Philos. Technol. (2011) 24:115–136
DOI 10.1007/s13347-010-0010-7

RESEARCH ARTICLE

Imaging Technology and the Philosophy of Causality

George Darby · Jon Williamson

Received: 8 October 2010 / Accepted: 6 December 2010 / Published online: 15 January 2011

Abstract Russo and Williamson (Int Stud Philos Sci 21(2):157–170, 2007) put forward the thesis that, at least in the health sciences, to establish the claim that C is a cause of E, one normally needs evidence of an underlying mechanism linking C and E as well as evidence that C makes a difference to E. This epistemological thesis poses a problem for most current analyses of causality which, in virtue of analysing causality in terms of just one of mechanisms or difference making, cannot account for the need for the other kind of evidence. Weber (Int Stud Philos Sci 23(2):277–295, 2009) has suggested to the contrary that Giere's probabilistic analysis of causality survives this criticism. In this paper, we look in detail at the case of medical imaging technology, which, we argue, supports the thesis of Russo and Williamson, and we respond to Weber's suggestion, arguing that Giere's account does not survive the criticism.

Keywords Causality · Causation · Difference making · Mechanism · Medical imaging

1 Introduction

Russo and Williamson (2007) put forward the thesis that, in the health sciences, in order to establish a causal claim C *is a cause of* E, one normally needs to establish two further claims: that C makes a difference to E and that there is

G. Darby (✉) · J. Williamson
Philosophy, SECL, University of Kent, Canterbury, CT2 7NF, UK
e-mail: G.A.Darby@kent.ac.uk

J. Williamson
e-mail: J.Williamson@kent.ac.uk

some mechanism linking *C* and *E* that explains this difference making. This epistemological thesis, which is referred to as RWT and which is outlined in more detail in Section 2, has been the object of some controversy (Weber 2009; Clarke 2010; Howick 2010; Illari 2011; Broadbent 2010; Gillies 2011; Russo and Williamson 2011), partly because of the metaphysical conclusions that it warrants: If both kinds of evidence are required, then no purely difference-making (respectively mechanistic) analysis of causality is viable because no such analysis can account for the need for mechanistic (respectively difference-making) evidence. In this paper, we shall show how imaging technology supports RWT: As we shall see in Sections 3 and 4, medical imaging technology is often used to provide evidence of the existence of a linking mechanism, and this evidence acts independently of difference-making evidence when establishing a causal claim. Finally, in Section 5, we shall rebut one objection, the suggestion of Weber (2009) that there is a difference-making analysis that successfully accounts for the need for mechanistic evidence, namely Giere's probabilistic analysis of causality.

2 The Epistemological Thesis

The epistemological thesis at the heart of this paper is the thesis that, in the health sciences, in order to establish a causal claim *C is a cause of E*, one normally needs to establish two further claims.

First, it should normally be the case that there is some mechanism linking *C* and *E* that can explain *E* in terms of *C*. This is for the following reason. Causal relationships are invoked as explainers. Thus, when asked 'Why did *E* occur?', one can answer 'Because *C* occurred and *C* caused *E*.' But it is widely acknowledged that in order to explain some phenomenon, one needs to point to the mechanism for that phenomenon—the relevant constitution and organisation of reality that is responsible for the phenomenon in question (see, e.g. Machamer et al. 2000). Hence, for causal relationships to be explanatory, they had better accord with underlying mechanistic explanations. This accounts for why evidence of the existence of a linking mechanism is required in order to establish a causal claim. Note that this is evidence of the *existence* of a mechanism that can play the explanatory role—one does not need to know the details of the mechanism itself for the causal claim to be explanatory.[1]

[1]Of course, the more of the mechanism that is known, the better the explanation that can be offered. But the existence of the mechanism is sufficient for the causal claim to play an explanatory role.

The details of the mechanism between *C* and *E* and those of surrounding mechanisms can be very useful in other ways in establishing a causal claim, e.g. they can shed light on the nature of possible confounders and on whether a causal claim can be extrapolated from animals to humans or from one time to another (see Section 5). Although these considerations are important, they are orthogonal to the particular role of mechanisms under consideration here.

Second, it should normally be the case that C makes a difference to E. This is for the following reason: Causal relationships are invoked to draw predictions (to predict E on the basis of C, or to diagnose C on evidence E) and to decide how to intervene on the world in order to achieve one's goals (e.g. to eliminate C in order to avoid outcome E). But prediction and control are only possible if cause and effect are evidence for each other and if changing the cause changes the effect—i.e. if the cause makes the appropriate sort of difference to the effect.

In sum, evidence of both a mechanism and of difference making is normally required to establish a causal claim.[2] This epistemological thesis is supported by practice in the health sciences, both in terms of methodological work in the health sciences and in terms of particular cases of causal inference (Russo and Williamson 2007). From the methodological point of view, note in particular that some of Bradford Hill's well-entrenched principles for establishing a causal claim require information about mechanisms and others require information about difference making. From the point of view of cases, consider two well-known examples. That difference-making evidence is not on its own normally sufficient to establish a causal claim is witnessed by the case of Semmelweis. In 1833, Semmelweis collected statistics in the Vienna maternity hospital showing that hand-washing makes a difference to the incidence of puerperal fever. But the causal claim was not accepted until, later on in the nineteenth century, the germ theory of disease, and hence an underlying mechanism, was accepted. On the other hand, that mechanistic evidence is not on its own normally sufficient for establishing a causal claim is witnessed by the case of Snow. In 1849, Snow identified the mechanism for cholera (a living organism that contaminates drinking water by proximity to sewage is responsible for cholera). But the causal claim was not accepted until he found in 1854 that the incidence of cholera was dependent on the source of water.

The epistemological thesis allows for exceptions though. There is no suggestion that *every* causal relationship charts an underlying mechanism: Indeed, cases involving causation between absences or involving double prevention are arguably cases in which there is causality without a corresponding mechanism (see, e.g. Williamson 2011). Consider, for example, the causal claim that my missing the flight at Gatwick caused the lack of my talk in Australia. Since the cause and effect are non-entities, there can hardly be a physical mechanism linking them. One might suggest that one substitute what was actually present for the absences, but then the relevance relationship is lost: Nothing about the

[2] As Illari (unpublished manuscript) notes, RWT points to a distinction between the *objects* of evidence (a difference-making relation, a mechanism) rather than between *items* of evidence: It is possible that the same item of evidence could be evidence both of a difference-making relation and of an underlying mechanism, in which case a single item of evidence could be sufficient to establish a causal claim.

boarding procedure at Gatwick can be said to have caused the extended coffee break at the conference in Australia.

