Logique & Analyse **206** (2009), 177–190

CONCEPTUAL ANALYSIS OF CAUSATION AND THEORETICAL UTILITY IN EVERYDAY CONTEXTS*

ERIK WEBER AND LEEN DE VREESE

Abstract
In this paper we elaborate Ned Hall's theoretical utility perspective for causation in everyday contexts. We do this by presenting some instances of it, thereby adding some flesh to the skeleton that Hall has provided. Our elaboration of the theoretical utility perspective also provides arguments for it: the instances we present show the fruitfulness of the approach. A question raised by Hall's proposal is: should we give up descriptive analysis of causation (and descriptive analysis in general) completely? We argue that, at least for causation, traditional descriptive conceptual analysis must be given up. However, we also argue that a more modest variant of descriptive conceptual analysis can be useful.

1. *Introduction*

Conceptual analysis of causation, as it is traditionally done, is concerned with our everyday causal intuitions, the way we think and reason about causation in commonsense situations, and the causal concepts we use when making everyday causal judgements. In general, a conceptual analysis can be descriptive or revisionist. This is also true for a conceptual analysis of causation. The aim of a descriptive conceptual analysis of causation is to develop a "correct definition" of causation that "fits the facts", i.e. fits the way the concept is used in everyday language. Dowe (2000) gives two quotes of philosophers who conceive their job this way. The first quote is by Curt Ducasse:

*We thank Tim De Mey and Anton Froeyman for their comments on earlier versions of this paper. The research for this paper was supported by the Research Fund of Ghent University through research project BOF2001/GOA/008 and by the Fund for Scientific Research – Flanders through research project G.0651.07.

> The problem of giving a 'correct' definition of the causal relation is that of making analytically explicit the meaning which the term 'cause' has in actual concrete phrases that our language intuition acknowledges as proper and typical cases of its use. (from Ducasse, 1926; quoted in Dowe, 2000, p. 2)

The second quote is by Hart and Honoré:

> The ordinary man has a quite adequate mastery of various concepts within the field of their day-to-day use, but, along with this practical mastery goes a need for explicit statement and clarification of the principles involved in the use of these concepts (from Hart & Honoré, 1985; quoted in Dowe, 2000, p. 3)

A revisionist conceptual analysis proposes changes to the way we use the word "cause". David Hume's analysis of causation is an example of revisionism: after rejecting the way we use the word in ordinary language on the grounds that this use is epistemologically suspect (necessary connections cannot be observed) he proposes to use "cause" in a different, acceptable way: cause = contiguity + time order + constant conjunction. Hume is well aware that this definition is not completely consistent with our daily intuitions; but it is free of epistemological problems and that justifies the revision.[1]

In De Vreese & Weber 2008 we have argued that revisionist conceptual analysis is problematic. A review of revisionist attempts shows that it is very difficult to be a revisionist in a non-arbitrary way. Revisionists have failed to provide good arguments for their revisions (sometimes they don't even try to give an argument).

What about descriptive conceptual analysis? Though Ned Hall practiced it himself till a couple of years ago (see e.g. Hall 2004) he has recently criticised the idea of descriptive conceptual analysis, which he characterizes as follows:

> ...the tradition that says that philosophers should focus attention on concepts of central philosophical interest, and pursue analyses of them by gathering up intuitions about hypothetical cases, gathering up a priori "platitudes" involving the given concept, and seeking

[1] According to the traditional interpretation of Hume, he denies that there are necessary connections in nature. According to the more recent "skeptical realist" interpretation, he does not deny the existence of necessary connections, but says that we cannot know anything about them. Both views warrant a revisionist attitude towards everyday causal language.

> a philosophical account of the concept that does the best job of systematizing all of this "data". (2006, p. 3)

In Hall's view, philosophy of causation is still almost completely in the grip of this tradition, while other areas of analytic philosophy have moved towards a more fruitful approach. He mentions the analyses of the concept of function in the philosophy of biology. According to Hall, these

> ... analyses earn their keep not by virtue of the way they fit the intuitive "data", but by virtue of their utility in explicating biological practice. (p. 3)

