

Generic versus Single-case Causality: the Case of Autopsy

Federica Russo & Jon Williamson

Draft of January 28, 2010

To appear in *European Journal for Philosophy of Science*

Abstract

This paper addresses questions about how the levels of causality (generic and single-case causality) are related. One question is epistemological: can relationships at one level be evidence for relationships at the other level? We present three kinds of answer to this question, categorised according to whether inference is top-down, bottom-up, or the levels are independent. A second question is metaphysical: can relationships at one level be reduced to relationships at the other level? We present three kinds of answer to this second question, categorised according to whether single-case relations are reduced to generic, generic relations are reduced to single-case, or the levels are independent.

We then explore causal inference in autopsy. This is an interesting case study, we argue, because it refutes all three epistemologies and all three metaphysics. We close by sketching an account of causality that survives autopsy—the epistemic theory.

Contents

| | | |
|----------|--------------------------------------|-----------|
| 1 | Introduction | 1 |
| 2 | Three Causal Epistemologies | 3 |
| 3 | Three Causal Metaphysics | 6 |
| 4 | Autopsy | 8 |
| 5 | Causal Epistemology Revisited | 13 |
| 6 | Causal Metaphysics Revisited | 14 |
| 7 | Conclusion | 19 |

1 Introduction

Intuitively, different causal claims operate at different levels. For instance, the claim that *smoking is a cause of lung cancer* appears to be at a different level to the claim that *Harry's smoking is a cause of his lung cancer*. In this paper

we investigate what this difference between levels amounts to. In particular, we shall look at various epistemological and metaphysical claims as to the relationship between the levels of causality, arguing against a range of views and in favour of one particular position.

There are various ways in which one can distinguish levels of causality. We will be interested in the *generic/single-case* distinction. A causal claim of the form *C is a cause of E* is *generic* if at least one of the causal relata *C* or *E* is generic, i.e., can be instantiated on more than one occasion. Otherwise the claim is *single-case*. Thus *Helicobacter Pylori is a cause of gastric ulcer is a generic claim*, while *the Helicobacter Pylori in Harry's stomach caused his gastric ulcer* is single-case. Usually, either both relata are generic or both relata are single-case—apparent cases of mixed causal claims are typically at root homogeneous. Thus, *the war is a cause of all my troubles* in fact has two single-case relata since *all my troubles* can only be instantiated once. However there are some genuine cases of mixed causal claims; for example, in a nested causal relationship such as *smoking causing cancer is a cause of restrictions on tobacco advertising*, the cause is single-case (a fact) while the effect is generic (Williamson and Gabbay, 2005). We will class causal claims with mixed-level relata as generic-level claims.

The generic/single-case distinction is similar in spirit to the *type/token* distinction common in philosophy (a *token* is a particular object while a *type* is a kind of object or entity). While causal claims sometimes appear to relate objects or kinds of object (e.g., *this government is a cause of the pauper*), causal relata are more properly construed as events, variables, facts, omissions and so on (e.g., *this government's actions are a cause of the recent increase in poverty*). Hence a distinction between type-level causal claims and token-level causal claims is, strictly speaking, rarely applicable. But 'type/token causality' is also often more loosely used just to invoke the generic/single-case distinction.

A more useful distinction between levels is the *population/individual* distinction. A causal relatum is *individual-level* if it concerns a particular individual or unit of observation, while it is population-level if it concerns a population or collection of units. Note that the population/individual distinction does *not* in general correspond to the generic/single-case distinction. Consider *the drop in birth-rate in 1990 is a cause of his financial successes*; here *the drop in birth-rate in 1990* is population-level but single-case, while *his financial successes* is individual-level but generic. The population/individual distinction will not play a role in this paper.

Common in the social sciences is the *aggregate-level/individual-level/multi-level* distinction. These descriptions are often used to describe a kind of model or the kind of variables that are used in a model. This is close to the population/individual distinction, but individual-level can here apply to a 'statistical individual', i.e., an arbitrary individual selected at random from a population, which is essentially population-level in the sense of the previous distinction (Russo, 2009, Chapter 6).

There are, of course, other ways in which one can distinguish between levels, including levels of organisation in mechanisms, levels of composition, and levels of size (Craver, 2007, Chapter 5). But those outlined above are perhaps the most important given our focus on causality. We will be concerned just with the generic/single-case distinction in this paper.

In this paper we ask how the levels of causality are related. We divide this into two questions. The first is epistemological: how can relationships at

one level be evidence for causal relations at the other level? The second is metaphysical: how can relationships at one level be reduced to relationships at the other level?

In §2 we present three different answers to the epistemological question. In §3 we present three different answers to metaphysical question. Then, in §4 we present in some detail a case study of causal inference in autopsy. We argue that the case of autopsy refutes all three answers to the epistemological question (§5) and all three answers to the metaphysical question (§6). But all is not lost—we just have to accept that causal epistemology and metaphysics are more complicated than might appear at first sight. In particular, as explained in §6, an epistemic theory of causality offers scope for giving a coherent account of causality at both levels.

2 Three Causal Epistemologies

Three basic kinds of epistemological view as to the relationship between generic-level and single-case causal claims have been predominant in the literature. We shall sketch these in this section.

Top down. One view is that is that epistemic access to causal relationships flows from the top, generic level, down to the bottom, single-case level: first we learn about generic causal relations, and only via this generic causal knowledge do we infer the truth of single-case causal claims. This is arguably Hume’s view:

It appears that, in single instances of the operations of bodies, we never can, by our utmost scrutiny, discover any thing but one event following another; without being able to comprehend any force or power by which the cause operates, or any connexion between it and its supposed effect. ...

Even after one instance or experiment, where we have observed a particular event to follow upon another, we are not entitled to form a general rule But when one particular species of event has always, in all instances, been conjoined with another, we make no longer any scruple of foretelling one upon the appearance of the other, and of employing that reasoning, which can alone assure us of any matter of fact or existence. We then call the one object, *Cause*; the other, *Effect*. ...

This is the sole difference between one instance, from which we can never receive the idea of connexion, and a number of similar instances, by which it is suggested. ... When we say, therefore, that one object is connected with another, we mean only that they have acquired a connexion in our thought, and give rise to this inference, by which they become proofs of each other’s existence. (Hume, 1748, pp. 73–76)

The top-down epistemology is an enduring view, advocated recently by Sober (1986) for instance. According to Sober, the likelihood of a *single-case* relation—e.g., *Harry’s smoking is a cause of his cancer*—depends on which *generic* relations—e.g., *smoking causes cancer*—hold in the population of reference.

To say that causal epistemology is top-down is not to say that generic causal relations provide the *only* evidence for single-case causal relations. There may also be experimental evidence or evidence of physical mechanisms, contiguity, temporal priority, and so on. The generic causal relations are themselves posited on the basis of non-causal single-case evidence. ‘Non-causal’ here means that while causal knowledge may be used in the inference to a generic causal relation of the form *A is a cause of B* (e.g., to identify the other known causes of *B*), there is no known causal relationship between *A* and *B* itself. In sum, then, to infer *striking the match is a cause of its lighting*, first one learns that *striking matches is a cause of their lighting* from repeated observations of particular strikings and particular lightings together with theoretical knowledge about the underlying mechanism, and then one particularises this generic causal claim to the single case by applying the further evidence that the match was struck and it lit, together with any further knowledge that bears on the mechanism involved in that particular striking.

