References

Cartwright, N. (1989) *Nature's Capacities and Their Measurement.* Oxford: Oxford University Press.

Cartwright, N. (1999) *The Dappled World: A Study of the Boundaries of Science.* Cambridge: Cambridge University Press.

Cartwright, N. (2000) *Measuring Causes: Invariance, Modularity and the Causal Markov Condition.* London: CPNSS Measurement in Physics and Economics Discussion Paper Series, LSE.

Cooley T. F. and LeRoy, S. F. (1985) A theoretical macroeconometrics. *Journal of Monetary Economics,* 16, 283-308.

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

Hausman, D. and Woodward, J. (1999) Independence, invariance and the causal Markov condition. *British Journal for the Philosophy of Science,* 50, 1-63.

Hoover (forthcoming) *Causality in Macroeconomics.* Cambridge: Cambridge University Press.

Macaulay, D. (1988) *The Way Things Work.* Boston: Houghton Mifflin.

Pearl, J. (2000) *Causality.* Cambridge: Cambridge University Press.

Simon, H. (1977) Causal ordering and identifiability. In H. Simon (Ed.) *Models of Discovery and Other Topics in the Methods of Science.* Boston: Reidel, 53-105.

Woodward, J. (1997) Explanation, invariance and intervention. *British Journal for the Philosophy of Science* 64, 26-41.

Woodward, J. (2000) Causation and manipulation. ms, California Institute of Technology.

5

Three Dogmas of Humean Causation

GÜROL IRZIK

Boğaziçi University

1 Introduction

Dogmas are beliefs held come what may. They survive despite contrary evidence and argument. As I see it, Humean causation is a dogmatic doctrine that consists of the following:

The First Dogma: Ontological Reduction. Exceptionless regularities or probabilistic-correlational complexes are constitutive of causal relations. The latter are reducible to the former.

The Second Dogma: Supervenience. Singular causal facts are true by virtue of generic causal facts (causal law statements).

The Third Dogma: Epistemological Reduction. Our knowledge of causal relations (generic or singular) is based on our knowledge of regularities, or probabilistic-correlational complexes, and of spatio-temporal relations. This is essentially an inductive, therefore, inferential kind of knowledge.

Leaving aside the part about probabilities and correlations, we owe these dogmas to David Hume (1739), so they have been with us for a long time. The qualifications concerning exceptionless regularities were added much later (in 1970, to be exact, by Patrick Suppes) in order to avoid some counter arguments and examples. This marks the beginning of the end of the Humean hegemony in matters of causality, as alternative approaches have proliferated

Stochastic Causality.
Maria Carla Galavotti, Patrick Suppes and Domenico Costantini (eds.)

since then. To name a few, we now have the process approach, the capacity approach, the DAG approach, the agency approach, and a number of probabilistic approaches to causality. This boom, I believe, is due in part to the rediscovery by philosophers of new tools like conditional probabilities and in part to a liberation from the first dogma; all of these new approaches employ probabilistic or statistical apparatuses one way or another, and all contain strongly anti-reductionist elements. One can even show that the two reasons are related in interesting ways, but that is not the purpose of my paper. What I instead would like to argue is that while recent accounts of causality have gone a long way toward emancipation from the grip of Humeanism, one last vestige still remains. This is the third dogma, which in effect says that no causal relation is observable and that none of our causal knowledge is innate. It has received surprisingly little critical attention from philosophers, even from those who reject the first two dogmas. Inspired by the pioneering work of Albert Michotte, psychologists have been for years carrying out experiments that strongly suggest that even very young children can perceive causal relations. Philosophers of diverse inclinations persistently ignored and continue to ignore such work. What better label than "dogma", then, can characterize the thesis of epistemological reduction?

The plan of my paper is as follows. In the next section I will summarize the main objections to the first dogma, that is, the thesis of ontological reduction. I will say very little about the second dogma, but the interested reader can consult Cartwright (1989) and Woodward (1990). Section 3 will be devoted to some aspects of the capacity approach, particularly in the form developed by Nancy Cartwright. I believe that the capacity approach provides the most promising framework to make sense of causation. However, it faces an hitherto unnoticed epistemological problem, a circularity concerning our knowledge of causal claims understood as ascription of capacities. I will introduce this problem and offer a way of solving it. Although my solution has restricted scope, it seems to me to be fruitful. For by involving precisely the rejection of the third dogma, it opens up new avenues for investigation. Naturally, the bulk of my paper will be devoted to defending, in sections 4 and 5, the view that at least some causal relations are observable. My defense will be based on both empirical and philosophical grounds. I will argue that the third dogma of Humean causation results from a very poor epistemological theory; that there exists a much better philosophical account of seeing and knowing, which allows us to say that causal relations can be observed; and, finally, that there is considerable empirical evidence supporting the observability of some causal relations.