We think that most literature on function ascriptions in biology is still written in the old tradition which Hall wants to do away with. But we agree that a new approach which puts "an emphasis on theoretical utility as the prime desideratum for an account of causation" (p. 5) would be very fruitful. In some recent papers we have been engaged in analysing which causal concepts social scientists, biomedical scientists, engineers, physicists, etc. need in order to do good research. These analyses certainly fit into the "theoretical utility perspective" which Hall proposes (see De Vreese 2006 and Weber 2007 for details). However, these papers are about causation in the sciences, not about causation in everyday contexts. The first aim of this paper is to elaborate the "theoretical utility perspective" for causation in everyday contexts.[2] We do this by presenting some instances of it, thereby adding some flesh to the skeleton that Hall has provided. Our elaboration of the theoretical utility perspective also provides arguments for it: the instances we present show the fruitfulness of the approach. Our examples in Sections 2 till 4 will highlight — much more than Hall himself does in his brief paper — that the theoretical utility perspective is a promising direction (the title of Hall's 2006 paper is "Philosophy of Causation: Blind Alleys Exposed; Promising Directions Highlighted").

A further question raised by Hall's proposal is: should we give up descriptive analysis of causation (and descriptive analysis in general) completely? In other words: should the "theoretical utility perspective" be complementary to descriptive conceptual analysis, or should it replace it? Hall does not really argue that traditional conceptual analysis is a blind alley; he merely asserts this. The second aim of our paper is to argue that, at least for causation, traditional descriptive conceptual analysis must be given up. However,

[2] With "causation in everyday contexts" we refer to causal talk outside scientific research. This is a very wide domain, as will become clear in our examples.

we will also argue that a more modest variant of descriptive conceptual analysis can be useful. These arguments will be elaborated in Sections 5 and 6. Section 7 will contain our final conclusions.

2. *Causation in the courtroom*

2.1 Consider judge J who is confronted with the following case:

> In breach of a statute forbidding the sale to an infant under the age of 16 of dangerous weapons, the defendant sold an air rifle and ammunition to a boy of 13. The boy's mother told the boy to return the weapon to the defendant and get a refund: on the defendant's refusal to take the rifle back, the boy's mother took it from the boy and hid it. Six months later the boy found it and allowed a playmate to use it, who shot and accidentally wounded the plaintiff, destroying the sight of one eye. (Example taken from Lehmann et al., 2006, pp. 286)

Judge J rules that the playmate's playing with the gun is the most proximate cause of the injuries, and therefore holds the playmate legally responsible. The judge admits that the selling of the gun to the boy and the mother's inadequate reaction are also causes, but they are more remote. So he decides that the seller and the mother are not legally responsible and should not be punished.

There are several explanation-seeking questions one might ask about this case. For instance, one can ask why the mother and the seller were acquitted, while the playmate was found guilty. The answer is that the judge used a proximity criterion. We can also ask why all three of them were accused and brought to trial. The answer then is that there is a certain type of causal relation (counterfactual dependence) between the acts of the accused and the victim's wounds. So a counterfactual concept of causation is, in this case, "theoretically useful", in the sense that it helps to answer why-questions (it is helpful in constructing explanations).

2.2 Our second example is also taken from Lehman et al. (2006):

> For instance, a truck driver transporting toxic substances has an accident and the toxic substances are spilled on the road. Subsequently, the toxic substances drain through the soil into a groundwater reservoir, thereby polluting the water and creating damage to a nearby village's water supply. Despite the clear factual connection

> between the truck driver's conduct and the polluting of the water, from a legal causal perspective one may argue that the truck driver did not cause the polluting, because he had no means of foreseeing that there was a groundwater reservoir right at the place where the accident took place. (p. 281)

Again, we can ask several why-questions about such a case. If we want to explain why the truck driver was accused, a counterfactual notion of causation is certainly useful: there is a counterfactual dependence of the damage on the accident. If we want to explain why the truck driver was acquitted, we have to mention the fact that — according to the judge — the driver could not foresee the damage and thus had no reason to be extra careful. If the truck driver respected the speed limits and kept enough distance, he has behaved as he was supposed to behave. In legal discourse, the term "legal cause" is often used to denote a "factual" cause that satisfies whatever additional criterion one considers necessary for declaring someone guilty (this is also done on the quote above). Again, the counterfactual conception is theoretically useful in the context of explanations of juridical outcomes.

