DRAFT VERSION FEBRUARY 2017

CAUSATION & TECHNICAL PROBLEM SOLVING
An analysis of causal knowledge underlying proposed solutions for technical problems

Inge De Bal - corresponding author
Erik Weber

Centre for Logic and Philosophy of Science
Ghent University (UGent)
Blandijnberg 2, B-9000 Gent, Belgium
Inge.DeBal@UGent.be
Erik.Weber@UGent.be

Abstract

In this paper we analyse the causal underpinnings of remedy claims found in technical problem solving instructions. For these claims to be successful, they need to be based on causal relations that hold in the world and that have certain properties. These required properties are the focus of our paper. We first introduce several examples from car- and bike repair manuals that demarcate our topic and function as illustrations throughout the paper. We then formulate three success criteria for problem solving manuals: the efficiency requirement, the no harm requirement and the maximal assistance ideal. These criteria determine the required strength and properties of the causal relation, and are used to frame our analysis. We start from theories of causation by Ronald Giere, Ellery Eells and John Mackie and develop a series of definitions to capture the properties of the aforementioned causal relations. We conclude that remedy claims need to be based on causal relations with the following property: positive causal factorhood with weak context-unanimity. Moreover, it is desirable to look for Mackie causes, viz. causes that are sufficient in maximally normal contexts. We finally show that our analysis is not limited to means of conveyance (this is the field in which our initial examples are located) by means of a case study from a different field (radio repair manuals).

CAUSATION & TECHNICAL PROBLEM SOLVING
An analysis of causal knowledge underlying proposed solutions for technical problems

1. Introduction

1.1 *Demarcation of the topic and some terminological conventions*

Everyone is familiar with what we call *technical problem solving instructions* (TPSI for short). Here are some examples that are taken from car - and bicycle repair manuals:

E1. Excessive fuel consumption: the air filter element is dirty or clogged. Remove the air filter, clean with compressed air and refit. (Mead & Legg, ref. 15, 1.13)

E1'. Excessive fuel consumption: the tyres are underinflated. Check and adjust pressures. (Mead & Legg, ref. 15, 0.14)

E2. The major cause of slow engine cranking or a "no-start" condition is battery terminals which are loose, dirty, or corroded. […] disconnect the battery and clean the terminals of both the battery and the cables. (Chilton, p.16)

E3. Difficulty engaging gears: worn or damaged gear linkage or gear cable. Replace the cable. (Mead & Legg, ref. 15)

E4. Indication: Engine sputters, may fail to start. Condition: water in the fuel. Remedy: […] For a layer of water, the tank must be drained, and the fuel lines blown out with compressed air. (Chilton, p. 288)

E5. Starter motor turns engine slowly: partially discharged battery. Recharge. (Strasman, p. 23)

E6. When pedaling forward, the cassette spins, but there is no drive to the bike: the freehub body is worn. Replace the freehub body. (Sidwells & Ballantine, p.37)

E7. The brakes are hard to apply, and/or sluggish to release: grit and dirt is inside the cable outers or the lubrication on the inner cables has dried. Strip down the cables, flush the outers, and clean the inner cables with degreaser, lubricate both, and reassemble. (ibid.)

We will use most of these examples further in this paper. They illustrate that a TPSI has three elements:

- a problem (e.g. excessive fuel consumption)
- a diagnosis (e.g. dirty air filter element)
- a remedy (e.g. remove, clean and refit the air filter).

In some cases (e.g. E7) there is more than one possible diagnosis.

1.2 Criteria of success for of repair manuals

This paper focuses on the third element: we will investigate the causal underpinnings of remedy claims. We start from three criteria that, in our view, a successful repair manuals must satisfy to be successful. We present and motivate these criteria in this section.

The first criterion is that repair manuals should avoid prescribing useless actions, i.e. actions whose result is irrelevant for solving the problem. In order to clarify this criterion, it is useful to distinguish between the *immediate result* of an action and possible *further consequences*. Suppose I experience the temperature in my office as too high. I decide to open the window because I want a cooler room (that is my aim). The open state of the window is the immediate result of my action. Whether this immediate result in its turn leads to the desired state (a cooler room) after some period of time depends on an additional factor, viz. the outside temperature. If the room cools down, then we call this a consequence of our action. The useful distinction between results and consequences of actions has a history in the philosophy of action as well as in the literature on causation.[1] With this terminology in place we now can formulate the general *efficiency requirement*:

> Including a remedy "Do X" is suitable only if there is a causal relation between R (the immediate result of the doing X) and the problem stated.

This general requirement has a specific instantiation for each TPSI. Let us give some examples. For E5 and E1', the specific efficiency requirements are, respectively:

> Including the remedy "recharge the battery" is suitable only if there is a causal relation between battery charge and the speed at which starter motors can make the main engine run.
> Including the remedy "inflate the tyres" is suitable only if there is a causal relation between tyre pressure and fuel consumption.

If a remedy claim does not satisfy the general efficiency requirement, it prescribes a useless action. Surely, making users waste time in performing useless acts is not a good idea.

The second criterion is the *no harm requirement*: executing the instructions in repair manuals should make the problem *worse*. In terminology above: the consequences of the actions should not be harmful.

The last criterion is that repair manuals should help the users as much as possible – within the range of actions they are capable of performing easily – in solving the problem(s) that the users experience. We call this the *maximal assistance ideal.*

1.3 Aims and s*tructure of the paper*

This paper can be seen as a stepwise analysis of how these criteria determine the nature of the *causal knowledge* underlying proposed solutions. As should be clear, remedy claims themselves are not causal claims: grammatically they are imperative clauses (see the examples: "remove', "recharge", "adjust", ….). However, they have to be based on certain causal relations that hold in the world (cf. the efficiency requirement). Hence, we do not analyse causation *in* remedy claims here (because there is no such thing) but causation *underlying* remedy claims.

Our first aim in this paper is to investigate what kind of causal knowledge is necessary to satisfy the efficiency requirement. We characterise the minimal strength of causal knowledge: if the kind of knowledge defined in this analysis is absent, the efficiency requirement is violated.

The other aims relate to how the other criteria in 1.2 (combined with the efficiency requirement) may demand the presence of stronger type of causal relations. Our second aim is to investigate whether and how the no harm requirement, if incorporated in the analysis on top of the efficiency requirement, motivates a stronger type of causal relation. Our third aim is to investigate how incorporating the maximal assistance ideal affects the kind of causal relations that are desirable.