Moreover, there is no suggestion that *every* causal relationship charts a difference-making relationship: Indeed, cases where the effect is overdetermined (i.e. would happen anyway) are arguably cases in which there is causality without a corresponding difference-making relationship (see, e.g. Williamson 2009). Such a case arises when $C_1, \ldots, C_n$ is a *partition* of causes of E, each operating by a distinct mechanism, such that none makes a difference to E. For example, state E of a particle is obtained by transition from state C_1 or C_2. If C_1 and C_2 are mutually exclusive and exhaustive, neither may make a difference to the incidence of E. One might suggest that were C_1 and C_2 *not* mutually exclusive and exhaustive then there would be difference making. But such a counterfactual is very hard to evaluate, and in any case, knowledge of the underlying physical mechanism provides the only grounds for making this claim. (The much-discussed failures of *faithfulness* offer other examples of causation without difference making, e.g. in some cases, difference making is intransitive; in other cases, difference making along a path from C to E can be cancelled out by a second path along which an equal and opposite difference is made; Simpson's paradox provides further examples in which difference-making relationships can be lost.)

One might restate the epistemological thesis as follows, then. In order to establish a claim of the form C is a cause of E, one needs two of the three following sorts of evidence:

1. Evidence that C makes a difference to E
2. Evidence that there is some physical mechanism linking C and E
3. Evidence that (1) or (2) is not appropriate in this case (e.g. evidence that C or E is an absence or evidence of overdetermination)

The epistemological thesis is of interest to philosophers because, if true, it has important consequences concerning the nature of causality. Currently, the principal accounts of causality interpret the causal relation either solely as a mechanistic relation, or solely as a difference-making relation, or as some pluralistic combination of the two (mechanistic in some cases and difference making in others). But all such accounts appear to fail if the epistemological thesis is true. An analysis of causality solely in terms of mechanisms cannot explain why, when the mechanism is known, evidence of difference making is also required. On the other hand, an analysis of causality solely in terms of difference making cannot explain why, when good difference-making evidence is available, evidence of a mechanism is also required. Finally, a pluralistic combination of the two inherits one of the above problems for any particular causal claim, according to whether that claim is given a mechanistic interpretation or given a difference-making interpretation. (Note that a conjunctive analysis, according to which a causal claim is true iff there is *both* a mechanism *and* difference making, is prone to both the counterexamples to the necessity of mechanisms—e.g. those stemming from absences—as well as to the counterexamples to the necessity of difference making—e.g. those stemming from

underdetermination.) Russo and Williamson (2007) argued that the *epistemic theory of causality* (according to which causal claims should be interpreted in terms of rational beliefs rather than in terms of worldly mechanistic and/or difference-making relations) is not prone to these problems.

Having explicated the epistemological thesis (RWT) of Russo and Williamson (2007), we will shortly see that imaging technology provides further support for RWT. First is a general introduction to imaging technology.

3 Imaging Technology

A philosophical discussion could be had about what is meant by 'imaging technology'. Without going into too much detail, it is worth thinking about the kind of thing that we have in mind. For the sake of the discussion and since the Russo–Williamson thesis is intended in the first instance as a thesis about causality in the biomedical sciences, it is primarily *medical* imaging that we are interested in. What naturally springs to mind is the striking image of the brain scan produced by functional MRI, which appears to show in real time the activation of different parts of the brain during cognitive processes. But in their different ways, ordinary optical microscopes, cameras, even spectacles, might also count as imaging technologies. What these have in common is that the image we receive is mediated by technology and is in most cases a different *kind* of image from that obtainable by the naked eye.

One way of drawing a line might be to co-opt, say, van Fraassen's account of the observable–unobservable distinction, according to which something is observable iff, were it present to us, we would observe it with our usual visual faculties (van Fraassen 1980, p. 16). Then we could take imaging technology to be that which makes the unobservable accessible to normal vision. This is not just a matter of scale: The minute variations in blood flow picked up by functional MRI are not observable in this sense, nor are the differences in density picked up by X-rays, even on a large scale.

Imaging technologies can be categorised in various ways. With some exceptions, such as ultrasound, the majority exploit electromagnetism in some way. In this section, we therefore briefly review some of the key properties of electromagnetism that make the technologies possible and then describe three basic kinds of imaging that rely on it. First we present the different kinds of microscopy, including the sophisticated biotechnological techniques that make it possible to image the individual components of cellular mechanisms. Second is the familiar X-ray, including its modern incarnation in the computerised tomography (CT) scan. Microscopy and X-ray technology have in common that they are based on electromagnetic radiation passing through a body. The third type of imaging technology is the MRI scan, which is of a somewhat different kind. For details of the physics underlying these technologies, see Hendee and Ritenour (2002); Cho et al. (1993); Barrett and Swindell (1981). For the history, see Webb (1990); Kelves (1997); Doby and Alker (1997).

3.1 Electromagnetism

Most imaging technologies, in their various ways, utilise aspects of *electromagnetism*. The aspects of electromagnetism that are important here are illustrated in Table 1.

At one end of the spectrum are the radio waves of large wavelength. At the other are the gamma rays of very short wavelength. As well as wavelength λ, the different types of electromagnetism can be categorised by frequency ν. Take the distance travelled by a wave in a second and divide by frequency, the number of oscillations per second. Then we have the wavelength—the length of one oscillation. So for electromagnetic waves travelling at the speed of light c, the frequency and wavelength are related by $\nu = c/\lambda$. The *energy* E of a photon (particle manifestation of electromagnetic radiation) is related to frequency via $E = h\nu$, where h is Planck's constant. The important point here is that the different portions of the electromagnetic spectrum can be categorised by either wavelength or frequency or energy (for example, X-rays have higher energy, higher frequency and shorter wavelength than visible light). The different wavelengths, frequencies and energies are important in the various imaging technologies.

One way in which these parameters are important is in determining the *attenuation* of a species of electromagnetism by a given material. Whether or not a wave (a radio wave, say) will pass through a gap depends on the relationship between the size of the gap and the length of the wave because this will determine whether *diffraction* occurs (this is why you can sometimes receive a good radio signal on one frequency but not another). The energy of a photon determines whether it will interact with a material to produce the photoelectric effect, whereby the photon is absorbed and an electron emitted (the photoelectric effect is also important in the X-ray detectors used in CT scans—see below), or whether Compton scattering will occur, whereby photons change direction and frequency in collision with an electron in an atom. Various characteristics of a given material interact in complicated ways to produce an *attenuation coefficient* that is specific to that material *and a given frequency* on the electromagnetic spectrum.

That the relationship between attenuation and position on the electromagnetic spectrum is not simple is illustrated by the fact that the Earth's atmosphere absorbs all of the gamma and X-rays from the Sun, much of the

Table 1 The electromagnetic spectrum

Wavelength (m)	10^{3}	10^{-2}	10^{-5}	10^{-6}	10^{-8}	10^{-10}	10^{-12}
Frequency (Hz)	10^{4}	10^{8}	10^{12}	10^{15}	10^{16}	10^{18}	10^{20}
Type	Radio waves	Microwaves	Infrared	Visible light ROYGBIV	Ultraviolet	X-rays	Gamma rays
Energy (eV)	10^{-12}	10^{-7}	10^{-1}	1	10^{2}	10^{4}	10^{6}

Energies are in electronvolts. Values to within an order of magnitude

ultraviolet light, little of the visible light, most of the infrared and none of the radio waves, except those with very long wavelength (absorption is not therefore simply a matter of energy—the high-energy gamma and X-rays are blocked, whereas the low-energy radio waves get through). The *technological* wizardry involves the manipulation and detection of electromagnetism using these various factors.

