. Uncertainty is intrinsic to all the phases of the encounter with the single patient, from diagnosis to prognosis to treatment. Some of these uncertainties must be accepted as such, but most of them need to be somehow qualified. Even the most experienced healthcare professional needs, in every single instance, to find an answer to the question:

What is the probability that this intervention will work *this* time, for *this* particular patient?

This question, we can say, is universal in the clinical encounter. However, there are different ways to approach it, conceptually and philosophically.

We are used to think about probability as *the* way to deal with uncertainty. Probability turns uncertainty into something more tangible by somehow giving it a qualification, or even a quantification. For instance, we say that the chances that a

E. Rocca (✉)

Centre for Applied Philosophy of Science, School of Economics and Business, Norwegian University of Life Sciences, Aas, Norway
e-mail: elena.rocca@nmbu.no

certain treatment will work in the single patient is 30%, fairly low or higher compared to another treatment. This will inform the clinical choice in a more satisfactory way than just acknowledging that “the outcome is uncertain”. Thinking in terms of probability is then a useful tool for dealing with uncertainties for the single patient. But what do these numbers mean? Can they be interpreted in only one way, or in many?

Thinking in terms of probability offers multiple ways, rather than only one, for dealing with uncertainty. Although we might assign similar probabilities to a certain event, there are different ways to interpret what such assigned values, or descriptions, mean. Think, for instance, about a healthcare professional who assures a patient that the probability of recovery with a certain intervention is ‘very high’. How should the patient understand such a statement? That the majority of the patients, previously treated with the same intervention, finally healed? That the professional has a high confidence in the positive outcome of this particular case? Or does it mean that the patient’s condition at this moment is optimal to respond to this particular treatment, considering its mode of action? Clearly, these are three different types of statement, and might lead to different clinical decisions.

In this chapter, we will first explore the basic assumptions that hide behind the usual conceptualisation of probability in evidence based approaches. One version of probability is objective and the other is subjective. We then explore a dispositionalist concept of probability, which we think is most relevant for the clinical context. We will show how this understanding of probability can help integrate evidence from healthcare guidelines with evidence from the single patient.

3.2 Probability from Statistics: Frequentism

To understand the frequentist approach to probability, imagine that you see a coin for the first time, and you notice that it is two-sided. You wonder what is the probability of the coin landing heads when tossed. One way to approach this question is to toss the coin many times, and count how frequently you get a head. After a sufficient number of repetitions you might calculate that about half the outcomes were heads, and therefore you can infer that the probability of getting a head at the next toss is $\frac{1}{2}$ or 50%. Crucially, you would not be confident in drawing a conclusion after only 3, 5, or 20 tosses. The more instances you have to base your calculation on, the more you can trust the result to be accurate.

We see that the frequentist approach calculates the probability that a certain event will happen by investigating how often it happened in the past. Philosophers call this type of approach ‘empirical’, meaning that it is exclusively based on observation. Recall that David Hume only trusted knowledge that could be observed through our senses, which was his empiricist starting point (see Anjum, Chap. 2, this book). To know how probable it is for the coin to land heads or tails, you would then not need to understand anything about the coin’s properties or hidden

dispositions. You simply toss the coin as many times as you can and count the distribution of outcomes.

In reality, the calculation of past events is usually much more elaborate than in the coin example, which is why scientists use statistical models and tools to calculate the relative frequency of a certain outcome in a sequence of events.

Note also that this approach was developed by thinking about games of chance, similar to coin tossing, and as such it needs two important premises to count as successful. First, the frequentist approach presupposes that there is the possibility (at least theoretically) of an infinite number of repetitions. Second, one needs to repeat many instances of the exact same conditions. As a consequence of this, it is not possible to calculate a frequentist probability for a single case if it cannot be repeated. From a single case, all we know is the actual outcome, not the proportion of such outcomes over a series of similar trials.

Simply put...

Frequentism means to calculate the objective probability that a certain event will happen based on the proportion of positive outcomes in a sequence of trials. To calculate probability, one then needs to observe how often the same type of event happened in previous, similar cases.

3.2.1 Frequentism and Evidence Based Approaches

Now imagine that, instead of the coin toss, we are talking about an intervention for the single patient. Clearly in this case, we are missing both premises that we saw are important for the frequentist approach to be reliable. We only have one instance of this particular patient meeting this particular intervention under some particular conditions. How can a frequentist notion of probability possibly be applied to the single patient? How can we say something about the probability of effect for this individual case?

There is a way out, and it is one that is widely used in evidence based approaches. By assuming that each single patient is a statistical average of a group of individuals that are similar enough to the patient in question, one can use the group as a representative for that patient. The probability that an intervention works for this patient can then be derived from the calculation of the statistical frequency of successful outcome in their patient group. This is the principle on which clinical studies are based. To say that a patient has 30% probability of recovery from a certain intervention based on clinical studies, means that 30% of sufficiently similar patients who tried that intervention under sufficiently similar conditions, recovered (at least for the patients who participated in those trials). These kinds of predictions are

common in evidence based medicine and practice, which takes the best evidence to be statistical evidence from clinical trials.

