

Chapter 2

An Introduction to Mechanisms



Abstract This chapter offers a brief summary of mechanisms, as including complex-system mechanisms (a complex arrangement of entities and activities, organised in such a way as to be regularly or predictably responsible for the phenomenon to be explained) and mechanistic processes (a spatio-temporal pathway along which certain features are propagated from the starting point to the end point). The chapter emphasises that EBM+ is concerned with evidence of mechanisms, not mere just-so stories, and summarises some key roles assessing evidence of mechanisms can play, particularly with respect to assessing efficacy and external validity.

This chapter introduces mechanisms and their use in the context of working with evidence in medicine. The first section gives an extremely short introduction to mechanisms that assumes no prior knowledge. Subsequent sections develop our account of mechanisms in more detail.


2.1 Mechanisms at a Glance

Mechanisms allow us to understand complex systems (e.g., physiological or social systems) and can help us to explain, predict, and intervene. An important subclass of mechanisms is characterised by the following working definition:

A complex-systems mechanism for a phenomenon consists of entities and activities organised in such a way that they are responsible for the phenomenon (Illari and Williamson 2012, 120).

For the example mechanism of Fig. 2.1, the phenomena are the effects of a drug, the drug and the receptor are the parts, and the interactions are the binding and triggering.

Fig. 2.1 A simple example of a mechanism diagram that shows the effects produced by a drug via receptor binding



Why do mechanisms matter? Mechanisms explain how things work. This makes them important in their own right, but also means that they are often used when designing clinical studies. For example, one might decide to use a biomarker to evaluate the effect of a drug, and that decision would rely on our knowledge of some mechanism that links the biomarker with the drug. Note that while mechanisms of drug action are an important kind of mechanism, they are not the only kinds of mechanism that we will consider here.

Caveats:

- We will be interested in *evidence* of mechanisms, not descriptions of mechanisms for which there is no evidence. To be useful, descriptions of mechanisms should be connected to high-quality research, and not just to either background knowledge or to what Pawson (2003) calls ‘programme theories’. Otherwise they are merely just-so stories. Descriptions of mechanisms need to be supported by evidence to be useful.
- Mechanistic studies are not normally sufficient on their own to justify treatment or policy decisions. Other supporting evidence (such as that arising from clinical studies) is normally required.
- As is the case with other kinds of evidence, evidence of mechanisms is not infallible.

Why should one scrutinise evidence of mechanisms in healthcare? As explained in Sect. 2.3 below, evidence of mechanisms can support or undermine judgements of efficacy and external validity. Therefore, using evidence of mechanisms in concert with other forms of evidence results in better healthcare decisions. (We use the analogy of reinforced concrete to explain this claim; see p. 92.) If this sort of mechanistic reasoning is not properly scrutinised, medical decisions may be adversely affected. For example, current tools for evaluating the quality of clinical research (such as GRADE) do not scrutinise assumptions about mechanisms that have been used to design clinical studies. Just as EBM improved clinical practice by scrutinising clinical studies, scrutinising evidence of mechanisms can lead to further improvements. We have provided some suitable tools for assisting such scrutiny in Chap. 4.

Mechanisms are the + in EBM+

2.2 What is a Mechanism?

Mechanisms are invoked to explain (Machamer et al. 2000; Gillies 2017b). Textbooks in the biomedical and social sciences are replete with diagrams and descriptions of mechanisms. These are used to explain the proper function of features of the human body, to explain diseases and their spread, to explain the functioning of medical devices, and to explain social aspects of health interventions, among other things.

One kind of mechanism, a *complex-systems mechanism*, is a complex arrangement of entities and activities, organised in such a way as to be regularly or predictably responsible for the phenomenon to be explained (Illari and Williamson 2012). In such mechanisms, spatio-temporal and hierarchical organisation tend to play a crucial explanatory role (Williamson 2018, Sect. 1).

Another kind of mechanism, a *mechanistic process*, consists in a spatio-temporal pathway along which certain features are propagated from the starting point to the end point (Salmon 1998). Examples include the motion of a billiard ball from cue to collision, and the trajectory of a molecule in the bloodstream from injection to metabolism. This sort of mechanism is often one-off, rather than operating in a regular and repeatable way. In the case of environmental causes of disease, the repercussions of these processes may take a long time to develop—e.g., they may be mediated by epigenetic changes.