2 The First Dogma: Virtues and Vices

It is not difficult to see why a Humean account of causality incorporating the three dogmas outlined above appeals to so many philosophers: it has an anti-metaphysical, deflationist-reductionist ontology and has an attractive "solution" to the epistemological problem. Regularities conveniently do both ontological and epistemological work. These virtues are nicely displayed by Hume's own definition: "We may define a cause to be an object precedent and contiguous to another, and where all the objects resembling the former are placed in like relations of precedency and contiguity to those objects that resemble the latter" (Hume 1739/1962, p. 221). To put it differently, an event A causes another event B if and only if A is earlier than B, A is contiguous to B, and whenever events of type A occur events of type B follow. Ontologically, this definition employs only spatiotemporal relations which exhibit a regular, law-like behavior. A world that contains regularities automatically contains causal relations as well simply because regularities are constitutive of causal relations. As causal relations are reducible to constant conjunctions, the Humean world has no ontological excess like powers, natures, or necessities for connecting causes to their effects. Such a world is presumably more easily accessible because all one needs in order to have knowledge of causes is knowledge of regularities, and that can be acquired by induction over observations of repeated instances of A and B.

Despite its virtues, however, Hume's definition has a number of well-known vices. It is worth recalling some of them. First, there are constantly conjoined pairs of events which are not causally related. EPR-type correlations are a case in point. To take a pedantic example from daily life, nights follow days regularly, but one is not the cause of the other. Hume's view has no way of distinguishing between genuine causal relations and 'spurious' constant conjunctions symptomatic of common causes. Second, not all causal relations imply exceptionless regularities; some are probabilistic, and Hume's definition has no resources to cope with them. Finally, there seems to be causal relations which do not display any regular behavior. Here is a typical example, which I believe I read in one of Elizabeth Anscombe's papers: my stepping on ice (call it C) made me fall down (call it E) the other day, though it is not the case that each time I step on ice I fall down. So, there does not seem to be any law that relates stepping on ice to falling. Donald Davidson's ingenious reply has become the standard one to all such counter-examples: The law does not have to connect the cause to its effect under those descriptions. Rather, C falls under a certain event kind or description c, E falls under a certain event kind or description e, and the law relates c to e. Davidson claims that such a law always exists: "In any case, in order to know

that a singular causal statement is true, it is not necessary to know the truth of the law; it is necessary only to know that some law covering the events at hand exists" (Davidson 1968, pp. 93-94). Now, what reason is there to think that some such law always exists? Davidson writes: "And very often, I think, our justification for accepting a singular causal statement is that we have reason to believe an appropriate causal law exists, though we do not know what it is." (Davidson 1967, p. 701). There are two difficulties with this sort of reasoning. First, covering laws are sparse. Second, the postulation of the existence of a covering law is not the only justification we have for accepting a singular causal statement. Often we rely on the elimination of other possible causes. Thus, if we doubt whether it was A that caused B, we carefully review the circumstances to see if some factor other than A could have produced B. Failure to find it justifies the conclusion that A caused B. The Humean account blinds us to the eliminationist strategy which is at the heart of all experimental design and reasoning.

It is mainly for these reasons that philosophers have turned to probability theory to rescue Hume's program. A notable attempt is Patrick Suppes' *A Probabilistic Theory of Causation* (1970). According to his account, effects need not follow their causes regularly; they just need to occur more often than not when their causes are present. This is usually expressed with the slogan that a cause must increase the probability of its effect. More precisely, an earlier event A is said to cause a later event B if and only if $P(B/A) > P(B)$ and there exists no event C earlier than both A and B, which screens off A from B. C is said to screen off A from B if and only if $P(B/A) > P(B)$ and $P(B/A.C) = P(B/C)$.

The idea of screening off turned out to be fruitful indeed, especially in capturing 'spurious' correlations arising from common causes: a common cause screens off the correlation between its effects. Furthermore, screening off is an asymmetric relation in the sense that if C screens off A from B, A does not necessarily screen off C from B. This was encouraging because causal relation is also asymmetric, so there was reason to hope that the notion of screening off, together with the idea that a cause raises the probability of its effect, may provide a reductive definition of causation after all.

Unfortunately, as we now know very well, not only common causes but also intermediate causes function as screeners off. The standard way to distinguish between these cases is to consider time. However, some philosophers are reluctant to appeal to temporal order and direction, which, they believe, arise from causal direction and order (see, for example, Mackie 1980, Papineau 1985), so for them reduction is even more elusive.

Intermediate causes are not the only kind of problem that stands in the way of reduction even when temporal considerations are taken into account. Simpson's 'paradox' is another. Suppose that A causes B and that there is no

further factor which screens off A from B. So we expect the probability of B given A to be greater than the probability of B in general. Assume furthermore that A also contributes to C, which prevents B. In that case, A is a positive causal factor for B through one path and a negative causal factor through another. If these two influences cancel each other out, we may not observe the expected probability increase.