From a healthcare perspective, however, there are some problems with this type of reasoning. By seeing the patient as an average of similar patients, one must be able to define what counts as ‘similar’. Which pieces of information are relevant in each case? Similar age, medical history, lifestyle or social status? We are normally not aware of which factors play a causal role in the single process, which is why it can be misleading to see a patient as an average of groups of other patients, even within the appropriate patient group. This issue is well acknowledged, and it is sometimes referred to as the ‘reference class’ problem. The reference class problem influences the interpretation of statistical data from population studies for the purposes of inferring the probability of an outcome for the single patient. Imagine for instance having to calculate the probability of a patient to respond to a certain class of anti-depressants. The patient has countless properties (woman, young, history of eating disorders, wealthy, diabetic, highly educated, hyperactive...), but which of these will have a role in her condition and in the therapeutic process? This is known only partially. Our patient is a member of many different classes (young wealthy women, diabetic patients, patients with eating disorders), for which the frequency of recovery from the anti-depressant differs. However, it is not obvious how the patient should be ‘classified’ in relation to her depression and its treatment, since we do not have a complete knowledge of which of her properties will play a causal role in her clinical development. A number of tactics have been suggested as a solution to this problem, and many of these consist in moving away from a purely frequentist approach to probability, and including different types of evidence in the calculation, such as mechanistic evidence (Clarke et al. 2013, 2014; Wallmann and Williamson 2017).

3.2.2 Randomisation, Inclusion Criteria and Exclusion Criteria in Population Trials

Let us now look at the case of clinical trials, of which randomised controlled trials (RCTs) are currently considered the most reliable. RCTs have the purpose to assess the frequency of recovery in a group of patients that received a treatment, compared to a group of patients that received only a control or placebo. Based on such frequency, one can predict the probability for the treatment to have a positive outcome for the single patient. This is a frequentist approach to probability, as we described above. Recall the example of the coin toss. There are two important premises for being able to infer the probabilities for the next toss, from the frequency of outcomes of previous coin tosses: first, the repetitions should be many, and second, they should all happen under similar conditions. How can these conditions be met in an RCT?

The first requirement is served by including a large number of patients in the trial. One criterion of quality for an RCT's design is that the more patients included, the better. This is however not sufficient. If we want to end up with many instances in which the same treatment is tried out in similar contexts, we also need to cancel out the influence of individual variations between different patients. This is pursued with two strategies at the same time.

The first strategy is *randomisation*. The reasoning goes like this: if one randomly assigns the patients to the two groups (the group of patients who receive the treatment and the group of patients who receive the placebo) and the groups are large enough, then there is greater probability that the relevant causal factors are distributed evenly across the groups. This way, the two groups can be considered homogeneous, or at least similar enough. This is important in case we detect a statistical difference in the outcome among the two groups, since it allows us to infer that any such difference is caused by the intervention we are testing, rather than by some other factor. Note that the difference spotted in an RCT is at the population level, not in single patients, since there will be many individual variations within both the test group and the control group.

There is an additional strategy for cancelling out the influence of individual variation on an RCT. This is to define in advance strict *inclusion and exclusion criteria* for selecting which patients qualify to take part in the study. Typically, the patients included in the study belong to a certain age range and have a certain medical profile.

When designing an RCT, it is important to define in advance inclusion and exclusion criteria, so that the sample included in the study is representative enough for the population we want to study. In other words, at the end of the study one wants to use the results observed in the selected sample in order to say something *general* about the population we wanted to study. If the aim of the RCT is to get some information about Norwegian women of fertile age, for instance, the study sample should not be imbalanced with respect to age, geography, income or health condition. As a result of this knowledge, exceptional cases, outliers, patients with conditions that could influence or confound the interpretation of the statistical results, or patients at higher risk of adverse effects might be excluded.

3.2.3 Internal and External Validity of Causal Claims from Randomised Controlled Trials

These two strategies (randomisation and predetermined inclusion and exclusion criteria) are aimed at increasing the reliability of the trial's results. We want the study to allow us to detect the effect of the intervention, and not to be confounded with other factors that could influence it. This is also called the *internal validity*, or *reliability*, of a causal claim based on a certain study design. A different matter, however, is to figure out what use we may have for such causal claim, once we know

they are valid for the experimental sample *on average*. How does the knowledge of how often something happened to the participants of an RCT apply when predicting what is going to happen in the single patient? This is the question of *external validity*, or *relevance*, of a causal claim based on the study for the case in question. The external validity of causal claims based on RCTs might be low when we are faced with marginal cases, or with multi-morbid, chronically sick patients, who rarely meet the inclusion criteria of an RCT. To what extent does the available evidence from clinical studies represent these patients? Scientists and philosophers of science have worried about this issue, when thinking about evidence based decisions and how they should be made. As we shall see, the inductive inference from ‘it worked there’ (in the study) to ‘it will work here’ (in my case) is not an easy one, and it is paved with challenges and pitfalls where the doctor’s expertise and knowledge of her patient seem to be indispensable ingredients (for a critical discussion of the external validity of RCTs, see Rothwell 2005, 2006; Cartwright and Hardie 2012).