3.2 Microscopy

No doubt the basic idea of the optical microscope is very familiar: Light reflected from (or passing through) a sample can be magnified many times to make it visible in an eye-piece. The basic mechanism of magnification is refraction—light beams will change direction on passing from a material with one refractive index to another; by altering the angle of incidence on the boundary, the angle of refraction can be altered; a lens is shaped so that the angle of incidence of beams of light spreading out from a microscopic source is such that they are brought back together. Modern microscopic technology is of course now very sophisticated, including optical techniques for improving contrast without physically staining the sample, stereomicroscopes for three-dimensional viewing, replacement of the eyepiece with a camera for connection to a computer, scanning microscopes, electron microscopes, scanning electron microscopes and so on. Here, though, we would like to focus on something that one might not think of as being an example of imaging technology—the use of *bio*technology as an aid to microscopy.

Since the structures in a cell are such that roughly the same amount of light passes through each of them and with colours equally represented, contrast in the image can be limited. This is where *staining* comes in, with the use of *contrast agents*, which will stain one particular part of the sample (a particular organelle, for example), filtering out all but one wavelength of light. A simple example would be the addition of iodine to make starch show up blue.

Modern biotechnology offers far more sophisticated variants on this idea, for example, in the use of *immunochemistry*. Immunohistochemistry is a tool used for visualising the location of proteins within a tissue sample, and immunocytochemistry is the analogous procedure on cells in culture isolated from the rest of the tissue such as the extracellular matrix.

For example, in the investigation of the mechanism of inflammatory response, it is already known that transcription factor NF-κB promotes the production of particular proteins involved in inflammation (a transcription factor binds to DNA and affects the extent to which different regions are transcribed). Since transcription takes place in the nucleus, the effect will only be found if the NF-κB is located in the nucleus. A certain inflammatory cytokine attaches to a receptor on the outside of the cell membrane and promotes the synthesis of NF-κB and its entry into the nucleus. The question now is whether a particular treatment inhibits the action of the cytokine specifically in regard to its effects on NF-κB. So a sample of cells is treated

with the cytokine and then with the cytokine plus the putative inhibitor, and immunocytochemistry is used to highlight the location of NF-κB within the cells.

The technique is this: Having decided which protein you are interested in, obtain an antibody that is specific to it, that is, that will bind to the protein in question and this protein only. In the simplest case, the antibody may have a florescent enzyme on it, which can be observed under the microscope. In practice, it is necessary to use a second antibody that binds to the first, but the principle is the same. The florescent enzyme must be exposed to electromagnetism of the right wavelength in order to view the results: When exposed to light of this wavelength, the florescent compound emits light of a *different* wavelength. The light used to *cause* the florescence can then be filtered out, and the remaining florescent light gives a clear image of the location of the desired protein in the sample (this technology comes from *immuno*chemistry because antibodies are produced by white blood cells. Bacteria contain proteins that are not found in the body, so an antibody binds specifically to a sequence of amino acids that forms part of one of these proteins so that it can serve as a marker in the immune response).

Figure 1 shows endothelial cells (the cells that line blood vessel walls) viewed under a fluorescent microscope. An antibody has been selected that binds to NF-κB and a second antibody applied that binds to the first. This second antibody has been tagged with florescent compounds that show up green when correctly stimulated. A blue stain has been used that binds to DNA; this identifies the nucleus, where the bulk of a cell's DNA is contained. A red stain that binds to actin, in the cytoskeleton, highlights the structure of the cells to complete the picture (by similar means, other organelles could be picked out if required). Figure 1a shows a control (NF-κB present in the cytoplasm but not in the nucleus): Fig. 1b shows cells after treatment with the inflammatory cytokine (NF-κB now present in the nucleus, where it will act to increase the production of inflammatory proteins); Fig. 1c is on treatment with the cytokine plus the inhibitor. This is an example of immunocytochemistry

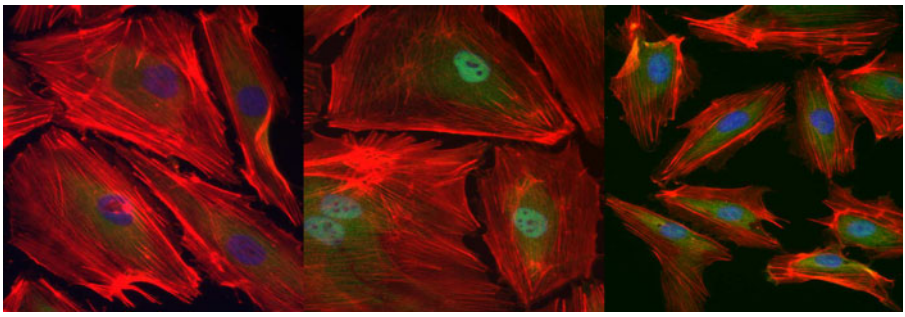

Fig. 1 Immunocytochemistry used to locate components of a mechanism

used to show where a component of a mechanism is acting and thereby to figure out how to inhibit the operation of the mechanism: In this case, the NF-κB is to be found in the cytoplasm and *not* in the nucleus after application of what is now known to be an effective inhibitor.

This kind of technology is used to establish *mechanisms* in medical research, by locating the various component molecules involved in the mechanism.

3.3 X-rays and CT

X-rays occupy the section of the electromagnetic spectrum from around 10^{-11} to 10^{-8} m (the distinction between X-rays and gamma rays is not firm, especially historically speaking). Conceptually, the X-ray is not so far from the ordinary mechanisms of vision. Electromagnetic radiation is passed through an object, some is absorbed by denser tissues and the remainder reaches a detector. In this case, since the radiation is along the spectrum from visible light, the detector cannot be the human eye. Photographic film was used early on; modern electronic detectors now do the job, not unlike a digital camera, which opens up the possibility of manipulating and combining the X-ray images.

It is the distance of X-rays from visible light on the spectrum that makes them medically useful, since they have different attenuation coefficients for a given material. For the tissues found in the human body, the attenuation is far less than for visible light, which allows them to pass through. The attenuation is different enough for the different tissues, on the other hand, that the X-rays are blocked by structures such as bones, so that images can be taken. (Their high energy is also responsible for X-rays' being a type of *ionising* radiation, which makes them less than perfect for medical use because of their own effects on health.)