2.3 Till now, we have provided examples which show the fruitfulness of a counterfactual definition of causation. In order to extend this to other conceptions, let us look at cases where there is a causal mechanism but no counterfactual dependence (because there is simultaneous overdetermination). Ned Hall gives examples of such cases. Suppose that Billy and Suzy are engaged in a competition to see who can shatter a target bottle first. Suppose further that Suzy throws her rock a split second before Billy. Suzy's throw is spatiotemporally connected to the shattering in the right way, but Billy's is not:

> Suzy's throw is a cause of the shattering, but Billy's is not. Indeed, every one of the events that constitute the trajectory of Suzy's rock on its way to the bottle is a cause of the shattering. But the shattering depends on none of these events, since had they not occurred the bottle would have scattered anyway, thanks to Billy's expert throw. (2004, p. 235)

Let us adapt the example a bit so that we can let a judge come in. We assume that Billy and Suzy are throwing rocks at a window of a building. Suzy's throw is spatiotemporally connected to the shattering of the window, Billy's throw is not. If the judge acquits both, this can be understood by the fact that he uses a counterfactual theory of causation (i.e. he thinks that no one can be guilty unless there is a counterfactual connection). If the judge

condemns Suzy and acquits Billy, this can be understood from the point of view of a causal mechanical conception of causation. If the judge condemns both, this can be understood if we assume that he uses a probabilistic notion of causation (if two people throw a rock at a window, the chance that it shatters is bigger than if only one person throws a rock). The three concepts are theoretically useful, because we can use them to explain differences in outcomes.

3. *Blaming people*

Let us now leave the courtroom and turn to moral judgments in everyday contexts (more specifically: blaming people for something that they have done or not done). In this section, we will borrow and adapt some examples of Longworth (2006). Longworth's paper is a classical descriptive conceptual analysis, which is sympathetic to, but also exposes some problems for, Ned Hall's dualism (as developed and defended in Hall 2004). Longworth uses a whole series of examples, including the following:

> Gardener: My plants died when I was away on vacation. If my gardener had watered them, as he was supposed to have done, they would not have died. (p. 50)
> Queen Elizabeth: My plants died when I was away on vacation. If Queen Elizabeth had watered them, they would not have died. (p. 61)

Consider the following adaptations of the examples:

> My plants died when I was away on vacation. If my gardener had watered them, as he was supposed to have done, they would not have died. So I blame him for this: I hold him morally responsible for the death of my plants.
> My plants died when I was away on vacation. If Queen Elizabeth had watered them, they would not have died. But I don't blame her for this: I don't consider her morally responsible for the death of my plants.

These examples show that we do not attribute moral responsibility on the basis of causal relations alone. There is an additional requirement: the behaviour should violate some norm or expectation. However, causation is certainly important: we will not hold our gardener morally responsible for events that happened during our vacation which are not causally connected

to his misbehaviour. So having a well-defined concept of causation helps us to understand why we blame people for some events, but not for others.

4. *War games*

According to Ned Hall (2004), causation "understood as a relation between events" comes in at least two basic and fundamental different varieties:

> Events can stand in one kind of causal relation — *dependence* — for the explication of which the counterfactual analysis is perfectly suited [namely, had c not occurred, e would not have occurred] (...). And they can stand in an entirely different kind of causal relation — *production* — which requires an entirely different kind of causal analysis [namely, c produces e] (...). (2004, p. 226, cf. pp. 252–257; emphasis added)

If we put this view into a definition of causation, we get:

> C causes E if and only if [(E counterfactually depends on C) or (there is a causal mechanism by which C produces E)].

The following example is used by Hall as a case where only dependence is required:

> Suzy and Billy have grown up, just in time to get involved in World War III. Suzy is piloting a bomber on a mission to blow up an enemy target, and Billy is piloting a fighter as her lone escort. Along comes an enemy fighter plane, piloted by Enemy. Sharp-eyed Billy spots Enemy, zooms in, pulls the trigger, and Enemy's plane goes down in flames. Suzy's mission is undisturbed, and the bombing takes place as planned. If Billy hadn't pulled the trigger, Enemy would have eluded him and shot down Suzy, and the bombing would not have happened. (2004, p. 241)

Billy's pulling the trigger *did not produce* the bombing, rather it neutralized a state-of-affairs that would have prevented the effect from occurring. The occurrence of the bombing *was dependent* on Billy's pulling the trigger, but not produced by it. In this example, the effect counterfactually depends on the cause, but there is no mechanism linking cause and effect.

In other cases there is a causal mechanism but no counterfactual dependence (cf. the rock throwing example in Section 2.3). An important aspect

of Hall's view is that in typical cases of causation, both relations are present and production and dependence coincide (ibid., p. 254, p. 265). So, though the relations are conceptually distinct, in the actual world their extensions overlap in most cases. Atypical cases of causation occur where there is a production relation without counterfactual dependence (e.g. overdetermination) or where there is a relation of counterfactual dependence without a production relation (e.g. double prevention).