3.3 Probability as Degree of Belief: Subjective Credence

Say you are evaluating whether to prescribe a painkiller to a 30 year old patient and want to predict the probability that the patient will get gastrointestinal side effects from the treatment. Imagine that you have also already prescribed painkillers of the same class to 100 patients, and of those patients, 30 patients experienced that side effect. If this is the only information you have on the matter, and you want to make a prediction about the probability of side effects for your patient, you might think that such probability is 0.3 or 30%.

Let us assume that after a conversation with your patient, you learn that she suffered from chronic gastritis for a number of years and only got better 5 years ago, while she still suffers from temporary relapses. With this information at hand, you might now change your belief on the likelihood of a gastrointestinal effect in the patient, from 0.3 to >0.3 . Another colleague in the same situation, however, might just have been to an information meeting with the manufacturer of the pain killer and learned that this particular painkiller acts through a different molecular pathway than the others in the same class, and does not interfere with gastrointestinal pathways. In light of this additional information, your colleague might have a different opinion than yours, and conclude that the probability of the patient getting a side effect is not that high after all, and at <0.3 or even less.

We see that, within this philosophical theory, we can have three different estimates of probability for the same patient in the same situation. The estimate will depend on which relevant facts we are aware of and how important we think these facts are for this particular patient. This suggests that the estimation of probability P is not objective or ontological, but subjective and epistemological: it concerns the information and knowledge that the healthcare professional has available at that time. The value assigned to P will then be the subjective measure of one’s own

degree of belief that an outcome O will happen, given the available evidence E . In mathematical language, this can be written in the following way:

$P(O|E)$ ('the probability P of the outcome O happening, given the evidence E ')

Every time we acquire new information, the evidence E changes and therefore our degree of belief in the outcome is updated. This might strike us as an intuitive and straightforward practice, but there are nevertheless some practical and philosophical issues to consider.

3.3.1 *Updating Belief*

The first issue is a practical one. How exactly should we update our degree of belief in light of new information? Who is apt to do it? And does 'subjective measure' of probability entail an uncontrolled subjectivism, by which anyone and anything goes? Certainly not. Proponents of subjective probability postulate that the way in which new evidence is used to update the degree of belief must follow some common rules. These are the rules of probability calculus. In other words, two clinicians might calculate a different probability of a certain treatment to work in a specific patient, but this will only happen because they have access to different information. But the way in which a piece of evidence updates the belief, should be the same for both clinicians.

Simply put...

Subjective probability (credence) is the degree of belief, or confidence, in a specific outcome given certain available evidence, as estimated by a suitable agent. A suitable agent is an agent that uses the rules of probability calculus in order to update its expectations, and the value of the probability is always expressed quantitatively.

One way to update beliefs or expectations in light of new evidence is given by the Bayesian formula. The Bayesian formula for calculating probability includes some *prior* probability (or belief), which in light of a new piece of evidence is then updated into *posterior* probability (or belief). One necessary assumption of the Bayesian formula is that we adopt a specific way to calculate probabilities that depends on the value of another probability. In the clinic, the need for such evaluations is quite common. For instance, we might want to know the probability that a therapy works well, given the patient's conditions. This is called 'conditional probability', which intuitively means 'the probability of an outcome given an intervention'. In probabilistic calculus, however, 'conditional probability' has a technical meaning and is calculated in a specific way. (For more details on the notion of conditional probability, see Anjum et al. [2018](#).)

3.3.2 Understanding the Basic Bayesian Formula

Bayesian calculations of probability can be rather complicated, and often they are made through a computational tool, such as a software. Many of the software available to the decision makers, not only in medicine and healthcare, but also more generally in the field of risk assessment, are based on Bayesianism. These software programs calculate the posterior probability of an outcome, every time new evidence is typed into the software. The disadvantage of making decisions based on software packages is that the user has to adopt the assumptions of the programmer, without the possibility of critical consideration. Although the programming of software based on Bayesian principles can be complicated, the principle on which the whole system is based is not difficult to grasp. Let us have a look at it.

The Bayesian formula postulates a way to derive posterior probabilities from the combination of prior probabilities, new evidence, and the likelihood of the event that constitutes the new evidence occurring if our prior hypothesis is correct. In its most basic form, the formula looks like this:

$$P(\text{Hypothesis} \mid \text{Evidence}) = P(\text{Hypothesis}) \times [P(\text{Evidence} \mid \text{Hypothesis}) / P(\text{Evidence})]$$

First of all, let us explain every term of the formula with an example from the clinic.

