*Ratio*, XXIV, 2, 1982

# CAUSAL SCEPTICISM OR INVISIBLE CEMENT

*David-Hillel Ruben*

A physical world sceptic may well believe that there is a physical world. Indeed, he may feel as confident of its existence as you or I. But what he will also believe is that, on the basis of his experience, he has no evidence that shows that his belief in the physical world is even reasonable, or justified, or probably true. Since he will deny that belief in the physical world can be analysed without remainder into beliefs about actual or possible experience, he will say that no amount of experiential evidence could ever deductively establish anything about the physical world. Moreover, he will deny that experiential evidence could lend any inductive support to our physical world beliefs. Some anti-sceptics argue that one's experience, when taken in conjunction with certain 'plausible' or 'modest' principles, will show how belief in the physical world is reasonable,[1] but to my mind this sort of tactic can only relocate the problem by raising the question of whether, on the basis of his experience, a person has any evidence whatever for thinking that the principles he wishes to employ are reasonable, or probably true.

I consider the causal sceptic to be in an analogous position. He believes:

(A) Constant conjunction between events of type $F$ and type $G$, supplemented by any spatio-temporal relation between the events whatever, is not enough to establish deductively that events of type $F$ and type $G$ are causally connected.[2]

and

(B) No amount of such evidence can render belief that there is a causal connection between the events of the two kinds reasonable, or justified, or probably true.

[1] See for example J.L. Mackie, *Problems from Locke*, Oxford University Press, 1976, pp. 62–7, and Michael Slote, *Reason and Scepticism*, Allen & Unwin, London, 1970, Chap. 2.

I wish to thank Mark Sainsbury, the late John Mackie, Ardon Lyon, John Watling, and Martin Hollis, who have helped make this paper better than it otherwise would have been.

[2] 'Causally connected' and 'a causal connection', here and elsewhere, are meant to cover the case in which $F$'s cause $G$'s, and the case in which they are joint effects of a single cause, or indeed bear any other sort of causal relation to one another, however remote and mediated. The contrast I wish to draw is between causal connection in this broad sense on the one hand and coincidentality on the other.

The causal sceptic thus parallels the physical world sceptic. In this paper, I wish to argue the case for scepticism concerning causal generalisations. I shall argue that no set of statements about constant conjunctions, supplemented by whatever spatio-temporal information about the relations between the events one might like, either deductively entails or inductively supports a generalisation that there is a causal relation between the conjoined events. In what follows, I use 'inductive argument' in a wide sense, to include any argument that is not deductively valid and is such that the truth of the premisses provides some reason for thinking that the conclusion is true.

Of the propositions held by the causal sceptic, (A) is certainly the least controversial. Almost every philosopher would accept that universal regularities of conjunction are insufficient for causation because it could always be a cosmic coincidence that, unrestrictedly, events of type F were constantly conjoined with events of type G. That there is this difference is indicated by our belief that if the constant conjunction is coincidental, it will not support a counterfactual claim. For example, it could have been a sheer cosmic accident that every case of any person's yawning was followed by that person's sneezing, without yawning and sneezing being in any way causally connected. Or imagine a universe of two objects, $a$ and $b$. Select quite at random two events, $e_1$ and $e_2$, in the history of $a$, and occurring at time $t$. Select two events, $e_3$ and $e_4$, in the history of $b$ and occurring at the next instant $t'$. Suppose $a$ and $b$ are contiguous. Then let an event be of type $F$ iff it is of a type of which, as a matter of fact, only $e_1$ is an instance, or of a type of which only $e_2$ is an instance, and let an event be of type $G$ iff it is of a type of which, as a matter of fact, only $e_3$ is an instance, or of a type of which only $e_4$ is an instance. In this possible world, 'Events of type $F$ are constantly conjoined with events of type $G$' is a truth, but one which lacks any causal force.

The actual world is only much more complicated than this. If God could see in a single glance the past, present and future of the actual world, he could formulate an indefinitely large number of true non-causal generalisations about the constant conjunction of events along the same lines of construction that I have suggested for the universe with two objects. But we find ourselves in a strange position with regard to whatever non-causal constant conjunction generalisations that there may actually be. As I have claimed, we have reason to believe that there are many such generalisations which are true. But the instances of such generalisations cannot give us any reason for thinking that they are true. Once the instances of a constant conjunction generalisation give us reason for taking the generalisation to be true, then the generalisation must be causal in character. Consider again the case of every person's yawning being followed by his sneezing. Suppose I knew by a complete enumeration that this has been true in the past. But if it is a true generalisation, it is true for the future, and if past instances are to provide me with grounds for believing that future yawnings will be followed by sneezings, I must be presupposing some causal tie, whether direct or indirect, between yawning and sneezing in the past, observed instances. If the instances of a constant conjunction generalisation provide me with evidence for the truth of the generalisation, then I cannot consistently take the constant conjunction generalisation to be non-causal in character.