As it is, however, this inherits one of the drawbacks of normal vision: it cannot be used to look inside a body at a hidden part, and if a series of distinct structures are stacked between the source and the detector, then they are combined in the image, like hands in a shadow puppet show. What we would like is a *slice* through the body, so as to see what is happening at a given depth (and ultimately the slices might be put together into a three-dimensional model). This is the aim of *tomography*. which has a history long before the computerised version. The basic, ingenious, idea is along these lines:[3]

Figure 2a shows an X-ray of a complex body consisting of a cylinder on top of a square. We did not want an image of the square—we wanted a section through the cylinder—but that is what we get because the whole lot is projected down onto the screen. To solve this, first allow the source and detection screen to move around in planes parallel to one another on either side of the subject.

[3]See Webb (op. cit.) for the origin and development of this idea.

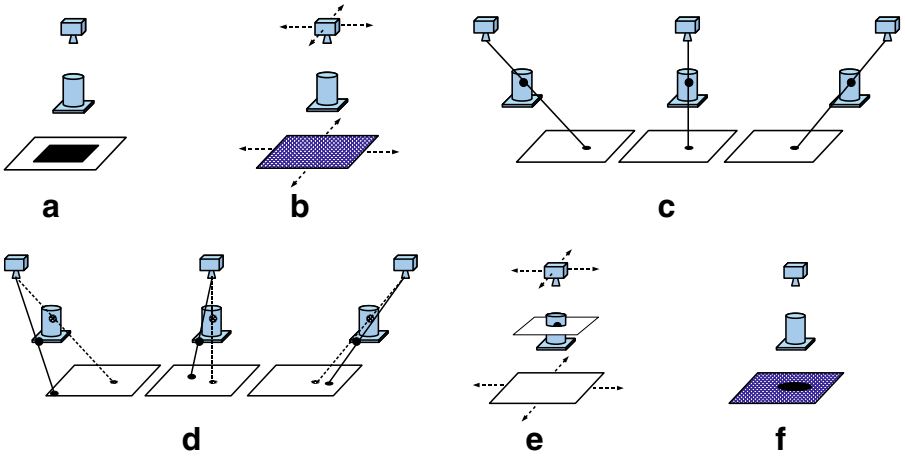

Fig. 2 An X-ray of a section through a complex body

If X-rays are emitted throughout this movement, then each point in the source will be projected onto different points on the screen, and the image will be a low-intensity smear across the screen (Fig. 2b). But now constrain the relative motion of source and screen so that one point in the subject is always projected onto the same point on the screen throughout the movement (Fig. 2c). When the screen and source are moved, this point will thus be sharply represented at high intensity on the image. Now consider some other point in the subject. This point will in general *not* be projected onto the screen at the same point in all positions (Fig. 2d). Its image will therefore be spread out at lower intensity across the screen. Now ask: which of the *other* points on the body *do* have this property (that they are projected onto the same point on the screen throughout the movement)? The answer is: those that lie in the plane that is parallel to the planes of movement and that contains our first point (Fig. 2e). The points on that plane will come out sharply on the image, and the rest will be blurred (Fig. 2f). Here, then, is a way of producing an X-ray of a section through a body.

Modern (X-ray) CT scanners replace the photographic film with detectors that produce a signal in response to arriving photons and employ enormous computing power to process the signals into a visual reconstruction.

As before, a source sends out X-rays, but this time, only a narrow beam ('pencil beam') is required. A typical detector might use a scintillation crystal (perhaps sodium iodide laced with thalium) which emits (visible) photons when struck by an X-ray photon; these are then sent to a photomultiplier which will then emit electrons to produce a current which feeds into the processing system (the emission of electrons here in response to the arrival of a photon is another application of the photoelectric effect). The photon beam arriving at the detector has of course been affected by passing through the imaging subject. Precisely *how* it has been affected depends on the attenuation

properties of the structures that make up the subject. Since the beam is only one-dimensional, all of this information is bound up into a single number, the effect of integrating the relevant coefficients along the path of the beam through the subject. By moving the beam across the sample, a one-dimensional map of the attenuation along one side is found. By rotating, the image along a second side is found and so on. It is to work back from these data to the reconstruction of the pattern of materials, with their different attenuation properties, in the cross section of the subject as a whole that sophisticated algorithms and enormous computing power are needed. More accurate scans are obtainable by scanning from more directions, and considerations of efficiency and practicality determine the optimum arrangement of source, detectors and their relative movement (see Cho et al. p. 160).

3.4 MRI

Different substances in the body are composed of varying amounts of different atoms—for example, lipids are composed primarily of carbon, hydrogen and oxygen, whereas proteins are composed primarily of carbon, hydrogen, oxygen and nitrogen. We might build up a picture of the composition and structure of a bodily structure, therefore, if there were some way of asking 'How much hydrogen is there at this point?' and 'How much carbon is there at that point?'. In fact, since different tissues contain different amounts of water—blood contains more water than fat; fat contains more water than muscle—it would suffice to know how much hydrogen there was at each point. This is the basis on which MRI is used to construct images of the body.

The basic mechanism is again electromagnetic, but this time what is important is the way in which the motion of charged particles interacts with a magnetic field. Moving electric charges produce a magnetic field, and a charge moving *in* a field feels a force. Even an elementary particle that was apparently not moving would produce a magnetic field because it can be thought of as possessing (intrinsic) angular momentum, *spin*. All magnetic fields have a direction, including that associated with the *proton*, the particle that we are most interested in (because the proton is the nucleus of the hydrogen atom).

When placed in a strong magnetic field (with a strength of, say, 1T), protons and also the nuclei of many atoms besides hydrogen do two things. First, they align their own little magnetic fields (and spin axes) with the main field. Second, the spin axis *precesses* (wobbles—as does the axis of rotation of a spinning top and of many planets), and it does this with a certain frequency. This frequency is determined by the type of particle and the strength of the magnetic field (for hydrogen 1T field, this is about 40 MHz).

This frequency is the key to imaging: The energy of these nuclei can be increased by means of radio waves (since 40 MHz just happens to be in the radio portion of the electromagnetic spectrum), so long as the frequency of the radio waves is equal to the frequency of precession—the *resonant* frequency (this is like pushing a child on a swing—you have to supply energy at the correct frequency). When the radio transmission is turned off, the nuclei

lose energy, and electromagnetic radiation is emitted (and can be detected). By switching the radio waves on and off, the response of the nuclei can be measured.

So far, then, this is a way of manipulating a desired type of atom in a sample, since the frequency of radio waves and the magnetic field that determines the resonant frequency can be precisely controlled. What is so far missing is a way of pinpointing the precise location of the atoms that are being manipulated. This is the point of introducing *gradients* into the magnetic field: If you know precisely how the magnetic field varies from a little over 1T on one side of the scanner to a little under 1T on the other, then you know precisely which slice you are interrogating with 40 MHz radio waves (since only those atoms where the field strength is precisely 1T will be excited at 40 MHz). By varying the frequency of the radio waves, you can interrogate a different slice. By varying the pattern of gradients in the magnetic field, you can explore different slices to build up a picture of the distribution of hydrogen and therefore of water in the sample. This, in essence, is how MRI works.