$P(\text{Hypothesis})$ is the prior probability, or the probability of a hypothesis being true prior to getting to see the new evidence. For instance, $P(\text{Hypothesis})$ could be the probability of a patient having hypertension. Let us say that the patient is young and has a healthy lifestyle. The prior probability in this case is low. Note that the problem of how to assign prior probabilities is an important one, and Bayesians disagree on the matter. We will come back to this later on.

$P(\text{Hypothesis} \mid \text{Evidence})$ is the posterior probability, or the probability that the Hypothesis is true given that we get to know the new Evidence. In our example, this could correspond to the probability that the same patient has hypertension after she came to consultation complaining about a severe headache (Evidence = headache).

$P(\text{Evidence} \mid \text{Hypothesis})$ is the likelihood that the new Evidence occurs given that our Hypothesis is true. In our example, it would be the likelihood that our patient has a headache given that she has hypertension.

$P(\text{Evidence})$ is the likelihood of the new Evidence happening at all. For instance, the likelihood of a young healthy person having a severe headache.

Now that we know what the terms mean, let us have a look at what the formula tells us.

In order to obtain the posterior probability, Bayes is telling us to multiply the prior probability by the factor $[(\text{likelihood of Evidence given Hypothesis})/(\text{likelihood of Evidence at all})]$. Why this?

The first observation is that the posterior probability is directly proportional to the likelihood of evidence happening given that the hypothesis is true. We can

understand this by thinking about our example. The likelihood of a hypertensive patient having a severe headache is high, therefore the posterior probability (i.e., our degree of confidence) of the patient having hypertension after knowing that he has headache, is higher than before knowing it. Let us imagine for a moment that the patient, instead of complaining about a headache, had complained about lower back pain. The probability of a hypertensive patient having lower back pain is not particularly high, therefore the value of posterior probability would not be higher than the probability prior to knowing that he has back pain.

The second observation is that the posterior probability is inversely proportional to the probability that the new evidence happens at all. In other words, the lower the probability of the new evidence, the higher is the updated belief on the hypothesis. Why? Think again about our example. $P(\text{Evidence})$ corresponds to the probability that a young healthy patient has severe headaches. This probability is low. Therefore, if it happens in our patient, it updates consistently the belief that we had in the hypothesis that he is hypertensive, before knowing the new piece of information. But let us now imagine that the patient, instead of complaining about a headache, complains about moderate fatigue. Moderate fatigue is a relatively frequent condition even in young and healthy people. Therefore, after knowing that the patient is often tired, our posterior belief in the hypothesis of hypertension is not that much higher than it was before knowing the new evidence.

We see, then, that the Bayesian formula is intuitive, as long as we bear in mind that prior and posterior probabilities are not intended as existing entities, but rather as *subjective* knowledge, or *degrees of belief*. Note that, ontologically speaking, it would not make so much sense to suggest that the probabilities of me having hypertension given that I have a headache somehow depends on the probability of a generic healthy woman having a headache. (For a dispositionalist discussion of the Bayesian formula, see Anjum and Mumford 2018, ch. 19 and 21.)

There is, however, something called *objective* Bayesianism, which might create some confusion. For our purposes here, it is sufficient to point out that what we said so far is general enough to apply both to subjective and objective Bayesian inference. The difference between these two is the strategy one might use to assign prior probabilities: what is the probability of a young healthy woman having hypertension, if we just know about her that she is young and otherwise healthy? How should one assign such probability? Objective Bayesian inference postulates that there needs to be a rational and agreed way to do this task (e.g. Williamson 2010). For instance, one could use the incidence of hypertension in the general population of young and healthy women.

3.3.3 *Uncertainty as Lack of Knowledge*

One important aspect of interpreting probability as degree of belief is to notice what it is exactly that generates the uncertainty. Given a certain patient and a certain treatment, why are we uncertain about the outcome? According to this philosophical

understanding of probability, any uncertainty in prediction comes from lack of knowledge. There are many sources of uncertainty in the clinic. We might lack information about our patient, her condition or about how the treatment works. We might base our expectation on clinical studies or lab tests that are flawed. Not least, we cannot know the complete set of possible outcomes, for instance whether the treatment might provoke some hitherto unknown side effects in this specific patient. If all the possible knowledge were available to us, there would be no uncertainty left.