The example that we will use here is a counterexample which Hall gives to his own theory:

> [T]here are certain kinds of cases that we have some inclination to call cases of causation, but that also elude classification in terms of production or dependence. Here is an example, a slight variation on the story of Billy, Suzy and Enemy: This time, there is a second fighter plane escorting Suzy. Billy shoots down Enemy exactly as before, but if he hadn't, the second escort would have. (ibid., p. 271)

This example is an instance of preempted double prevention; in this case it is no longer true that the bombing wouldn't have happened if Billy hadn't pulled the trigger. Hall admits that one will nonetheless be inclined to grant Billy some causal responsibility for the success of the bombing, just as when there was no second escort. Hall sees this as "a piece of unfinished business that affects my account of causation" (ibid., p. 272).

Let us consider a group of air strike planners, who have to decide on the number of escorts on future bombing raids. If these experts claim that the presence of Second Escort contributed to the success of Suzy's raid, this can be understood from a probabilistic view of causation. They might even have evidence for it like the following. They ask a group of experienced army pilots to simulate some fights (like they would do during ordinary training sessions). First, we have 500 control trials in which there is only a pilot playing the role of Suzy, and one playing the role of Enemy. Suzy is virtually shot down 450 times before she can drop her bomb (so her success rate is 10%). Then we have a first set of 500 experimental trials, with a Billy-pilot also present. Enemy is shot down in time 450 times. He shoots at Suzy only 50 times, and succeeds in shooting her down 45 times. Suzy's success rate is much higher now: 91%. In the second series of 500 experimental trials we also add Second Escort. Billy misses Enemy 50 times, Second Escort shoots 50 times and misses 5 times. Enemy shoots only 5 times, and doesn't miss. Suzy's success rate is 99%. If we assume that the pilots are assigned randomly to the different trials and roles, this experiment is good evidence for the claim that the presence of Second Escort is a positive probabilistic cause of Suzy's success rate (the same goes for Billy's presence).

Though the example is fictitious, we can find many real-life cases in which people make decisions about policy: politicians and economic agents make them all the time. Well-defined concepts of causation can help us to understand the decisions that are made: in order to understand the decisions, we have to know how they use "cause" and related terms. The philosophical concepts of causation can function here as analytic tools.

5. *Giving up traditional descriptive conceptual analysis*

In the previous sections, we analysed the use of the concept of "cause" in different everyday (non-scientific) contexts. We learned that the use and the interpretation of "cause" differ in different contexts. The moral, legal, policy,... contexts in which one has to decide about which factor is to be labelled as the cause influence themselves the depicted interpretation of the term "cause". In other words, the everyday meaning of the term "cause" is highly contextual. Sometimes, the underlying concept guiding our deliberation is the same (e.g. a counterfactual interpretation of "cause" in case of the legal and moral examples), but different additional criteria are used to select one factor from a constellation of causes as *the* cause (e.g. a proximity criterion in the case of legislation versus a normative criterion in the case of decisions on moral blame). In other cases, (also) the underlying interpretation of the notion of "cause" differs (e.g. a probabilistic concept of "cause" in the case of policy decision versus the counterfactual concept in the moral and legal case). To systematize our everyday causal talk completely — as we should be able to do according to the traditional goals of a descriptive conceptual analysis — seems highly problematic given the influence of very different everyday contexts in which the term "cause" is used.

Hence, we will argue that traditional descriptive conceptual analysis must be given up. Actually, John Stuart Mill gave the basis of an argument against such a conceptual analysis before it was even invented by analytic philosophers. Mill defines causation as constant conjunction plus temporal priority:

> To certain facts, certain facts always do, and, as we believe, will continue to, succeed. The invariable antecedent is termed the cause; the invariable consequent, the effect. (1973, p. 327)

Mill adds a criterion (unconditionality of the constant conjunction) to make sure that non-simultaneous effects of a common cause (e.g. day and night) do not come out as cause-effect pairs. Mill is also aware that in most cases there is not a single antecedent:

> It is seldom, if ever, between a consequent and a single antecedent, that this invariable sequence subsists. It is usually between a consequent and the sum of several antecedents; the occurrence of all of them being requisite to produce, that is, to be certain of being followed by, the consequent. (1973, p. 327)

Note that Mill requires non-redundancy: each of the antecedents is required to produce the cause; if you leave one out, there is no invariable succession anymore (cf. Mackie's INUS conditions (1974)).