In the health sciences, mechanistic explanations often involve a combination of these two sorts of mechanism. For example, an explanation of a certain cancer may appeal to the mechanistic processes that bring environmental factors into the human body, the eventual failure of the body's complex-systems mechanisms for preventing damage, and the resulting mechanistic processes that lead to disease, including the propagation of tumours (Russo and Williamson 2012).

We shall use 'mechanism' to refer to a complex-systems mechanism or a mechanistic process or some combination of the two. We should emphasise that mechanisms in medicine and public health may be social as well as biological (see Chap. 9 and Clarke and Russo 2017), and, in the case of medical devices, for instance, they may also include technological components.

A clinical study is the usual method for establishing that two variables are correlated:

A **clinical study** for the claim that A is a cause of B repeatedly measures the values of a set of measured variables that includes the variables A and B . These values are recorded in a **dataset**. In an **experimental study**, the measurements are made after an experimental intervention. If no intervention is performed, the study is an **observational study**: a **cohort study** follows a group of people over time; a **case control study** divides the study population into those who have a disease and those who do not and surveys each cohort; a **case series** is a study that tracks patients who received a similar treatment or exposure. An **n-of-1 study** consists of repeated measurements of a single individual; other studies measure several individuals. Clinical studies are crucial for estimating any correlation between A and B , and they indirectly provide evidence relevant to the claim that A is a cause of B (see Fig. 3.1).

On the other hand, a much wider variety of methods can provide good evidence of mechanisms—including direct manipulation (e.g., in vitro experiments), direct observation (e.g., biomedical imaging, autopsy), clinical studies (e.g., RCTs, cohort studies, case control studies, case series), confirmed theory (e.g., established immunological theory), analogy (e.g., animal experiments) and simulation (e.g., agent-based models) (Clarke et al. 2014; Williamson 2018). A mechanistic study is a study which provides evidence of the details of a mechanism:

A **mechanistic study** for the claim that A is a cause of B is a study which provides evidence of features of the mechanism by which A is hypothesised to cause B . Mechanistic studies can be produced by means of in vitro experiments, biomedical imaging, autopsy, established theory, animal experiments and simulations, for instance. Moreover, consider a clinical study for the claim that A is a cause of C , where C is an intermediate variable on the path from A to B —e.g., a surrogate outcome. Such a study is also a mechanistic study because it provides evidence of certain details of the mechanism from A to B . A clinical study for the claim that A is a cause of B is not normally a mechanistic study for the claim that A is a cause of B because, although it can provide indirect evidence that there exists some mechanism linking A and B , it does not normally provide evidence of the structure or features of that mechanism. Similarly, a mechanistic study for the claim that A is a cause of B is not normally a clinical study for the claim that A is a cause of B , because it does not repeatedly measure values of A and B together. A study will be called a **mixed study** if it is both a clinical study and a mechanistic study—i.e., if it both measures values of A and B together and provides evidence of features of the mechanism linking A and B . To avoid confusion, the terminology **clinical study** and **mechanistic study** will be used to refer only to non-mixed studies.

2.3 Why Consider Evidence of Mechanisms?

There are various reasons for taking evidence of mechanisms into account when assessing claims in medicine. In general, when evidence is limited, the more evidence one can take into account, and the more varied this evidence is, the more reliable the resulting assessments (Claveau 2013). Moreover, when deciding whether to approve a new health intervention, or whether a chemical is carcinogenic, for example, it can take a very long time to gather enough evidence if the only evidence one considers is clinical study evidence. By considering evidence of mechanisms in conjunction with clinical study evidence, decisions can be made earlier: one can reduce the time taken for a drug to reach market (Gibbs 2000), and reduce the time taken to restrict exposure to carcinogens, for instance.

There are also reasons for considering evidence of mechanisms that are particular to the task at hand. While evidence of mechanisms can inform a variety of tasks (see below), in this book we focus on its use for evaluating efficacy and external validity. Williamson (2018) provides a detailed justification of the need for evidence of mechanisms when performing these two tasks. Here we shall briefly sketch the main considerations.