Within the subjective Bayesian notion of probability, uncertainty is treated as an epistemological matter: a matter of what we can possibly know. In this case, probability is not understood as an ontological matter. The uncertainty or degree of belief that we estimate using the subjective Bayesian inference should then not be understood as something that exists outside of us, in the world. In other terms, given a certain set of initial conditions (a patient, an illness, a stage of illness, a treatment, a context of treatment) there would be no inherent uncertainty about the outcome. The possible outcome is only one: the trouble is that it is impossible to know it for sure unless we were omniscient beings. This suggests that the *credence* notion of probability assumes that causality is a deterministic matter: of all (probability 1) or nothing (probability 0). Probability itself, then, does not come in degrees. Probabilistic claims do not, therefore, express something about the causal strength of an intervention, but about the limits of our knowledge and our confidence in a particular outcome. We will now look at a third notion of probability that instead sees probabilities as ontological, dispositional and intrinsic.

3.4 Probabilities as Dispositional and Intrinsic: Propensities

We have seen two possible interpretations of probability: probability as the frequency of outcome in a relevant population, and probability as subjective degree of belief, or credence. Both these approaches, however, might seem somehow inadequate or unsatisfactory for the clinic. Frequentism because it must treat the patient as a statistical average of their relevant sub-group, and credence because it takes probabilities to be entirely subjective. Indeed, clinicians are likely to think of the probability that their patient will recover as ontological: something real, physically existing in the world, but also to some extent as something that is intrinsic to the patient. Regardless of whether I have a limited knowledge about a certain condition and its prognosis, there is an actual, existing probability that the patient will recover. And this is entirely independent from my own subjective belief. Such physical probability that the patient will recover is produced by the patient's own pathophysiological and contextual situation and can be called a *propensity*.

There are many different understandings of propensity in literature. Common to all these definitions, however, is that propensities refer to the single event. Karl Popper, one early proponent of the propensity theory of probability, describes propensities as dispositional properties of singular events (Popper 1959). Propensities, we might say, are an explanatory understanding of probabilities. The probability of

me falling asleep after an injection of morphine is explained and qualified by the sedative power of morphine, by my degree of habituation and by other properties of the whole situation, which in total we can call an overall propensity. A further example might help to clarify this idea.

An elderly person affected by influenza, for instance, has a certain probability of recovery, which is generated by the type of viral infection, the state of the patient's immune system, his general health and his context and lifestyle. This set of dispositions of the whole single situation generates a certain propensity of the patient to heal. Certainly, a clinician might draw insights into this individual case from the previous experience of similar cases. However, the probability of recovery of this particular patient is affected only by the physical properties, or dispositions, in place (type of virus, patient's immune system, et cetera). Ontologically, this probability is independent of the outcomes in other similar cases. Epistemologically, what happens in similar cases can be an indication for the probability of outcome in the single patient, but not necessarily. To use a brutal example, if a bus full of patients with a rare condition crashes on its way to a medical conference, their collective death does not affect the propensity of the patient who missed the bus to survive the rare condition. On a frequentist account, however, that patient's probability of survival will have changed with the bus accident (without it being a good epistemological indicator either, in this case).

Simply put...

The individual propensities of a single patient, treatment or context are given by its unique combination of dispositions and the dispositions' degree of tendency toward certain outcomes *in that individual case*.

3.4.1 *Individual Propensities Are Not Always Seen Through Frequencies*

A dispositionalist understanding of causality, as described in Chap. 2, fits best with a view of probabilities as single *propensities*. Although propensity interpretations of probability are less known than frequentism and credence, it has been defended by philosophers (e.g. Popper 1959, 1990; Mellor 1971; Gillies 2000, 2018) and scientists (e.g. Bohm 1957) since the early 1900s. Propensities seem particularly relevant in the clinic.

Think about an epileptic child who is about to start a therapy with valproic acid, a widely used anticonvulsant. What are the probabilities that the patient develops liver toxicity as an undesired effect of the drug? One possible answer is that the child will most probably be unhurt, thus with probability close to zero, on the basis of the outcome for the majority of children. However, this evaluation might be met as superficial, or unsatisfactory, given that children *can* be hurt by valproic acid, sometimes even fatally.

A better clinical approach to the question might be: what are the *propensities* of *this* child to develop liver toxicity from valproic acid? If the child has the particular physical and contextual onset that makes her sensitive to the drug's undesired effect, then the toxic outcome, no matter how rare, will nevertheless be very likely for her. The frequency of the toxic outcome in other children can in some cases be indicative of the individual propensity to the outcome (for instance, if we consider closely the properties of the patient in comparison with the properties of the harmed children), but is not necessarily so.

3.4.2 *Propensities as Qualities*

How should propensities be expressed? Can a number between 0 and 1 be estimated for this purpose, as is generally done in probability theory? Philosophers have different views on this matter, depending on their understanding of propensity. In the CauseHealth project, however, we favour a singular and qualitative, rather than a numerical, description of propensities. (For an approach to propensities that is more compatible with frequentism, see Gillies 2018.)