Finally, Mill insists on calling the complete set of antecedents the cause:

> In such cases it is very common to single one only of the antecedents under de denomination of Cause, calling the others merely Conditions. (1973, p. 327).
> The real Cause, is the whole of these antecedents; and we have philosophically speaking, no right to give the name of cause to one of them, exclusively of the others (1973, p. 328).

In Mill's view, singling out one of the antecedents as the cause (and thus distinguishing the cause from mere conditions) may be useful in everyday life, but we should not do that in philosophical treatises, nor in scientific investigations of nature. This latter point is not important. What is important here is that we single out certain antecedents as cause. John Mackie also stresses this phenomenon of singling out:

> And there is no doubt that we tend to be a bit selective, to be more ready to call some kinds of factors causes rather than others. There are no firm rules governing this selection, but there are some fairly systematic tendencies. (1974, p. 34)

For example, in the case of the gardener and Queen Elizabeth, all of the following antecedents are parts of the cause of the death of my plants: my plants need water to survive, the earth of my plants dried out during my vacation, it did not rain during my vacation, I did not water my plants since I was not at home, the gardener did not water my plants during my vacation although I told him to do it, Queen Elizabeth did not come over to water my plants, nobody else came to water my plants,... Although all these antecedents together resulted in the death of my plants, and the effect is counterfactually dependent on all of these antecedents, we are inclined to single out the gardener's negligence as the cause of the death of my plants, since he was the one who was supposed to water my plants (the additional normative criterion).

The problem is not that this selection occurs in everyday causal talk. But for a traditional descriptive analysis to succeed there should be firm rules for this selection. Suppose that we accept, for the sake of argument, that nothing is called a cause unless it is an INUS condition of its effect. A descriptive conceptual analysis can only cope with selection if it contains a set of strict selection rules, e.g.:

> American defence lawyers in defending their clients always use selection criterion X to single out causes from the set of INUS conditions.
> Belgian judges in writing their sentences always use selection criterion Y to single out causes from the set of INUS conditions.

Without a complete set of rules like this (covering all types of persons and contexts) a descriptive conceptual analysis cannot do what it promises to do. Moreover, also for the selection of a basic concept of "cause" depending on the context of use (e.g. a probabilistic or a counterfactual interpretation of "cause"), there does not seem to be firm rules. Our use of the word "cause" varies in a complex and subtle way. Even if it can be systematised (this is not certain: there might be no rules at all, just chaotic behaviour) the system will be very complex, and thus very difficult to construct.

Consequently, the instances we have provided in Section 2 till 4 seem to show that Hall's "theoretical utility perspective" provides the only possible way to analyse "cause" in everyday contexts. An important characteristic, which the cases have in common, is that the philosophical analyses of the concepts of causation are also the analysis of *something else*: a moral claim, a legal claim, a policy decision, ... Given that the context of use is of crucial importance for the selection of the applied concept of "cause" and for the selection of *the* cause from a series of antecedent conditions of the effect, we are not able to abstract from these concrete contexts of use. Hence, we are unable to provide the results of a generalizing descriptive analysis stipulating the sufficient conditions for the meaning of "cause" in any everyday context.

As becomes clear from De Vreese (2006) and Weber (2007), the same holds true for the use of the notion of "cause" in the sciences. Also in a scientific research context, the meaning of the notion of "cause" cannot be separated from its use in a certain scientific context with its specific goals and methodologies. In conclusion, analysing causation from a "theoretical utility perspective" seems to us the most fruitful, making clear how the notion of "cause" gets different useful interpretations in different contexts of practice.

6. *Completely giving up descriptive analysis?*

Is the inevitable consequence that descriptive conceptual analysis should be completely given up? It is if one requires from such an analysis that it provides sufficient conditions. However, if philosophers want to stick to descriptive conceptual analysis, they can give up the biconditional. Instead of trying to give a set of jointly sufficient conditions, they should then try to find necessary conditions. This is how Mackie conceives of his INUS account of causation. If we apply this to Hall's dualistic conceptual analysis, the result is this:

> If C causes E, then [(E counterfactually depends on C) or (there is a causal mechanism by which C produces E)].

The aleatory nature of causal attributions in everyday contexts does not preclude such an analysis. It is strange that analytic philosophers (even Ned Hall) have struggled so long to do the (sheer) impossible, instead of settling for the more modest aim.