The structure of this paper is as follows. In Section 2 we present a tool that we will use throughout this paper: the concept of physical setup. Section 2 also contains some clarifications of how we work in the rest of the paper. In sections 3 till 5 we gradually develop our analysis: each of these sections is devoted to one of the three aims. Each section contains many examples to clarify our points. In Section 6 we summarise what we can learn from Sections 3 till 5. Finally, in Section 7 we provide an elaborate illustration that comes from a different domain than the previous examples.

1.4 *Motivation*

---

[1] The distinction between results and consequences as we use it here stems from the work of Georg Henrik von Wright (1971, pp. 66-67).

Analysing which types of causal relations are the basis for successful remedy claims provides an epistemological framework for reflection on technical problem solving instructions. The goal of TPSIs is to help the users of the technical artefacts covered by the manual in which they occur. It is therefore important that there are ways of evaluating the remedy claims from the TPSIs and ensuring that they satisfy the success criteria presented in 1.2. Our analysis provides an epistemological framework for such an evaluation. Before putting instructions in a manual, we can reflect on (1) whether they are based on appropriate causal relations and (2) whether we have sufficient evidence to back up our claims about these relations.

Remedy claims are parts of TPSIs. The latter also contain problem descriptions and diagnostic claims. There are probably interesting philosophical issues about these other parts too, but they fall outside the scope of this paper.

2. Preliminaries

2.1 *Physical setups*

In the biomedical sciences, the validity and meaning of causal claims are intrinsically linked with the population that the causal claim is about. This can be shown by an example that we borrow from Daniel Steel (2008, p. 82). The following causal claims are true:

Aflatoxin B_1 causes liver cancer in rats.
Aflatoxin B_1 causes liver cancer in humans.

However, the following claim is false:

Aflatoxin B_1 causes liver cancer in mice.

Since the population we talk about makes a difference (not only in this case, but general), it is important that we always explicitly mention the intended population when making a causal claim. If we only say

Aflatoxin B_1 causes liver cancer

it is not clear which population we are talking about. And depending on which population we intend to talk about, the truth value differs.

It is important to do the same kind of specification regarding physical causal claims. Without specifying what our causal claim is about, we cannot understand or evaluate it. Consider the following examples of general physical causal claims:

(G_1) Increasing the temperature of a gas causes an increase in volume occupied by the gas.
(G_2) Increasing the temperature of a gas causes an increase in volume occupied by the gas in rigid, closed containers.
(G_3) Increasing the temperature of a gas causes an increase in volume occupied by the gas in non-rigid or open containers.

G_2 is incorrect. Gas in a rigid, closed container cannot expand. So in rigid, closed containers, an increase in temperature of the gas will not lead to an increase in volume occupied by the gas. G_3, on the other hand, is a correct causal claim. We therefore cannot evaluate G_1 as such: it is underspecified. We need a concept that can play the same role as ‘population’ does in causal analysis of the biomedical sciences, viz. specifying what the causal claim is about.

The concept we propose for this purpose is called ‘physical setup’. It is defined as follows:

> A physical setup is a whole comprising at least two physical objects, located in space and time, with each having at least one variable feature. [2]

Consider an example of a physical setup based on example E_4 above:

(S_1) (1) The engine (with the variable feature whether it sputters or not) and (2) the fuel-system (with the variable feature whether it contains water or not), with (1) and (2) organised such that fuel from the fuel tank is pumped into the engine via the fuel system and located in time.

Physical causal claims can be seen as referring to physical setups in the following way: general (type level) physical causal claims are about collections (types) of physical setups; token physical causal claims (which we do not discuss in the paper) are about individual setups.

By means of the concept of physical setup we can formulate the causal claim underlying the remedy claim from E_4 as follows:

(C_1) For all physical setups of the type S_1: the value of the variable W (whether it contains water) of the fuel-system immediately before *t* influences/has an effect on the value of the variable *S* (whether it sputters) of the engine immediately after *t*.

Time *t* refers here to the moment at which the user of the car attempts to start it (i.e. turns the ignition key). The concept of physical setup now allows us to delineate what the causal claim is about. It is a tool to get a grip on the scope of causal claims. But we have to analyse the meaning of this and similar causal claims more thoroughly: how can we further characterize “influences/has an effect on” ? In Section 3-5 we gradually develop our answer and illustrate it with examples.

2.2 *Properties of our definitions*

We will work with a series of definitions that capture different types of causation. The definitions in Section 3 are the weakest ones, because they are motivated by the efficiency requirement only. The definitions in later sections are stronger, because they are motivated by the efficiency requirement plus the no harm requirement (Section 4) or by the efficiency requirement plus the maximal assistance ideal (Section 5).

The series of definitions we develop is based on various existing theories of causation: the comparative model of Ronald Giere (1997), the context unanimity theory of Ellery Eells (1990) and the INUS theory of John Mackie (1974). We will extract interesting ideas from these theories and incorporate them into definitions that have a certain standard format that we use. One reason to use the standard format is that it renders the relations between our definitions (how and why one is stronger than the other) crystal clear. Our standard format also has several other advantages:

(I) The relevant type of physical setup is explicitly mentioned (so the claim is not underspecified in the aforementioned sense).

(II) A definition in the standard format is “purely consequential”, i.e. it specifies what *follows from* causal beliefs without making any assumptions about how causal claims are confirmed. In this way the definitions are compatible with many ways in which evidence for causal claims can be gathered.

(III) A definition in the standard format makes clear why causal knowledge is – in principle – practically useful (in the case of TPSIs: solving problems, evaluating the proposed procedures, …).

---

[2] “Physical object” is to be interpreted here in a pragmatic way. It bears no metaphysical implications.

(IV) The standard format takes into account the fact that the truth of causal claims often depends on the alternative cause we have in mind. We illustrate this with an example from Peter Menzies (2007, p. 204-206). Consider three options for administering a drug to a patient: no dose, a moderate 100 mg dose, or a strong 200 mg dose. Suppose we give the patient a moderate dose and he recovers. It is possible that both of the following claims are true:

> Taking the moderate dose (as opposed to no dose) was a cause of the patient's recovery.
> Taking the moderate dose (as opposed to the strong dose) was *not* a cause of the patient's recovery.

If we leave out the specification "as opposed to", we get two claims that seem contradictory but maybe are compatible. To rule out this kind of confusion, definitions should explicitly mention the alternative cause. The same holds for causal claims regarding TPSIs. Suppose that we experience difficulty braking when driving our bike. When we are referring to the type of physical setup S_2:

(S_2) (1) The brakes (with the variable feature whether they are sluggish to release) and (2) the cable (with the variable feature whether it is gritty, clean or lubricated), in all cable operated bicycle brake systems[3].

both causal claims can be true:

> The cable being clean (as opposed to lubricated) is a cause of the braking difficulties.
> The cable being clean (as opposed to gritty) is not a cause of the braking difficulties.