Propensities, we said, are generated by dispositions. These are intrinsic qualities of things. But a propensity also depends on the disposition's magnitude or intensity, which, although being in some sense quantitative, cannot be directly derived from statistical frequencies. The presence (or absence) of mutual manifestation partners for a certain disposition, and the presence (or absence) of possible causal mechanisms which might result in an outcome, affect the propensity of such an outcome to happen. Evaluating the propensity for an outcome therefore requires describing and understanding, at least partially, *how* that certain outcome could or could not happen. Numbers or scores cannot completely fulfil this purpose, at least not alone, since what happens statistically at a population level might not reflect what happens in each individual case.

On average, for instance, a population might seem to have a weak propensity toward cirrhosis, but this average only represents the sum of all the manifestations of individual propensities toward cirrhosis. Individual propensities will depend on a number of dispositions related to age, gender, genetics, lifestyle, diet and medical history. All of these will be different from one individual to another, so we should not expect two people to have exactly the same combination of dispositions, or dispositions with exactly the same magnitude.

3.4.3 *Propensities and Prediction*

There is a further important consequence of propensities being generated by dispositions. We saw in Chap. 2 that dispositions can exist unmanifested. A patient might carry a certain genetic mutation that disposes toward an allergic reaction from

penicillin, for instance. However, we are likely to remain unaware of such a disposition until it meets its proper mutual manifestation partner. In other words, it might be difficult to correctly evaluate the propensity toward this toxic reaction until the patient uses penicillin. In some cases, we might even get it totally wrong. As a consequence, evaluations of propensity ought to be carried out with some degree of epistemic humility. That means, there might be dispositions and interactions in this particular case, which we were not aware of at the moment of evaluation. Any evaluation or prediction must therefore be interpreted with some caution.

Although this attitude is valid for any interpretation of probability, it is especially important when we think about propensities and dispositions. Accordingly, probabilities as propensities are a matter of qualitative evaluation, theoretical knowledge and practical expertise, and cannot be generated by an algorithm as a definite number. Notice that this does not mean that propensities are more fallible than other ways to calculate probability. Rather, it might just make us more aware of the fallibility of prediction.

If we accept the dispositionalist notion of probability, a question remains: how should clinical inquiry (as well as research) be organised in order to uncover propensities for the single patient?

3.5 Propensities and the Clinic

A clinician who adopts the propensity approach to probability also adopts a certain specific approach to clinical inquiry. In this section, we list some of the methodological and epistemological implications of the propensity view of probability. Note that these follow from the particular version of propensities here presented and dispositionalism presented in Chap. 2, and that other versions of propensity theory might have other implications.

3.5.1 *The Importance of Local Knowledge*

The more one knows about the dispositions and interactions in place in the particular case of interest, the more reliably one can evaluate the propensity toward one specific outcome in that case. This might sound like nothing particularly new. The frequentist approach, indeed, also requires us to know as much as possible about the case of inquiry, so that the most relevant sub-population can be found, and more reliable statistics performed. So what is particular about the propensity approach? The difference is in the *type* of knowledge required. Uncovering propensities requires knowledge about local processes and interactions, rather than knowledge of mere values and parameters. Typically, local processes and interactions need to be observed in their own context and as aspects of a whole, while values and parameters can be picked and chosen, and analysed in isolation. Local knowledge of a

patient, ideally, would not be reduced to knowing the value of his biomarkers or genetic onset, but requires that we have as much knowledge as possible about his unique context, including history, lifestyle, reactions and interactions.

3.5.2 Person Centered Clinical Analysis

Knowing about the patient's local context requires, first of all, time. While parameters and values can be collected through tests and orthodox clinical enquiry, processes and interactions need to be understood through person centered dialogue. Person centered dialogue is a type of interaction in which the patient is met as a whole person: her biology, her biography, her history and her narrative are taken as equally important information for the purpose of the clinical inquiry (see also Anjum and Rocca, Chap. 4 and Low, Chap. 8, this book).

3.5.3 Focus on Theories of Causal Mechanism

In order to evaluate the propensity of an outcome for a single case, it is necessary to have an insight into how and why such an outcome might be generated. This requires a certain degree of general, theoretical knowledge about the process at hand. It is impossible, for instance, to evaluate the propensity of someone to develop diabetes without having an idea about the biological mechanisms underlying the onset of the illness. This requires that the clinician cultivates a high level of theoretical, pathophysiological knowledge along with statistical evidence. At the same time, as mentioned above, local causal mechanisms are of considerable relevance for the propensity approach. These include biological symptoms of broad interest that the clinician might notice in the patient, besides the symptoms of relevance for the targeted examination. But such local mechanisms also include higher level, socio-psychological mechanisms which might influence a patient's physiological conditions.

3.5.4 Multidisciplinarity and Networking

By adopting the propensity approach to probability, the clinician makes use of a wide number of scientific and theoretical insights, besides statistical and population studies. For instance, in order to maximise the propensity for recovery in a depressed patient, a clinician must be updated about scientific insights on the various causal mechanisms influencing the onset of depression. But the connection between research and the clinic is not a one-way street. Since the clinical search for propensities is focused on local processes and interactions, it potentially becomes a reliable source of new general scientific hypotheses about the mechanisms of healing and

disease. This is particularly the case with unexpected clinical observations, such as side effects of drugs. We can illustrate this with an example.