Having been convinced that these difficulties are insurmountable, a number of philosophers rejected the first dogma; that is, they gave up the project of reducing causal relations to regularities or probabilistic dependencies and turned to the less ambitious but perhaps more fruitful task of establishing the right kind of connection between them. Salmon's process approach and Cartwright's capacity approach are well-known examples. More recently, Spirtes, Glymour and Scheines (1993), Pearl and Verma (1994), who are the main founders of the DAG approach, have joined the non-reductionist camp.

However, some philosophers, most notably David Papineau (1991, 1993) and Wolfgang Spohn (this volume), believe that the DAG approach can be used for reductive purposes. They employ very similar strategies, the major difference being that Papineau attempts to define not only causal relatedness but also causal directionality in terms of the resources of the DAG approach. I have argued against his version in detail elsewhere (see Irzik 1996). So, let me turn to Spohn's. The key idea is that causal dependence relations among a given set of variables U (called the frame) can be represented by a directed acyclic graph (called DAG). A probability measure can be assigned over the members of U, so there are probabilistic dependence and independence relations among them. Basically, three conditions specify the kind of relations that hold between the frame U represented by a DAG and the associated probability measure. The first is *the Markov condition* that says that each variable is probabilistically independent of all its non-descendants conditional on its parents. Screening off relation defined above is a special case of this. The second condition says that no proper subgraph of the DAG satisfies the Markov condition. This is called *the minimality condition*. The third condition is known as *the faithfulness condition*. Its exact formulation is a bit complicated, but its intuitive content is clear: no conditional independence relation (such as screening off or a zero partial correlation) in U results from accidental situations regarding the parameter values of U. In other words, every such relation is determined by the structure of the DAG alone. A DAG that satisfies all three conditions is called a *Bayesian net*. Spohn then asserts that every causal graph is a Bayesian net and defines causal dependence in terms of it: B directly causally depends on A if and

only if there is a causal path from A to B in a Bayesian net (Spohn, this volume). In other words, Bayesian nets are all there is to causal relatedness.

But are these conditions universally satisfied? Take the faithfulness condition, for example. All cases of Simpson's 'paradox' are violations of this condition; whenever we have a cause that acts as a contributor through one route and that acts as a preventer through another in such the way that the two influences cancel each other out, the faithfulness condition will be violated. Admittedly, such cases will be rare, so methodologically it may be all right to impose it on our models, but it will not do if the aim is ontological reduction. Violations of the faithfulness condition, no matter how rare they are, form a serious obstacle for reductionist accounts such as Spohn's and Papineau's.

The Markov condition does not fare better in this regard either. For one thing, there is the EPR paradox that violates it. For another, as Nancy Cartwright has shown, the condition also fails when a cause acts in a truly probabilistic way. To see this, consider the following abstract example provided by her (Cartwright 1999, p. 109; see p. 108 for a concrete example). Suppose there is a cause C that has two separate yes-no effects, X and Y. We then have four possible outcomes: +X+Y, -X+Y, +X-Y, -X-Y. Assuming that C occurs, there are four joint probabilities that must be fixed by the world: Prob(+X+Y), Prob(-X+Y), Prob(+X-Y), Prob(-X-Y). As Cartwright puts it, "nothing in the concept of causality, or of probabilistic causality, constrains how Nature must proceed" (ibid.). The Markov condition holds only in one very special case, where Prob(+X+Y).Prob(-X-Y) = Prob(+X-Y).Prob(-X+Y). In all other cases, it will fail. Short of determinism, then, there is no reason to expect that the Markov condition will be satisfied universally. As before, this does not mean that Markov condition is useless methodologically, but it does imply that ontological reduction does not succeed.

The typical reductionist response to such objections is to argue that it is always possible to find a larger, more refined graph-theoretic frame in which the Markov, the minimality, and the faithfulness conditions are satisfied. This is the strategy pursued by David Papineau, for example. He claims that a world of cosmic conspiracies in which the probabilistic dependencies and independencies conceal the true causal structure of the world forever is conceivable but metaphysically impossible (Papineau 1993, p. 246). In a similar vein, Spohn claims that "in the final analysis it is the all-embracive Bayesian net representing the whole of reality which decides about how the causal dependencies actually are" (Spohn, this volume). Spohn denies that we possess an independent concept of causality to rely on, independent, that is, from Bayesian nets.

Such strong claims need some reason for believing them! I think it is plain false that we do not have a notion of causal relatedness independent of Bayesian nets. Our basic notion of causality is simply that causes make their

effects happen, and I claim that we acquire this notion primarily by manipulating objects; in other words, it is ultimately tied to the fact that we are agents. I will say more about this in section 6. Let us first look at what happens to the reductionist project with Papineau's and Spohn's moves. That project was attractive to empiricist philosophers because it was thought to be anti-metaphysical and epistemologically feasible. Now, it has itself become metaphysical and epistemologically elusive. To be consistent, reductionists *must* deny the possibility of a "non-Bayesian" or "conspiratorial" world. They appeal to probabilistic relations because they believe that they are empirically more easily accessible, and they reject notions such as power and capacity because they find them metaphysical. To accept that probabilities may never reveal the true causal structures which they constitute would be of course self-defeating. But the commitment to the existence of a "non-conspiratorial" causal world (Papineau) or an "all-embracive Bayesian net representing the whole of reality" (Spohn) surely transcends all possible experience since only God can have this sort of knowledge. Consequently, such a commitment is itself metaphysical, and the first dogma survives only within this deadly dialectic.