Again, if we do not specify what state we consider as the alternative, we get two claims that seem contradictory.[4]

---

[3] Cable operated brake systems encompass all brake systems where pulling the cable engages the brakes, viz. makes the bike brake. This includes e.g. rim brakes and roller brakes.

[4] Note that "clean" here is to be interpreted strictly, as the state in which there are no substances present on the surface of the cable inside the sleeve. A lubricated cable has greasy material on its surface and thus is not clean in this sense.

3. Remedy claims and positive/ negative causal factors

In this section we investigate what kind of causal knowledge is necessary to satisfy the efficiency requirement. We do this in four steps. Section 3.1 contains definitions of two concepts : positive causal factors (PCF) and negative causal factors (NCF). These definitions are based on the work of Ronald Giere on the meaning of causal claims. In Section 3.2 we clarify how our definitions relate to Giere's comparative model of causation. In Section 3.3 we show that our definitions have the properties specified in 2.2. Finally, in 3.4 we argue that, in order to satisfy the efficiency requirement, remedy claims must be based on either positive causal factors or negative causal factors as defined in 3.1.

3.1 *Positive and negative causal factors*

Here is the first definition, 'PCF' (Positive Causal Factor):

(PCF) C (as opposed to C*) is a positive causal factor for E in the collection of setups U if and only if $\mathbf{P}_X(E)$ is greater than $\mathbf{P}_K(E)$.

C and C* are mutually exclusive but not necessarily jointly exhaustive (for instance: a clean cable as opposed to a lubricated cable; there are other possibilities such as a gritty cable). The collection of setups (U) will consist of all individual setups that belong to a certain type (S). X is the hypothetical collection of setups which is obtained by changing, for every individual setup in U and does not exhibit the value C, the value into C. K is the analogous hypothetical collection of setups in which all individual setups that do not exhibit C* are changed into C*. $\mathbf{P}_X(E)$ and $\mathbf{P}_K(E)$ are the probability of E in X and K respectively.

An example might clarify this. If we claim that water in the fuel system (C) as opposed to only fuel in the fuel system (C*) is a positive causal factor for a sputtering engine (E) in the physical setups of U (setups of type S_1), this amounts to claiming that if we poured water into every fuel system that is part of a physical setup in U there would be more sputtering engines than if every fuel system of the setups in U was completely filled with fuel.

Our second definition (Negative Causal Factor) is the negative counterpart of the first:

(NCF) C (as opposed to C*) is a negative causal factor for E in the collection of setups U if an only if $\mathbf{P}_X(E)$ is less than $\mathbf{P}_K(E)$.

The four remaining definitions (section 4 and 5) are reinforcements of the first. It is possible to construct reinforcements of the above definition of negative causal factorhood. We will not do this, because we do not need these additional definitions.

3.2 *Our use of Giere's comparative model*

As mentioned above definitions (PCF) and (NCF) are based on the comparative model which Giere developed mainly to analyse the meaning causal claims in the biomedical sciences. The following definitions constitute the core of this model:

> C is a *positive causal factor* for E in the population U whenever $\mathbf{P}_X(E)$ *is greater than* $\mathbf{P}_K(E)$.
> C is a *negative causal factor* for E in the population U whenever $\mathbf{P}_X(E)$ *is less than* $\mathbf{P}_K(E)$.
> **C** is *causally irrelevant* for **E** in the population U whenever $\mathbf{P}_X(E)$ *is equal to* $\mathbf{P}_K(E)$. (1997, p. 204)

Most of Giere's examples come from the biomedical sciences. So the population U mostly is a subclass of human beings, e.g. all Americans or all women in Germany. He considers only binary variables. This is an important difference with our definitions, to which we come back below. In

Giere's definitions, **C** is a variable with two values (C and not-C); the same for **E** (values E and not-E). For Giere, X is the hypothetical population which is obtained by changing, for every member of U that exhibits the value not-C, the value into C. K is the analogous hypothetical population in which all individuals that exhibit C are changed into not-C. $\mathbf{P}_X(E)$ and $\mathbf{P}_K(E)$ are the probability of E in respectively X and K. Probabilities are defined as relative frequencies (Giere takes U to be finite, i.e. causal claims are about finite populations).

Let us give an example. If we claim that smoking (**C**) is a positive causal factor for lung cancer (**E**) in the Belgian population (U), this amounts to claiming that if every inhabitant of Belgium were forced to smoke there would be more lung cancers in Belgium than if everyone were forbidden to smoke. Conversely for the claim that smoking is a negative causal factor. Causal irrelevance is a relation between variables (represented in bold) rather than a relation between values of a variable (like the first two relations). If we claim that "smoking behaviour" (**C**) is causally irrelevant for "the incidence of lung cancer" (**E**) this means that we believe that in the two hypothetical populations the incidence of lung cancer is equally high.

It will be clear that the idea of hypothetical sets that we have used in 3.1 stems from Giere. We also preserve the notation he uses for the relevant probabilities that are compared in the definitions: $\mathbf{P}_X(E)$ and $\mathbf{P}_K(E)$. $\mathbf{P}_X(E)$ should not be confused with $\mathbf{P}_U(E|C)$. The latter is the relative frequency of E in the subclass C of the *actual* population U, while the former is the relative frequency of E in the hypothetical population X (which is defined starting from U but certainly not identical to U or to $U \cap C$). In our smoking example, $\mathbf{P}_U(E|C)$ would be relative frequency of lung cancer in the subclass of actual smokers in the real population of Belgium. $\mathbf{P}_X(E)$ is the relative frequency of lung cancer in the hypothetical population in which every inhabitant of Belgium were forced to smoke. Put differently: a difference between $\mathbf{P}_U(E|C)$ and $\mathbf{P}_U(E|\neg C)$ entails that the variables are correlated. In order to have causation we need something else: a difference between $\mathbf{P}_X(E)$ and $\mathbf{P}_K(E)$.