This sketch is of course crude, and many refinements are possible. Probably the most prominent example is blood oxygen level-dependent ***functional*** MRI (BOLD fMRI), wherein different types of brain activity are associated with cognitive or behavioural functions. As well as measuring blood flow, this exploits the fact that the magnetic properties of haemoglobin vary depending on whether it is oxygenated or deoxygenated to determine whether an increase in blood flow is accompanied by oxygen usage. The methodology is controversial, especially when data about brain activity are tied to particular cognitive functions—it is not quite as simple as watching a certain region of the brain light up when thinking of chocolate; see especially Bechtel and Richardson (2010). However, this technology does at least offer some insight into neural mechanisms (see Raichle and Mintun 2006).

3.5 Summary

What is imaged is different in each case: in microscopy the capacity to reflect or transmit light of different wavelengths, in immunochemistry the presence and location of an antigen (via the location of an antibody, via microscopy), in CT and X-ray images the attenuation coefficient and in MRI (typically) the concentration of hydrogen nuclei (or less directly the water content). But in all cases, the technologies are used to generate visual images of the mechanisms and their components.

4 The Role of Imaging in Establishing Causal Claims

In this section, we describe case studies which support the RWT: cases from current research in the health sciences in which the identification of a mechanism with the aid of the various types of imaging technology is

an important component of the establishment of a claim of causality. (We take it as uncontroversial that evidence of difference making is required to establish a causal claim in the health sciences: This is clear, for example, from developments such as the evidence hierarchy of the evidence-based medicine (EBM) movement.)

Obviously in many cases, the further investigation of mechanisms is important irrespective of whether causality has been established. But here we are specifically interested in the use of mechanisms to help establish causal claims. First, the Russo–Williamson thesis shows up in scientists' own reports of their work: Some of the cases below are ones in which the claim of causality is made on the basis of mechanistic evidence as well as difference-making evidence, where it could not be made just on the basis of difference-making evidence. Second, the thesis shows up *implicitly* in the motivations offered for investigation of mechanisms: In some of the cases, we can see causality being hypothesised on the basis of mechanistic evidence gathered from just a few individuals. Such a small sample size would not be sufficient to establish probabilistic difference making with any certainty, and in any case, different methods would be used to support a claim of difference making. Thus, mechanistic evidence is acting independently of difference-making evidence in supporting causal claims.

The challenge is to determine whether mechanistic evidence was considered key to establishing causality, or a desirable addition. In the published results, it is not always made explicit that *causality* was the target at all. This is amplified by the fact that, as noted in Russo and Williamson (2007, p. 1), it is sometimes hard to disentangle the various uses of 'cause' versus 'covary', 'association', 'risk factor' and the like. Notice also that even when 'cause' and 'effect' are used, it is not straightforward to infer the precise intention. For example, 'effect' can be ambiguous between an effect in our sense and a mere phenomenon (see, for example, the end of the abstract of Roberts and Garavan (2010)). It is therefore sometimes necessary to extrapolate information about the intention behind a study. This is to be expected given the vexed nature of the concept of causation in the health sciences, as witness the attempts at clarification found in the literature of the sciences themselves referenced in Russo and Williamson (2007, p. 2). Nevertheless, the following sampling does indicate that mechanistic evidence is required in addition to difference-making evidence when making claims of causality. (Note of course that it is besides the point whether these claims of causality are *true* or not—no doubt they are controversial. What matters is how they are established. Given the controversy over what is actually shown in fMRI, for example, we ought to be very cautious about claims of causality for the mechanisms that are investigated by it. But that does not change the fact that the imaging results are used to establish a causal link between the stimuli and behaviour that they accompany.)

In Section 4.1, we look at cases where imaging is used in establishing long-term effects of recreational drugs. In Section 4.2, we look at the case of the alleged benefits of alcohol for the heart. In Section 4.3, we note that recent

work by Brendan Clarke on the RWT once again highlights the importance of imaging technology.

As it happens, we might note in passing that imaging technology played a role in the ultimate acceptance of the mechanism, an essential piece of evidence in establishing causality, in the Semmelweis case too: Microscopes had been used to observe bacteria (but not viruses) for around 100 years by the 1870s, when Robert Koch used microscopy, together with staining techniques that he had developed, in the identification of the anthrax bacterium. This, together with the work of Pasteur and others on microorganisms in brewing, which also made use of microscopy, was among the factors that led to the acceptance of the germ theory of disease.

4.1 Case 1: the Effects of Recreational Drug Use

There has been fairly high-profile discussion of the long-term psychiatric effects in regular users of recreational drugs, such as the dangers associated with, and so the proper legal classification of, marijuana. Bound up with these policy discussions is the question of whether a causal link, from marijuana use to the psychiatric features in question, can really be established. This prompts studies including psychiatric evaluations and also MRI imaging of the brains of users. What is interesting from our point of view is this: If you wanted to establish just difference making, then psychiatric evaluation would be enough—establish the right kind of statistical relationship between the relevant behaviour and the relevant drug use and the case is made ('Behaviour' here is meant broadly, to include performance on memory tasks and the like). But that is not what happens. With the statistical relationship fairly well established (anyone who has smoked a joint could have told you about certain difference making relationships), there is still unwillingness to accept a causal link. What really *does* impress at this point—what seals the case for *causality*—are these MRI studies showing effects such as brain shrinkage, altered activation in response to stimuli etc. The pursuit of mechanistic evidence in this case seems to us to demonstrate that it is required to establish causality.

In many of the studies into the health effects of recreational drugs, evidence may roughly be divided into mechanistic evidence from fMRI scans, in which the brain mechanisms underlying behaviour can be observed (modulo the caveats mentioned earlier) and difference-making evidence gathered by observing the behaviour itself in various cognitive tasks.

Roberts and Garavan (2010) studied 20 subjects for performance on cognitive tasks and monitored their brain activity using fMRI during the tasks. In fact, no performance effects were evident (the authors note that in fact the absence of performance effects can be helpful in avoiding confounding factors such as frustration), but there *were* differences in brain activity between users and non-users. The absence of evidence of covariation of drug use and performance deficits does not concern the authors, since that covariation is *already* taken to be plausible—the purpose of this study is to identify the mechanism by which covariation happens. Gathering mechanistic data is indeed taken

to support a causal claim, although the authors do not even commit on the *direction* of causation. Whichever way it does go, the mechanistic evidence is taken to support the relevant causal claim which was not already established by statistical evidence.

In Duchene et al. (2010), a report of vascular events (myocardial infarction and stroke) in one individual following cannabis use, part of the background is that there is *not* an accepted causal link between the drug use and the events in question. Of course this is not statistical data constituting difference-making evidence because there is only one individual in the study. Nevertheless, it is considered to make plausible a claim of causality. Imaging technology (in this case angiography, a technique for imaging the vascular system based on a technology such as X-ray CT together with markers injected into the blood stream) provided negative evidence about the mechanism, by ruling out the more usual causes of myocardial infarction and stroke. Russo and Williamson (2011) present further evidence that the attribution of causality in single cases (autopsies, in that study) supports the RWT.