Bottom up. An alternative view is that epistemic access to causal relationships flows from the bottom, single-case level, up to the top, generic level: first we learn about single-case causal relations, and only via this single-case causal knowledge do we infer the truth of generic causal claims.

This view is natural for those who, contra Hume, maintain that causal relations are directly perceivable. Ducasse (1926, 1968) is a key philosophical proponent of this position:

Hume’s view that no connection between a cause and its effect is objectively observable would be correct only under the assumption that a ‘connection’ is an entity of the same sort as the terms themselves between which it holds, that is, for Hume and his followers, a sense impression. For it is true that neither a color, nor an odor, nor a sound, nor a taste, nor any other sense impression, ‘connecting’ the cause and the effect, is observable between them. Indeed, we must even add that if a sense impression were present between those said to constitute the cause and the effect, it would, from its very nature as a sense impression, be quite incapable of doing any connecting and would itself but constitute one more of the entities to be connected. This is true in particular of the feeling of expectation which Hume would have us believe is what the words ‘necessary connection’ ultimately denote.

But there is fortunately no need for us to attempt to persuade ourselves that whenever people during the past centuries have talked of objective connection they thus have not really meant it at all. For the fact is that a causal connection is not a sensation at all, but a relation. The nature of that relation has already been minutely described above. It is, as we have seen, a relation which has individual concrete events for its terms; and, as analyzed by us, its presence among such events is to be observed every day. We observe it whenever we perceive that a certain change is the only one to have taken place immediately before, in the immediate environment of another. (Ducasse, 1926, pp. 63–64)

Michotte (1962) conducted psychological research and concluded that humans directly perceive causal relationships. This view has subsequently been developed by other psychologists—e.g., Leslie and Keeble (1987); Cohen et al. (1998); Schlottmann and Surian (1999); Twardy and Bingham (2002); Muentener and Carey (2006).

Independent levels. A third view is that epistemic access to causal relationships is neither top-down nor bottom-up but proceeds independently at the single-case and generic levels. Thus we learn about single-case causal relations, not by particularising generic causal knowledge but directly from single-case data; we learn about generic causal relations not by generalising single-case causal knowledge but directly from generic data. Thus Eells, for example, says:

I argue that (1) very little (if anything) about what happens on the token level can be inferred from type-level probabilistic causal claims, and that (2) very little (if anything) about type-level probabilistic causal relations can be inferred from token-level probabilistic causal claims. (Eells, 1991, p. 6)

Eells' position was, as we will see in §3, motivated by metaphysical considerations.

The independent-levels view might also be motivated by observing that both the top-down and bottom-up epistemologies have their share of problems. Top-down inference is subject to an analogue of the ecological fallacy. The ecological fallacy—a case of Simpson's paradox—was put forward by Robinson (1950). On the basis of evidence of the 1930 US census, Robinson showed that (generically) the literacy rate and the proportion of immigrants are positively correlated. But it is fallacious to infer from this that an immigrant will be more likely to be literate than a non-immigrant—in fact, analysing individual-level data, these variables¹ are negatively correlated, the explanation being that immigrants tended to settle in states in which the native population was more literate. Likewise, it is fallacious to infer, from the fact that hitting your golf ball into a tree is a generic preventative (i.e., negative cause) of it getting into the hole, the claim that in a particular single case hitting the tree did not positively cause it to enter the hole (Salmon, 1980). These fallacies suggest a problem with the top-down epistemology. On the other hand, bottom-up inference is itself subject to an analogue of the atomistic fallacy (Ní Bhrolcháin and Dyson, 2007, p. 5). Entering a raffle with a single monetary prize is, for all but one individual, a cause of financial loss. But entering the raffle may not be deemed to be a cause of financial loss at the generic level if the prize is sufficiently large (if the expected gain is positive). Hence one might have similar reservations about bottom-up inference to those that were had about top-down inference. If both top-down inference and bottom-up inference are problematic, the view that the levels are epistemologically independent has some support.

¹Note the change from population-level variables—*literacy rate* and *proportion of immigrants*—to individual-level variables—*being literate*, *being an immigrant*.

3 Three Causal Metaphysics

There are three basic kinds of metaphysical view as to the relationship between generic-level and single-case causal claims. We shall sketch these in this section.

Generic relationships primitive. One view is that single-case causal relationships are reducible in part to generic causal relationships: generic causal relationships are primitive and help determine single-case causal relationships. This was Hume’s view:

We say, for instance, that the vibration of this string is the cause of this particular sound. But what do we mean by that affirmation? We either mean *that this vibration is followed by this sound, and that all similar vibrations have been followed by similar sounds: Or, that this vibration is followed by this sound, and that upon the appearance of one the mind anticipates the senses, and forms immediately an idea of the other.* (Hume, 1748, p. 77)

This regular succession and the corresponding transition of mind from cause to effect constitute Hume’s definitions of generic causation: ‘Similar objects are always conjoined with similar’ (Hume, 1748, p. 76), and ‘an object followed by another, and whose appearance always conveys the thought to that other’ (Hume, 1748, p. 77). Hence, for Hume, causation in the single case of a vibration of a string is analysed in terms of generic causation.

The view that generic causal relationships are primitive is a natural one for many regularity, probabilistic and agency theories of causality. If causality is analysed as a *difference-making* relationship—i.e., if the defining feature of causality is that a cause makes a difference to its effects—and this difference-making only shows up in a succession of cases, then causality is, at root, a generic relation. A single-case relation can be construed derivatively as an application of a generic relation to a single-case scenario.

For example, one kind of probabilistic analysis of causality cashes out the causal connection in terms of frequency-raising: a cause raises the frequency of its effects. But frequency itself is a generic relation, in the sense that frequencies are defined on repeatably instantiatable attributes and undefined on single cases.² Hence, under such a probabilistic account, the required sense of difference making can only be manifested in the general case, and causality must be fundamentally a generic relation. (As Hume indicated, one can go on to make sense of a single-case causal claim by taking it to be subsumed by a corresponding generic relation.) The reader is referred to Williamson (2009) for an introduction to probabilistic theories of causality.

²This fact—that frequency is at root a generic relation—also motivates some accounts where levels are metaphysically independent, a position discussed below. In fact, it is not difficult to find cases where at the generic level the causal relatum is, say, a negative causal factor, and yet there exist single cases where an instantiation of the same causal relatum is instead a positive causal factor. This situation has led some authors to argue that the probabilistic account is not undermined by such cases, provided that distinct concepts of cause at work at distinct levels be defined. A stock example in the philosophical literature is the golf ball example mentioned in §2. See, for instance, Sober (1984), Sober (1986) Salmon (1980), Salmon (1984), Eells (1991), and (Russo, 2009, ch. 2 and 6) for a discussion.

Single-case relationships primitive. A contrasting view holds that generic causal relationships are reducible in part to single-case relationships: single-case relationships are primitive and help determine generic relationships.

