Available online at www.sciencedirect.com
ScienceDirect
Stud. Hist. Phil. Sci. 39 (2008) 112–123
Studies in History and Philosophy of Science
www.elsevier.com/locate/shpsa

Causally productive activities

Jim Bogen

Department of History and Philosophy of Science, University of Pittsburgh, 1017 Cathedral of Learning, Pittsburgh PA 15260, USA

Received 31 May 2006; received in revised form 11 January 2007

Abstract

This paper suggests and discusses an answer to the following question: What distinguishes causal from non-causal or coincidental co-occurrences? The answer derives from Elizabeth Anscombe's idea that *causality* is a highly abstract concept whose meaning derives from our understanding of specific causally productive activities (e.g., pulling, scraping, burning), and from her rejection of the assumption that causality can be informatively understood in terms of actual or counterfactual regularities.

Keywords: Elizabeth Anscombe; Causality; Explanation; Inhibition

1. Hume is dead

Elizabeth Anscombe's 1970 Cambridge inaugural address dismissed Humean and other philosophical theories that try to reduce the notion of causality to notions of regularity. Causality is one thing, she said, and regularity, another. If A is caused by B,

> this does not imply that every A–like thing comes from some B–like thing or set up, or that every B–like thing or set up has an A–like thing coming from it … Any of these may be true, but if any is, that will be an additional fact, not comprised in A's coming from B … (Anscombe, 1981, p. 136)

Causes that operate in regular ways to produce an effect, or to make its occurrence highly probable, are typically of more theoretical and practical interest to us. In many cases we learn about causes from empirically established regularities. But even so, causality is one thing and regularity is another. Today many philosophers would agree: although some of them would do so only on the proviso that if tokens of B sometimes fail to produce tokens of A, then there must be regularities that conspire behind the scenes to account for the failures. Anscombe also rejected the idea that to cause is to necessitate. The common core of our ideas about causality, she said, is that effects 'derive' or 'arise' or 'come' from their causes, not that given a cause its effect could not possibly fail to occur (ibid.). If a batter hit a home run over the centerfield fence, it is immediate that nothing interfered to prevent the ball from clearing the fence after it was hit. But as Anscombe would insist, and as most people would agree, it does not follow that it was *impossible* for anything to interfere. Thus hitting the ball caused the ball to clear the fence without necessitating it.[1] Furthermore, many philosophers and scientists believe causes operate probabilistically, at least at the quantum

E-mail address: rtjbog@comcast.net
[1] Cf. Russell (1971), p. 135 n.

0039-3681/$ - see front matter © 2007 Elsevier Ltd. All rights reserved.
doi:10.1016/j.shpsa.2007.11.009

level, and some believe that even some macroscopic causal interactions are genuinely indeterministic.[2]

More controversially, Anscombe claimed that one cannot explain what it is for one thing to cause another without appeal to concepts that are themselves irreducibly causal. Suppose the tent you are camping in collapses when a wind comes up. It could be that the wind caused the tent to collapse. It could be that the tent collapsed for other reasons and it was only a coincidence that the wind was blowing when it collapsed. Anscombe would say the difference between these possibilities cannot be explained without invoking irreducibly causal concepts. This was far too strong a claim for her 1970 audience, but now Cartwright (1983), Hausman (1998), Machamer et al. (2000), Pearl (2000), Spirtes et al. (2000), Woodward (2003) and some others agree.

So much for agreement. As I understand her, Anscombe would say that if the tent's collapse was caused, rather than coincidentally accompanied, by the wind, then that is because the wind blew it down. Banal as that sounds, it raises issues that are as controversial now as they were in 1970. *Blowing something down*, like *scraping*, *pushing*, *wetting*, *carrying*, *eating*, *burning*, *knocking over*, *keeping off*, *squashing*, *hurting*, *making noises*, and *making paper boats* (Anscombe's own examples) are highly specific causal concepts (Anscombe, 1981, p. 137). Anscombe held that general causal concepts like *causing*, *bringing about*, *making happen*, *intervening*, *resulting from*, and the like, are placeholders that derive their meanings from specific causal concepts.[3] If we understand words that signify *scraping*, *squashing*, etc., we can understand general terms like 'cause'. But if a language had no words for specific causes, it could not have, and we could not introduce into it, a word that means what we mean by 'cause' (ibid.).

If blowing down the tent is what distinguishes a causal co-occurrence of wind blowing and tent collapse from a non-causal or coincidental co-occurrence, what makes what the wind does a causally productive activity in the case in which it causes the collapse? Anscombe said that

> causality consists in the derivativeness of an effect from its cause. This is the core, the common feature, of causality in its various kinds. Effects derive *from*, arise out of, come of [sic] their causes. (Ibid., p. 136)

But not every derivation, arising from, or coming from, is causal: premises don't cause the mathematical results mathematicians derive from them; cars that come from Japan are not caused by Japan; people are not caused by the beds they arise from in the morning. You might want to say that derivations, arisings, and comings from are causal only when what is derived, etc., is an effect which is *produced* by what it derives from, arises from, or comes of. But *production* is every bit as abstract a causal notion as *cause*. If our understanding of what it is for one thing to produce another derives from our understanding of specific causally productive activities, then it's uninformative to say, for example, that what makes X a causally productive activity is that an effect derives from it and that 'Y derives from X' means that X *produced* Y. What makes the idea that causally productive activities distinguish causal from non-causal sequences of events controversial is that there is no informative general condition which discriminates causally productive activities from goings on which are not causally productive of the effect of interest.

To make this vivid, consider how it applies to a counterexample Wesley Salmon used to argue against Hempelian accounts of causal explanation. Hempel believed that to explain an effect is to deduce a canonical description of it from law-like generalizations and descriptions of initial conditions. Salmon's counterexample is a deduction whose conclusion is the explanandum sentence 'John Jones avoided becoming pregnant', and whose premises are the initial condition sentence 'John Jones has taken his wife's birth control pills regularly' and the law-like generalization, 'every male who regularly takes birth control pills fails to become pregnant'. Salmon's objection to Hempel is that this deduction is valid but non-explanatory (Salmon, 1971, p. 34).

An Anscombian response to Salmon's example would appeal to the empirical fact that in the male body the chemicals in the pills Jones took cannot engage in any activities that prevent pregnancy. Oral contraceptives typically include several different chemicals, each of which interferes with a different biological system. One chemical suppresses hormones whose activity is required for the release of eggs from the ovary. Another acts on the cervical mucous to thicken it so that sperm cannot move through it. Another acts on the lining of the uterus to make it unsuitable for the implantation of a fertilized egg. Another kills sperm. The reason the pills don't prevent pregnancy in men is that men have no ovaries, or uteri for the chemicals to act on, and no eggs for sperm to fertilize. The reason men do

[2] Epidemiologists, psychologists, social scientists, and others find it convenient to use deterministic models to study causal systems even when they believe they operate probabilistically. These models treat failures of co-occurrence among probabilistically dependent events as due to the influence of unmeasured exogenous factors. This commits the investigator to the usefulness of deterministic models, but not to deterministic causal relations among the items of the systems they are used to investigate. For an illustration of how genuinely indeterministic models can be constructed see Steel (2005), p. 13.