In van Hell et al. (2010), the authors note in the conclusion that theirs is the first study to show 'an effect of chronic cannabis use on reward processing in humans'. Between putative cause—cannabis smoking—and effect—altered behaviour—there is a complex mechanism to be discovered before cannabis use can be said to cause specific behavioural changes. During reward tasks, blood flow in various regions of the brain implicated in reward processing was different among cannabis users, regular smokers and non-smokers, and there are plausible possibilities for further details of the mechanism, such as the effects of cannabis and nicotine on transmission of neurotransmitters and sensitivity of receptors (p. 160). What we find interesting about this is that, again, two types of evidence—behavioural and fMRI—were gathered. Suppose that the goal were simply to show difference making, with the thought that that would be enough to establish causality; then there would be no need for the MRI study. But the authors clearly think that the MRI is necessary and the RWT concurs: It is necessary because the fMRI study illuminates the mechanism. Theirs is the first study to show *causality* because theirs is the first to provide this mechanistic evidence. (One might also find the way this is presented in the abstract telling: 'Our findings imply that chronic cannabis use as well as nicotine, may cause an altered brain response to rewarding stimuli'. There was already *behavioural* data, and unless the authors think that there is some *other* plausible explanation for the behavioural data (not, however, replicated in this study), which there is not, one might think that that was enough to show that there must be an altered brain response. But it is only having done the mechanistic study that causation is implied.)

4.2 Case 2: the Effects of Alcohol on Heart Disease

There is perhaps even more high-profile discussion of the miraculous health benefits of a daily glass or two of red wine. The growing consensus *seems* to be that moderate intake has some positive effect. Anecdotally, this starts

with the observation that the French, who have a diet that is relatively high in saturated fat, nevertheless have a relatively low incidence of heart disease. The explanation for this is said to be that they also regularly drink moderate amounts of red wine. Of course, there are many possible confounding factors such as lifestyle, and although early studies appeared to show that there was a definite benefit, this became controversial—see, for example, Davidson (1989, p. 436). Recent studies, however, have re-established the epidemiological evidence and moreover appear to show that it is alcohol itself, rather than wine specifically, that is beneficial (Rimm et al. 1996). For our purposes, perhaps the most interesting opinion is that of Steinberg et al. (1991), who, despite acknowledging the epidemiological evidence, argue that the case has not at all been made for recommending even moderate alcohol consumption on the basis of its apparent benefits, as long as the protective *mechanism* is unclear.

Given the RWT, what one would expect now are investigations of the possible mechanisms underlying this effect (if it *is* an example of cause and effect). Of course, statistical studies and in particular meta-analyses are still important, and so we would not expect them to disappear. But the mechanisms would become important, not just because of the intrinsic utility of intervention but because they are key to the claim of causality. And indeed they do.

Demirovic et al. (1993, p. 2787) note the inverse association between moderate alcohol assumption and coronary heart disease, but stop short of claiming a causal link. They investigated a possible mechanism, as follows: Given this mechanism, one would *also* expect an increase in *carotid* atherosclerosis. So they measured artery wall thickness using ultrasound in a population. As it happens, their results were negative, but this still goes to show that mechanistic evidence, arrived at by the (albeit indirect) use of imaging technology, would support a currently contentious claim of causality in a way that the currently promising difference-making evidence does not.

Of course, we should be wary of over-interpreting the precise words that scientists use, but the RWT appears to be mirrored in the approach of Lakshman et al. (2010, p. 113). The authors note the epidemiological evidence: There is a U- or J-shaped curve, with incidence of coronary artery disease (CAD) decreasing from abstinence to moderate alcohol consumption, then increasing as alcohol use becomes heavier. They then say that '[T]he possibility that lighter alcohol use protects against CAD is supported by plausible hypothetical mechanisms' and go on to list a number of studies, using various types of imaging technology, in support of the plausibility of these mechanisms. For example, Vliegenthart et al. (2004) investigated the onset of atherosclerosis, a key factor in coronary heart disease, in individuals with varying levels of alcohol consumption. Using a form of CT scanning to determine the level of calcification of the arteries, they showed that in individuals who had so far no symptoms of heart disease there was nevertheless a variation in degree of atherosclerosis with alcohol consumption similar to that between alcohol and heart disease itself. Thus, the imaging supports the claim of causality by providing evidence of a highly plausible mechanism. The investigation continues

at the level of basic science and the investigation of cellular mechanisms, with the attendant use of biological imaging technologies.

4.3 Case 3: Viruses as Causes of Cancer

Brendan Clarke documents in detail the work that led to the identification of the Epstein–Barr virus (EBV) as cause of Burkitt's lymphoma and the viral causes of cervical cancer, arguing that the two cases ultimately support a version of the RWT (Clarke 2010). For our purposes, two things are salient.

First, imaging technology was a key to establishing causality here. For example, Clarke (pp. 26ff) traces the search for the cause of EBV from speculation to observation of virus-like particles (p. 41), discovery and description of the virus itself (p. 46) and the cementing of its causal role (pp. 47–48), during which a wide variety of techniques were used. For example, the immunofluorescent techniques described above were used to establish the effects of viral antigens on the lymphoma, and this combined importantly with the use of electron microscopy to observe the viruses (or virus-like particles) themselves in similar situations (p. 47). In general, this case is one in which electron microscopy had to play an important role because the virus was not the kind of thing that could be seen with the naked eye, nor even with an optical microscope: The Epstein–Barr virus is in the region of 10^{-7} m across, and the resolving power of an optical microscope barely goes down that order of magnitude.

Second, both mechanistic and difference-making evidence were key in demonstrating that EBV really was the cause of Burkitt's lymphoma, in the face of obstacles, competing hypotheses and so on (pp. 86–88). This case was more complicated than simply having a putative C and E and needing a mechanism (perhaps investigated by imaging technology) to demonstrate that causality was really present, since the cause had not been properly identified. Nevertheless, Clarke demonstrates that mechanistic evidence was tied up in the search for the virus. In fact, imaging technology provided the mechanistic evidence *before* difference making was established through a large-scale epidemiological study—a respect in which this differs from other examples discussed above. Having established the presence of EBV within the mechanism for Burkitt's lymphoma and having proved that levels of antibodies to the Epstein–Barr virus (an indicator of exposure to the virus itself) are *correlated* with development of Burkitt's lymphoma, the causal claim was taken to be established (p. 49). Meanwhile, research continued on further elucidating the details of the underlying mechanism.

5 Can Giere-Causality Account for Mechanistic Evidence?

While Weber (2009) agrees with the epistemological thesis of Russo and Williamson (2007) and agrees that several difference-making analyses of causality—including those of Suppes, Eells and Humphreys—are inferior

because they cannot account for the use of mechanistic evidence to establish causal claims, he argues that there is one difference-making analysis of causality that can, namely Ronald Giere's probabilistic theory of causality.