To avoid misunderstanding, we have to provide some further clarification regarding the hypothetical populations. The idea underlying X and K is that they differ only with respect to the value of **C** and the values of all variables causally downstream of **C**. This is related to the way we 'change' the value of **C** in order to obtain the hypothetical populations. Giere does not specify what this change entails, but we define it as a surgical change following Jim Woodward:

> Let X and *Y* be a variables, with the different values of X and *Y* representing different and incompatible properties possessed by the unit *u*, the intent being to determine whether some intervention on X produces changes in *Y*. Then I is an intervention variable for X with respect to *Y* if and only if I meets the following conditions:
>
> (**IV**)
> I1. I causes X.
> I2. I acts as a switch for all the other variables that cause X. That is, certain values of I are such that when I attains those values. X ceases to depend on the values of other variables that cause X and instead depends only on the value taken by I.
> I3. Any directed path from I to *Y* goes through X. That is, I does not directly cause *Y* and is not a cause of any causes of *Y* that are distinct from X except, of course, for those causes of *Y,* if any, that are built into the *I-X-Y* connection itself; that is, except for (a) any causes of *Y* that are effects of X (i.e. variables that are causally between X and *Y)* and (b) any causes of *Y* that are between I and X and have no effect on *Y* independently of X.
> I4. I is (statistically) independent of any variable Z that causes *Y* and that is on a directed path that does not go through X.
> ("Cause" in this characterization always means "contributing cause" rather than "total cause. ")
> Given the notion of an intervention variable, an *intervention* may be defined as follows:

> (**IN**) I's assuming some value $I = z_i$, is an intervention on X with respect to *Y* if and only if I is an intervention variable for X with respect to *Y* and $I = z_i$ is an actual cause of the value taken by X.
> (Woodward 2003, p. 98)

If we define 'change' by means of Woodward's notion of intervention, the hypothetical populations X and K do indeed only differ with respect to the value of **C** and all variables causally downstream of **C**.

The differences between our definitions and Giere's are the following:

(a) We use "if and only if" while Giere uses "whenever" which is a conditional in one direction only. Because of the complementarity of Giere's three definitions (i.e. the fact that the three relations are jointly exhaustive and mutually exclusive) our biconditional formulation is in fact equivalent to the original formulation. Given that we will focus on positive causal factorhood, it is better for us to use a biconditional formulation. This formulation avoids confusion when one definition is used in isolation from the other two.[5]

(b) Giere refers to populations, which is due to his interest in biomedical sciences, while our definitions are phrased in terms of collections of setups. The latter is of course due to our interest in technical problem solving instructions.

(c) The third difference is (philosophically) the most important. Giere does not have explicit contrasts in the cause variable. C is, by default, compared with not-C. Our definitions are more fine-grained because different substates of not-C may be used to make causal claims. Recall our example at the end of in Section 2.2, for which the following claims may be true:

> The cable being clean (as opposed to lubricated) is a positive causal factor of braking difficulties.
> The cable being clean (as opposed to gritty) is not a positive causal factor of braking difficulties.

It is important to notice that by introducing contrasts we do not risk any form of relativism. Contrasts allow us to make more precise causal claims. For instance, with respect to smoking we can claim:

> Smoking 20 cigarettes a day (as opposed to none) is a positive causal factor for lung cancer.
> Smoking 20 cigarettes a day (as opposed to 100) is a negative causal factor for lung cancer.

These have very different meanings. The first means that $\mathbf{P}_X(E) > \mathbf{P}_{K\text{-}0}(E)$, where X is the hypothetical population where everyone smokes 20 cigarettes a day and K-0 the hypothetical population where no one smokes cigarettes. The second means that $\mathbf{P}_X(E) < \mathbf{P}_{K\text{-}100}(E)$, where X is in the first claim while K-100 is the hypothetical population where everyone smokes 100 cigarettes a day. The two claims have different truth conditions, but the truth condition of each claims is mind-independent.

3.3 *Properties of definitions (PCF) and (NCF)*

It is immediately clear that definitions (PCF) and (NCF) have the characteristics mentioned in (I), (II) and (IV) in Section 2.2. What about (III)? As we have seen, an important feature of Giere's original model is that he defines causation in terms of what would happen in two *hypothetical* populations. In this way the policy relevance of biomedical causal claims becomes clear. Why should policy makers

---

[5] If you believe that C is a positive causal factor for E in a population U, you cannot believe at the same time that C is a negative causal factor for E in U. Then (according to Giere's one-sided definition of negative causal factorhood) you have also to reject that $\mathbf{P}_X(E) < \mathbf{P}_K(E)$. A similar line of reasoning based on the definition of causal irrelevance leads to the rejection of $\mathbf{P}_X(E) = \mathbf{P}_K(E)$. Hence, you are forced to accept that $\mathbf{P}_X(E) > \mathbf{P}_K(E)$: this is the only option left. In this way it can be shown that Giere's definitions that use "whenever" jointly entail that in each case the other direction is also valid.

want causal knowledge? The hypothetical populations X and K correspond to populations a policy maker may create by means of some direct intervention (e.g. a ban on smoking, a mandatory inoculation, ...)[6]. This feature is preserved in our definitions: they define causation in terms of what would happen in hypothetical sets of setups. Contrary to the policy maker, the TPSI user does not want to intervene on the entire set of setups. He is interested in the individual level: he wants to fix his specific setup, e.g. (a part of) his car. Knowledge about positive causal factors can still guide him in this quest. A person following the TPSI may create a member of the hypothetical sets X, by means of some direct intervention (e.g. removing the water from the fuel tank).

3.4 *Implications of the efficiency requirement*

In example E4, the instruction is to drain the tank and blow out the fuel lines with compressed air. Suppose we find this instruction in a manual, while on the other hand we have good reasons to believe that the following claim is false:

> Containing only fuel in the fuel tank (C) as opposed to fuel and water (C*) is a positive causal factor for a normally running engine, in collections of setups of type S_1.

In such case, the manual proposes a useless intervention and thus the efficiency requirement is violated. In order for the instruction “drain the tank” to be suitable for inclusion (given the aims of the manual) this claim about positive causal factorhood must be true.

Similarly, in E5 the instruction to recharge the battery is suitable for inclusion only if the following claim is true:

> The battery being fully charged (C) as opposed to only partially charged (C*) is a positive causal factor for a starter motor turning the engine at normal speed, in collections of setups of type S_3.

Where S_3 refers to:

(S_3) (1) the battery (with the variable feature whether it is charged) and (2) the starter motor (with the variable feature the speed with which it turns), in all cars.

Some remedy instructions rely on negative causal factorhood. We give two examples. In E1 the instruction to clean the air filter is suitable for inclusion if the following is true:

> A clean air filter (C) as opposed to a dirty one (C*) is a negative causal factor for excessive fuel consumption, in collections of setups of type S_1.