[3] I thank Ted McGuire for pointing out that Austin made a similar suggestion with regard to the phrase 'doing an action' as philosophers used it in his time (Austin, 1970, p. 178).

not become pregnant is that the male reproductive system is composed of parts which cannot engage in the causally productive activities they would have to engage in to produce pregnancy.[4] This treatment of Salmon's example assumes there is no need to explain, for example, why thickening cervical mucous causes rather than non-causally accompanies a change in its consistency.

Birth control chemicals do engage in *some* causally productive activities in men who ingest them. They dissolve in the stomach, pass through the small intestine and travel in the blood stream into the tissues it irrigates, interacting in various ways with molecules they encounter along the way. Why don't those activities qualify as causes of John Jones's non-pregnancy? According to Jim Woodward's counterfactual dependence theory of causality the answer depends upon the fact that non-pregnancy in males is not counterfactually dependent upon the activities the birth control chemicals engage in when males ingest them. By contrast, there is a counterfactual dependency between non-pregnancy in women and the activities of birth control chemicals they ingest. In both cases, the relevant dependencies obtain only if non-pregnancy *would have* resulted in idealized versions of John and his fertile wife Joan *if* ideal interventions had occurred to promote or damp the activities of the chemicals in their bodies (Woodward, 2004, pp. 48 n.).

The Anscombian alternative appeals not to counterfactuals but to facts about the chemicals and the causally productive activities they engage in. We might know enough about the chemicals and their activities to draw counterfactual conclusions about what would have resulted from ideal interventions on them. But if we know enough to evaluate the counterfactuals Woodward proposes as conditions for causal relevance, we should be able to implement an Anscombian treatment of Salmon's case without invoking them. If you appreciate the fact that John Jones has no uterus, ovaries, or eggs you needn't appeal to counterfactuals to explain why the chemicals which suppress hormones, thicken cervical mucous, and kill sperm aren't responsible for his non-pregnancy.[5]

According to Anscombe, it is a brute fact, for example, that the activity of pulling on a door can open it rather than coincidentally accompanying a door opening, that scraping a carrot removes rather than coincidentally accompanies the removal of its skin, that wetting something does not merely accompany its getting wet, and so on. If she's right, then it's misguided to think we need a general account of causality like Hume's, Hempel's, or Woodward's to explain why any specific causally productive activity is causally relevant to the production of an effect. Instead of suggesting a general, uniformly applicable, answer to the question of what differentiates causes from non-causes, what I take to be Anscombe's view calls for piecemeal treatments of questions about how specific effects are produced. Accordingly, whether a given factor made a causal contribution to the production of a given effect depends upon contingent, empirically checkable facts about how the relevant things behave, and what results from their behavior.

To some, this may seem tantamount to conceding that there is nothing philosophically interesting to be said about causality. I think that philosophical or not,[6] the Anscombian treatment of Salmon's example exemplifies a perfectly satisfactory response to the philosophical question of what differentiates causes from non-causal factors that coincidentally accompany an effect.

2. Some questions about Anscombian causality

There are some recognizably philosophical questions for an Anscombian to answer. The first have to do with the concept of a causally productive activity. Anyone who rejects the received general conditions for distinguishing causes from non-causal factors had better be able to supply something to take their place. Unless there are principled constraints on what can legitimately be called a causally productive activity it is vacuous to say that causes are things that engage in causally productive activities. If (as said) general notions like *productivity* don't illuminate the difference between activities which are causally productive of a given effect and goings on which are not, how then can there be principled decisions about whether something should be considered a causally productive activity?

[4] Anscombe's views are congenial to the Mechanist account of causal explanation advocated by Machamer, Darden, and Craver (MDC). What MDC mean by a mechanism is an organized collection of things, each of which contributes to the production of an effect by engaging in a causally productive activity. In order for the mechanism to produce an effect, its parts and their activities must be spatially and temporally organized in such a way that the mechanism can proceed in an orderly way from its start up condition through one or more stages until it reaches an end state marked by the occurrence of the effect. Mechanists hold that adequate causal explanations help us understand why an effect occurred by answering questions about the parts of the mechanism that produced it, the causally productive activities in which they engage, and how they contribute to the production of the effect of interest. Their examples of causally productive activities (invoked to explain neuronal signaling) include such things as releasing neuro-transmitters, chemical binding, depolarizing neuronal membranes, electrical attraction and repulsion, ion channel gating, pumping, rotating proteins to change their positions relative to other components of ion channels, moving structures in and out of a channel pore, and so on (Machamer, Darden, & Craver, 2000, pp. 3, 5–6, 12). It is apparent from these examples that MDC's examples of causally productive activities are the ontological counterparts of the causal concepts in terms of which Anscombe said we were to understand causality. But not all Mechanists accept Anscombe's view—at least as I have interpreted it. Stuart Glennan holds that an interaction, A, is causally productive of a change of property, E, only if a lawful counterfactual regularity obtains among A-type and E-type events (Glennan, 2002, pp. 344–345).

[5] Mark Schroeder pointed out in discussion that counterfactual conditionals might belong to the truth conditions of a causal claim even though we don't evaluate the former in order to evaluate the latter. But that's not to say we have any good reason to think they do.

[6] The phrase 'philosophical or not' is misleading to the extent that it suggests that there is a fixed boundary between science and philosophy. Historically the boundary has been both indistinct and highly variable.

A related issue is whether *causally productive activity* is a unified concept, or an arbitrary catchall. If causally productive activities aren't all distinguished by a single, non-disjunctive feature (of no more than manageable complexity), then what makes the concept of a causally productive activity any more interesting than a randomly assembled list?

The second question has to do with what kinds of things we can perceive.

Seventeenth- and eighteenth-century empiricists believed the content of an idea can include nothing beyond what is given in introspective or perceptual experience. Locke thought we get the idea that things of one kind cause things of another by perceiving instances of the one causing instances of the other (Locke, 1988, pp. 324–325). But Hume thought we can't observe any such thing: watch any causal interaction as closely as you can, he said, and whatever you see, it will not be the connection between events which makes one a cause of another (Hume, 1990, pp. 155–172; cf. Anscombe, 1981, pp. 137–138). It follows, Hume thought, that in order to have content our idea of causality must reduce to ideas of other things that we *can* perceive.

Tell people you believe in causally productive activities that don't reduce to regularities, or other non-causal sorts of things, and they're likely to give you a funny look and ask if you really think we can perceive causality. Hume's ghost is whispering in their ears, telling them that we can't form an acceptably clear idea about a causally productive activity unless causality is or reduces to something we can see, hear, taste, smell, feel, or introspect. He is trying to ask whether what Anscombe says about causal concepts is psychologically realistic.