Giere (1979) holds that C causes E in population U if and only if C makes a difference to the frequency of E in that population, in the sense that $P_X(E) \neq P_K(E)$, where X is the hypothetical population in which C holds for each member of U, K is the hypothetical population in which $\neg C$ holds for each member of U and the probabilities $P_X(E)$ and $P_K(E)$ are the relative frequencies of E in populations X and K, respectively. To take Weber's example, smoking is a cause of lung cancer in Belgium iff forcing every inhabitant of Belgium to smoke would make a difference to the incidence of lung cancer in Belgium, in comparison with the case in which all inhabitants were forced not to smoke.

Weber (2009, Section 4) argues that 'Giere's theory can account for the use of mechanistic evidence in contexts in which no *prima facie* relevant probabilistic evidence is available' (p. 283), on the grounds that in this case we need to rely on thought experiments in order to estimate whether $P_X(E) \neq P_K(E)$, and since knowledge of the underlying mechanisms is the only clue as to what would happen in populations X and K in this case, mechanistic evidence needs to be taken into account. Note though that, as Weber acknowledges, this argument does not impinge on the claim made in Russo and Williamson (2007, Section 5) that probabilistic accounts of causality fail to explain why, *when good probabilistic evidence is available*, mechanistic evidence is typically also required because in this case good probabilistic evidence is not available.

Weber (2009, Section 5) addresses this last claim more directly, arguing that Giere's theory 'can also account for the use of mechanistic evidence when relevant probabilistic associations have been established' (p. 286). The argument proceeds as follows: Randomised experiments are the ideal way to establish causal relations on Giere's account: randomly select a large sample from U, randomly assign members of the sample to two groups, a C group and a $\neg C$ group, measure the incidence of E in each group and use the results to estimate $P_X(E)$ and $P_K(E)$. But randomised experiments on humans are usually impossible for ethical reasons. However, one can perform non-random experiments with humans, or one can perform randomised experiments with animals. In each case, mechanistic evidence is important. With non-random experiments, mechanistic evidence can help with the problem of confounding: If there is no plausible mechanism linking two correlated variables, one can conclude that the correlation is spurious, but if there is a plausible mechanism linking two correlated variables, one can conclude that one is a cause of the other (Weber 2009, Section 5.1). On the other hand, with random experiments on animals, mechanistic evidence can help with the problem of extrapolation to humans: If a causal connection is found in an animal population and it can be demonstrated that the underlying mechanism is similar in a human population, one can conclude that the causal connection obtains in the human population (Weber 2009, Section 5.2). Hence, in either case mechanistic evidence is important.

But this argument also lacks bite, for similar reasons. The claim under scrutiny is the claim that 'the proponent of the probabilistic theory can't account for the fact that mechanisms are required even when appropriate probabilistic associations are well established' (Russo and Williamson 2007, p. 164). In this case, the required probabilistic associations are *not* well established. With non-random human experiments and with animal experiments (Section 5), the probabilistic evidence that $P_X(E) \neq P_K(E)$ in a human population is actually rather weak. While on Giere's account the required probabilities $P_X(E)$ and $P_K(E)$ are not directly measurable, Weber holds that a randomised experiment on a human population would count as good probabilistic evidence for $P_X(E) \neq P_K(E)$ and hence for the corresponding causal claim. And if that sort of difference-making evidence were available, it seems apparent that on Giere's account there would be no need for mechanistic evidence: No mechanistic evidence is strictly required in order to perform such a randomised experiment and such an experiment would lead to direct estimates of $P_X(E)$ and $P_K(E)$, rendering any further mechanistic evidence redundant.

Weber admits as much:

> Let us now imagine that we live in the ideal world for experimenters. In this world, time travel is possible (so we can do randomised experiments in the past ...) and there are no ethical restrictions (so there is no need to do animal experiments or prospective or retrospective studies ...). In such a world, one could say, a scientist who wants to make a Gierean population claim does not have any use for mechanistic evidence. Even if that were true, that would not count as an argument against Giere's theory: Giere wants to deal with real science and scientists in the real world. (Weber 2009, p. 290)

We disagree—that would count as an argument against Giere's theory. If Giere's theory cannot in principle account for the need for mechanistic evidence when there is good difference-making evidence available, then it cannot account for real science and scientists in the real world because it is clear that sometimes real scientists *do* manage to perform randomised experiments and *do* obtain excellent difference-making evidence, yet *also* require mechanistic evidence. So much the worse for Giere's theory, then.

Weber continues, though:

> Moreover, it can be argued that even in the ideal world for exerimenters, mechanistic evidence would play a role from a Gierean point of view (see Weber (2007) for this; the argument is connected with the stability of causal generalizations over time). (Weber 2009, p. 290)

Now Weber (2007, Section 3) argues that knowledge of the stability of mechanisms is required to extrapolate a causal claim, which has been determined from evidence gathered at one time, to the time at which a public policy intervention is implemented, if the causal claim is to be used to justify that policy. This seems quite right—in general, a policy will be implemented at a different time to that at which the evidence is gathered, and the causal relationships

might change in the meantime. But not all causal inference has policy in mind. Weber himself is interested in causal claims in historiography; such claims tend to be of interest for their own sake rather than as a justification for policy decisions. One example that Weber gives is the question of whether large-scale Japanese irrigation projects in 1930s Taiwan caused the subsequent move from extended families to nuclear families, since fewer hands were needed per household to ensure reliable harvests (Weber 2007, p. 358). An answer to this question is of great interest, but irrelevant for public policy interventions these days, and stability of the underlying mechanisms over time is consequently of little concern.

It should be clear then that counterexamples to Giere's theory occur at least when (a) there is excellent evidence for difference making in Giere's sense, (b) there appears to be no underlying mechanism to explain this difference making and (c) mechanistic evidence is not required for other purposes such as extrapolation over time. For example, consider a randomised experiment that showed an association between the recovery from bloodstream infection of patients in 1990–1996 and the saying of prayers in 2000 to ask for the recovery of those same patients (Leibovici 2001). Giere's theory would surely be forced to deem the saying of prayers in 2000 to be a cause of the recovery in 1990–1996. However, most observers would be inclined to say that such a causal claim has not been established because our evidence points to there being no underlying mechanism: As the author of the study acknowledges, 'No mechanism known today can account for the effects of remote, retroactive intercessory prayer said for a group of patients with a bloodstream infection' (Leibovici 2001, p. 1451). The same can be said for randomised experiments used as evidence for many causal claims in astrology and homeopathy; sceptics are right to remain sceptical of these claims, as long as the evidence points to no underlying mechanism that could account for any association. If this is so then Giere's theory of causality cannot be correct.

6 Discussion

The case studies considered in this paper investigate causal claims about humans. As Weber points out, such studies do not normally appeal to randomised trials on humans. Nevertheless, there is often good difference-making evidence in such cases (notably, not normally derived from randomised trials on animals). And as we hope to have shown, imaging technology is used to provide mechanistic evidence, over and above difference-making evidence, to establish a causal claim. We have argued that, contra Weber, this aspect of scientific practice is hard to reconcile with Giere's theory of causality.