Given that a more modest type of descriptive conceptual analysis is possible, the question is whether it is worthwhile. We think the answer is positive, for two reasons. First, the modest kind of descriptive conceptual analysis can explain and predict certain aspects of our causal talk: it can explain why (in some cases) and predict when (if certain initial conditions are satisfied) we will *not* call one event a cause of another. It gives us a partial grip on the principles involved in ordinary causal talk. The second reason is that this kind of analysis leads to the development of new causal concepts in order to account for counterexamples which available concepts face[3]. Recall the example of pre-empted double prevention that we used in Section 4:

> [T]here are certain kinds of cases that we have some inclination to call cases of causation, but that also elude classification in terms of production or dependence. Here is an example, a slight variation on the story of Billy, Suzy and Enemy: This time, there is a second fighter plane escorting Suzy. Billy shoots down Enemy exactly as before, but if he hadn't, the second escort would have. (Hall 2004, p. 271)

[3] Although it might be very hard to distinguish between counterexamples that should be accounted for by reference to an additional criterion used for causal attribution on the basis of a known causal concept, and counterexamples that should be accounted for by reference to a new causal concept used for causal attribution.

For a descriptive conceptual analyst involved in an attempt to find necessary conditions for causation, this counterexample constitutes a challenge which might only be met by developing a third concept which either complements or replaces Hall's two concepts. This additional concept can turn out to be theoretically useful. In other words: descriptive conceptual analysis can have a heuristic value. As long as one does not make it too ambitious (by desiring completeness) there is no harm in it.

7. *Conclusion*

In this paper we have elaborated Ned Hall's theoretical utility perspective for causation in everyday contexts. We have presented instances of it, showing how various concepts of causation can help us to explain how people make causal attributions (in the context of practical decisions, condemning people, assigning responsibility to people). What the cases discussed in Sections 2–4 have in common, is that the philosophical analyses of the concepts of causation are at the same time the analysis of *something else*: a moral claim, a legal claim, a policy decision, ... Given that the context of use is of crucial importance for the selection of the applied concept of "cause" and for the selection of *the* cause from a series of antecedent conditions of the effect, we are not able to abstract from these concrete contexts of use in order to give a generalizing descriptive analysis of the meaning of "cause" in any everyday context. Because of this characteristic, we have argued that traditional descriptive conceptual analysis is no viable option. However, in the final section we have argued that a modest variant of descriptive conceptual analysis, which does not aim at completeness, can be heuristically useful. It can also have a "negative" explanatory and predictive value with respect to everyday causal attributions.

Centre for Logic and Philosophy of Science
Ghent University (UGent)
Blandijnberg 2
B-9000 Gent
Belgium
E-mail:
Erik.Weber@UGent.be
Leen.DeVreese@UGent.be

REFERENCES

De Vreese, Leen (2006), "Causal Pluralism and Scientific Knowledge: an underexposed problem," *Philosophica* 77, pp. 125–150.

De Vreese, Leen & Weber, Erik (2008), "Confusion and Bad Arguments in the Conceptual Analysis of Causation", in *Logique et Analyse* 201, pp. 81–99.

Dowe, Phil (2000), *Physical Causation*, Cambridge, UK: Cambridge University Press.

Ducasse, Curt (1926), "On the Nature and the Observability of the Causal Relation," *Journal of Philosophy*, 23, pp. 57–68.

Hall, Ned (2004), "Two Concepts of Causation," in: Hall, N., Collins J. and Paul, L.A. (eds.), *Causation and Counterfactuals*, MIT Press, Cambridge, pp. 225–276.

Hall, Ned (2006), "Philosophy of Causation: Blind Alleys Exposed; Promising Directions Highlighted", in *Philosophy Compass* 1 (PS 002), pp. 1–9.

Hart, H. & Honoré, T. (1985), *Causation in the Law*, Oxford: Clarendon.

Lehmann, Jos, Breuker, Joost and Brouwer, Bob (2004), "Causation in AI and Law," *Artificial Intelligence and Law* 12, pp. 279–315.

Longworth, Francis (2006), "Causation, Pluralism and Responsibility," *Philosophica* 77, pp. 45–68.

Mackie, John Leslie (1974), *The Cement of the Universe. A study of causation*, Oxford: Clarendon Press.

Mill John Stuart (1973), *Collected Works, volume VII.* Toronto: University of Toronto Press [Edited by J.M. Robson].

Weber, Erik (2007), "Conceptual Tools for Causal Analysis in the Social Sciences," in: Russo, Federica and Williamson, Jon (eds.), *Causality and Probability in the Sciences*, College Publications, London, 2007, pp. 197–213.