3. The question whether we can perceive causality is a red herring

I'll consider the perceptual question first because it's the easiest of the two. It assumes, as Locke and Hume believed, that concepts derive their content from introspective and perceptual experience in such a way that they can have no content over and above what is contained in ideas of the experiences that give rise to them. But why should we believe that? We don't have to think grammar is identical or reduces to anything we can see, hear, or otherwise experience in order accept the fact that at least some people have well developed concepts of the grammar of the languages they know. We don't have to think truth is identical or reduces to anything perceptible to accept that people who can distinguish true claims from false ones have a concept of truth. There are lots of interesting unanswered questions about how people acquire their ideas of grammar (truth) but 'Can grammar (truth) be perceived?' is not one of them. By the same token, you don't have to believe that causality or causal connections are perceivable or introspectible to believe we have concepts of causally productive activities.[7]

A related objection holds that the notion of a causally productive activity is too anthropomorphic to ground our understanding of causality. The objection assumes that concepts of causally productive activities derive from our experiences of manipulating things and, as before, that a concept can have no more content than the experiences which give rise to it. This is supposed to show that that we can't think an inanimate object (e.g., a charged particle) engages in a causally productive activity (e.g., repelling another charged particle) unless we conceive of it as being something like a human agent that manipulates things the way humans do. I don't know of any decisive arguments for the first assumption. The second assumption is implausible. If it were true, picking up one's idea of a wave from experiences of waves that travel through air or water should debar one from understanding wave propagation without a medium. But the most serious troubles people have when they try to understand wave propagation don't come from gaps in their introspective and perceptual experience. They come from the difficulty of the mathematics used in wave physics. And the fact that people do learn the physics argues that the empiricist assumption about the limits of conceptual content is false.

Alison Gopnik and her associates have impressive experimental evidence that, contrary to what Hume's and other regularity accounts of causality would predict, very young children engage in causal reasoning and acquire causal concepts that go well beyond generalizations of co-occurrences or regularities among co-occurrences they have observed. Their subjects could pick out the cause of a result from factors that accompanied it with the same relative fre-

[7] A referee objected that I haven't done justice to the following motivation for empiricism: 'If the causal facts outrun the observable facts, then ... two worlds could be identical with respect to all the observable facts and differ with respect to the causal facts, leaving the causal facts (that is, facts about activities) hopelessly underdetermined by any possible evidence'. All I have space to say about this is that it lends little if any support to the empiricist idea that causal claims cannot be tested and our idea of causality can have no content unless causal concepts can be analyzed without remainder into concepts of things we can perceive. Not all testable claims are tested against, and not all applications of concepts are legitimized by, appeal to evidence which humans can perceive or introspect. For example solar neutrino fluxes are detected by analyzing data obtained from Geiger counters and other mechanisms that are sensitive to things we cannot perceive.

More importantly, many claims are confirmed, and applications of concepts legitimized, by evidence that neither belongs to nor is represented (linguistically or non-linguistically) in their content. For example jagged lines moving up and down as they travel across a computer screen are used to test claims how axonal responses to artificial electrical stimulation vary with the concentration of sodium in the solution that surrounds the axon. But there is no reason to think the evidence belongs to the content of the claims or the relevant concepts. The claims and concepts were well understood before the tests were devised and the computer technology was invented. Both are now understood by laypersons who have no idea of what the computer displays look like.

Thus the rejection of the empiricist position does not imply that causal claims can't be tested or that specific causal concepts have no content.

quency (Gopnik & Schulz, 2004, pp. 372–373). They could distinguish, and appeared to understand the difference between direct and indirect causes of the same effect (Gopnik et al., 2002, p. 82). And they could figure out how to intervene to stop an ongoing effect by doing something they had not previously observed (ibid., p. 74).

This argues against the view that the content of the children's causal concepts must be limited to the co-occurences and other non-causal perceptibles the children have experienced. But it doesn't tell us whether all of the concepts the children use in their causal reasoning are as general as *making something happen* or *making something stop*, or whether at least some of their reasoning uses causal concepts as specific as the ones Anscombe said were fundamental. P. L. Harris et al. (1996) describe experiments which make it plausible that very young children understand and employ notions of highly specific causally productive activities. Their three and four year old subjects understood that walking across a floor in muddy shoes gets it dirty, and that one can keep the floor clean by taking off one's shoes before entering (ibid., p. 238).[8] They understood that pens can get your hands dirty by leaking ink, that you can't avoid the ink stains by choosing a different color, but that you can keep your hands clean by using a pencil instead of a pen (ibid., p. 243 n.). They understood that the pen isn't the only thing that can dirty their hands; if one touches the picture you draw it can smudge them (ibid., p. 246). They knew that if someone dirtied the floor by painting it, it wouldn't have kept it clean to use fingers instead of a brush (ibid., p. 241). They knew that a blower can blow out a candle, but only if it's turned on, pointed the right direction, and there's no barrier to block the wind (ibid., p. 253). They knew that a light source can illuminate a wall, but a blower cannot. They knew a light source cannot blow out a candle. Having seen what one blower could and could not do, they weren't surprised that other blowers could and could not do the same thing (ibid., p. 253).

Anscombe's ideas are psychologically adequate if children can engage in causal reasoning that uses concepts of specific causally productive activities. The moral I draw from Gopnik's, Harris's, and their co-authors' experiments is that whether or not children do whatever Hume and others may mean by 'perceiving causality', they can and do pick up and use non-anthropomorphic concepts of specific causally productive activities.[9]

4. Distinguishing genuine from spurious causally productive activities

Scientists do draw principled distinctions between genuine and spurious causally productive activities. But the constraints they rely on to do this are too various, and for the most part too local, to be captured informatively by any single account of causality as general as what the standard philosophical literature tries to supply. As I illustrate in Section 6, below, significantly different considerations constrain investigations into the causes of different kinds of effects. Furthermore, the constraints develop over time: investigators who look for the causes of one and the same effect during different historical periods typically do not operate under all of the same constraints. Moreover, different groups of investigators studying the same effect during the same historical period may not acknowledge the same constraints. That is not to say that scientific beliefs and practices are necessary or sufficient to determine whether any given activity actually occurs, and whether and what it contributes to the production of a given effect.[10] What changes over time are the kinds of activities scientists can recognize as responsible for the effects they study. For example, chemical and biological thinking has changed so much since the nineteenth century that without some familiarity with the relevant history it's difficult for us to grasp, let alone take seriously, all of the considerations which constrained nineteenth-century investigations of fermentation. What has changed since vitalists like Pasteur and anti-vitalists like Liebig argued about whether non-living things can convert sugar into alcohol is not the process of fermentation, but the way scientists investigate and understand it. The same holds for nineteenth and early twentieth-century debates about neuronal inhibition. Even though neuroscience has changed dramatically since the turn of the nineteenth century, neurons engaged then in the same inhibitory activities they engage in now.