In E1’ the instruction to adjust tyre pressure relies on the truth of the following claim:

> Properly pressured tyres (C) as opposed to underinflated tyres (C*) is a negative causal factor for excessive fuel consumption, in collections of setups of type S_1.

These four examples support the thesis that, in order to suitable for inclusion, remedy instructions need to be backed up by a true positive causal factor claim (where ‘positive causal factor’ is defined as in PCF) or a true negative causal factor claim (where ‘negative causal factor’ is defined as in NCF). If these claims are false, the remedy instructions are inadequate.

The constraint that is imposed here on remedy instructions is rather weak, because PCF and NCF are definitions that are not very demanding. There is a prima facie reason to look for stronger

---

[6] Of course, there are often ethical and practical limitations here.

constraints: the no harm requirement. The truth of type level claims as defined by PCF and NCF is compatible with adverse effects in certain subsets. Sticking to the engine examples: the truth of the positive or negative causal factorhood that backs up a remedy instruction is compatible with a situation in which you cause serious damage to your car's engine by executing the instruction. This is the main topic of Section 4.

4. Remedy claims and context unanimity

We started with definitions based on Giere's because this results in rather weak definitions. Stronger definitions (that correspond to crucial ideas of Eells and Mackie) can be easily obtained later by adding constraints. In this section we us an important idea of Eells, viz. context unanimity, to clarify how the no harm requirement can be satisfied. In 4.1 we explain what context unanimity is and why it is not present in Giere's model. In 4.2. we incorporate the idea of context unanimity in our analysis by means of two definition which are reinforcements of (PCF). In 4.3 and 4.4 we relate these two definitions to the no harm requirement.

4.1 *Average effect versus context unanimity definitions of causation*

Giere's original comparative model and our adaptation can be characterised as "average effect" definitions of causation in sets of physical setups and biological populations. Let us look at an example from the biomedical sciences to clarify what this means. Consider a dangerous virus, which threatens a population of humans (H). Some people are immune to the disease (I), but there is no way to find out who is and who is not. It is possible to vaccinate people before they become sick (V). We assume the following probabilities in the hypothetical populations (S stands for survival):

$P_V(S|I) = 0.9$
$P_V(S|\neg I) = 0.8$
$P_{\neg V}(S|I) = 1$
$P_{\neg V}(S|\neg I) = 0$

Furthermore, we assume that 50% of the population is immune, so we also have:

$P_V(S) = 0.85 \qquad (0.8\times 0.5 + 0.9\times 0.5)$
$P_{\neg V}(S) = 0.5 \qquad (1\times 0.5 + 0\times 0.5)$

Note that in the subpopulation I of people that are immune, vaccination is a negative causal factor: 10% of this subpopulation would not survive vaccination. In subpopulation ¬I and in H as a whole, vaccination is a positive causal factor. How is this possible? In I there is a group of people (10% of I) whose residual state is such that they die if vaccinated. For the others, vaccination is causally irrelevant at the individual level. Combined, this gives negative causal relevance at the level of subpopulation I. In subpopulation ¬I we have a large group (80%) whose residual state is such that vaccination is positively causally relevant at the individual level. For the others, it is irrelevant (their residual state is such that they die anyway). The combination of this gives positive causal relevance at the level of subpopulation ¬I. The population H contains a group of 5% (10% of the 50% immune) for whom vaccination has negative causal relevance at the individual level. It also contains a group of 40% (80% of the 50% non-immune) for whom vaccination has positive causal relevance at the individual level. For the others, vaccination is causally irrelevant at the individual level (vaccination makes no difference for them: they survive anyway because they are immune, or they die anyway because the vaccination does not work for them). Because the group with positive relevance is larger (40% as compared to 5%) the result is a positive causal relevance at the level of population H.

The vaccination example illustrates that, according to Giere’s definitions, causal relevance can be reversed or annihilated in subpopulations: if C is a positive causal factor for E in population U, it can be a negative causal factor or be causally irrelevant in subpopulations of U. The same holds for negative causal factors. Theories of causation which have this property are called “average effect theories”: whether there is a causal relation in a population depends – according to these theories – on the average effect in the population, no matter what happens in its subpopulations.

An alternative to average effect theories are the so-called “context unanimity theories”. The first context unanimity theory can be found in Cartwright (1979). A more recent version can be found in Eells (1991). In chapter 2 of his book, Eells gives the following example:

> To use an example of Cartwright’s (1979), ingesting an acid poison (*X*) is causally positive for death (*Y*) when no alkali poison has been ingested (~*F*), but when an alkali poison has been ingested (*F*), the ingestion of an acid poison is causally negative for death. I will argue that in a case like this it is best to deny that *X* is a positive causal factor for *Y*, even if, overall (for the population as a whole), the probability of death when an acid poison has been ingested is greater than the probability of death when no acid poison has been ingested (that is, even if *Pr*(*Y*/*X*) > *Pr*(*Y*/~*X*)). I will argue that it is best in this case to say that *X* is causally *mixed* for *Y*, and despite the *overall* or *average* probability increase, *X* is nevertheless not a positive causal factor for *Y* in the population as a whole. (p. 58)

The characteristic property of causes in the sense of context unanimity theories is that the causal tendency cannot be reversed (from positive to negative) or annihilated (from positive or negative to causally neutral) in a subpopulation. Note that we only borrow the idea of context-unanimity from Eells. This is because Eells considers *actual* probability distributions in defining causation. This is already clear from the quote above: he defines a causal factor and being causally mixed in terms of the actual probability of the occurrence of the effect given presence or absence of the cause (*Pr*(*Y*/*X*) and *Pr*(*Y*/~*X*) respectively). Because we want to preserve the advantages of Giere’s definitions in terms of hypothetical populations (see 3.3) we only borrow Eells’ idea of context-unanimity.

4.2 *Two additional definitions*

The idea of context unanimity can give rise to two reinforcements of PCF: PCF-WU and PCF-SU (where WU stands for weak unanimity and SU for strong unanimity):

(PCF-WU) C (as opposed to C*) is a weakly unanimous positive causal factor for E in the collection of setups U if an only if
(1) $\mathbf{P}_X(E)$ is greater than $\mathbf{P}_K(E)$; and
(2) there are no subgroups S of U for which $\mathbf{P}_{X'}(E)$ is less than $\mathbf{P}_{K'}(E)$.

X’ and K’ are defined in the same way as X and K above, but starting from the subset of setups S instead of the whole set of setups of a certain type U.