Evaluating efficacy. As noted above, establishing effectiveness can be broken down into two steps: establishing efficacy and establishing external validity. Establishing efficacy, i.e., that *A* is a cause of *B* in the study population, in turn requires establishing two things. First, *A* and *B* need to be appropriately correlated. Second, this correlation needs to be attributable to *A* causing *B*, rather than some other explanation, such as bias, confounding or some connection other than a causal connection (Williamson 2018, Sect. 1).

If it is genuinely the case that *A* is a cause of *B*, then there is some combination of mechanisms that explains instances of *B* by invoking instances of *A* and that can account for the magnitude of the observed correlation. As a mechanism of action may only be present in some individuals but not others, it needs to be credible that the mechanism of action operates in enough individuals to explain the size of the observed correlation in the study population. Just finding a mechanism of action in some individuals is insufficient. Thus, in order to establish efficacy one needs to establish both the existence of an appropriate correlation in the study population and the existence of an appropriate mechanism that can explain that correlation. We shall refer to this latter claim—that there is a mechanism that can explain that correlation—as the general mechanistic claim for efficacy:

General mechanistic claim. In the case of efficacy, the general mechanistic claim takes the form: there exists a mechanism linking the putative cause *A* to the putative effect *B*, which explains instances of *B* in terms of instances of *A* and which can account for the observed correlation between *A* and *B*. In the case of external validity, the general mechanistic claim is: the mechanism responsible for *B* in the target populations is sufficiently similar to that responsible for *B* in the study population.

More generally, evidence of mechanisms can help rule in or out various explanations of a correlation. For example, it can help to determine the direction of causation, which variables are potential confounders, whether a treatment regime is likely to lead to performance bias, and whether measured variables are likely to exhibit temporal trends.

Some alternative explanations of a correlation can be rendered less credible by choosing a particular study design. Adjusting for known confounders and randomisation can lower the probability of confounding. Blinding can reduce the probability of performance and detection bias. Larger trials can reduce the probability of chance correlations. Selecting variables *A* and *B* that do not exhibit significant temporal trends and that are spatio-temporally disjoint can reduce the probability of some other explanations.

In certain cases, clinical studies alone might establish that an observed correlation is causal (Williamson 2018, Sect. 2.1). However, establishing a causal claim in the absence of evidence of the details of the underlying mechanisms requires several independent studies of sufficient size and quality of design and implementation which consistently exhibit a sufficiently large correlation (aka ‘effect size’), so as to rule out explanations of the correlation other than causation. This situation is rare: evidence from clinical studies is typically more equivocal. Therefore, evidence of mechanisms obtained from sources other than clinical studies can play a crucial role in deciding efficacy. Considering this other evidence is likely to lead to more reliable causal conclusions. Where this evidence needs to be considered, its quality should be evaluated in ways such as those set out in this book.

Evaluating external validity. Having established efficacy, i.e., that a causal relationship obtains in the study population, one needs to establish external validity—that the causal relationship can be extrapolated to the target population of interest.

As noted above, establishing that *A* is a cause of *B* requires establishing both that *A* and *B* are correlated and that there is some mechanism that can account for this correlation. Having established these facts in the study population, one can infer causation in the target population with some confidence if one can establish that:

1. In the target population, there is a mechanism that is sufficiently similar to the mechanism of action in the study population, and
2. Any mechanisms in the target population which counteract this mechanism do not mask the effect of the mechanism of action to such an extent that a net correlation in the target population could not be explained mechanistically.

Evaluating external validity, then, requires evaluating whether the complex of relevant mechanisms in the target population is *sufficiently similar* to that in the study population, in the sense of (1) and (2) holding. Evidence of mechanisms is therefore crucial to this mode of inference.

This form of inference can be especially challenging when the study population is an animal study and the target population is a human population (Wilde and Parkkinen 2017). This is because, despite important similarities between several physiological mechanisms in certain animals and those in humans, many differences also exist. This form of inference can also be challenging when both the study and the target population are human populations. This is because human behaviour is often a component of an intervention mechanism and may in fact hinder the effectiveness of the intervention. We discuss this in Chap. 9. Some well-known examples of behaviour modifying effectiveness include the Tamil Nadu Integrated Nutrition Project (India) and the North Karelia Project (Finland), both discussed by Clarke et al. (2014).