If different scientists can accept very different kinds of things as causally productive activities because they operate under different constraints and base their decisions on different considerations, why isn't *causally productive activity* just an open ended catch-all for items that don't have enough of interest in common with one another to fall under a single concept? Why isn't it a cognitive analogue of what Aristotle called heaps as opposed to unified wholes? What makes *causally productive activity* a concept is more like what makes *game* and *organism* concepts than what makes *triangle* a concept. That may remind you of Morton O. Beckner's account of polytypic biological concepts, but there are differences and it may help to begin by noting them.

The concept of a triangle is monotypic, which is to say that the possession of a non-disjunctive property (that is, the property of being a closed, plane figure bounded by three straight line) is necessary and sufficient for something to qualify as a triangle. By contrast, Beckner would say, *ruby crowned kinglet* is a polytypic concept. That means that there is a collection, G, of properties such that every bird that falls under the concept possesses a 'large (but

[8] I am indebted to Jim Woodward for telling me about Harris's work.

[9] In invoking Gopnik, Glymour, Harris, and their collaborators in support of the psychological realism of Anscombe's views about general and specific causal concepts, I don't mean to imply that all or any of them subscribe to this paper's views.

[10] I'm ignoring cases where scientists develop new ways to produce an effect.

unspecified) number of the properties in G', each property in G 'is possessed by large numbers' of birds that qualify as ruby crowned kinglets, and 'no [property] in G is possessed by every individual in the aggregate'[11] (Beckner, 1968, p. 22). *Causally productive activity* is not polytypic in Beckner's sense because there is no G whose members are non-disjunctive properties possessed in large numbers by every instance of scraping, pushing, wetting, carrying, eating, burning, knocking over, keeping off, squashing, making noises, making paper boats (that was Anscombe's list), or opening the pore of an ion channel, pushing or rotating a helical component of a protein, depolarizing an axonal membrane, electro-chemical attracting and repelling, bonding, releasing, diffusing, dissolving,[12] and every instance of every other causally productive activity. Furthermore, causally productive activities are far, far too diverse to have any chance of meeting Beckner's requirement of a G such that each one of its members are possessed by a large number of causally productive activities.[13]

By contrast, Aristotle said that some perfectly useful and intelligible concepts cannot be captured by necessary and sufficient condition style definitions and must be explained instead by appeal to analogies. To understand such a concept is to grasp connections between analogies among things that justify its application to them (Aristotle, 1995, 1048a/30–b/8). Aristotle's idea fits the concept of a causally productive activity better than Beckner's. All that the many different activities that fall under the concept have to do with one another is that

> **1.** Each one is similar or analogous to one or more of the others with regard to features that are causally relevant to the production of their effects.

The similarities and analogies need not be the same for any appreciable number of causally productive activities. The features with respect to which they hold must be 'causally relevant', but if no single, universally applicable criterion distinguishes causes from causally extraneous factors, there should be no single, universally applicable criterion for causal relevance. The historical examples I sketch in Section 6 illustrate some of the many different kinds of considerations that can determine whether an activity is causally relevant to the production of an effect.

5. Is *causally productive activity* a unified concept of?

Something is a triangle just in case it is a three-sided closed plane figure. *Triangle* is a concept rather than a cognitive heap because nothing except closed plane figures fall under it. Thus what makes *triangle* a concept is also what determines whether or not any given shape is a triangle. Similarly if something could be a causally productive activity only if it is related to its effects in the same way that all other causally productive activities are related to theirs, then *causally productive activity* would be unified concept rather than a cognitive heap by virtue of the very same thing that determines which activities fall under it. But as I interpret her, things aren't so simple for Anscombe. She can say that *causally productive activity* qualifies as a concept because of similarities and analogies among the activities that fall under it. But what unifies the concept is not the same thing as what makes any given activity fall under it. Although any given causally productive activity will be similar to some and analogous to others, an activity is causally productive by virtue of facts about the activity, the things that engage in it, and what results from their doing so. These facts—different ones for different activities—give rise to, but do not reduce to, the similarities and analogies that unify the concept.

The similarities and analogies that unify the concept are also indispensable to scientists in their efforts to discover the causes of an effect, especially in the early stages of their research. But to repeat, those similarities and points of analogy are not to be confused with the facts that determine whether or how an activity makes a causal contribution to the production of an effect. For example, diphtheria bacteria inflame the throat by synthesizing and releasing a toxin whose influence on throat tissue depends upon facts about its molecular makeup (including the relative positions and sizes of the molecules, the charges they carry, their energetics, etc.). It depends also on the molecular makeup and physiology of the tissue they act upon on, the activities the toxic molecules engage in, and the ways in which various components of the throat tissue responds. Such facts give rise to, but do not consist of or reduce to, resemblances or analogies among the activities engaged in by diphtheria bacteria and toxin, and other biological agents. This means that Anscombe cannot tell as simple a story as her opposition.[14] The opposition thinks that one and the same regularity condition not only unifies the concept of a cause, but also determines what falls under it, and constrains decisions about what to classify as a causally productive activity. In the next section I sketch two case histories, which argue, to the contrary, that the constraints on ideas about causally productive activities, the ways in which similarities and analogies bear on the discovery of causal influences, and the kinds of facts that determine

[11] This last requirement serves to rule out G sets consisting of properties (like the property every individual has by virtue of flying if it flies) that can be trivially ascribed to everything. For the case of causally productive activities, the requirement that no property in G is possessed by everything that falls under the concept would rule out properties like 'occurring when it occurs' and 'producing what it produces'.

[12] See note 4. These examples come from (Machamer, Darden, & Craver, 2000, p. 14).

[13] Beckner saw his definition of 'polytypical' as an explication of what Wittgenstein said about concepts like *game* (Wittgenstein, 1968, §64–67). Because Wittgenstein was too averse to theorizing to provide a positive account of such concepts, this paper does not consider his views.

[14] The last ten lines owe a great deal to objections and suggestions from Peter Machamer.

whether an activity contributes to the production of an effect vary from case to case.

6. Two historical examples: inhibition and fermentation

During the nineteenth century, neuroscientists came to agree that in addition to exciting muscular activity, nerves are involved in some sort of process which inhibits it, and that their doing so is essential to the organization of voluntary as well as reflex movements. But there was no agreement about the nature of the inhibitory process or what nerves do to contribute to it.