3 An Alternative Approach: Causes as Stable Capacities

In my view, the starting point for an alternative approach should be to accept without embarrassment that causation is a primitive relation that cannot be defined in terms of non-causal relations. The notion of cause makes sense only within a semantic field of related notions such as capacity, nature, change, manipulation, and invariance. I do not believe that these terms can be defined independently from each other, nor do I believe that they can be learned in isolation from one another; they must be acquired together in clusters. Thus to know the meaning of 'cause' is to know the network of relations in the semantic field to which it belongs. I am, in other words, advocating a semantic holism with respect to the concept of cause.

Take a causal claim like "smoking causes lung cancer". Following Cartwright (1989, 1999) and Harre and Madden (1975) I suggest that it should be understood as an ascription of a capacity or power to smoking to produce lung cancer. To say that X has the capacity or power to Y means that "X can do Y, in the appropriate conditions, *in virtue of its intrinsic nature*" (Harre and Madden 1975, p. 86; emphasis original). Thus, there is something in the nature of smoking, probably certain chemicals, the inhalation of which can produce lung cancer under certain conditions. Often, these conditions must be created artificially, in labs. The capacity or power in question is a more or less stable one that is manifested in different situations, "a capacity

which if the circumstances are right reveals itself producing a regularity, but which is just as surely seen in one good single case" (Cartwright 1989, p. 3). In this account regularities are no more than empirical manifestations of underlying nature of things under very special circumstances.

The notion of stable capacity is tied to the notion of invariance. Causal claims must satisfy an invariance condition in the sense that the relation between a cause and its effect should continue to hold under some specified class of changes in various conditions including initial and background ones. In the case of smoking-lung cancer, for example, the causal relation should continue to hold under changes of age, sex and the like.

It may also be worth noting that 'capacity' is clearly an Aristotelian notion. Aristotle's term was potency, by which he meant "the source of change or movement in another thing, being moved by another thing or by itself qua other" (Metaphysica, Book Δ, ch. 12, 1019a). So, the notion of change is built into the notion of cause. If A and B are causally related, manipulating A results in a change in B. That is why quitting smoking, but not cleaning one's yellow fingers, is an effective strategy to reduce the chance of having lung cancer, despite the fact that having yellow fingers and having lung cancer are highly correlated.

Roughly, this is the semantic holism I have had in mind. Although I have not defended the capacity approach, I believe that it provides us with the right framework to make sense of causal claims both in the scientific and ordinary contexts. The interested reader should consult the excellent works of Cartwright (1989, 1999). What I want to do in this section is to draw attention to a problem that faces the capacity approach, a problem that threatens it with an epistemological circularity.

To see this, begin by noting the distinction between causal capacity claims and singular causal claims. For the latter to be true, the cause and the effect must actually occur. But a causal capacity claim can be true even if the cause and the effect in question never occur. For example, after the devastating earthquake in the Marmara Region of Turkey in 1999 so much stress has been put on the fault line near Prince Islands that it now has a capacity to produce a major earthquake in Istanbul, even though it has not happened yet. There is a straightforward relationship between singular causal and causal capacity claims: for a singular causal claim to be true, it is necessary that a corresponding causal capacity claim must also be true. By contrast, the truth of a singular causal claim is sufficient for the truth of a capacity claim.

Let us then ask how we know the truth of a singular claim of the form "A caused B". Well, to know this we must know that A is the sort of thing that can cause B; if A does not have the right nature and therefore the right capacity, it cannot produce B whatever the circumstances are. But how do we know that A has that sort of capacity? If the answer is, only by knowing that

it or things similar to it have exercised that kind of capacity in various situations, i.e., by knowing the truth of some appropriate singular causal claims, then we are running in a circle.

The circle may not be a vicious one. We know that the fault line under Prince Islands has the power to produce a major earthquake, though it has not exercised its power yet. We know this because we know that that fault line is an extension of the North Anatolian fault line, which did produce similar earthquakes elsewhere in the past. Although what we are facing here is not a vicious circle, it is nevertheless worrisome. For the problem still remains: how does our knowledge of causes understood in terms of capacities get off the ground? After all, token causal relations are the result of exercising of capacities, and capacities are not open to direct inspection.

It seems to me that a profitable way out of this circle involves the denial of the third dogma that our knowledge of causal relations is always inferentially obtained by induction. This opens up two possibilities: some of our causal knowledge may be innate, or it could be observational. Recent work in developmental psychology suggests both. Let me start with innateness.