A difference, then, between causal and non-causal constant conjunction generalisations is that the former but not the latter are inductively supportable by their instances. This is a difference that Davidson has noted: 'Lawlike statements are general statements that . . . are supported by their instances'.[3] Before we can discuss this view, I think we need to distinguish between the following three claims: (1) There is a constant conjunction between $F$-type events and $G$-type events; (2) There is a causal conjunction between $F$-type events and $G$-type events; (3) It is coincidental that there is a constant conjunction between $F$-type events and $G$-type events. The conjunction of (2) and (3) is logically inconsistent. (1), which merely states the universal generalisation, is consistent with the truth of either (2) or (3). Indeed, both (2) and (3) entail (1). Davidson's position is that (2), and claims like it, are supportable by their instances; (3), and claims like it, are not. A possible ambiguity is whether or not (1), and claims like it, are similarly supportable by their instances.[4]

[3] Donald Davidson, 'Mental Events', in Lawrence Foster and J.W. Swanson, eds., *Experience and Theory*, Duckworth, 1970, p. 92. Davidson speaks of lawlike statements, not causal ones. I do not conflate the causal and the nomological, but I do assume that Davidson would allow the same account for the relation between a causal generalisation and its instances.

[4] One might argue in this way for the inductive supportability by instances of claims like (1). Suppose we have some inductive evidence for the belief that there is a causal connection between $F$-type events and $G$-type events. Since 'There is a causal connection between $F$-type events and $G$-type events' entails 'There is a constant conjunction between $F$-type events and $G$-type events', it follows that we can have inductive evidence for the belief that there is a constant conjunction between $F$-type events and $G$-type events. Suppose further that, although we had evidence for 'There is a causal connection between $F$-type events and $G$-type events', it happens that this belief is false, that in fact

John Mackie makes a claim superficially similar to Davidson's: A causal law is a universal generalisation which, *inter alia*, is 'supported by what we take to be good inductive evidence',[5] Davidson speaks of supportability by instances; Mackie, of inductive supportability *tout court*. Do these different formulations make any difference? Can we imagine a (3)-type claim being inductively supportable, but not by its instances? Can there be any inductive evidence for the belief that it is a coincidence that all events of type $F$ are followed by events of type $G$, even if such inductive evidence does not include the instances of the generalisation?

We first need to distinguish between a wider and narrower sense of 'It is a coincidence that . . .' Suppose that all the yawners that there ever were, are, or will be, had colds that lead them to near constant sneezing. There is, then, a sense in which a reason has been given for why yawnings were followed by sneezings, namely that the yawners had colds. If one takes 'it is a coincidence that' in a very wide sense, as roughly equivalent to 'there is no reason whatever for . . .', then it will follow quite trivially that (3)-type claims cannot have inductive support of any kind. But taking 'it is a coincidence that . . .' in this wide sense will also insure that no determinist will think that there are any coincidences, and I do not believe that this was an intended consequence of the contrast between its being coincidental and its being causal that there is a constant conjunction between events of two types. On the other hand, if we take 'it is a coincidence that . . .' more restrictedly to mean that it is a coincidence only as far as the events of two specified types are concerned, that the types do not themselves supply the reason, directly or indirectly,[6] for the constant conjunction, then it is not at all obvious why there cannot be inductive evidence for a claim like 'It is a coincidence that all yawnings are

the constant conjunction is coincidental. Even if the causal connection claim is false, it may still be supported inductively by instances, and so 'There is a constant conjunction between F-type events and G-type events', which is entailed by the causal claim, will be inductively supported by instances. This means that the constant conjunction claim, (1), can be inductively supported by instances even when in fact it might only be true in virtue of a coincidence.

This argument relies on this principle: If a proposition $p$ inductively supports $q$, and if $q$ deductively entails $r$, then $p$ inductively supports $r$. This principle is controversial (but not, I think, obviously wrong, as some assume.).

[5] John Mackie, *Truth, Probability and Paradox*, Oxford University Press, 1973, p. 118.