This view is natural for those difference-making theories of causality in which the difference-making relation is manifested in the single case rather than the general case. A probabilistic theory which appeals to a notion of single-case chance is just such a theory, as is a counterfactual theory according to which possible worlds yield the truthmakers of counterfactuals.³

This view is also natural under the principal alternative to the difference-making approach, namely the *mechanistic* approach to causality. Thus if causal relationships are analysed in terms of physical processes that manifest conserved quantities (Dowe, 2000; Salmon, 1998), or in terms of connections in complex-systems mechanisms (Glennan, 1996), or in terms of dispositional capacities in the context of a nomological machine (Cartwright, 1999, Chapter 3), then causal relationships are analysed in terms of physical processes or mechanisms connecting single cases, and so it is single-case causal relations that are being analysed (Cartwright, 1989, Chapter 3). Derivatively, a generic causal relation can be thought of as a generalisation of underlying single-case relations.⁴

Independent levels. A third view has it that neither is the single-case reducible to the generic nor is the generic reducible to the single-case; rather, each level is to be analysed independently.

This view is appropriate for those who advocate a particular kind of metaphysical pluralism concerning causal claims. According to this kind of pluralism, single-case causal relations are to be analysed one way, but generic relations another. Eells (1991), for example, argues that two probabilistic analyses of causality are required: a *probability-comparison* analysis at the generic level and a quite different *probability-trajectory* analysis at the single-case level:

The problems of type and token probabilistic causation turn out to be quite distinct. Causal claims made on one of the two levels of causation turn out to be quite independent of claims made on the

³It is worth noting that in the account of Lewis (1973), truth conditions of counterfactuals are given in terms of similarity of worlds, and although similarity is not fully explicated in Lewis but taken as a primitive notion, it remains possible that the single-case causal relation depends on generic facts.

⁴Cartwright's capacity view of causality is close to the complex-systems mechanism view in important respects: complex-systems mechanisms and nomological machines are both conceived of as the underlying physical structures that are responsible for regularities, and they are responsible for those regularities in virtue of the way their components are organised and the behaviour that these components can engage in (Glennan, 1996; Bechtel and Abrahamsen, 2005; Machamer et al., 2000; Cartwright, 1999). But there are differences—for example, the behaviour of the components is variously underwritten by capacities (which are a special sort of dispositional property), activities (which differ from capacities in that they can be relational rather than monadic, and are manifestations rather than dispositions), or laws. The capacity theory maintains that an effect can only be regularly produced by a cause in the context of a nomological machine, so the theory may be classed as a mechanistic theory of causality. It cannot be classed as a difference-making theory because an entity can have a capacity to make no difference: a homeostatic mechanism has the capacity to maintain the status quo; a teenager has the capacity to do nothing. To discover a causal regularity we need to establish the underlying mechanism rather than establish that the cause makes a difference to its effects: 'we have to establish the arrangement and capacities of mechanical elements and the right shielding conditions that keep the machine running properly' (Cartwright, 1999, p. 50).




Figure 1: Template for describing causes of death D on a UK death certificate.




Figure 2: Template for describing causes of death D on a US death certificate.

other. And the two levels of probabilistic causation require quite different kinds of theories. (Eells, 1991, p. 6)

Weber (2007) offers the same kind of pluralism but with a different combination of theories: a difference-making (probabilistic) theory at the generic level, and a mechanistic theory for the single case.

4 Autopsy

In this section we describe in detail the kinds of inference in autopsy. Autopsy, we shall argue in §5 and §6, is an interesting case because it refutes all three epistemologies of §2 and all three metaphysics of §3.

Autopsy (or post mortem) is the examination of one or more corpses, the goal of which is to establish cause, mechanism and manner of death, and to individuate any possible injury or disease. Three types of autopsy are usually distinguished: the coroner’s autopsy, the forensic autopsy, and the clinical or academic autopsy:

Coroner. In the UK, as in many countries, when someone dies a death certificate needs to be completed. Usually (in about 55% of cases in England and Wales) a general medical practitioner signs this certificate, but if the cause of death is uncertain or unnatural, the death is referred to a coroner for further investigation. The coroner may have enough information to issue a certificate, or may arrange for a pathologist to carry out an autopsy. The coroner may also arrange an inquest into death. The purpose of the coroner’s autopsy is solely to determine a cause of death, and no consent from next of kin is required. In 1996, 22% of people who died in England and Wales were subject to a coroner’s autopsy; more than 95% of autopsies conducted in England, Wales and Northern Ireland are coroners’ autopsies (NCEPOD, 2006). Identifying the cause of death without autopsy is error-prone: in the UK and US, at least a third of all death certificates are incorrect (Roulson et al., 2005); it has been claimed that in Nigeria this figure may be as high as 71% (Izegbu et al., 2006).

Regulations of coroners’ autopsies vary according to the country’s legal system. In England, Wales and Northern Ireland, for instance, death is reported to the coroner if:

[there is] evident trauma; cause of death unknown; a doctor has not seen the deceased with known disease within the 14 days preceding the death; the patient died or was certified in an Accident & Emergency department; death occurred 14–28 days following surgery (individual coroner dependent); death was related to a mishap in hospital; industrial disease or acute poisoning. (Cooper et al., 2007)

In the ‘Death Certification Training Pack’ of the Office for National Statistics, nine more cases are mentioned, among which: death occurred in a violent or unnatural way, or there are suspicious circumstances. The death certificate must give

Ia: the disease or condition directly leading to death and, optionally,

Ib: another disease or condition leading to *Ia*,

Ic: another disease or condition leading to *Ib*,

II: other significant conditions contributing to the death but not related to the disease or condition (*Ic–Ia*) causing it.

This causal account is depicted in Figure 1. The corresponding template for death certificates in the US is given in Figure 2. In practice, deciding which event is the principal and direct cause of death *Ia* is not always easy, given the uncertainties involved:

If a man with lung cancer dies of a brain haemorrhage the cancer will not appear on the death certificate as it did not contribute to death. The cause of death given is the most probable one, and acute events are usually deemed more likely to have caused death than chronic ones, so a man with a recent myocardial infarction and chronic bronchitis will be deemed to have died of the former. (Alan Bates, personal communication)

See NCEPOD (2006) for an extended discussion of the coroner’s autopsy.

Forensic. Forensic autopsies (which are rare in the UK but more common elsewhere) are performed for medico-legal purposes, for instance when a corpse is found on a crime scene and also in civil law cases in some jurisdictions. Forensic autopsies are performed by forensic pathologists and do not require family consent. Often, forensic pathologists not only have to determine a cause of death but also the identity of the cadaver being examined. In determining the cause of death, the pathologist is also asked to determine the *manner* in which death occurred by identifying the series of events which led to death. Most jurisdictions include the following ‘manners of death’: homicide, accidental, natural, suicide, and undetermined. In forensic autopsies, examination may also serve to collect further evidence relevant to the crime investigated (Matthews, 2008).

Clinical/academic. Clinical or academic autopsies are carried out for medical and research purposes. This kind of autopsy is used to determine more information about the circumstances of death in a particular individual, or to

draw general conclusions about the causes of death in a population. In the former case, which for clarity of exposition we shall refer to as *clinical autopsy*, a death certificate is required before the autopsy starts. In the latter case, which we shall call *academic autopsy*, the cause of death in each of a series of individuals is established and general conclusions are drawn. These kinds of autopsy need family consent before they are carried out; they play a very important role in medicine as they aim to contribute to a better understanding of disease—see for instance Ayoub and Chow (2008), Dieboled (2007), Aho (2003), Sando et al. (1999).