According to what is called "the theory theory", "the processes of cognitive development in children are similar to, indeed perhaps even identical with, the processes of cognitive development of scientists" (Gopnik and Meltzhoff 1997, p. 3). Just as scientists have theories about the world, which they revise according to evidence they gather, so do infants. To put it differently, human beings are born with certain substantive principles about the world and mechanisms for revising them on the basis of their experiences. For example, there is considerable evidence that infants have an innate knowledge that the world is three-dimensional and that objects move in certain trajectories (ibid., chs. 4 and 5). In a similar vein, Gopnik et al. (forthcoming) argue that the process of intuitive 'theory' formulation and revision also involves a type of representation that they call a causal map. They define a causal map as an abstract representation of the causal relationships among events in the world. They propose that certain cognitive devices were designed by evolution to recover causal information, and that animals, including human beings, may have some hard-wired expectations about possible causes and causal behavior. These constrain the attention of the infants to only certain aspects of the world. In the language of the capacity approach, this means that human beings from birth have some knowledge about capacities of objects. Needless to say, this is a very limited, general and abstract sort of knowledge, but, contra Humeans, it is innate. Experience serves only to confirm or revise it. This is a non-inductive, a top-down kind of knowledge, so to speak, and breaks the circle we mentioned earlier: some

knowledge of capacities is given innately without prior knowledge of any singular causal truths.

This is the ground base from which all our knowledge of causes takes off, a base which also makes possible a bottom-up kind of causal knowledge when coupled with observation. In other words, once we have knowledge of capacities, in certain circumstances we can observe one event causing another, and that gives us knowledge of singular causal truths without appealing to any knowledge of corresponding regularities. As we shall see, developmental psychology provides considerable empirical evidence for this as well. I do not mean to suggest that all our causal knowledge is acquired in this way, nor do I deny that many casual relations are unobservable. What I claim is that if we can see tables, fires, and so on, we can also see a hammer's smashing a tomato, a boy's pushing another boy, and so on. I will present the findings of the developmental psychologists in section 5. Let me now discuss some likely objections and then provide an analysis of what it is to see a causal relation.

4 Observability of Causal Relations

The idea that causal relations can be observed would be appalling to a Humean. I suppose he would object as follows. For one thing, he would say that all we observe is one event being followed by another: we observe first the blow of the hammer and then the smashed tomato, but never one causing the other, much less the relevant capacity. For another, he would ask us how we know that it is the blow of the hammer that smashed the tomato without establishing a regular occurrence between the two. Indeed, the two objections are related in that an archaic skepticism lurks behind them. To the first objection I reply: We do not see the powers, we infer them. They come with the properties that enter into causal relations. Just like there are observable objects, events and properties, there are observable relations in the world as well. If being taller than is an observable relation, I do not see why smashing, pulling, pushing, grabbing and a host of similar relations cannot be observed as well. From this viewpoint, there is nothing peculiar about causal relations. Perceptual skepticism might have made sense at the time of Hume who believed that all we perceive are our immediate sensations of sight, touch, pain, etc., but it does not make sense today. My reply to the second objection is that of course we can be mistaken about our causal judgments based on perception, but that this sort of fallibility pervades all other contexts: my observations can fail me about any other objective relation. Checking regular occurrence is surely one way of avoiding making mistakes, but it is not the only one. As we saw earlier, elimination of other possible causes is another.

Dispelling the objections of Humeans is not enough, of course. I owe a positive account of what it is to see a causal relation. Luckily, such an account has already been developed and defended in great detail by Fred Dretske more than forty years ago, and it is a pity that it has not been appreciated sufficiently by philosophers of science. In his *Seeing and Knowing* Dretske (1969) distinguishes between epistemic and non-epistemic senses of seeing. He defines the latter as follows: "S $sees_n$ D = D is visually differentiated from its immediate environment by S" (Dretske 1969, p. 20). Here, S is a subject, D is an object or event, and the subscript indicates the non-epistemic sense of seeing. 'Visually differentiated' involves, among other things, "S's differentiation of D is constituted by D's looking some way to S, and, moreover, looking different than its immediate environment" (ibid.). Non-epistemic seeing is a kind of seeing which is devoid of any positive belief content. A person can see an object in this sense without knowing what it is, without even knowing that it is an object. Thus, non-epistemic seeing is different from, and indeed possible without, 'seeing as'.

Although we can see things around us without recognizing what they are, our vision has also epistemic import. We can see, for instance, not only water but also that it is water. Dretske states four conditions that must be met in order for seeing to be epistemically informative: S sees that b is P in the epistemic way if (1) "b is P", (2) "S $sees_n$ b", (3) "The conditions under which S $sees_n$ b are such that b would not look, L, the way it now looks to S unless it was P", (4) "S, believing the conditions are as described in (3), takes b to be P" (ibid., pp. 79-88).