We think that the way imaging technology is used to identify mechanisms and the fact that such mechanisms have to be identified before a claim of causality is well established show that there is more to the proper analysis of causality than simply difference making. Since difference making is also typically required (modulo qualifications to do with overdetermination etc.),

this shows that causality must involve both difference making and mechanistic aspects. This is a metaphysical as well as epistemological position and might therefore be seen as further supporting the claims of philosophers of science such as Ladyman and Ross (2007) that close attention must be paid to the practices of actual science in formulating metaphysical theories.

Imaging technology is philosophically interesting in its own right. Its use poses questions from the epistemological (e.g. concerning the demarcation between the observable and the unobservable) to the ethical (see Farrell (2010), for example, on the use of fMRI scans as evidence of the intentions of the defendant in criminal trials). But, as we hope to have illustrated for the methodology of the health sciences and the metaphysics of causality, it also sheds light on wider philosophical concerns.

Acknowledgements We are very grateful to Sarah Heathfield, Phyllis McKay Illari, Federica Russo, Erik Weber and two anonymous referees for helpful discussion and comments, to the Leverhulme Trust for supporting George Darby's research and to the British Academy for supporting Jon Williamson's research.

References

Barrett, H. H., & Swindell, W. (1981). *Radiological imaging (Vol. 2)*. New York: Academic.

Bechtel, W., & Richardson, R. (2010). Neuroimaging as a tool for functionally decomposing cognitive processes. In: Hanson, S.J. & Bunzl, M. (Eds.), *Foundational issues in human brain mapping* (pp. 241–262). Cambridge: MIT.

Broadbent, A. (2010). Inferring causation in epidemiology: Mechanisms, black boxes, and contrasts. In: Illari, P. M., Russo, F., & Williamson, J. (Eds.), *Causality in the sciences* (pp. 45–69). Oxford: Oxford University Press.

Cho, Z., Jones, J., & Singh, M. (1993). *Foundations Of medical imaging*. New York: Wiley-Interscience.

Clarke, B. (2010). *Causality in medicine with particular reference to the viral causation of cancers*. Ph.D. thesis, Department of Science and Technology Studies. London: University College London.

Davidson, D. (1989). Cardiovascular effects of alcohol. *Western Journal of Medicine, 151*, 430–439.

Demirovic, J., Nabulsi, A., Folsom, A. R., Carpenter, M. A., Szklo, M., Sorlie, P. D., et al. (1993). Alcohol consumption and ultrasonographically assessed carotid artery wall thickness and distensibility. *Circulation, 88*, 2787–2793.

Doby, T., & Alker, G. (1997). *Origins and development of medical imaging*. Carbondale: Southern Illinois University Press.

Duchene, C., Olindo, S., Chausson, N., Jeannin, S., Cohen-Tenoudji, P., & Smadja, D. (2010). Cannabis-induced cerebral and myocardial infarction in a young woman. *Revue neurologique, 166*(4), 438–442.

Farrell, B. (2010). Can't get you out of my head: The human rights implications of using brain scans as criminal evidence. *Interdisciplinary Journal of Human Rights Law, 4*, 89–95.

Giere, R. (1979). *Understanding scientific reasoning*. Fort Worth: Harcourt Brace, fourth (1997) edition.

Gillies, D. A. (2011). The Russo–Williamson thesis and the question of whether smoking causes heart disease. In: Illari, P. M., Russo, F., & Williamson, J. (Eds.), *Causality in the sciences* (pp. 110–125). Oxford: Oxford University Press.

Hendee, W., & Ritenour, E. R. (2002). *Medical imaging physics*. New York: Wiley-Liss, fourth (2002) edition.

Howick, J. (2010). Exposing the vanities—and a qualified defence—of mechanistic reasoning in clinical decision-making (unpublished manuscript).

Illari, P. M. (2011). Disambiguating the Russo-Williamson thesis. *International Studies in the Philosophy of Science* (in press).
Kelves, B. (1997). *Naked to the bone: Medical imaging in the twentieth century*. New Brunswick: Rutgers University Press.
Ladyman, J., & Ross, D. (2007). *Every thing must go*. Oxford: Oxford University Press.
Lakshman, R., Garige, M., Gong, M., Leckey, L., Varatharajalu, R., & Zakhari, S. (2010). Is alcohol beneficial or harmful for cardioprotection? *Genes and Nutrition, 5*, 111–120.
Leibovici, L. (2001). Effects of remote, retroactive intercessory prayer on outcomes in patients with bloodstream infection: Randomised controlled trial. *British Medical Journal, 323*, 1450–1451.
Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science, 67*, 1–25.
Raichle, M., & Mintun, M. (2006). Brain work and brain imaging. *Annual Review of Neuroscience, 29*, 449–476.
Rimm, E. B., Klatsky, A., Grobbee, D., & Stampfer, M. J. (1996). Review of moderate alcohol consumption and reduced risk of coronary heart disease: is the effect due to beer, wine, or spirits? *British Medical Journal, 312*, 731–736.
Roberts, G. M., & Garavan, H. (2010). Evidence of increased activation underlying cognitive control in ecstasy and cannabis users. *NeuroImage, 52*(2), 429–35.
Russo, F. & Williamson, J. (2007). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science, 21*(2), 157–170.
Russo, F., & Williamson, J. (2011). Generic versus single-case causality: the case of autopsy. *European Journal for Philosophy of Science* (in press).
Steinberg, D., Pearson, T. A., & Kuller, L. H. (1991). Alcohol and atherosclerosis. *Annals of Internal Medicine, 114*, 967–976.
van Fraassen, B. (1980). *The scientific image*. Oxford: Oxford University Press.
van Hell, H., Vink, M., Ossewaarde, L., Jager, G., Kahn, R., & Ramsey, N. (2010). Chronic effects of cannabis use on the human reward system: An fMRI study. *European neuropsychopharmacology, 20*(3), 153–63.
Vliegenthart, R., Oei, H.-H. S., van den Elzen, A. P. M., van Rooij, F. J. A., Hofman, A., Oudkerk, M., et al. (2004). Alcohol consumption and coronary calcification in a general population. *Archives of Internal Medicine, 164*, 2355–2360.
Webb, S. (1990). *From the watching of shadows: The origins of radiological tomography*. Bristol: Institute of Physics.
Weber, E. (2007). Social mechanisms, causal inference, and the policy relevance of social science. *Philosophy of the Social Sciences, 30*(3), 348–359.
Weber, E. (2009). How probabilistic causation can account for the use of mechanistic evidence. *International Studies in the Philosophy of Science, 23*(3), 277–295.
Williamson, J. (2009). Probabilistic theories. In Beebee, H., Hitchcock, C., and Menzies, P., (Eds.), *The Oxford handbook of causation* (pp. 185–212). Oxford: Oxford University Press.
Williamson, J. (2011). *Mechanistic theories of causality*. Philosophy Compass (in press).