Although procedures for carrying out post-mortem examination can be very similar in the different types of autopsies, it is worth noting that the goals, and therefore the types of inferences made, are quite distinct. Our interest lies exactly in the *kinds of inference* made in coroners', forensic, clinical and academic autopsies and in the *kinds of evidence* needed. In the following, we will show, first, that in autopsy, generic-to-single-case (i.e., top-down) as well as single-case-to-generic (i.e., bottom-up) causal inferences are performed, and, second, that at each level, causal claims depend on evidence both of difference-making and of mechanisms.

Generic-to-Single-Case. In some kinds of autopsy, single-case non-causal evidence is used in conjunction with generic causal knowledge to establish a single-case causal claim. This is true of the coroner's autopsy and forensic autopsy, which aim to determine the cause of death in a particular individual; it is also true of a clinical autopsy, which seeks to further elucidate the network of causal relationships leading to death in an individual.

Inference in these three kinds of autopsy fits the top-down causal epistemology of §2 well: in providing a single-case explanation of death much generic medical knowledge is taken for granted and indeed used to make inferences to the single-case.

Consider for instance how injuries or evidence of injuries are examined. In forensic medicine, the following types of injuries are distinguished: bruise, abrasion, laceration, incised wound, puncture or stab wound, gunshot wound, burn, fracture. Injuries, during autopsy, have to be described according to the location, size, shape and colour. Suppose the body shows a deep stab wound in the aorta blood vessel. Well established generic medical knowledge, notably about the functioning of the circulatory system, will be used to support the claim that haemorrhage caused death.

In fact, inference in these kinds of autopsy requires two types of evidence: mechanistic and difference-making. Reconstructing what happened clearly involves mechanistic considerations, whereas clarifying legally relevant circumstances and contexts of death clearly involves difference-making considerations. Consider again the injury example. The mechanism will be given by describing the injury; for instance, the dagger entered from above and with a certain inclination, thus provoking a wound of that size and shape, etc. (which might also provide evidence about the assailant: taller than the victim, for instance). But there are also difference-making considerations. Suppose the victim is found with other injuries as well, say fractures, but only light. It will be reasonable to infer, *ceteris paribus*, that the stab wound caused death rather than the assault that caused the fractures.

In general it is uncontroversial that in the single case the cause of death must be connected to the death of the individual through both difference-making and mechanistic considerations. Cause of death has been defined as follows:

the cause of death determined by the certifying physician [is] defined as (a) the disease or injury that initiated the train of morbid events leading directly to death, or (b) the circumstances of the accident or violence that produced the fatal injury. (NCHS, 2003, pp. 9–10)

That there needs to be difference-making is witnessed by the use of ‘initiated’ and ‘produced’: the cause of death must have changed things in a way that made death more likely. That there needs to be a mechanism is witnessed by the ‘train of morbid events’ and the causal chains of Figure 1 and Figure 2: there needs to be a clear physical connection between the initiator and the death. Hence it should be no surprise that, in order to establish the cause of death, medical practitioners require evidence not only of a difference-making relation between the supposed cause and death, but also of a mechanism or physical connection that explains the difference-making relation from a theoretical perspective. Typically, of course, this evidence is encapsulated as prior knowledge: the medical practitioner needs only identify the event (e.g., a heart attack) and needs no prompting to recognise the presence of both a (generic) difference-making connection and a (generic) mechanistic connection to death, applying this medical training to the single case in question. But in some cases, the pathologist might need to appeal to scientific research papers to provide the appropriate mechanistic and difference-making evidence. That mechanistic evidence guides autopsy inference is made clear in the ‘Guidelines for reports by autopsy pathologists’ (Adams, 2008, p. v):

Just as a sound knowledge of anatomy is essential to making the observations that go into the autopsy protocol, a sound knowledge of the mechanisms of death, that is, pathophysiological derangements, is necessary to make opinions concerning the cause, mechanism, and manner of death.

Single-Case-to-Generic. The mode of inference is rather different in the case of academic autopsy. Here, as before, single-case non-causal evidence is used in conjunction with generic knowledge of mechanisms, theoretical background knowledge and generic causal knowledge to establish single-case causal claims. But these single-case causal relations are then generalised by induction to a *new* generic causal claim. Academic autopsies thus contribute, unlike the other kinds of autopsy, to generic medical knowledge.

To give a flavour of this kind of inference, consider a few examples of academic autopsy:

- Tsokos et al. (2000), *Pressure sores: epidemiology, medico-legal implications and forensic argumentation concerning causality*, examined 10222 corpses in order to identify the frequencies and grading of pressure sores in various groups (age, gender, contributing cause of death, etc.) and to establish a causal relationship between pressure sores and fatal infectious complications or septicemia.

- Terrabuio Junior et al. (2007), *Autopsy-proven causes of death in lungs of patients immunocompromised by secondary interstitial pneumonia*, looked at 558 autopsies of patients who died during a stay in hospital after secondary interstitial pneumonia. In each case the authors examined the cause of the pneumonia (and thence death), and then generalised to determine which potential causes of pneumonia actually made a significant difference to the risk of death and could thus be counted as a generic-level cause. The aim was to use the results of this study to identify and treat those hospital patients who have a high risk of death from pneumonia. ‘It is thus important to know the main causes of death in this population to establish correct prophylactic actions, which are the cheaper and more intelligent ways to prevent secondary interstitial pneumonia’ (p.70).
- Abdel-Karim et al. (2007), *Causes of death at autopsy in an inpatient hospice program*, looked at 48 autopsies to ascertain the direct cause of death of individual hospice inpatients. While there were a variety of direct causes of death at the single-case level, the authors identified pneumonia as the significant direct cause of death at the generic level.
- Taggart and Craver (2006), *Causes of death, determined by autopsy, in previously healthy (or near-healthy) children presenting to a children’s hospital*, looked at 572 autopsies, determining the cause of death in each case, in order to identify the generic-level causes. They identified bacterial infection and neoplasia to be the key causes at the generic level.

Some remarks are in order. *First*, inferences in academic autopsy are neither simply top-down nor simply bottom-up. In fact, single-case autopsy results have to be established first, and then these are used as a basis for inferring a generic causal claim. What is used, then, is not a top-down or bottom-up causal epistemology but a subtle blend of the two.

Second, while one might think that one can trivially infer from *A is a cause of B* in the single case to *A is a cause of B* in general, this kind of inference is not appropriate to academic autopsy. At the generic level, population-wide difference-making is important—a cause should make a significant *generic* difference to its effects. This is because the generic causal claims are used primarily for prediction and control; but if a cause makes no overall difference to its effect then one can neither use the cause to predict the effect nor can one control the effect by controlling the cause. (Single-case causal claims, on the other hand, are primarily used for explanation and it only matters that the cause made a difference to the individual, not to the population as a whole.) Suppose, to take a toy example, that an individual had a peculiar genetic make-up that ensured that imbibing a glass of red wine each evening caused his death; a coroner’s autopsy rightly revealed this to be the cause of death in this single case. One might try to generalise and say that imbibing a glass of red wine each evening is a cause of death at the generic level. But, at the generic level, having a glass of red wine each evening is a preventative of death: for all other individuals in the population, their genetic make-up ensures that red wine in moderation prevents heart disease, and, on balance, lowers the generic probability of an early demise. It is this latter causal inference that is analogous to that of academic autopsy.