Other questions. Apart from when evaluating efficacy and external validity, evidence of mechanisms can also be helpful when:

- Drawing inferences about a single individual, for treatment and personalised medicine (Wallmann and Williamson 2017);
- Commissioning new research and devising new research funding proposals;
- Justifying the use of clinical studies, designing them, and interpreting their results (Clarke et al. 2014);
- Suggesting and analysing adverse drug effects—see Gillies (2017a), who argues that consideration of evidence of mechanisms would have been necessary to avoid the thalidomide disaster;
- Designing drugs and new devices;
- Building economic models in order to ascertain cost effectiveness of a health intervention;
- Deciding how surrogate outcomes are related to outcomes of interest.

Example. How evidence of mechanisms can help with the analysis of adverse drug effects: abacavir hypersensitivity syndrome.

Abacavir is a nucleoside analog reverse transcriptase inhibitor, widely used as part of combination antiretroviral therapy for HIV/AIDS, that received an FDA licence in 1998. However, its use was initially complicated by a severe, life-threatening, hypersensitivity reaction that occurred in approximately 5% of users (precise estimates vary; Clay (2002) gives a range of 2.3–9%). There was confusion regarding the cause of this reaction, and it was thought that ‘it is not possible to characterize those patients most likely to develop the HSR’ on the basis of reports of the syndrome (Clay 2002, 2505).

This changed with the discovery that the hypersensitivity syndrome only occurred in individuals with the HLA-B*5701 allele (Mallal et al. 2002). This discovery arose from evidence of mechanisms. These authors noted that there were similarities between the mechanisms of several hypersensitivity syndromes—by ‘evidence that the pathogenesis of several similar multisystem drug hypersensitivity reactions involves MHC-restricted presentation of drug or drug metabolites, with direct binding of these non-peptide antigens to MHC molecules or haptenation to endogenous proteins before T-cell presentation’ (Mallal et al. 2002, 727). Patients are now genetically screened for the HLA-B*5701 allele, and this has greatly reduced the incidence of the hypersensitivity syndrome (Rauch et al. 2006).

In this book, we focus largely on the use of evidence of mechanisms to help establish efficacy and external validity. The problem of drawing inferences about a single individual is briefly discussed in Chap. 10.

Importance of considering evidence of mechanisms. Recall that in certain cases clinical studies on their own suffice to establish efficacy and there is no need for a detailed evaluation of other evidence of mechanisms. In other cases, however, evidence of mechanisms arising from sources other than clinical studies can be decisive. In such cases, it is important to scrutinise and evaluate this evidence, just as it is important to scrutinise and evaluate clinical studies.

Situations in which it is particularly important to critically assess evidence of mechanisms arising from sources other than clinical studies include:

- Where clinical studies give conflicting results, are of limited quality, or otherwise exhibit uncertainty about the effect size;
- Where randomised clinical studies are not possible, for practical or ethical reasons, in the population of interest (e.g., evaluating putative environmental causes of cancer in humans; evaluating the action of drugs in children and pregnant women);
- Where clinical studies are underpowered with respect to the outcomes of interest (e.g., when assessing adverse reactions to drugs by means of studies designed to test the efficacy of the drug);

- Any question of external validity where clinical studies in the target population are limited or inconclusive;
- Assessing the effectiveness of a public health action or a social care intervention, where a thorough understanding of the relevant social mechanisms is important;
- Assessing the effectiveness of a medical device, where the mechanism of the device and its interaction with biological mechanisms may not be immediately obvious.