Once we are given a definition of epistemic seeing for objects and events, it can easily be extended to cover any relation at all. Let me apply the definition Dretske provides for relations to the relation of causality (see Dretske 1969, p. 141):

S sees that A causes (or better: is causing) B in the epistemic way if

(i) A is causing B

(ii) S $sees_n$ A and S $sees_n$ B

(iii) The conditions are such that A and B would not look the way they do, L, relative to one another to S unless A were causing B,

(iv) S, believing the conditions are as described in (iii), takes A to be the cause of B.

Go back now to our example of the hammer and the tomato: We hit a tomato with a hammer (A) and then see the tomato being smashed (B). Now, do we also see that A caused B? The answer is an unambiguous 'yes'. Conditions (i)

and (iv) pose no problem. Condition ii is satisfied because both the hammer blow and the smashing of the tomato are observable in the sense that we can distinguish them from their environment by the way they look to us. Whether condition iii is satisfied or not is an empirical matter; if the blow and the smash would not look to us the way they do were it not the case that one was causing the other, then the condition would be fulfilled. Note that it is perfectly possible that some other event C can cause the same effect without changing the way A and B would look to us. This would be the case, for instance, if a small bomb implanted inside the tomato exploded at the right moment and in the right way. In that case, condition iii would be violated, and we would not be able to see that A was causing B (cf. Dretske 1969, p. 231).

I submit that this analysis of epistemic and non-epistemic seeing provides a much more satisfactory framework than Hume's simplistic theory of impressions and ideas. Moreover, with some qualifications, it can be extended to cover touching, smelling, and so on although I lack the space here to do so. Its chief virtues for the purposes of the present paper are that it dispenses with the third dogma of Humean causation and that it is consistent with the psychologists' empirical findings to which I now return.

5 Evidence for Infants' Perception of Causal Relations

Recent research in developmental-cognitive psychology has provided considerable evidence to the effect that a causal relation can be directly observed as being distinct from the spatiotemporal properties of events. The pioneer in this field is Albert Michotte (1963) who suggested that even infants might have a direct perception of some cause-effect relationships as a kind of perceptual gestalt. Since then a number of psychologists followed up Michotte's suggestion and devised various experiments to test it. Here I will summarize the striking results of a series of such experiments conducted by Alan Leslie (1982, 1984), and Alan Leslie and Stephanie Keeble (1987). Like most others, theirs too use the habituation-dishabituation of looking technique. Six-month old infants are shown films of a red object colliding with a green object in a variety of ways. Habituation involves subjecting the infants to the image of the same motion many times. This causes a decrease in the infant's attention, which can then be measured in terms of the time of looking. Dishabituation involves presenting a contrasting image just once and then observing the infant's response in terms of the recovery of attention manifested as a longer period of looking. The idea behind the habituation-dishabituation technique is that because we are surprised at the events that violate our expectations, we look at them for a longer period of time. If this logic is correct, then subjects should be surprised at a violation on the first trial (dishabituation) and should look at this violation for longer than presentation of an event that is not in violation.

An initial experiment suggested that infants can distinguish between direct launching and similar but discontinuous events, where direct launching means that an object (here, a red one) moves continuously in a straight line to the right and collides with another object (a green one), causing it to move in the same manner (Leslie 1982). This is a perfect example of seeing in the non-epistemic sense. The experiment provides nice evidence that not only adults but also infants as young as six-month old can see_{n} certain events.

In another experiment, the effect of direct launching was compared to the continuous motion of a single object (Leslie 1984). One group of infants was habituated to direct launching and then subjected to its reversal once, i.e., dishabituated. Another group was habituated to continuous motion of a single object and then dishabituated in the same way. The result was that the first group recovered their looking more than the second group. This suggested that the infants could detect even an internal structure and parse the submovements in direct launching.

A third experiment was designed to show that the infants could perceive something more than the spatiotemporal dimension in the connection between the events that make up direct launching (Leslie and Keeble 1987). To this end, a group of infants were habituated to direct launching and then subjected to the image of delayed reaction-without collision. Another group was habituated to launching-without-collision and then dishabituated with an image of delayed reaction. In both cases, spatial and temporal contrast between the image of habituation and that of dishabituation was equal. However, the infants in the first group showed a greater degree of recovery of attention, suggesting that they had perceived a greater change in the replacement of direct launching, which is the only apparently causal sequence, with delayed reaction-without-collision, which appears non-causal.

In a final and more conclusive experiment, Leslie and Keeble compared the reversal of direct launching with the reversal of delayed reaction (ibid.). The first group of infants was habituated to a sequence in which a red object directly launched a green object by colliding with it in a rightwards direction. They were then shown a reversal of this motion where the green object came back and directly launched the red object in a leftward direction. The second group was subjected to a similar reversal in the case of delayed reaction. Again, the first group displayed a significantly greater recovery of attention, indicating that the infants in this group were responding to the reversal of causal direction as well as that of spatiotemporal direction while the infants in the second group were responding only to the spatiotemporal reversal.