So while, for causal inference in autopsy, difference-making is important at both levels, a different kind of difference-making is required at each level. At

the generic level the difference-making is population-wide. For instance, [Tsokos et al. \(2000, p. 283\)](#) say that ‘[...] a significant correlation was found between the prevalence of pressure sores and certain underlying diseases, e.g. trauma, senile dementia, neurological diseases, apoplexy and marasmus;’ [Terrabuio Junior et al. \(2007, p. 69\)](#) evaluate the risks of various factors and make such claims as the following: ‘in patients with lung cancer or cachexia, the risk of death by bronchopneumonia as determined histologically was high (OR = 3.6; OR=2.6, respectively);’ [Abdel-Karim et al. \(2007, p. 894\)](#) identify pneumonia as the significant generic cause of death of hospice patients because ‘Pneumonia was present in 79% of the autopsied patients (n = 38) and appeared to be the major cause of death in 44% of patients (n = 21). It contributed to death in an additional 17% (n = 8) and was found incidentally in 19% of patients (n = 9). Twenty-one percent of patients (n = 10) had no evidence of pneumonia on autopsy.’

Third, as in the case of other kinds of autopsy, causal inferences in academic autopsy require evidence of mechanisms as well as evidence of difference-making. A causal claim cannot be considered established at the generic level without a plausible mechanism linking the cause and the effect. Consider for instance ([Fornaro et al., 2008, p. 408](#)), who conducted a postmortem study on patients dead of cardiac diseases. The goal of the study was to ‘assess the incidence of cardiac rupture and compare it with the other causes of cardiac death.’ The authors collected a fair amount of statistics pointing to heart rupture in a number of different cases. Their results, however, rest at a very conjectural level as ‘the responsible mechanism is not yet known.’

Error measurement. Clinical autopsy is sometimes used to monitor diagnostic or treatment error: in what proportion of cases is the diagnosis that is made while the patient is alive the correct diagnosis? This is yet another form of autopsy inference. Although it doesn’t squarely fall under any of the strategies discussed in §2, it is an important form of inference, as results of clinical autopsies may inform, e.g., medical practice or public health. See for instance [Shojania et al. \(2003\)](#), [Cummings et al. \(2007\)](#), [Jönssona et al. \(2007\)](#), [Ullah and Alamgir \(2006\)](#).

5 Causal Epistemology Revisited

It is time to reconsider the three epistemological answers to the question of the relation of the levels, as outlined in §2. The three views are: causal inference is top-down; causal inference is bottom-up; causal inference is independent at each level. (These are not the only possible answers to the epistemological question, but they are the answers that are prevalent in the literature.)

As the discussion of the two types of inference in autopsy in §4 shows, these three standard views all fail. In the case of autopsy, causal inference is certainly not independent at each level: each kind of autopsy infers causal claims at one level from those at the other. But neither is exclusively top-down or exclusively bottom up: some inferences are top-down (in coroner’s, forensic and clinical autopsy); others are a mixture of top-down and bottom-up (in academic autopsy). Thus none of the answers of §2 fits the case of autopsy.

Hence the epistemological relation between the levels is more subtle than that captured by any of the three standard views. Any viable causal epistemology must be able to handle the complex interplay between top-down and bottom-up causal inference that takes place in autopsy, for example.

6 Causal Metaphysics Revisited

Let us now turn to the three responses of §3 to the question of the metaphysical connection between the levels. These responses were: generic causal relations are primitive; single-case causal relations are primitive; the levels of causation are metaphysically independent.

The argument. In this section we will argue that these three views do not hold water. But our argument will be less direct than that of the last section. It is not that autopsy provides a straightforward counterexample to each of these three views—as it does to each of the epistemological views of §2—but, rather, that autopsy creates a problem for each of the metaphysical theories of causality that support these three views. We saw in §3 that, broadly speaking, the standard metaphysical theories of causality can be classed as either difference-making or mechanistic. Mechanistic theories of causality tend to support the view that single-case causal relations are primitive, while difference-making theories support the view that generic-causal relations are primitive if the difference-making in question is generic, or that single-case relations are primitive if the difference-making invoked by the theory is single-case. Pluralistic theories of causality can be used to support the view that each level admits of its own analysis (Eells adopts a different difference-making theory at each level; Weber advocates a mechanistic theory for the single case and a difference-making theory at the generic level).

The key point to note is that most metaphysical theories of causality—including pluralistic combinations of these theories (Williamson, 2006a)—analyse any particular causal relation either in terms of difference-making relations or in terms of mechanistic relations, but not both.⁵ Now in the health sciences in general (Russo and Williamson, 2007) and in the case of autopsy in particular (§4), evidence *both* of mechanisms *and* of difference-making is normally required to establish a causal claim. To determine that *Alf's heart attack was a cause of his death*, the medical practitioner needs to have evidence *both* that there is viable biological mechanism linking heart-attack and death *and* that the heart

⁵There are exceptions such as Glennan (2002), discussed below. Note however that several accounts that apparently appeal to both aspects simultaneously do not on closer inspection. While apparently appealing to both mechanisms and counterfactuals, Glennan (2009) is in fact pluralist, while for Craver counterfactuals are used as tests. Glennan (2009) distinguishes two concepts: (i) causal relevance and (ii) causal production. On the one hand, a property, in his account, is causally relevant to an event if the event counterfactually depends on the property; on the other hand, an event causally produces another if they are connected by a chain of causal processes. (In fact for Glennan, mechanisms are truthmakers for both kinds of causal claim.) In Craver's account, counterfactuals do not enter the *definition* of 'mechanism' but are rather used to evaluate causal claims. Indeed, Craver (2007, ch. 4) does not mention any counterfactual element in the *features* of mechanisms; on the contrary, he stresses—e.g., Craver (2007, p. 64, 65, 86)—that counterfactuals and manipulations serve to evaluate causal mechanisms; that is to say, counterfactuals are tests, not definers.

attack made a difference to his death (brought it about sooner than might otherwise be expected). At the generic level, in order to establish that *pneumonia is a cause of death in hospital patients*, those conducting an academic autopsy need to be aware of evidence *both* of a mechanism linking pneumonia and death, *and* that pneumonia makes a significant difference to death in the population in question.

But it is very hard for standard metaphysical theories of causality to account for this epistemic practice. If a particular causal claim is to be given a mechanistic analysis then surely mechanistic evidence would suffice to establish that claim. The proponent of the mechanistic analysis cannot explain why, in cases where there is excellent mechanistic evidence, evidence of difference-making is also required.⁶ On the other hand, if the causal claim in question is to be given a difference-making analysis then surely difference-making evidence would suffice to establish it. The proponent of the difference-making analysis cannot explain why, in cases where excellent difference-making evidence is not wanting, evidence of a viable biological mechanism is also required.

In sum, each of the standard metaphysical theories of causality fails to account for the epistemic practice of autopsy inference. Hence these metaphysical theories cannot be considered viable as they stand. But the three standard answers to the metaphysical connection between the levels were motivated by these metaphysical theories. This leaves the three views of §3 unmotivated. There are no grounds for preferring one of the views over any of the others, nor indeed for thinking that any one of them is correct.