Zolpidem is a hypnotic drug used to treat short-term insomnia. Clinicians reported a variety of anecdotal undesired and beneficial effects in Zolpidem users: sleepwalking, sleep-eating and sleep-driving, compulsory behaviours followed by amnesia, but also speech recovery after stroke, recovery of mobility after brain injury, and recovery from posttraumatic semi-unconscious state. These insightful clinical observations resulted in new hypotheses for basic research, for instance about the mechanism of recovery after brain injury. Notice that these effects are very rare and sometimes unique, therefore they would not count as particularly relevant evidence for a frequentist.

These considerations highlight that clinicians should ideally work in a multidisciplinary network with researchers, so that information can be easily shared among diverse experts.

3.5.5 The Potential of Clinical Experience for Advancing Medical Knowledge

We often think of the perfect medical research and healthcare system as a system that places patient care as the final aim of a long process. In a way, this is hardly controversial: patients' interests must be prioritised over commercial or other economic interests, for instance. Research hypotheses, funding, and experimental designs ought to be developed with a special consideration that they are meant to be primarily useful for the patient. Important steps are being taken in this direction, and bioethics has this as a key principle of both healthcare and research.

This conception, however, must be somehow adjusted. There is nothing “final” about the clinical meeting between practitioner and patient. Quite the contrary: each of such encounters is potentially the beginning of a new hypothesis, a challenge for established paradigms, and the springboard for broadening medical knowledge. This is not difficult to believe if we think about the history of medicine.

Many have already emphasised the value of patient centered medicine and healthcare for the final purpose of improved clinical decision making, patient care, and clinical ethics. But few have talked about the fact that a patient centered clinical approach also has a significant epistemological value: it is the best available opportunity for advancing causal knowledge in research. Expansion of knowledge does not happen in a straight line, with the patient at the end of it. It is a continuous circle of trial and success or failure, where evidence from clinical cases are looping back to pre-clinical and clinical research. The more attention that is given to the clinical cases, therefore, the more opportunities we have to improve research (Rocca 2017).

From a practical point of view, what does this entail? First of all, the clinical interview takes on a crucial role, not only for the patient's wellbeing, but also for the whole healthcare community. This important process of gathering clinical evidence should not be left to individual skills and improvisation (see Hagen, Chap. 10, this

book). Medical schools should teach patient centered models of clinical communication, and should stress their key value (see Broom, Chap. 14, this book). Second, clinical evidence should be collected in databases and networked within the broad medical community (see Copeland, Chap. 6, this book). Third, researchers should recognise the primary role of patient centered evidence for the corroboration, challenge, and advance of causal knowledge.

3.5.6 *What Does $N = 1$ Mean, Within the CauseHealth Project?*

We have seen that a propensity approach to probability requires that theoretical understanding of physiology and of illness is prioritised. Statistical knowledge can sometimes be a useful tool to gain such knowledge, but it is certainly not the only type of evidence one needs in order to understand the how and why of medical phenomena. The single patient represents a major part of the causally relevant information for understanding the illness and choosing the best treatment. In CauseHealth, we sometimes summarise this central concept through the slogan “ $N = 1$ ”, which has been a source of philosophical debate among healthcare practitioners. We give a specific meaning to $N = 1$, which is distinct from the traditional meaning of $N = 1$ trials in medical research. In the following, physiotherapist Roger Kerry provides a full explanation of the slogan’s meaning:

“ $N = 1$ ” is a slogan used to publicise a core purpose of the CauseHealth project. $N = 1$ refers to a project which is focussed on understanding causally important variables which may exist at an individual level, but which are not necessarily represented or understood through scientific inquiry at a population level. There is an assumption that causal variables are essentially context-sensitive, and as such although population data may by symptomatic of causal association, they do not constitute causation.

The project seeks to develop existing scientific methods to try and better understand individual variations. In this sense, $N = 1$ has nothing at all to do with acquiescing to “what the patient wants”, or any other similar fabricated straw-man characterisations of the notion which might emerge during discussions about this notion.

In Evidence Based Medicine terms, of course, an $N = 1$ trial is a randomised controlled trial involving a single subject with a random allocation of the temporal sequence of interventions. Such a trial has traditionally sat at the very top of evidential hierarchies because it offers the best scientifically controlled conditions. CauseHealth is sympathetic to such a methodology, although the clinical notion of $N = 1$ means much more than just this method.

$N = 1$ is both an ontological claim, about the causal singularism (this means that causation is something intrinsic to the person and the situation, and does not have to be repeated in exactly the same way elsewhere to count as causation) and possibility/plausibility of the situation that each causal setting is unique.