Around the middle of the century Eduard and Heinrich Weber demonstrated that one can damp heart contractions by stimulating the vagus nerve (Fearing, 1970, p. 188 n.; Smith, 1992, p. 80 n.). In 1863 Sechenov published evidence that regions in the midbrain facilitate reflex responses (e.g., the frog's scratching reflex) by inhibiting opposing muscles whose contraction would interfere with them. This could be because inhibition is a causally productive activity, that is, something nerves do to muscles to relax them. That's what Charles Bell proposed several decades earlier when he wrote that contrary to the received view that nerves are exclusively instruments for stimulating muscles nerves can relax muscles too. In addition to their previously recognized contribution to coordination by which the simultaneous excitation of several muscles produces a combined response they also provide coordination 'between classes of muscles by which the one relaxes and the other contracts' (quoted in Sherrington, 1989, p. 287). Near the end of the nineteenth century Sherrington argued that the coordination required for knee jerk and other reflexes is accomplished by 'inhibito-motor spinal reflexes [that] occur quite habitually and concurrently with many of the excito-motor [spinal reflexes]' (quoted in Swazey, 1969, p. 87). Sherrington's studies of reflex responses led him to think that one nerve can inhibit another, and that inhibitory nerves relax skeletal muscles by damping the activity of the excitatory nerves which innervate them (Eccles et al., 1979, p. 209).

Well established as it is now, the view that nerves engage in inhibitory as well as excitatory activities had opponents as late as 1938 when B.F. Skinner said that if the term 'inhibition' signifies anything that can be studied experimentally, it can only refer to something like a decrease in the frequency of observable behavioral responses to a stimulus. Nothing 'was to be gained by contrasting excitation and inhibition (*pace* Sherrington) since [those terms] ... in fact referred to a continuum of degrees of reflex strength' (quoted in Smith, 1992, p. 161). The predominant view among neuroscientists and psychologists in the nineteenth and early twentieth centuries was that apart from nutritional and metabolic processes required to maintain them and support their functions, nerves are anatomically and physiologically incapable of engaging in anything except excitatory activities. Accordingly investigators were constrained to understand inhibition as a byproduct of the excitatory neuronal activity. T. L. Brunton believed inhibition occurs when excitatory impulses transmitted through two or more nerves interfere with one another 'in much the same way as two rays of light interfered with one another in Newton's rings' (Fearing, 1970, p. 195). G. H. Lewes concurred, arguing that theorists who posit special inhibitory centers to explain the damping of muscular or neuronal activity must resort to the desperate Ptolemaic maneuver of positing additional centers, which inhibit the inhibitors to explain how inhibited tissue can resume its activity; 'just as epicycles are heaped upon cycles, so nerve centers are being added to nerve centers' (Lewes, 1877, p. 237). But inhibition can't work the way Brunton and Lewes supposed unless nerves transmit excitatory influences which resemble light or sound waves with regard to factors required for the production of interference effects. Following Wundt and others, Sherrington rejected such theories on the basis of physiological evidence that

> though in a certain sense of the word the nerve impulse can be described as wave-like, it is not an undulatory disturbance at all in the sense in which those reactions are which show physical interference. (Sherrington, 1989, p. 193–194)

That amounts to saying that neural impulses do not resemble and are not analogous to light and sound waves as they would have to be in order to interfere with one another as Brunton supposed.

William McDougall (1902, passim) advocated an alternative account of inhibition based on an analogy between the excitatory activities of nerve cells and the discharge of water from one channel of a hydraulic system into another. To exert an excitatory influence, he thought, a nerve must transmit some sort of neuronal force into whatever it excites. In doing so it drains force away from one or more other nerves just as expelling water from one channel drains another and lowers its pressure. Just as one channel can reduce the flow from another by lowering its pressure, one nerve can reduce another's excitatory activity by draining force out of it. (See Smith, 1992, pp. 133–134, and also James, 1950, p. 584 n.). Drainage accounts were abandoned for the same kind of reasons as interference accounts: the details of the hydraulic analogy could not be developed to any satisfactory degree of empirical adequacy.

During the 1880s, Sherrington's teacher, W. H. Gaskell, published conclusions based on anatomical studies of the nervous system, and observations of the reactions of non-skeletal muscles in frogs, turtles, and other animals to electrical stimulation of the nerves that innervate them. Like the Webers, Gaskell found that stimulating vagus fibers weakens and slows heart contractions. Furthermore, Gaskell could make the heart beat faster and more strongly by stimulating what he called sympathetic nerve fibers (Gaskell, 1886, p. 42). This, and similar evidence from experiments on eye and sphincter muscles, convinced him that

neuronal activity can inhibit as well as excite non-skeletal muscles and, furthermore, that no single nerve can perform both functions. Accordingly he proposed that in addition to the motor nerves there must be inhibitory nerves, whose anatomy and physiology must be understood before investigators 'can attempt to understand the part played by the nervous system in the regulation of the different vital processes' (ibid., p. 40). As a first step toward finding out what activities nerves engage in to inhibit muscles, Gaskell investigated the effects their activity produces. One way to weaken or slow muscular contractions is to exhaust the muscle, but Gaskell's experiments demonstrated that inhibition and exhaustion are different effects. For example while exhausted heart muscles cannot become fully active until they have rested for some time, no rest is required for the resumption of full activity after inhibition (ibid., pp. 49–50). In order to explain how the vagus can inhibit the heart without exhausting it, Gaskell drew an analogy between inhibition and one of the metabolic effects Herring's opposed process theory of color sensation invokes to explain experienced differences in hue and saturation. According to Herring, experienced hue depends upon the activity of a 'red–green substance' and a 'yellow–blue' substance in the visual system. Saturation depends upon what goes on in a third, black–white substance. The hue and saturation one experiences while sensing a color depends upon the degree to which assimilative (anabolic or A) and dissimilative (catabolic or D) type metabolic processes predominate in each of the three substances.[15] To explain the inhibition and excitation of muscular activity, Gaskell appealed to metabolic activities analogous to Herring's, claiming that a muscle is inhibited when its metabolic activity is anabolic (it is building up energy) and excited when its metabolic activity is catabolic (it is using energy) (ibid., p. 50). Like the interference and drainage theories just mentioned, Gaskell's metabolic account foundered on physiological evidence that argued against the analogy it depends upon.

Sherrington extended Gaskell's idea that inhibitory as well as excitatory nerves are required to control and coordinate muscular activity to the explanation of skeletal muscle reflexes. He also proposed that inhibitory nerves can relax muscles by damping the excitatory nerves which innervate them. Most interestingly for our purposes he endorsed a theory devised by J. S. MacDonald about how one nerve inhibits the activity of another. MacDonald's theorizing is an example of how the acceptance of the claim that something is a causally productive activity is constrained, and what must be done to meet the constraints.