Scope of the argument. We should emphasise that the requirement for evidence both of mechanisms and of difference-making is not confined to autopsy. That difference-making evidence is required is rarely controversial; but mechanisms are widely called for too. For instance, according to Hill (1965, p. 10), a classic in the health science literature, a causal claim should either (i) be biologically plausible or (ii) be the only possible explanation of the phenomenon in question.⁷ A similar point is made, more recently, by Swaen and van Amelsvoort (2008). In the social sciences, Gerring (2005, p. 179) claims that ‘Even in the most perfect experimental context where a causal relationship between two entities seems validated, one is unlikely to feel that an argument is clinched until a causal mechanism has been identified.’ And analogously, in their survey paper on causation in demography Ní Bhrolcháin and Dyson (2007, p. 25) claim that ‘Specifying and providing evidence of the mechanism involved is essential.’ In philosophy, Psillos (2004) put forward the idea that we can get a glimpse of the ‘secret connexion’ between cause and effect by using both mechanisms and counterfactuals (i.e., difference-making).

Gillies (2010, §2) poses some questions for the thesis, put forward in Russo and Williamson (2007), that typically evidence both of mechanisms and of difference-making is required to establish a causal claim. First, what kind of

⁶The proponent of the mechanistic analysis might argue that difference-making evidence can help establish the existence of a mechanism. But it still remains the case that according to this sort of position, once the mechanism is established no further difference-making evidence is required. It is this last point that we contest (Russo and Williamson, 2007).

⁷It is worth noting that Hill never considered his points as ‘criteria’ but just as guidelines. However, this is not relevant to our argument, which is just to point out what ‘elements’ are required to establish causal claims.

difference-making evidence is required? As Gillies suggests, statistical evidence is appropriate in the health sciences. We would agree with this, and point out that in the health sciences, and in particular in autopsy inference, it is usual to provide evidence that the cause and effect are probabilistically dependent conditional on possible confounding factors (i.e., conditional on the effect's other causes). We leave it open as to whether causal claims in other domains require other kinds of difference-making evidence: e.g., in the law, evidence of a counterfactual dependence between cause and effect may be required; in policy science, evidence of a dependence under intervention may be required.

Second, how plausible does the mechanism need to be in order for a causal claim to be established? Gillies distinguishes between a *plausible* mechanism and a *confirmed* mechanism, and suggests that it may suffice to show that there is a plausible mechanism linking the cause to the effect. We would suggest that the answer to this question depends on the epistemic context. Consider again the use of autopsy to determine the cause of death. If the results of the autopsy are to be presented in a civil legal case, then it may be sufficient that the cause of death be established on the balance of probabilities. On the other hand if the autopsy results are to be presented in a criminal case, then the cause of death may need to be established beyond reasonable doubt. In the latter situation the causal claim needs to be more credible, and correspondingly the evidence for the causal claim needs to be more credible. In a civil case, it can be enough that it is plausible that there is an underlying mechanism, but in a criminal case the mechanism itself may need to be established. (The epistemic context can also vary with type of autopsy: clinical and academic autopsies can be more speculative than coroners' and forensic autopsies in that while the latter are meant to conclusively determine the principal cause of death, the former are intended to pick up on a wider range of contributing causes. Hence evidence of underlying mechanisms may be more speculative in the case of clinical and academic autopsy.)

The way forward. Where does our argument leave our understanding of causality? Although the standard metaphysical theories of causality do not hold water for the reasons outlined above, we do not conclude that all attempts to analyse causality are doomed. There is at least one viable option as we shall now see.

One might suggest that one should appeal to *both* mechanisms *and* difference-making to analyse causality—arguably this is the approach of Glennan (2002), for instance. But one should beware of a simple-minded conjunctive analysis of the form: *A* is a cause of *B* if and only if (i) there is an appropriate mechanism linking *A* and *B*, *and* (ii) *A* stands in the right difference-making relation to *B*. This conjunctive analysis is unappealing for reasons that are independent of the problems outlined above. We alluded above to an epistemological problem facing standard metaphysical theories of causality: their inability to cope with multifarious evidence. But there are other problems too. Most notably, they are prone to counterexamples. There are well known cases in which causality is accompanied by mechanisms but no difference-making (Williamson, 2009), and there are also cases, often involving absences as causal relata, in which causal relations are accompanied by difference-making but no physical mechanisms. (These counterexamples led Hall (2004) to advocate pluralism, but, as

mentioned above, pluralism also succumbs to the epistemological problem—see [Williamson \(2006a\)](#) on this point.) While these counterexamples give the standard analyses of causality a rough time, they will give a conjunctive analysis an even rougher time: such an analysis will be prone to *all* the counterexamples, not just one sort or another. For this reason, a conjunctive analysis is, we think, not viable.

There is another option, though, that has better prospects: an *epistemic* theory of causality. Here is not the place to offer a detailed description or a sustained defence of the epistemic theory—see [Williamson \(2005, 2006a,b, 2007, 2009\)](#) and [Russo and Williamson \(2007\)](#) for more detailed discussion of the epistemic theory and its application to the health science domain. We shall merely give a flavour of the theory; enough to suggest that it has the potential to avoid the problems discussed in this paper and is worthy of further investigation.

According to the epistemic approach, a metaphysics for causality can be derived from an epistemology of causality. This is not a bad thing, since the metaphysics of causality is notoriously controversial, while there is substantial agreement (at least amongst practicing scientists) as to the kinds of way in which we can come to know about causal relations. The epistemic theory proceeds as follows. A causal epistemology can be thought of as a mapping from sets of possible evidence to sets of causal relationships: given one's evidence, a causal epistemology yields one or more sets of causal relationships that are compatible with that evidence. Now imagine that one had total evidence—one knew everything there was to know about the fundamental furniture of the world: the fundamental objects and the pattern of instantiation of the fundamental properties and the fundamental relations. Presumably, then, the correct causal epistemology—call it mapping ϵ —would yield the correct set of causal relationships. (This ought to be the case whatever one's metaphysical views about causality.) The epistemic theory takes this to be all there is to causality: according to this view, the causal relation is not one of the fundamental relations to be taken as basic in an ontology, nor does it supervene on the fundamental stuff by being definable in terms of difference-making or mechanisms; rather, the correct set of causal relationships is just the set of relationships that would result from applying mapping ϵ to total evidence. The epistemic theory, then, analyses the causal relation in terms of causal epistemology.

The usual practice, of course, is to go from metaphysics to epistemology: to construe the correct causal epistemology as the mapping that, when applied to total evidence, would yield the causal relation. In order to turn this around one needs an independent handle on the causal epistemology. What makes ϵ the correct *causal* epistemology, rather than an arbitrary mapping? According to the epistemic theory, the uses of our causal claims determine the causal epistemology. Causal claims are put to a variety of uses, inferential (e.g., prediction and control) as well as explanatory, and the correct causal epistemology is the mapping that yields claims which best fulfill these uses—i.e., which lead to the best inferences and explanations. In sum, then, the epistemic theory analyses the causal relation in terms of the mapping ϵ which is in turn characterised in terms of the inferential and explanatory role of the claims that we call 'causal'. Causality is the relation that best fulfills the role demanded of it, nothing more.

It is worth emphasising that the epistemic theory is a metaphysical, rather than epistemological, theory—it is an account of what causality *is*. According to this theory, causal claims, be they generic or single-case, may be construed as

claims about rational belief (Williamson, 2005): C is a cause of E iff an agent with total evidence, the ideal causal epistemology, and the capacity to apply the latter to the former, ought to hold a particular sort of directed belief from C to E ; such beliefs licence certain inferences that we call ‘causal’ inferences—these include inferences about interventions, as well as some predictive and explanatory inferences—and it is the appropriateness of these inferences that determine the ideal causal epistemology and hence that determine the set of beliefs deemed rational.⁸

Although the epistemic account is a metaphysical theory, the metaphysical and the epistemological aspects of causality are very tightly meshed. This has its own advantages: if causal relationships are what the epistemic theory says they are, there is no mystery as to how one can come to know about them.