Some commentators have argued that one should disregard evidence of mechanisms, largely on the grounds that mechanistic reasoning has sometimes proved dangerous in the past. An infamous example concerns advice on baby sleeping position in order to prevent sudden infant death syndrome (Evans 2002, 13–14). On the basis of seemingly plausible mechanistic considerations, it was recommended that babies be put to sleep on their fronts, since putting a baby to sleep on its back seemed to increase the likelihood of sudden infant death caused by choking on vomit. However, comparative clinical studies later made clear that this advice had led to tens of thousands of avoidable cot deaths (Gilbert et al. 2005). There are several other examples of harmful or ineffective interventions recommended on the basis of mechanistic reasoning (Howick 2011, 154–157). As a result, it has been argued that relying on evidence of mechanisms can do more harm than good.

In many of these cases, however, the proposed evidence of mechanisms was not explicitly evaluated: often, there was little more than a psychologically compelling story about a mechanism (Clarke et al. 2014, 350). In such cases, making the evidence explicit and explicitly evaluating that evidence would have been enormously beneficial. Thus there is a difference between *mechanistic reasoning*, which in some cases is based on rather little evidence and can be problematic, and evaluating *mechanistic evidence*, which is almost always helpful. The case of anti-arrhythmic drugs may help to illustrate this distinction. Arguably, anti-arrhythmic drugs were recommended on the basis of ill-founded mechanistic reasoning (Howick 2011). The story goes as follows. After a heart attack, patients are at a higher risk of sudden death. Those patients are also more likely to experience arrhythmia. On the basis of some mechanistic reasoning, it was thought likely that there was a mechanism linking arrhythmia to heart attacks. Anti-arrhythmic drugs were, as a result, prescribed in an attempt to indirectly prevent heart attacks by directly preventing arrhythmia. It was later discovered on the basis of the Cardiac Arrhythmia Suppression Trial (CAST) that, unfortunately, the drugs led to an increase in mortality (Echt et al. 1991). See also Furberg (1983). However, at least in retrospect, it looks as though insufficient attention had been paid to mechanistic evidence. In particular, there was little reason to think that reducing arrhythmia was a good surrogate outcome for reducing mortality due to heart attacks. Indeed Holman (2017) argues that pharmaceutical company influence was largely responsible for that choice of surrogate outcome. In this case, properly considering the mechanistic evidence may have led to not recommending anti-arrhythmic drugs.

A critic of the use of evidence of mechanisms might respond that even when there exists good evidence of mechanisms, many biomedical processes are so complex that it remains difficult to establish causal claims on the basis of evidence

of mechanisms (Howick 2011, 136–143). For example, there was arguably some good mechanistic evidence in favour of the claim that dalcetrapib lowers the risk of developing coronary heart disease by increasing the ratio of HDL:LDL. However, a randomised controlled trial showed that risk of coronary heart disease was not significantly affected (Schwartz et al. 2012). A possible explanation for this failure was identified by Tardif et al. (2015), who identified two genetic subgroups of patients. While one subgroup appeared to benefit from dalcetrapib, the second genetic subgroup was harmed. Here, while further work was required to understand the mechanisms in play at the stage of the dalcetrapib clinical trial, it appears as if a credible conclusion has now been reached.

More generally, it is widely accepted that the complexity of biomedical processes presents a significant hurdle for establishing causal claims solely on the basis of evidence of mechanisms. But this is exactly why this book recommends explicitly evaluating evidence of mechanisms *alongside* evidence of correlation. Evidence of mechanisms is not sufficient for good clinical decision making—but neither is evidence of mere correlation.