It is also a methodological claim, arguing against the idea that the individual can best be captured by searching for the relevant sub-population. Which group should represent Rani?

Women between 40 and 50 years of age? Mixed ethnic background? Norwegians? Educational status? Etc. Say we find such a group, then why assume that this group is Rani's 'twin population'? There might be all sorts of causally relevant factors that they cannot represent but that are important in Rani's individual case. $N = I$ is about starting from the expectation that everyone is different, rather than from the assumption that everyone is statistically average. This is a fundamentally important scientific shift in how research should be operationalised and interpreted.

Despite the above, $N = I$ thinking is not at all dismissive of population studies, and sees them as critical tools which are well suited to signalling to where causal activity may well lie. However, the above limitations of population studies related to individual clinical decision making are highlighted within the $N = I$ notion.

Paradoxically, as we gain more data, experience, and maturity with our population research programmes, contextual analysis of such data starts to reveal that there is indeed no "one size fits all" approach to the management of much burdensome disease, for example low back pain. Such analyses are exemplars of how $N = I$ and population data work together.

$N = I$ is about contextualising the individual human within population data. It moves beyond a level of thinking which says "here is a patient with low back pain, let me see what evidence based interventions are available for low back pain". Rather, it is committed to understanding the human in front of us and the causally relevant factors which will influence that person's return to a desired functional level. Some of those factors will have been represented in population data, many will not have been.

Roger Kerry, 'What does CauseHealth mean by $N = I$?', CauseHealth blog (<https://causehealthblog.wordpress.com/2017/06/22>)

3.6 To Sum Up...

This chapter outlined three different interpretations of the concept of probability and explained why causal dispositionalism supports an understanding of probability as propensities, and how this influences clinical decision making and medical investigations in general. For a final illustration of the difference between the three perspectives presented above, imagine a situation in which we are going to cross a bridge with a heavy truck, and we want to evaluate the probability that the bridge will endure the weight of the truck (and consequently the risk of an accident). The frequentist approach would face this challenge by looking at how often similar bridges collapsed under the weight of similar trucks. The Bayesian approach would treat the probability as a subjective matter that changes depending on the information we have about the bridge and the truck, and would treat the probability as the value of how certain we are that an accident will (or will not) happen. The measure of such certainty will be updated every time we gain a new piece of information. The propensity approach would describe the probability using the qualities of the bridge, the truck, and the

whole situation, and trying to understand the intrinsic disposition of the bridge to collapse under a certain weight. Such intrinsicity will be evaluated by investigating the properties at hand (height, length, solidity, material) and by understanding the causal and physical processes involved. All these perspectives – frequencies, uncertainty and propensities – offer something that can be useful for expanding our causal knowledge. The philosophical question is which we take to be basic.

References and Further Readings

- Anjum RL, Mumford S (2018) Causation in science and the methods of scientific discovery. Oxford University Press, Oxford
- Anjum RL, Mumford S, Myrstad JA (2018) Conditional probability from an ontological point of view. In: Anjum RL, Mumford S (eds) What tends to be. The philosophy of dispositional modality. Routledge, London, pp 101–114
- Bohm D (1957) Causality and chance in modern physics. Routledge, London
- Cartwright N, Hardie J (2012) Evidence-based policy. A guide to do it better. Oxford University Press, Oxford
- Clarke B, Gillies D, Illari P, Russo F, Williamson J (2013) The evidence that evidence-based medicine omits. *Prev Med* 57:745–747
- Clarke B, Gillies D, Illari P, Russo F, Williamson J (2014) Mechanisms and the evidence hierarchy. *Topoi* 33:339–360
- Gillies D (2000) Varieties of propensity. *Br J Philos Sci* 51:807–835
- Gillies D (2018) Causation, probability, and medicine. Oxford University Press, Oxford
- Mellor DH (1971) The matter of chance. Cambridge University Press, London
- Popper K (1959) The propensity interpretation of probability. *Br J Philos Sci* 10:25–42
- Popper K (1990) A world of propensities. Thoemmes, Bristol
- Rocca E (2017) Bridging the boundaries between scientists and clinicians. Mechanistic hypotheses and patient stories in risk assessment of drugs. *J Eval Clin Pract* 23:114–120
- Rothwell PM (2005) External validity of randomised controlled trials: “to whom do the results of this trial apply?”. *Lancet* 365:82–93
- Rothwell PM (2006) Factors that can affect the external validity of randomised controlled trials. *PLoS Clin Trials*. <https://doi.org/10.1371/journal.pctr.0010009>
- Wallmann C, Williamson J (2017) Four approaches to the reference class problem. In: Hofer-Szabó G, Wroński L (eds) Making it formally explicit, European studies in philosophy of science, vol 6. Springer, Dordrecht
- Williamson J (2010) In defence of objective Bayesianism. Oxford University Press, Oxford

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.