(PCF-SU) C (as opposed to C*) is a strongly unanimous positive causal factor for E in the collection of setups U if and only if
(1) $\mathbf{P}_X(E)$ is greater than $\mathbf{P}_K(E)$; and
(2) there is no subgroup S of U for which $\mathbf{P}_{X'}(E)$ is less than *or equal to* $\mathbf{P}_{K'}(E)$.

The difference between the two definitions is that weak context unanimity allows that a causal tendency in the whole population is annihilated in a subpopulation. It only prohibits tendency reversal. Strong context unanimity prohibits both tendency annihilation and reversal.

4.3 *Remedy claims and weak context unanimity*

There is a connection between the no harm requirement and weak context unanimity. In order to argue for this, we imagine a sloppy manual that contains the following TPSI:

> Problem: engine sputters and then stops running.
> Diagnosis: empty fuel tank.
> Remedy: put gasoline in the fuel tank.

The sloppiness consists in the fact that this instruction is given for varieties of the given car type (let us call that S_4), regardless of whether they have a diesel engine or a gasoline engine. Let us assume that 80% of the cars that belong to type S_4 have a gasoline engine. Then the following causal claim is correct:

> The fuel tank containing gasoline (C) as opposed to diesel (C*) is a positive causal factor for a normally running engine, in collections of setups of type S_4.

Indeed, if all fuel tanks in S_4 cars contain diesel, only 20% runs normally. If all fuel tanks contain gasoline, 80% runs normally. Let S_5 be all the cars of type S_4 with a diesel engine. Then the following claim is true:

> The fuel tank containing gasoline (C) as opposed to diesel (C*) is a negative causal factor for a normally running engine, in collections of setups of type S_5.

So S_4 violates weak context unanimity. The following claims is false:

> The fuel tank containing gasoline (C) as opposed to diesel (C*) is a weakly unanimous positive causal factor for a normally running engine, in collections of setups of type S_4.

The fact that there is no weak context unanimity entails that there are contexts (or subgroups of setups - in casu: cars with diesel engines) where the action "put gasoline in the fuel tank" worsens the problem rather than helping to solve it. This is why the proposed remedy "put gasoline in the fuel" is not suitable. It should not occur (in this form, given the context we have sketched here) in manuals because it violates the no harm requirement.

4.4 *Remedy claims and strong context unanimity*

Can an analogous case be made for strong context unanimity? We do not think so. In order to argue for this, we look at two examples of section 3.4 in which there is no strong context unanimity. The first example is:

> The battery being fully charged (C) as opposed to only partially (C*) is a positive causal factor for a starter motor turning the engine at normal speed, in collections of setups of type S_3.

There are subsets in which the positive difference is annihilated, e.g. in the set of cars with heavily corroded battery terminals. A fully charged battery makes no difference for the behaviour of the engine in setups of that type. So there is no strong context unanimity in this case. The following claim is false:

> The battery being fully charged (C) as opposed to only partially charged (C*) is a strongly unanimous positive causal factor for a starter motor turning the engine at normal speed, in collections of setups of type S_3.

A similar observation can be made with respect to the following example:

> Containing only fuel in the fuel tank (C) as opposed to fuel and water (C*) is a positive causal factor for a normally running engine, in collections of setups of type S_1.

Again there are subsets in which this positive difference is annihilated, e.g. in the set of cars with defective fuel pumps. A properly filled fuel tank makes no difference for the behaviour of the engine in setups of that type. This means that, again, there is no strong context unanimity. The following claim is false:

> Containing only fuel in the fuel tank (C) as opposed to fuel and water (C*) is a strongly unanimous positive causal factor for a normally running engine, in collections of setups of type S_1.

The crucial question now is: is this absence of strong unanimity a problem? No, because no harm can be done if there is weak context unanimity. If there is weak context unanimity, there are no contexts in which the remedy instruction worsens the problem. Moreover, TPSI's are typically part of a larger set of instructions. If one instruction does not lead to a functioning artefact, there often is a follow-up instruction that prescribes another action aimed at solving the problem. In this way, a TPSI not need not be self-contained, but functions as part of a bigger whole.

5. Remedy instructions and sufficient causation

In this section we investigate the implications of the maximal assistance ideal. Remedy claims that are based on true causal claims in the sense of (PCF-WU) do not guarantee that the problem is solved. So maybe remedy claims should be based on knowledge about sufficient causes? In 5.1 we argue that this is not a good idea. In 5.2 and 5.3 we develop an alternative based on Mackie's concept of INUS condition.

5.1 *Sufficient causation?*

In order to guarantee that the problem is solved if the prescribed action is performed, a stronger causal relations has to be present in the world. A concept of causation that captures this stronger causal relation can easily be obtained by adding a sufficiency clause to definition (PCF). This results in the following definition:

(SC) C (as opposed to C*) is a sufficient cause for E in the set of physical setups U if and only if
(1) $\mathbf{P}_X(E)$ is greater than $\mathbf{P}_K(E)$; and
(2) $\mathbf{P}_X(E)=1$.

Note that, according to this definition, a sufficient cause is always a positive causal factor. The sufficiency condition implies that there is weak context unanimity. The reason for this is that, if $\mathbf{P}_X(E)=1$ as required by the second clause in (SC), it will also be the case that $\mathbf{P}_{X'}(E)=1$ for all subsets of U. Put differently: if all objects in X have a property E, then it is also the case, for all subsets of X, that all their members have E.

Should we advise producers of repair manuals to include only remedy claims that are based on true sufficient cause claims? It seems not. The reason is that there are hardly any true claims of this kind available. For instance, the following claims are all false:

> The battery being fully charged (C) as opposed to only partially charged (C*) is a sufficient cause for a starter motor turning the engine at normal speed, in collections of setups of type S_3.
>
> Containing only fuel in the fuel tank (C) as opposed to fuel and water (C*) is a sufficient cause for a normally running engine, in collections of setups of type S_1.

These claims are false because the problems that are to be solved may have multiple causes. The first claim is false because the battery terminal may be corroded. The second claim is false because the fuel pump may be defective.

Because of this scarcity, the requirement that remedy claims are to be based on sufficient causation relations as defined by (SC), is untenable: repair manuals would hardly contain any instructions. And this would be at odds with the maximal assistance ideal.

So we have to allow remedy claims to be based on non-deterministic causal relations. But given the maximal assistance ideal it is useful to investigate whether some weaker alternative is possible. We think there is such an alternative, and we call it "sufficiency in maximally normal contexts". In Section 5.2 we clarify what we mean by this and why it is a desirable property of TPSIs. In section 5.3 we show that the idea can be captured by adding John Mackie's concept of INUS condition to definition (PCF).