MacDonald concluded his Royal Society communication on the structure and function of nerve fibers with a methodological remark about his disinclination to use the term 'excitation'. In addition to carrying 'many conceptions foreign to my meaning and possessing no relation to the facts discovered' he said that talk of excitation couldn't help explain the phenomena he was investigating until someone provided (as he hoped he had been able to begin doing) a physiological account of how the nerve performs this function (MacDonald, 1905, p. 350). I suppose he would have said the same about the word 'inhibition'.[16] This illustrates that (as Carl Craver put it in correspondence) 'filler terms such as "activate", "inhibit", "encode", "cause", "produce", or "process" are often used to gesture at a kind of activity in a mechanism without providing any detail about exactly what activity is to fill that role'. Such terms can do no more than mark explanatory gaps until they are supplemented by enumerations of facts about the relevant causally productive activities, and accounts of how and under what circumstances those activities contribute to effects. In the course of filling the gaps, the analogies and similarities which suggested to investigators how an effect might possibly be produced must eventually give way to the enumeration of facts about causal factors, what they do, and how they do it. For example, we'll see that Macdonald tried (without complete success) to replace the suggestive analogy between inhibiting neuronal activity and impeding the flow of a current with an account of facts about ion currents and factors which influence the conductance of the fluids which carry them through the axon.

To establish that inhibition is an activity by which one nerve can retard signaling in another, MacDonald first had to explain what a nerve signal is. He recognized that the signals by which a nerve excites a muscle or another nerve (he called them 'action currents') are electrical in nature, and identified them with ion currents. He understood that neuronal axons and dendrites are filled with fluid and proposed that the fluid's electrical conductance depends on how conducive it is to ion flow. It follows that one nerve can excite or inhibit the production of action currents in another if it can act on the fluid it contains to render it more or less conducive to ion flow. To say what kind of activity could accomplish that, MacDonald had to identify features that influence ion flow. To this end he proposed that the fluid that carries the action current contains colloidal particles and inorganic salts, and that the ease with which charged ions move through it varies with the concentration of salts in solution. This in turn depends partly on osmotic pressure. Conductivity changes as changes in osmotic pressure allow salt to flow outward

[15] When A and D activities balance each other in the red–green substance, the hue of the resulting sensation derives exclusively from yellow and blue. When they balance each other in the yellow–green substance, the hue derives exclusively from red and green. Balance in the white–black substance produces a medium grey. If the hue of the experienced color contains red or green, yellow or blue, the degree to which it does so depends upon the degree to which D or A activity predominates in the two hue substances, and similarly for degree of saturation (Kries, 2000, p. 435).

[16] I think the reason he did not say so explicitly is that his research was focused on the production, rather than the inhibition, of the injury current (i.e., electrical activity he observed in cut nerves) and by extension, the 'action currents' of normally functioning nerves.

or inward through the nerve membrane. The colloidal particles suspended inside the nerve also influence conductivity by trapping salts on their surfaces and taking them out of solution to reduce conductivity. If things actually worked that way, one nerve could regulate another nerve's signal propagation by influencing its osmotic pressure or the sizes of the colloidal particles. MacDonald considered the latter possibility. He proposed that (for reasons he didn't try to explain) the colloids congeal in the presence of negative electrical charges, and break up into smaller bits in the presence of positive charges. When they become smaller, their total surface area increases, more salts are trapped, and fewer electrolytes are available to carry ion currents. When they congeal, their surface area decreases, releasing some salts into the solution to increase its conductivity (ibid., pp. 340–342). If so, the introduction of negative charges can promote neuronal signaling, and the introduction of positive charges can inhibit it. Accordingly, one should 'expect the possibility of finding nerve fibres capable of communicating a positive charge in place of the more usual negative charge, to the tissues innervated by them' (ibid., p. 348; cf. Sherrington, 1989, p. 198 n.). Inhibition would be the activity by which one nerve communicates a positive charge to another. On MacDonald's theory the acceptance of inhibition as a causally productive neuronal activity would require the identification of the features and behaviors by virtue of which inhibitory nerves are capable of communicating positive charges.

MacDonald speculated briefly and unsuccessfully about how a nerve might discharge the positive current he considered necessary for inhibition (MacDonald, 1905, p. 349). As we know, the theory of inhibition didn't develop the way he would have expected.[17] But, unsatisfactory as it was, his story about colloids and inorganic salts could not be replaced with a more adequate account until decades later when neuroscientists learned (among other things) about ion channels, neurotransmitters, neuroreceptors, and their contributions to the propagation and damping of action potentials. Needless to say, it also had to await the development of new experimental technologies. As Eccles said, there was no chance of developing a satisfactory theory of inhibition before it became possible to record intra-cellular potential in motor neurons (Eccles, 1956, p. 161). Without the data the new technologies provided, investigators could not develop accounts of inhibitory activities that were much better than the ones that were available to Sherrington.[18] The account that developed and became irreversibly well established (in outline at least) during the last half of the twentieth century is very different from MacDonald's.[19] But in order to secure the credentials of inhibition as a causally productive neuronal activity its twentieth-century proponents had to provide empirically adequate answers to questions about what kind of activity inhibition is, how, and under what conditions it operates, and what makes nerves anatomically and physiologically capable of engaging in it.

The examples of Skinner, Brunton, MacDonald, and the others illustrate some of the many different kinds of constraints that can bear on decisions about what is and what is not a causally productive activity. The goal of Skinner's behaviorist research program was to find ways to predict and modify behavior without appealing to unobservable causal factors internal to the mind or brain. His unwillingness to treat inhibition as a causally productive activity comes from the methodological constraints his brand of behaviorism imposed, together with the limits of the observational technology available to him. G. H. Lewes's aversion to Ptolemaic clutter is a general aesthetic constraint, which favors theories that come closer than their competitors to meeting certain standards for simplicity. Although these are highly general a priori requirements, they are nothing like the general requirements of lawful regularity, counterfactual dependence, and so forth, which philosophers typically invoke to distinguish causes from causally irrelevant factors. For the most part, the constraints which kept Brunton, McDougall, and others from thinking of inhibition as something nerves do directly (instead of something that results from their excitatory activity) were local, rather than general. They depended for their authority on empirically testable ideas from anatomy, physiology, psychology, and physics. The difference between these constraints and such philosophical principles as 'causes and effects must instance law-like regularities' will be obvious.

When he studied fermentation in the 1850s, Pasteur tried to explain why far more alcohol is produced when yeast is immersed in a sugar solution where no free oxygen can reach it, than when it grows exposed to the air above the sugar solution. His explanation assumed that in both conditions yeast must obtain oxygen with which to derive nourishment from the sugar. In the aerobic condition, he proposed, it engages in a respiratory activity to absorb free oxygen from the air, and a distinct nutritive activity to ingest nutrients. Neither of these activities produces any appreciable amount of alcohol. In the anaerobic condition, he thought, the yeast extracts the oxygen it needs from the sugar. In order to do this, it does something different from, but analogous to aerobic respiration. Yeast produces alcohol by 'performing its respiratory functions somehow or other with oxygen existing combined in sugar' (Pasteur, 1969, p. 259). Pasteur insisted that the kind of respiration

[17] It would be Whiggish to suggest that either MacDonald or Sherrington anticipated or directly influenced the development of theories of neuronal signaling during the second half of the twentieth century and on to the present. Although MacDonald is mentioned by some historians, I didn't find any mention of his work, or Sherrington's description of it in any of the textbooks I spot checked.