[6] I mean by 'directly or indirectly' to cover all the sorts of possibilities mentioned in fn. 2.



followed by the yawner's sneezing', even if it is not the instances of the constant conjunction that provide the evidence. We shall be in a better position to evaluate this suggestion after we get clearer on what the form of an inductive argument for a causal generalisation looks like.

The question I first want to answer is this: what is the form of the simplest inductive argument that yields justified belief in a causal generalisation? That is, we are looking for the minimum number of evidential premisses which are required to (non-deductively) support or justify a general causal conclusion to any degree whatever. It might be thought that such an argument had this form:

A: (P) All observed events of type $F$ have been followed by events of type $G$.

---

(C) Events of type $F$ are causally connected with events of type $G$.

Those who think that A is the form of such an argument would concede that the support (P) gives to (C) is weak. My point is that this premiss on its own can give *no* reason whatever, however weak, for thinking that the conclusion is true. At best, A could only represent an enthymeme of some acceptable inductive argument for a causal conclusion like (2). There is no supportability of claims like (2) by their instances alone.

My argument for this is as follows: If A gave the form of an inductive argument for general causal belief, with no missing premisses, we would always be in a position, on the basis of observed constant conjunctions, to conclude that the causal generalisation conclusion is more probable than not. But this is not so. We know that there are an indefinitely large number of coincidentally true generalisations about the constant conjunction of events of two types, even if we never knew which they were. Moreover, we can conjecture that, relative to the number of these coincidental constant conjunctions, the number of causal constant conjunctions must be small. Moreover, it is plausible to believe that we in fact have observed many of the instances of those coincidental constant conjunctions, whether knowingly or not. Of those observed constant conjunctions that figure into our experience, I cannot see that it is more likely that they should be the instances of the causal conjunctions than of the coincidental

ones, although this might be true of that proper subset of observed constant conjunctions in which we take any real interest. So the observed conjunctions ought to be better evidence for the belief that the correlation is not causal than for the belief that it is causal. The premisss (P) should lead to the probable truth of *not*-(C), rather than to the probable truth of (C)! A cannot be the form of the inductive argument that we seek.

Someone who holds that causal generalisations, unlike non-causal ones, are inductively supportable by their instances need not believe that there are no other differences between them. So we might inquire about the minimum number of premisses we would need to add to (P), in order that the conjunction of the premisses is able to support or justify the general causal conclusion to any degree whatever. If there are other differences between the causal and non-causal constant conjunctions, then perhaps those differences might provide us with the clue we need in finding whatever additional premisses we need.

At this point, there is what I regard as an inevitable lacuna in my paper. I believe that at least one of the additional premisses necessary for the support of the causal conclusion in A is itself about causal connection. I do not believe that the addition of any set of premisses to (P) in A will permit the premisses to support, to any degree whatever, the causal conclusion unless at least one of those premisses is itself about, or presupposes the existence of, causal connections. I cannot prove that this is so. The best I can do is to remind the reader of the other suggested candidates in the philosophical literature for the role of the missing premiss(es) in the argument, candidates against which I take it that there are decisive objections. These alternative candidates are espoused by: realism, empiricism, and humanism.[7] To pursue the ways in which

[7] First, there is the thesis championed by realists that a difference between causal and non-causal conjunctions is the presence in the former but not in the latter of underlying mechanisms or structures that connect the events of the two types: 'For a generalisation to qualify as a law the properties specified by its contained predicates must not only be universally co-instantiated, they must be connected in some way . . . [by] the presence of some underlying generative mechanism or structure . . .' (Colin McGinn, *Aristotelian Society Sup. Volume*, LII, 1978, p. 205). This account must either accept that there is some level of reality at which there are causal constant conjunctions unconnected by any further mechanism of any sort, or must accept a vicious ontological regress of levels.

Second, there is the thesis championed typically by empiricists such as Braithwaite that would seek to distinguish between statements of causal and non-causal constant conjunction by the derivability of the former from higher level generalisations. Such an account leaves unexplained the causal character of whatever highest level, underived causal laws there may be, except by resort to *ad hoc* considerations (e.g., that theoretical terms occur in them) which look in any event to be begging the question at issue, cf., R.B. Braithwaite, *Scientific Explanation*, Cambridge University Press, 1964, pp. 293–318.