5.2 S*ufficiency in maximally normal contexts*

The issue we want to bring up here can be illustrated by means of the water-in-the-fuel-system example. Let us compare two TPSIs. The first is:

> Problem: engine sputters and fails to start.
> Diagnosis: water in fuel supply system.
> Remedy: drain the fuel tank.

The second is:

> Problem: engine sputters and fails to start.
> Diagnosis: water in fuel supply system.
> Remedy: drain the fuel tank and blow out the fuel lines with compressed air.

Which instruction is the best one? In general, draining the fuel tank is not sufficient for solving the problem because there may also be water in the fuel lines. So the second TPSI is more adequate. This indicates that we want the remedies to be complete in some sense: they have to be sufficient for solving the problem *provided that there is nothing wrong with the artefact on top of what is stated in the diagnosis*. In other words, the remedies have to be complete relative to the diagnosis that is part of the TPSI. This is what we call the "maximally normal context": the context where all components – except the one addressed in the diagnostic claim – of the artefact function normally (i.e. as conceived by the designers, and can be legitimately expected to function in a newly produced artefact that has passed the manufacturer's quality control).

5.3 *Mackie causation and its application to TPSIs*

Mackie's theory of causation is situated at the token level, i.e. it is about causal relations between particular events. We are dealing here with causal relations between variables. We first give a brief presentation of Mackie's account and then clarify how we propose to use the crucial INUS concept at the type level in combination with our definition PCF.

In *The Cement of the Universe* Mackie claims that a cause is always an INUS condition for its effect (1974, p. 64). He defines INUS conditions as follows:

> … an *insufficient* but *non-redundant* part of an *unnecessary* but *sufficient* condition: it will be convenient to call this (using the first letters of the italicized words) an *inus* condition. (p. 62)

So in Mackie's view, a cause in itself is not sufficient for its effect, but combined with a set of other factors, it is. Furthermore, a cause is not necessary for its effect, since multiple sets of factors can produce the same effect. Consider one of his examples:

> The short-circuit caused the fire.

As Mackie points out, the short-circuit (*c*) in itself is not sufficient for a given fire (*e*): you also need oxygen and combustible material (A). The factors in (A) are themselves also not sufficient for the fire (*e*), so the short-circuit is a non-redundant part of the condition (A+*c*). Combined, the short-circuit and the factors in A are sufficient for the fire. Yet (A+*c*) is not necessary for a fire to happen, a fire can also occur because of (A+ lightning), (A+ a lit match), etc.

In this paper we are interested in type level causation, while Mackie deals with the token level. However, we can use his crucial INUS idea by adding it as a constraint to definition (PCF) and linking it to maximally normal contexts. This results in our sixth definition MC (Mackie cause):

(MC) C (as opposed to C*) is a Mackie cause for E in the collection of setups U if an only if
(1) $\mathbf{P}_X(E)$ is greater than $\mathbf{P}_K(E)$; and
(2) if A characterises the maximally normal context, it is the case that (a) all members of U that have A and C, also have E, and (b) not all A have E.

Let us give an example based on the following TPSI:

> Problem: starter motor turns engine slowly.
> Diagnosis: partially discharged battery.
> Remedy: recharge.

Let C stand for "fully charged battery" and E for "starter motor turns engine rapidly". C is not a sufficient cause for E, because there may be other malfunctioning components (e.g. corroded battery terminals, bad starter motor, …). But C in combination with A (where A is a description of the maximally normal context, and includes a.o. the proviso that the battery terminals are not corroded and that the starter motor is not broken) is sufficient for E. A manual that aims at instructions that are maximally useful for its users, preferably contains this kind of instructions.

6. Synthesis and further reflections

In this paper we have performed a stepwise analysis of how the success criteria determine the nature of the causal knowledge underlying remedy claims. In order to satisfy the efficiency requirement and the no harm requirement, it is legitimate to expect underlying causal relations causes in the following sense:

(PCF-WU) C (as opposed to C*) is a weakly unanimous positive causal factor for E in the collection of setups U if an only if
(1) $\mathbf{P}_X(E)$ is greater than $\mathbf{P}_K(E)$; and
(2) there are no subgroups S of U for which $\mathbf{P}_{X'}(E)$ is smaller than $\mathbf{P}_{K'}(E)$.

This has been argued in in Sections 3 and 4. In Section 5 we have argued that the maximal assistance ideal motivates a preference for sufficiency in maximally normal contexts:

(MC) C (as opposed to C*) is a Mackie cause for E in the collection of setups U if an only if
(1) $\mathbf{P}_X(E)$ is greater than $\mathbf{P}_K(E)$; and
(2) if A characterises the maximally normal context, it is the case that (a) all members of U that have A and C, also have E, and (b) not all A have E.

In other words: we can legitimately expect positive causal factors that are weakly context-unanimous and Mackie causes are to be preferred. Requiring only positive causal factorhood as defined in (PCF) is not enough. We need something stronger in order to satisfy the success criteria.

Note that both concepts are needed. It is obvious that, if a purported cause satisfies (PCF-WU), it may fail to satisfy (MC). Though less obvious, the opposite may also be the case; state C that in the maximally normal context A is sufficient for solving a problem, may be harmful in other contexts (i.e. if there are other problems with the artefact besides the one mentioned in the diagnosis).

In light of the maximal assistance ideal, stronger requirements or preferences cannot be justified. We have argued in Section 5.1 that including only remedy claims based on sufficient causation (as defined in (SC)) is not an optimal strategy. The manuals would be almost empty. A similar argument can be made with respect to strong context-unanimity (as defined in (PCF-SU)). In Section 4.4 we have argued already that strong context-unanimity it is not necessary in order to avoid harm. However, we can go one step further: if composers of repair manuals would restrict themselves to instructions based on strongly unanimous causal relations, they would not provide maximal assistance to users because many adequate remedy claims would be left out. For instance, from the perspective of maximal assistance, remedies such as "recharge the battery" or "drain the fuel tank" should be included (see 4.4).

In the introduction we have labelled some criteria "requirements" while one was called an "ideal". Now that we have used these in our analysis, it should be clear why we have used these different labels. The efficiency requirement and the no harm requirement are either satisfied or violated. The maximal assistance ideal, on the other hand, is a regulative principle which is important for the optimisation of the content.