How does the epistemic theory overcome the difficulties that we have encountered in this paper? First, it takes causal epistemology as basic and does not force this epistemology into any of the straightjacket views of §2. There is no claim that we only learn causal relations top-down; nor bottom-up; nor independently at each level. Thus the epistemic theory withstands the subtlety of causal inference in autopsy, for instance, which admits both top-down and bottom-up causal discovery. Second, the epistemic theory neither analyses causality solely in terms of mechanisms nor solely in terms of difference-making, and so it neither asserts the supremacy of evidence of mechanisms nor of evidence of difference-making. It thus avoids the epistemological trap into which the views of §3 fall, and copes with another key aspect of inference in autopsy. In reality the correct causal epistemology will yield causal relationships on the basis of a variety of kinds of evidence—including that of mechanisms and that of difference-making—and there is no reason to suspect that an easy analysis of causality in terms of one or other or both of these indicators is possible. Third, the epistemic theory avoids the counterexamples to mechanistic and difference-making accounts (and hence conjunctive accounts). Indeed, by analysing causality in terms of the correct causal epistemology, no counterexample is possible. A counterexample takes the following form: analysis X of causality deems A to be a cause of B (or, not a cause), but clearly A is not a cause of B (respectively, is a cause); hence analysis X is false. But, if A is clearly not a cause of B then presumably the ideal causal epistemology would not yield the claim ‘ A is a cause of B ’ when applied to total evidence, in which case the epistemic theory could not deem A to be a cause of B —no counterexample is possible.

Hence the epistemic theory withstands this rather substantial array of problems that beset other theories of causality.

⁸Of course, as to how to characterise the ideal causal epistemology is open to some speculation. In Williamson (2005) a hybrid of hypothetico-deductive and inductive epistemology is proposed and in Russo (2009) it is argued that the epistemic theory is compatible with a variational epistemology. A variational epistemology says, simply put, that reasoning underlying the discovery and confirmation of causal relations hinges on the notion of *variation*. A variational epistemology is not confined to either level—generic or single-case. Quite to the contrary, it can be shown that the concept of variation plays a role in causal discovery and confirmation at the generic level as well as in the single-case (Russo, 2009, ch. 7). This means that a ‘level-monism’ does not enter from the back door—the epistemic theory does not require an epistemology that is ‘single-level’, so to speak.

7 Conclusion

In this paper we have addressed the question of how the levels of causation are related, splitting the question into two: (i) whether causal relationships at one level can be evidence for causal relationships at the other level (epistemology), and (ii) whether causal relationships at one level can be reduced to causal relationships at the other level (metaphysics). Concerning the first, epistemological question, in §2 we discussed three kinds of answer, categorised according to whether inference is top-down, bottom-up, or the levels are independent. As for the second, metaphysical question, in §3 we discussed three kinds of answer, categorised according to whether single-case relations are reduced to generic, generic relations are reduced to single-case, or the levels are independent.

We then argued that autopsy (§4) is an interesting test case because none of the standard epistemologies (§5) nor metaphysics (§6) squarely fits it. After introducing the different types of autopsy, we showed that inferences are neither universally top-down nor universally bottom-up, but instead a blend of the two. In autopsy, moreover, causal relations are neither reducible to difference-making relationships nor to mechanisms, but rather both kinds of evidence are required in order to establish a causal claim.

Those considerations led us to suggest a different view of causality—the epistemic theory—that has the potential to overcome the difficulties raised by standard epistemologies and metaphysics. Our claim is that we need an account of causality that treats single-case and generic causal claims in an egalitarian way and that allows one to go back and forth from one to the other—none of the standard epistemologies are adequate in this respect. Also, we need an account of causality that treats mechanisms and difference-making in an egalitarian way—none of the standard metaphysics are adequate in this respect.⁹

This does not mean that we are pluralists. We do not endorse the view that causation can be neatly divided into a set of cases where a mechanistic relation is the salient relation and another set of cases where difference-making is the salient relation (Hall, 2004). Rather, we need a single metaphysical account that puts different kinds of evidence, mechanistic and difference-making, on an equal footing. Put otherwise, evidential pluralism (which we endorse) ought not to be confused with metaphysical pluralism (which we reject).

Acknowledgments

This research was supported by funds from the British Academy and the Belgian Fonds Nationale de Recherche Scientifique. We are grateful to Alan Bates, Nancy Cartwright, Brendan Clarke, Carl Hoefer and an anonymous referee for very helpful comments.

⁹Note that, in the light of the epistemic theory, the standard classification of theories of causality into difference-making and mechanistic theories is in fact somewhat blurred. The mechanistic theory is neither classifiable as purely mechanistic nor as purely difference-making. Hence the two categories are not mutually exclusive and exhaustive.

References

- Abdel-Karim, I. A., Sammel, R. B., and Prange, M. A. (2007). Causes of death at autopsy in an inpatient hospice program. *Journal of Palliative Medicine*, 10(4):894–898.
- Adams, V. I. (2008). *Guidelines for Reports by Autopsy Pathologists*. Humana Press.
- Aho, H. (2003). The roles of autopsies in the development of medicine. *Duodecim*, 119(13):1255–1264. Article in Finnish.
- Ayoub, T. and Chow, J. (2008). The conventional autopsy in modern medicine. *Journal of the Royal Society of Medicine*, 101:177–181.
- Bechtel, W. and Abrahamsen, A. (2005). Explanation: a mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36:421–441.
- Cartwright, N. (1989). *Nature’s capacities and their measurement*. Clarendon Press, Oxford.
- Cartwright, N. (1999). *The dappled world: a study of the boundaries of science*. Cambridge University Press, Cambridge.
- Cohen, L. B., Amsel, G., Redford, M. A., and Casasola, M. (1998). The development of infant causal perception. In Slater, A., editor, *Perceptual development: Visual, Auditory, and Speech Perception in Infancy*, pages 167–209. Psychology Press, East Sussex.
- Cooper, H., Leigh, M. A., Lucas, S., and Martin, I. (2007). The coroner’s autopsy. The final say in establishing cause of death? *The Medico Legal Journal*, 75:114–119.
- Craver, C. F. (2007). *Explaining the brain*. Oxford University Press, Oxford.
- Cummings, P. M., Le, B. H., and Lopes, M. B. S. (2007). Postmortem findings in morbidly obese individuals dying after gastric bypass procedures. *Human Pathology*, 38(4):593–597.
- Dieboled, J. (2007). Clinical autopsy—its role in modern medicine. *Praxis (Bern 1994)*, 96(43):1667–1671. Article in German.
- Dowe, P. (2000). *Physical causation*. Cambridge University Press, Cambridge.
- Ducasse, C. J. (1926). On the nature and the observability of the causal relation. *The Journal of Philosophy*, 23(3):57–68.
- Ducasse, C. J. (1968). *Truth, Knowledge and Causation*. Humanities, New York.
- Eells, E. (1991). *Probabilistic causality*. Cambridge University Press, Cambridge.
- Fornaro, G., Lazzerio, M., Giacalone, A., Aralda, D., Prando, M. D., Rossi, L., and Bongo, A. S. (2008). Death after ST-elevation myocardial infarction in patients treated with fibrinolytic therapy, angioplasty, or conventional therapy. A post-mortem study to verify cardiac rupture as a cause of death. *Giornale italiano di cardiologia*, 9(6):408–420. Article in Italian.
- Gerring, J. (2005). Causation: A unified framework for the social sciences. *Journal of Theoretical Politics*, 17:163–198.
- Gillies, D. A. (2010). The Russo-Williamson thesis and the question of whether smoking causes heart disease. In McKay Illari, P., Russo, F., and Williamson, J., editors, *Causality in the Sciences*. Oxford University Press, Oxford.
- Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science. Supplement: Proceedings of the 2000 Biennial Meeting of the Philosophy of*