Finally, there is the humanist thesis put forward by those who consider that testing, experiment, or human intervention afford us a wholly new sort of evidence for causation than that which arises from observation. Von Wright is an example of this way of thinking (G.H. Von Wright, 'On the Logic and Epistemology of the Causal Relation', reprinted in *Causation and Conditionals*, ed., Ernest Sosa, Oxford University Press, 1975, quote. from p. 105). Suppose we have observed may cases in which *F*'s have been followed by *G*'s, and no cases in which they have not. I then wait for an occasion on which no event of type *F* is present. I then bring such an event into existence, and notice that an event of type *G* occurs. 'This operation will . . . "impress" us strongly and confirm the surmise . . . that the regular sequence of *p* and *q* in the past was no mere accident but signified a causal tie between the two factors'. But it is hard to see how this is relevant. If it were a true non-causal generalisation that events of type *F* and type *G* are constantly conjoined, it would be as true for cases in which the *F*-event is artificially produced as it is for cases in which the *F* event occurs naturally. All that interventions can do is to add to the number of cases in which observed *F*'s have been followed by *G*'s; so if a causal conclusion does not follow from the evidence of observed correlations in which the antecedent event occurs without human intervention, it will not follow when the observed correlations are expanded to include cases in which the antecedent event is made to occur by a deliberate human act.



these alternatives do fail would involve another paper on its own. Rather, I wish here to simply put forward what I regard as the full form of the weakest possible inductive argument for a causal generalisation. I use 'inductive argument' in the widest sense to include any argument such that the premisses give some sort of non-deductive support for the conclusion.

It is obvious, I think, that the argument for which we are looking cannot be an enumerative inductive argument as traditionally conceived. The conclusion, that events of type *F* are causally connected with events of type *G*, is not a simple generalisation of the evidence given in the premisses, that all observed events of type *F* have been followed by events of type *G*, in the way in which, for example, all rubies are red is a simple generalisation on the evidence that all observed rubies have been red. The conclusion in A is not that all (observed and unobserved) events of type *F* are followed by events of type *G* – which would be the simple generalisation – but rather that events of type *F* are causally connected with events of type *G*. If this argument is not an enumerative inductive one, what kind of inductive argument, in the wide sense of 'inductive argument', is it? I believe that it is a theoretical inductive argument, or an inference to the best explanation, as that has been discussed in the literature in other connexions.[8] Such an argument has this form:

[8] A simple account of this kind of non-deductive inference is given by Wesley Salmon, *Logic*, Prentice-Hall, 1963, pp. 76–88, and discussed more fully in his contribution to *Explanation*, ed., S. Korner, Blackwell, 1975. It has also been discussed by N.R. Hanson,

B ($P_1$) If an hypothesis $H$ is true, then our observations will be of a certain sort.
($P_2$) Our observations are of that sort.
($P_3$) Hypothesis $H$ is a better explanation of our observations being of that sort than is any alternative hypothesis.

---

(C) Hypothesis $H$ is true.

The basic idea behind construing arguments to general causal conclusions to be of this form is that the hypothesis that events of the constantly conjoined types are causally connected is sometimes a better explanation for the fact that the events of the two types have been observed to be constantly conjoined than is any alternative hypothesis equally consistent with the same evidence – for example, better than the hypothesis that the observed constant conjunction is a sheer cosmic coincidence. The argument we need, with form B, looks like this:

($P_1$) If the hypothesis that events of type $F$ are causally connected with events of type $G$ is true, then we will observe a constant conjunction of $F$-type events and $G$-type events.
($P_2$) All observed events of type $F$ have been followed by events of type $G$.
($P_3$) The causal hypothesis is a better explanation of the observed constant conjunction than is any alternative hypothesis.

---

(C) The hypothesis that events of type $F$ are causally connected with events of type $G$ is true.

What makes the causal hypothesis sometimes better than the hypothesis that the constant conjunction is coincidental? Which

*Patterns of Discovery*, Cambridge University Press, 1975, pp. 85&ff., and Gilbert Harman, 'The Inference to the Best Explanation', *Philosophical Review*, 74, 1965, pp. 88–95. The type of inference is sometimes called 'abduction', 'retroduction', and 'hypothetical argument', as well as 'theoretical inference' and 'inference to the best explanation'. My claim is only that there is a type of non-deductive inference with the form I describe. I do not need to accept Harman's stronger claim that 'all warranted inferences which may be described as instances of enumerative induction must also be described as instances of the inference to the best explanation'. Peirce, for example, held that induction and abduction 'are utterly irreducible', although I do not accept his account of an abductive argument as one whose conclusion states that 'something may be' (rather than probably is so).