[18] For a relatively early survey of some of experimental technologies which were available by the middle of the twentieth century, the kinds of evidence they yielded, and the arguments which established how neurons could engage in inhibitory activities, see Eccles (1964), passim.

[19] See Kandel (2000), Pt. III, pp. 175–295.

that produces alcohol from sugar is a vital activity that only living organisms can engage in. He supported this claim in part by appeal to evidence that amyl alcohol has two components, one of whose crystals are optically active. According to Pasteur such optical activity is indicative of a molecular asymmetry that 'plays a considerable role in the most intimate laws of the organization of living organisms and intervenes in the most intimate of their physiological processes' (quoted in Geison, 1995, p. 96). The sugars from which amyl alcohol is fermented are themselves optically active, but Pasteur argued on empirical grounds that this could not explain the crystals' optical activity because non-vital influences cannot decompose optically active substances into new optically active substances. '[L]ife alone is capable of creating full fledged new [molecular] asymmetries' (ibid., p. 100). If what holds for amyl alcohol production holds for the production of other alcohols as well, fermentation must be an activity (like respiration) which non-living things cannot engage in.

Liebig was an antivitalist. His investigations of fermentation were constrained by the assumption that every physiological process results from physical and chemical activities which non-living things can engage in (Dubos, 1960, p. 154). He held that fermentation is a mechanical activity (in the seventeenth-century sense of that term) of lifeless bits of decomposed yeast. When the yeast in a sugar solution stops growing and begins to decompose

> the bond which unites the constituent parts of its cellular contents is loosened, and it is through the motions produced therein that the cells of yeast bring about a disarrangement of the elements of the sugar into other organic molecules. (Quoted in Kries, 2000, p. 327)

According to Liebig bits of decomposed dead yeast set up motions in sugar molecules by bumping into them. In effect, they shake them until alcohol and other molecules separate out.

Liebig was constrained by the principle that only non-vital activities may be invoked to explain physiological phenomena, while Pasteur was constrained by the principle that only vital activities could produce molecularly asymmetrical fermentation products. But what is a vital activity? Pasteur offers no clear, general characterization of the difference between vital and non-vital activities.[20] What made it possible for him to argue with Liebig was their mutual acceptance of a stock of examples that vitalists and non-vitalists alike considered crucial. They agreed, for example, that yeast is an organism, that living yeast behaves differently from dead yeast, and that dead yeast decomposes into non-living particles whose motions and collisions do not qualify as vital activities. The constraints at work in the disagreement between Pasteur and Liebig depended for their application on the fact that all parties to the dispute agreed that respiration and nutrition were the kinds of activities anti-vitalists were obliged to reduce, and that moving and shaking were good examples of non-living activities to which they were obliged to reduce them. What twentieth-century biologists learned about the molecular makeup and physiology of yeast made the constraints Liebig invoked against Pasteur irrelevant to the understanding of fermentation, and resulted in the development of new criteria to determine what kinds of activities could be responsible for fermentation. Fermentation turned out to be a physico–chemical process, as Liebig hoped it would, even though it is not the kind of process he was in any position to envisage. Eventually people were able to agree, as Pasteur thought they should, that fermentation results from activities characteristic of living things, but they were able to do so without committing themselves to the consequences that made vitalism unacceptable to its nineteenth-century opponents.

7. Concluding remarks

I conclude by considering very briefly one reservation and two objections you may have to what I've been saying.

[a] Philosophers who look for a single, universally applicable, exception-free condition (such as 'causes necessitate their effects', 'effects depend counterfactually on their causes', 'causes and effects are instances of law-like regularities') to discriminate causes from non-causal factors that accompany the production of an effect won't be at all satisfied with Anscombe's view as I've developed it. Nor will philosophers who think the distinction between causes and non-causes must turn on non-contingent principles to be discerned a priori by analyzing general concepts like *cause* or the meanings of general terms like 'cause'. As I understand her, Anscombe thinks that whether something makes a causal contribution to the production of an effect depends upon whether it engages in a causally productive activity. The examples of inhibition and fermentation argue that principled decisions about what qualifies as a causally productive activity which can contribute to the production of the effect of interest depend largely on local constraints that rest upon contingent, empirically testable beliefs, which develop historically. The decisions may be influenced by general methodological preferences like Skinner's and Lewes's. Furthermore, in addition to background beliefs as limited as Pasteur's ideas about optically active substances and MacDonald's ideas about the influence of neuronal colloids and salts on ion flows, investigators were constrained not to violate general mathematical principles most philosophers believe are necessary and a priori. They were also constrained by highly general physics principles that many philosophers think are necessary, universally applicable, and exceptionless. But these are not the princi-

[20] Since there were disputes about which things qualified as living organisms, it didn't help much to say (as Pasteur did) that vital activities are the ones that only living things can engage in.

ples Hume and Anscombe's other opponents include in their analyses of causality, and they are not by themselves sufficient to settle questions about causally productive activities.

In order to extract a general characterization of causally productive activities that Anscombe could accept from the examples we've looked at one would have to resort to characterizations as abstract as **1**. (the Aristotelian characterization in Section 4 above) with its talk of similarities or analogies with respect to 'causally significant' features. This characterization is acceptable to the extent that it recognizes that different things can qualify as causally productive activities for different reasons, and hints (albeit with Aristotelian vagueness) at what kinds of considerations decide whether what purports to be a causally productive activity qualifies as such. But no such characterization can pretend to be informative in the way that Hume's and other traditional accounts were thought to be. Anscombe's approach emphasizes the details of individual causal explanations at the expense of the generality philosophers typically prefer. Whether this is a serious drawback depends upon the feasibility of the traditional philosophical attempt to find general necessary and sufficient conditions for causality. I don't think that approach has a good enough track record to compare at all favorably to Anscombe's.

[b] We've seen that many of the constraints on nineteenth-century investigations of inhibition and fermentation were based on ideas which later researchers rejected. Why then should we have any confidence in the constraints present day investigators rely on? And what is to prevent investigators from arbitrarily rejecting constraints they find inconvenient? The answer to this kind of objection is that the old constraints did not give way to new ones without argument, and the arguments for their rejection were no different in kind or degree of cogency from arguments scientist use to make up their minds about just about everything they consider. The best of these arguments are so rigorous and so well supported by empirical evidence that only an extreme skeptic could doubt them. The anatomical, physiological, and physical evidence that eroded the constraints under which Brunton and the other interference theorists labored is a case in point. If there wasn't enough evidence to falsify the assumptions those constraints were based on, it's hard to see how any experimental evidence could be adequate for any scientific purpose.