- Science Association. Part II: Symposia Papers (Sep., 2002)*, 69(3):S342–S353.
- Glennan, S. (2009). Mechanisms, causes and the layered model of the world. *Philosophy & Phenomenological Research*, forthcoming.
- Glennan, S. S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, 44:49–71.
- Hall, N. (2004). Two concepts of causation. In Collins, J., Hall, N., and Paul, L., editors, *Causation and counterfactuals*, pages 225–276. The MIT Press.
- Hill, B. (1965). The environment and disease: association or causation? *Proceedings of the Royal Society of Medicine*, 58:295–300.
- Hume, D. (1748). Enquiry into the human understanding. In *Enquiries concerning human understanding and concerning the principles of morals*. Clarendon Press, Oxford, 1777 edition.
- Izegbu, M. C., Agboola, A. O. J., Shittu, L. A. J., and Akiode, O. (2006). Medical certification of death and indications for medico-legal autopsies: The need for inclusion in continue medical education in Nigeria. *Scientific Research and Essay*, 1(3):61–64.
- Jönssona, A. K., Holmgrenb, P., Druid, H., and Ahlner, J. (2007). Cause of death and drug use pattern in deceased drug addicts in Sweden, 2002–2003. *Forensic Science International*, 169(2–3):101–107.
- Leslie, A. M. and Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, 25:265–288.
- Lewis, D. K. (1973). Causation. In *Philosophical papers*, volume 2, pages 159–213. Oxford University Press (1986), Oxford.
- Machamer, P., Darden, L., and Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, 67:1–25.
- Matthews, P. (2008). Coroners’ law resource. <http://www.kcl.ac.uk/depsta/law/research/coroners/>.
- Michotte, A. (1962). *The perception of causality*. Methuen, Andover, MA.
- Muentener, P. and Carey, S. (2006). What is the domain of causal perception? investigating causal perception of motion and non-motion state change events in infancy. *Journal of Vision*, 6:965a.
- NCEPOD (2006). The coroner’s autopsy: do we deserve better? <http://www.ncepod.org.uk/2006.htm>. National Confidential Enquiry into Patient Outcome and Death.
- NCHS (2003). *Physicians’ handbook on medical certification of death*. National Center for Health Statistics, Hyattsville, Maryland.
- Ní Bhrolcháin, M. and Dyson, T. (2007). On causation in demography: Issues and illustrations. *Population and Development Review*, 33(1):1–36.
- Psillos, S. (2004). A glimpse of the secret connexion: harmonising mechanisms with counterfactuals. *Perspectives on Science*, 12(3):288–319.
- Robinson, W. S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review*, 15:351–357.
- Roulson, J., Benbow, E. W., and Hasleton, P. S. (2005). Discrepancies between clinical and autopsy diagnosis and the value of post mortem histology; a meta-analysis and review. *Histopathology*, 47:551–559.
- Russo, F. (2009). *Causality and causal modelling in the social sciences. Measuring variations*. Springer, New York.
- Russo, F. and Williamson, J. (2007). Interpreting causality in the health sciences. *International Studies in Philosophy of Science*, 21(2):157–170.
- Salmon, W. C. (1980). Probabilistic causality. In *Causality and explanation*,

- pages 208–232. Oxford University Press (1988), Oxford.
- Salmon, W. C. (1984). *Scientific explanation and the causal structure of the world*. Princeton University Press.
- Salmon, W. C. (1998). *Causality and explanation*. Oxford University Press, Oxford.
- Sando, Z., Borisch, B., and Bosman, F. (1999). Autopsies: from the past to the future. *Rev Med Suisse Romande*, 119(10):763–767. Article in French.
- Schlottmann, A. and Surian, L. (1999). Do 9-month-olds perceive causation-at-a-distance? *Perception*, 28:1105–1113.
- Shojania, K. G., Burton, E. C., McDonald, K. M., and Goldman, L. (2003). Changes in rates of autopsy-detected diagnostic errors over time: a systematic review. *Journal of the American Medical Association*, 289(21):2849–2856.
- Sober, E. (1984). Two concepts of cause. *PSA: Proceedings of the Biannual Meeting of the Philosophy of Science Association (1982)*, vol. 2:405–424.
- Sober, E. (1986). Causal factors, causal inference, causal explanation. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 60:97–113.
- Swaen, G. and van Amelsvoort, L. (2008). A weight of evidence approach to causal inference. *Journal of Clinical Epidemiology*.
- Taggart, M. W. and Craver, R. (2006). Causes of death, determined by autopsy, in previously healthy (or near-healthy) children presenting to a children’s hospital. *Archives of Pathology and Laboratory Medicine*, 130(12):1780–1785.
- Terrabuio Junior, A. A., Parra, E. R., Farhat, C., and Capelozzi, V. L. (2007). Autopsy-proven causes of death in lungs of patients immunocompromised by secondary interstitial pneumonia. *Clinical Sciences*, 62(1):69–76.
- Tsokos, M., Heinemann, A., and Püschel, K. (2000). Pressure sores: epidemiology, medico-legal implications and forensic argumentation concerning causality. *International Journal of Legal Medicine*, 113(5):283–287.
- Twardy, C. R. and Bingham, G. P. (2002). Causation, causal perception, and conservation laws. *Perception & Psychophysics*, 64:956–968.
- Ullah, K. and Alamgir, W. (2006). Comparison of clinical causes of death with autopsy diagnosis using discrepancy classification. *Journal of College of Physicians and Surgeons Pakistan*, 16(12):768–772.
- Weber, E. (2007). Conceptual tools for causal analysis in the social sciences. In Russo, F. and Williamson, J., editors, *Causality and probability in the sciences*, pages 197–213. College Publications, London.
- Williamson, J. (2005). *Bayesian nets and causality: philosophical and computational foundations*. Oxford University Press, Oxford.
- Williamson, J. (2006a). Causal pluralism versus epistemic causality. *Philosophica*, 77:69–96.
- Williamson, J. (2006b). Dispositional versus epistemic causality. *Minds and Machines*, 16:259–276.
- Williamson, J. (2007). Causality. In Gabbay, D. and Guenther, F., editors, *Handbook of Philosophical Logic*, volume 14, pages 95–126. Springer.
- Williamson, J. (2009). Probabilistic theories of causality. In Beebe, H., Hitchcock, C., and Menzies, P., editors, *The Oxford Handbook of Causation*, pages 185–212. Oxford University Press, Oxford.
- Williamson, J. and Gabbay, D. (2005). Recursive causality in Bayesian networks and self-fibring networks. In Gillies, D., editor, *Laws and models in the sciences*, pages 173–221. King’s College Publications, London. With comments pp. 223–245.