These experiments provide evidence that even very young children can perceive some of the causal relations in the world. In all likelihood, there is a causal percept factor at work in infants' observations of motion. Leslie and Keeble hypothesize that there is a visual mechanism, already operating at the age of six months, which is responsible for organizing a causal percept. They conclude that "instead of causality being entirely a result of the gradual development of thought (Piaget) or of prolonged experience (Hume), an important and perhaps crucial contribution is made by the operation of a fairly low level perceptual mechanism" (ibid., p. 285).

This is perhaps to be expected. After all, causal information is vitally important for survival and adaptation. It is therefore not surprising that human beings (and also some animals) have biologically evolved in such a way that they are endowed with certain mechanisms that specifically target the causal structure of the world.

Now, how does all this fit into the analytic framework of seeing that we have discussed earlier? I think it fits perfectly well. First of all, the experimental evidence suggests that even infants can see certain objects, events and causal relations. Of course, initially, this must be a seeing in the non-epistemic sense. Later, when they are old enough, their visual experiences begin to acquire epistemic import. By the time they are ten months old, for example, they realize the significance of spatial contact for causal efficacy. In other words, they begin to see that they can move an object only by coming into spatial contact with it and use this sort of information to make predictions (Gopnik and Meltzoff 1997, p. 138). It is probably from this point onward that non-epistemic seeing becomes epistemically clothed.

6 Concluding Remarks

Since Hume's publication of *Treatise* in 1739, the regularity account of causation has dominated the philosophical scene despite repeated failures to reduce causation to exceptionless regularities, probabilistic dependencies or correlational complexes. In the last several decades, however, what I have called the first dogma of Humean causation lost its grip, and, as a result, a number of non-reductive approaches have emerged. Some philosophers of science still hope to rescue the reductionist project, but I have argued that their attempts too fail, this time in a self-defeating manner. The rescue attempt itself becomes metaphysical, inconsistent with the original intent.

Although regularities are no more believed to be ontologically privileged, they continue to play an epistemologically privileged role in many accounts. Probabilistic or otherwise, they are often taken to be our only source of causal information, our only justification for causal claims. While I do not deny their epistemological role, I have argued that that role is limited.

We have other means of epistemological access to causes, such as the method of elimination and direct observation. I have borrowed Dretske's analysis of seeing to specify the conditions under which a person can be said to observe a causal relation. This is a far better account than Hume's which makes the observability of causal relations impossible. I have then drawn attention to recent empirical studies that provide evidence against the third dogma. These studies also suggest that, contra Humeans, some of our causal knowledge is inborn. I have used these alternative sources of causal knowledge to solve the circularity problem faced by the capacity approach.

Let me conclude by saying a few words about our concept of cause and its acquisition. As Cartwright has pointed out, 'cause' does not seem to be a unitary notion (Cartwright, this volume; see also her 1999, pp. 118-121). It is highly abstract, unspecific and varied: causes may be "standing conditions, auxiliary conditions, precipitating conditions, agents, interventions, contraventions, modifications, contributory factors, ...etc." (ibid.). This may explain why our attempts to define causation fail; there is simply no single notion of cause to be defined. Perhaps that is also why we must be semantic holists in the sense I explained in section 3. To know the meaning of the term 'cause', we must know the meanings of a cluster of related terms like 'change', 'produce', 'manipulation', 'intervention', 'prevention', 'power', and the like.

How do we then acquire such a non-unitary notion? The answer, I believe, is by extension from the primordial cases. By primordial cases I mean our earliest experiences as agents. Again, developmental psychology provides some clues (see Leslie and Keeble 1997, ch. 5). Apparently, infants can distinguish people from inanimate objects very early on and become aware of the effects of their actions on such objects when they are only two or three months old. That is to say, they seem to notice the causal connection between their actions and their effects. The paradigmatic examples are grabbing a toy, sucking a pacifier, pulling a napkin, pushing a cup, and so on. This awareness in turn enables them to make predictions about which events will be followed by which actions. At this early stage, however, infants' attention is directed to the temporal dimension of the link between their actions and their consequences. Around the age of ten months, they begin to realize the importance of the spatial contact between the two. For instance, at this age they know not only that they can move objects by kicking them, but also that stronger kicks result in greater displacement.

A second primordial source of our notion of cause is the link between mental states and actions. This is established primarily through imitation. There is overwhelming evidence that infants who are only a few weeks old can imitate facial gestures such as tongue protrusion and mouth opening.

Since such imitations "require a mapping from visually perceived physical movements to internally felt kinesthetic sensations, the most fundamental of action representations," and since such sensations are closely related to mental states such as pain, there seems to be an innately established connection between these internal sensations and facial expressions (ibid., p. 130). "Similarly", write Leslie and Keeble, "simple motor plans, like the intention to move your tongue, seem to be a primitive kind of mental state, and we seem to map these plans onto perceived actions" (ibid., p. 131).