References

- Clarke, B., & Russo, F. (2017). Mechanisms and biomedicine. In S. Glennan & P. Illari (Eds.), *The Routledge handbook of mechanisms and mechanical philosophy*. Chap. 24 (pp. 319–331). London: Routledge.
- Clarke, B., Gillies, D., Illari, P., Russo, F., & Williamson, J. (2014). Mechanisms and the evidence hierarchy. *Topoi*, 33(2), 339–360.
- Claveau, F. (2013). The independence condition in the variety-of-evidence thesis. *Philosophy of Science*, 80(1), 94–118.
- Clay, P. (2002). The abacavir hypersensitivity reaction: A review. *Clinical Therapeutics*, 24(10), 1502–14.
- Echt, D. S., Liebson, P. R., Mitchell, L. B., Peters, R. W., Obias-Manno, D., & Barker, A. H., et al. (1991). Mortality and morbidity in patients receiving encainide, flecainide, or placebo. *New England Journal of Medicine*, 324(12), 781–788.
- Evans, D. (2002). Database searches for qualitative research. *Journal of the Medical Library Association*, 90, 290–3.
- Furberg, C. (1983). Effect of antiarrhythmic drugs on mortality after myocardial infarction. *The American Journal of Cardiology*, 52, C32–C36.
- Gibbs, J. B. (2000). Mechanism-based target identification and drug discovery in cancer research. *Science*, 287(5460), 1969–1973.
- Gilbert, R., Salanti, G., Harden, M., & See, S. (2005). Infant sleeping position and the sudden infant death syndrome: systematic review of observational studies and historical review of recommendations from 1940 to 2002. *International Journal of Epidemiology*, 34(4), 874–887.
- Gillies, D. (2017a). Evidence of mechanism in the evaluation of streptomycin and thalidomide. *Studies in the History and Philosophy of Biological and Biomedical Sciences*, 66, 55–62.
- Gillies, D. (2017b). Mechanisms in medicine. *Axiomathes*, 27(6), 621–634.
- Holman, B. (2017). Philosophers on drugs. *Synthese*. <https://doi.org/10.1007/s11229-017-1642-2>.
- Howick, J. (2011). *The philosophy of evidence-based medicine*. Chichester: Wiley-Blackwell.
- Illari, P. M., & Williamson, J. (2012). What is a mechanism? Thinking about mechanisms across the sciences. *European Journal for Philosophy of Science*, 2, 119–135.

- Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1–25.
- Mallal, S., Nolan, D., Witt, C., Masel, G., Martin, A., Moore, C., et al. (2002). Association between presence of HLA-B*5701, HLA-DR7, and HLA-DQ3 and hypersensitivity to HIV-1 reverse-transcriptase inhibitor abacavir. *The Lancet*, 359(9308), 727–732.
- Pawson, R. (2003). Nothing as practical as a good theory. *Evaluation*, 9(4), 471–490.
- Rauch, A., Nolan, D., Martin, A., McKinnon, E., Almeida, C., & Mallal, S. (2006). Prospective genetic screening decreases the incidence of abacavir hypersensitivity reactions in the Western Australian HIV cohort study. *Clinical Infectious Diseases*, 43(1), 99–102.
- Russo, F., & Williamson, J. (2012). EnviroGenomarkers: The interplay between mechanisms and difference making in establishing causal claims. *Medicine Studies: International Journal for the History, Philosophy and Ethics of Medicine and Allied Sciences*, 3, 249–262.
- Salmon, W. C. (1998). *Causality and explanation*. Oxford: Oxford University Press.
- Schwartz, G. G., Olsson, A. G., Abt, M., Ballantyne, C. M., Barter, P. J., Brumm, J., et al. (2012). Effects of dalcetrapib in patients with a recent acute coronary syndrome. *New England Journal of Medicine*, 367(22), 2089–2099.
- Tardif, J.-C., Rhéaume, E., Perreault, L.-P. L., Grégoire, J. C., Zada, Y. F., Asselin, G., Provost, S., Barhdadi, A., Rhinds, D., L'Allier, P. L., Ibrahim, R., Upmanyu, R., Niesor, E. J., Benghozi, R., Suchankova, G., Laghrissi-Thode, F., Guertin, M.-C., Olsson, A. G., Mongrain, I., Schwartz, G. G., & Dubé, M.-P. (2015). Pharmacogenomic determinants of the cardiovascular effects of dalcetrapib. *Circulation: Genomic and Precision Medicine*, 8(2), 372–382.
- Wallmann, C., & Williamson, J. (2017). Four approaches to the reference class problem. In G. Hofer-Szabo & L. Wroński., editor, *Making it formally explicit: Probability, Causality, and Determinism*. Dordrecht: Springer.
- Wilde, M., & Parkkinen, V.-P. (2017). Extrapolation and the Russo–Williamson thesis. *Synthese*. <https://doi.org/10.1007/s11229-017-1573-y>.
- Williamson, J. (2018, in press). Establishing causal claims in medicine. *International Studies in the Philosophy of Science*. <http://blogs.kent.ac.uk/jonw/>.

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.