Optimising the content of a manual requires expert judgement. As we have argued, Mackie causes are desirable but there are other factors that may influence the decisions about which remedy claims are included and how precise they should be. Since repair manuals target a non-expert, non-professional audience, feasibility may be such a factor. Manuals typically prescribe actions which are relatively easy to perform (like recharging a battery, cleaning, draining, blowing, ….). Another trade-off may be specificity versus intelligibility: by specifying, the success rate of an instruction may rise, but at the same time, it might make the instruction harder to understand. Our analysis can therefore be interpreted as a means to make explicit *some* of the criteria used by writers of repair manuals.

7. An additional example: radios

The examples we gave in Sections 3 till 5 were about cars and bicycles. Though there is no reason to suspect that our analysis is somehow limited to these means of conveyance, it is good to have some positive evidence for the validity of our analysis in other technical domains. The domain we have chosen is that of radios. In *Old Time Radios! Restoration and Repair* (Carr 1990), we find instructions like:

> If the […] cone is badly torn or warped […] then the sound produced by the speaker will be distorted. Replace the loudspeaker with a new speaker, or recone the old one if a new model is not easily available. (p. 203)

Our framework can help to understand and represent what exactly is being said in this example as well. We can isolate a specific TPSI:

Problem: the sound produced by the speaker is distorted.
Diagnosis: the cone of the loudspeaker is badly torn or warped.
Remedy: replace the cone with an intact one or recone the old one.

According to our analysis at the very minimum the following causal claim has to be true to make the remedy claim warranted:

(C_2) An unwarped, untorn cone (C) instead of a warped or torn cone (C*) is a weakly unanimous positive causal factor for a loudspeaker that produces a clear sound (E) for the setups of type S_6.

where

(S_6) (1) The cone (with the variable features whether it is (a) torn and (b) warped) and (2) the loudspeaker (with the variable feature whether it produces distorted sound), in all radios with conic loudspeakers.

This is because, according to our analysis, remedy claims have to be based on weakly unanimous positive causal factors. Definition (PCF-WU) applied to this example entails that the following conditions must be satisfied:

(1) $\mathbf{P}_X(E)$ is greater than $\mathbf{P}_K(E)$; and
(2) there is no subgroup S of U for which $\mathbf{P}_{X'}(E)$ is smaller than $\mathbf{P}_{K'}(E)$.

In these conditions X is the set of setups of type S_6 with unwarped, untorn cones, K is the set of S_6-setups with warped or torn cones and U the set of all S_6-setups.

Each of conditions (1) and (2) corresponds to a potential reason for removing or revising the considered instruction:
(1’) If replacing the cone with a new intact one would not increase the possibility of a properly functioning radio, the instruction is not adequate
(2’) If, for *some* radios, replacing the cone actually makes the situation worse e.g. produces a more distorted sound or no sound at all, the instruction is not adequate.

According to our analysis remedy claims are preferably based on Mackie causes. So ideally, a third condition is satisfied:

(3) if A characterises the maximally normal context, it is the case that (a) all members of U that have A and C, also have E, and (b) not all A have E.

Applied to the example, this third condition comes down to requiring that if the radio is functioning properly on all fronts except the torn or warped cone, replacing the cone (without any further additional actions) will ensure that the sound is no longer distorted. So if the radio does not suffer from other problems (e.g. dirty tube socket connections, open capacitors,…) that can hinder its functioning, the ideal outcome is that, after our intervention, the radio functions properly. Otherwise, the instruction is not optimal: we are still left with a malfunctioning radio, despite the fact that there are no other problems diagnosed. In that case there is better, more elaborate instruction that should be included instead of the current one.

8. Concluding remarks

In this paper, we first studied the required properties of causal relations underlying remedy claims. In order for remedy claims to be adequate, they have to be based on causal relations that (i) hold in the world and (ii) have certain properties: positive or negative causal factorhood with weak context unanimity. That these properties are required follows from two success criteria: the efficiency requirement and the no harm requirement. Positive and negative causal factorhood were explicated in definitions (PCF) and (NCF), which were inspired by Giere's comparative model of causation. Weak context unanimity was explicated in definition (PCF-WU) hast its origins the idea of context unanimity as put forward by Cartwright and Eells.

In the second part of this paper we studied the desirable properties of causal relations underlying remedy claims. We introduced the idea of sufficiency in maximally normal contexts and explicated it in definition (MC) which is the result of adding Mackie's INUS idea to Giere's comparative model. We used the criterion of maximal assistance to argue that Mackie causation is a desirable property for causal relations underlying remedy claims.

There are many other interesting topics related to causation and engineering, but they cannot be discussed in this article because a full treatment of them requires a separate papers. For instance, causal cognition plays an crucial role in failure analysis, the branch of engineering sciences that deals with artefact failure and its causes. This topic is discussed in one of our other papers [Author 1, article A]. Causation also plays an important role in engineering design (e.g. reverse engineering). This topic is discussed in [Author 2, Article B].

References

Carr Joseph (1990) *Old Time Radios! Restoration and Repair*. McGraw Hill Professional.

Cartwright Nancy (1979), 'Causal Laws and Effective Strategies', in *Noûs* **13**, pp. 419-437.

Chilton Enterprises (1986), *Chilton Book Company Repair & Tune-up Guide: All U.S. and Canadian 305, 327, 350, 396, 427, 454 Models, Including Fuel Injection. Corvette 1963 to 1983.* Radnor: Chilton Book Company.

Eells Ellery (1991), *Probabilistic Causality.* Cambridge: Cambridge University Press.

Giere Ronald (1997), *Understanding Scientific Reasoning* (4th edition). Fort Worth: Harcourt Brace College Publishers.

Lumley John (1999), *Engines: An Introduction.* Cambridge University Press.

Mackie John (1974), *The Cement of the Universe: a Study of Causation.* Oxford: Clarendon Press.

Mead John and A. K. Legg (1997), *Toyota Carina E Service and Repair Manual*. Haynes Service and Repair Manual Series: Haynes Manuals Inc.

Menzies Peter (2007), 'Causation in Context', in H. Price & R. Corry (eds.), *Causation, Physics and the Constitution of Reality. Russell's Republic Revisited.* Oxford: Oxford University Press, pp. 191-223.

Sidwells Chris and Richard Ballantine (2004), *Bike Repair Manual.* London: Dorling Kindersley Limited.

Steel Daniel (2008), *Across the Boundaries. Extrapolation in Biology and Social Science*. New York: Oxford University Press.

Strasman Peter (1988), *Rover 213 & 216 Owners Workshop Manual.* Haynes Service and Repair Manual Series: Haynes Manuals Inc.

Von Wright Georg Henrik (1971), *Explanation and Understanding.* Ithaca, New York: Cornell University Press.

[2 references removed for anonymity]