It seems to me that it is these kinds of experiences that constitute the primordial sources from which we acquire our initial concept of cause. Since it essentially derives from our actions as agents, it involves 'change', 'make happen' and 'bring about', which is the standard dictionary meaning. Once we have this 'core' meaning, we extend it to situations that are sufficiently similar to the original ones: a bee's sting, objects colliding, wind blowing leaves, fire burning paper, and so on. How we extrapolate to such cases from the 'core', which aspects of objects, their properties and relations we pick out as similar, of course, requires careful empirical scrutiny. This is all the more so in scientific contexts. One rule of thumb for extension seems to be manipulability. Anything that can be used to manipulate anything else for a certain purpose can be classified as a cause. Perhaps regularity and counterfactual dependence constitute other such rules.

This account accords well with the agency view that our concept of cause derives from our experiences as agents. Huw Price (1996) used that view to argue that the direction of causation is a projection of the direction of the means-ends relation. I, on the other hand, am a realist about powers and natures, so I believe that if it is not in the nature of a thing to do X, it cannot be manipulated to achieve X. This may create a tension between the capacity approach and Price's perspectivalism, but the discussion of this issue is better left for another occasion.

References

Aristotle (1928) *Metaphysica*. Translated by W. D. Ross. Oxford: Clarendon Press.

Cartwright, N. (1989) *Nature's Capacities and Their Measurement*. Oxford: Clarendon Press.

Cartwright, N. (1999) *The Dappled World*. Cambridge: Cambridge University Press.

Cartwright, N. (This volume) Modularity: It can—and generally does—fail. In D. Costantini, M. C. Galavotti, and P. Suppes, *Stochastic Causality*. Stanford: CSLI Publications.

Davidson, D. (1967) Causal relations. *The Journal of Philosophy*, 64, 691-703.

Davidson, D. (1968) Actions, reasons, and causes. In A. White (Ed.), *The Philosophy of Action*. Oxford: Oxford University Press, 79-94.

Dretske, F. (1969) *Seeing and Knowing*. London: Routledge & Kegan Paul.

Gopnik, A. and Meltzoff, A. (1997) *Words, Thoughts, and Theories*. Cambridge: The MIT Press.

Gopnik, A., Glymour, C. and Sobel, D. (forthcoming). Causal maps: A cognitive and computational account of theory formation. In P. Gardenfors, K. Kijania-Placek, and J. Wolenski (Eds.) *Logic, Methdology and Philosophy of Science XI*. Dordrecht: Kluwer Academic Publishers.

Harre, R. and Madden, E. (1975) *Causal Powers*. Oxford: Oxford University Press.

Hume, D. (1739/1962) *A Treatise of Human Nature*. Glasgow: Fontana.

Irzik, G. (1996). Can causes be reduced to correlations? *British Journal of Philosophy* 36, 273-289.

Leslie, A. (1982) The perception of causality in infants. *Perception*, 11, 173-186.

Leslie, A. (1984) Spatiotemporal continuity and the perception of causality in infants. *Perception*, 13, 287-305.

Leslie, A. and Keeble, S. (1987) Do six-month-old infants perceive causality? *Cognition*, 25, 265-288.

Mackie, J. L. (1980) *The Cement of the Universe*. Oxford: Clarendon Press.

Michotte, A. (1963) *The Perception of Causality*. London: Methuen.

Papineau, D. (1985) Causal asymmetry. *British Journal of Philosophy*, 36, 273-289.

Papineau, D. (1991) Correlations and causes. *British Journal of Philosophy*, 42, 397-412.

Papineau, D. (1993) Can we reduce causal direction to probabilities? In D. Hull, M. Forbes and K. Okruhlik (Eds.), *PSA 1992*, Vol. 2, East Lansing: Philosophy of Science Association, 238-252.

Pearl, J. and Verma, T. (1994). A theory of inferred causation. In D. Prawitz, B. Skyrms, and D. Westerstahl (Eds.), *Logic, Methodology, and Philosophy of Science IX*. Amsterdam: Elsevier, 789-811.

Price, H. (1996) *Time's Arrow and Archimedes' Point*. New York: Oxford University Press.

Spirtes, P., Glymour, C. and Scheines, R. (1993) *Causation, Prediction and Search*. New York: Springer-Verlag.

Spohn, W. (This volume) Bayesian nets are all there is to causal dependence. In D. Costantini, M. C. Galavotti, and P. Suppes, *Stochastic Causality*.

Suppes, P. (1970) *A Probabilistic Theory of Causality*. Amsterdam: North Holland.

Woodward, J. (1990) Supervenience and singular causal statements. In D. Knowles (Ed.), *Explanation and its Limits*. Cambridge: Cambridge University Press, 211-246.