Similar considerations apply to the objection that even if our ideas about causally productive activities are heavily constrained, and even if the constraints are based in large part on empirically testable considerations, that's no reason to believe that there really are causally productive activities, or if there are, that they operate to produce effects in the ways we think they do. This objection should be answered, case by case, by appealing to the ways in which claims about causally productive activities are tested. In some cases, there is room for doubt. In others, there is next to none. For example, even though plenty of questions about details of neuronal inhibition remain to be answered, an overwhelming body of empirical evidence has been marshaled to argue for the existence of inhibitory neurotransmitter and neuro-receptor molecules, some of the activities through which they are produced and released at the synapse, and some of the activities they engage in to damp neuronal activity by opening and closing ion channels in postsynaptic cell membranes. The evidence leaves no room for reasonable doubt about the occurrence of the causally productive activities or the correctness of at least some widely accepted qualitative claims about how they do their work.

[c] The experimentalists I've mentioned repeated their experiments in hopes of obtaining uniform results. Physiologists repeated their experiments in hopes of obtaining similar effects from similar experimental setups and manipulations. Anatomists repeated their dissections in hopes of observing the same causally relevant structural features. And so on. Someone might think that suggests, contrary to Anscombe, that causes must contribute to their effects in accordance with law-like regularities of some sort. Anscombe can cheerfully agree that looking for regularities and irregularities is epistemically important to the investigation of causes, that knowledge of regularities is crucial to practical applications of causal knowledge, and that there are obvious reasons for theoreticians to want to find out how regularly and under what conditions a cause succeeds in contributing to the production of an effect.[21] But it does not follow from the fact that regularities are of great epistemic, theoretical and practical interest that actual or counterfactual regularities are constitutive of the ontological or conceptual difference between the causes of an effect and the non-causal items that accompany its production.

Acknowledgements

I'm indebted to the Eastern Pennsylvania Philosophical Society for inviting me to give the talk this paper comes from, and for helpful discussion, to Lindley Darden for extremely helpful comments on a first draft, to Peter Machamer for help with this paper and for decades of conversation, to Carl Craver for several helpful suggestions, and to the audience who heard a version at the University of Maryland.

References

Anscombe, G. E. M. (1981). Causality and determination. In idem, *The collected philosophical papers of G. E. M. Anscome, vol. 2. Metaphysics and the philosophy of mind* (J. Barnes, Ed.) (pp. 133–147). Minneapolis: University of Minnesota Press.

[21] I discuss this at some length in Bogen (2005).

Aristotle. (1995). *Metaphysics theta*. In idem, *The complete works of Aristotle, Vol. 2* (J. Barnes, Ed.) (pp. 1552–1728). Princeton: Princeton University Press.

Austin, J. L. (1970). A plea for excuses. In idem, *Philosophical papers* (3rd ed.) (pp. 175–204). Oxford: Oxford University Press.

Beckner, M. O. (1968). *The biological way of thought*. Berkeley: University of California Press.

Bogen, J. (2005). Regularities and causality: Generalizations and causal explanations. *Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*, 397–420.

Cartwright, N. (1983). *How the laws of physics lie*. Oxford: Oxford University Press.

Dubos, R. (1960). *Louis Pasteur, free lance of science*. New York: Da Capo Press.

Eccles, J. C. (1956). *The neurophysical basis of mind*. Oxford, Oxford University Press.

Eccles, J. C. (1964). *The physiology of synapses*. New York: Academic Press.

Eccles, J. C., & Gibson, W. C. (1979). *Sherrington: His life and thought*. Berlin: Springer–Verlag.

Fearing, F. (1970). *Reflex action: A study in the history of physiological psychology*. Cambridge, MA: MIT Press.

Gaskell, W. H. (1886). On the structure, distribution, and function of the nerves which innervate the visceral and vascular systems. *Journal of Physiology*, *7*, 1–80.

Geison, G. L. (1995). *The private science of Louis Pasteur*. Princeton: Princeton University Press.

Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, *69*(3) (Suppl., Pt. II), S42–S53.

Gopnik, A., & Schulz, L. (2004). Mechanisms of theory formation in young children. *Trends in Cognitive Sciences*, 372–376.

Harris, P. L., German, T., & Mills, P. (1996). Children's use of counterfactual thinking in causal reasoning. *Cognition*, *61*, 233–259.

Hausman, D. M. (1998). *Causal asymmetries*. Cambridge: Cambridge University Press.

Hume, D. (1990). *A treatise of human nature* (2nd ed.) (L. A. Selby-Bigge, Ed.; rev. by P. H. Nidditch). Oxford: Oxford University Press. (First published 1781; this ed. first published 1978)

James, W. (1950). *The principles of psychology*. New York: Dover Publications.

Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (2000). *Principles of neural science* (4th ed.). New York: McGraw Hill.

Kries, J. (2000). Historical and critical survey. In H. Helmholtz, *Treatise on physiological optics, Vol. 2* (pp. 435–462). Bristol: Thoemmes Press.

Lewes, G. H. (1877). *The physical basis of mind*. Boston: James R. Osgood & Company.

Locke, J. (1988). *An essay concerning human understanding*. (P. H. Nidditch, Ed.). Oxford: Oxford University Press. (First published 1689; this ed. first published 1975)

MacDonald, J. S. (1905). The structure and function of nerve fibers: Preliminary communication. *Proceedings of the Royal Society of London: Series B, Containing Papers of a Biological Character*, *76*(510), 322–350.

Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, *67*(1), 1–25.

McDougall, W. (1902). The physiological factors of the attention-process (I). *Mind, New Series*, *11*(43) 316–351.

Pasteur, L. (1969) *Studies on fermentation, the diseases of beer, their causes, and the means of preventing them* (F. Faulkner, & D. C. Robb, Trans.). New York: Kraus Reprint.

Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge: Cambridge University Press.

Russell, B. (1971). On the notion of cause. In idem, *Mysticism and logic* (pp. 132–151). New York: Barnes and Noble.

Salmon, W. C. (1971). Statistical explanation. In idem, *Statistical explanation and statistical relevance* (pp. 29–27). Pittsburgh: University of Pittsburgh Press.

Sherrington, C. S. (1989). *The integrative action of the nervous system*. Birmingham, AL: Classics of Medicine Library.

Smith, R. (1992). *Inhibition, history and meaning in the sciences of mind and brain*. Berkeley: University of California Press.

Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction, and search* (2nd ed.). Cambridge, MA: MIT Press.

Steel, D. (2005). Indeterminism and the causal Markov condition. *British Journal for the Philosophy of Science*, *56*(1), 3–26.

Swazey, J. P. (1969). *Reflexes and motor integration: Sherrington's concept of integrative action*. Cambridge, MA: Harvard University Press.

Wittgenstein, L. (1968). *Philosophical investigations* (3rd ed.). New York: Macmillan.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.

Woodward, J. (2004). Counterfactuals and causal explanation. *International Studies in the History and Philosophy of Science*, *18*(1), 41–72.