hypothesis is better may depend on numerous factors, but one such factor that plays a key role is whether the hypothesis 'fits in' with our other causal beliefs and hypotheses. Some have argued, and others denied, that the events of the two types must be spatially and temporally contiguous, but even if this were necessary, it could not be sufficient, since there can be cosmic coincidences between events of two types which stand in this, or any other, spatial and temporal relation to one another. Sometimes too considerations such as simplicity are mentioned, in the sense that it is somehow simpler to assume that the observed constant conjunction is a result of a causal connection than it is to assume that it results from a cosmic coincidence. If this use of 'simplicity' means that the causal hypothesis fits more naturally into the web of our other causal beliefs, I agree. But it should not be thought that causal generalisations are themselves 'simpler' than coincidentally true generalisations. It is true that, in the example of the imaginary universe we considered, 'All events of type $F$ are followed by events of type $G$' came out as a coincidental truth only because of the seemingly artifical way in which 'being an event of type $F$' and 'being an event of type $G$' were defined, and if we wanted to extend such examples to cover cases in which there were many events of the two kinds in question, the definitions of what it is to be an event of the one type or other will become very complicated indeed, artificially composed (or so it will seem to us) of a great number of disjuncts. But on the other hand, we must not understate how complicated the types will be whose instances are unrestrictedly causally connected. They will certainly not be of the simple form, 'All strikings of matches are followed by lightings of matches'; for once we put in all the qualifications necessary in order to obtain an unrestrictedly true causal generalisation, we shall find ourselves with something complicated indeed. In terms of simplicity of the generalisation itself, there is not liable to be much to choose between.

Rather, I claim that a major consideration in favour of moving from an observed constant conjunction to a causal hypothesis as the best explanation of the observed conjunction, in those cases in which we are warranted in so doing, is that such an hypothesis fits with our other causal beliefs. Of course, I do not claim that 'fit' with our other causal beliefs is a sufficient condition by itself for warranted belief in a causal generalisation; my contention is only that it is necessary part of a sufficient condition which also includes the observed constant conjunction. In this way, I think

my position escapes some of the more obvious 'relativistic' conclusions it might be thought to have.[9]

Immediately after arising each weekday morning, I hear the post being pushed through my letter box. Suppose that I sleep immediately adjacent to the letter box, so that the events are spatially as well as temporally contiguous. I do not take this observed constant conjunction as a support for a causal generalisation, because I judge, in the light of my other causal beliefs, that the best explanation for this observed correlation is not a causal hypothesis but a non-causal one. I explain the observed conjunction as arising from a coincidence, due to the coincidence between the time at which I arise and the time of the delivery of the first post. Notice that this judgment does not involve my finding or failing to find a connecting mechanism between my waking and the delivery of the post. It is perfectly true that one of the considerations relating to the 'fit' of the hypothesis with our other causal beliefs may include our beliefs as to whether it is likely that there are connecting causal mechanisms. So when and where relevant, the question of 'fit' can include the question of connecting mechanisms, but the question of 'fit' is the more basic, and permits the logical possibility of there being genuine causal action at a distance, as well as permitting us to deal with the vexing question of there being a level of reality ultimately basic, in the sense that between the events occurring at this level, no further connecting mechanisms or events could be found.

We are now in a position to answer the question we raised earlier: can we have inductive evidence for a belief that it is coincidental that events of two types are constantly conjoined? First, it is evident that we can have evidence that an observed constant conjunction is coincidental, because sometimes the hypothesis that it is coincidental fits better with our other causal suppositions, as in the case of the post-through-the-letter-box. Can I have reasons for 'projecting' this observed constant conjunction to unobserved, and hence all, cases? I think that this could arise in at least two ways; I might have reasons to think that the observed cases exhaust all the cases there ever will be; or I might have reasons for thinking that future cases of the occurrence of the antecedent event will also be coincidentally followed by the consequent event. In the first sort of case, I might have reasons to

[9] For instance, entrenchment of a causal belief, without support of an observed constant conjunction, is insufficient to provide support for some other causal belief that 'fits' with it.



expect my own death this afternoon, and hence reason to think that 'All cases of my waking are followed by post being pushed through the letter box' is an unrestrictedly true generalisation, since there will be no further cases of the antecedent event. In the second sort of case, I might have reasons for thinking that my future weekday wakings will occur at the same time as they have in the past, and good reasons for thinking that the delivery of post will occur at the same time as in the past. If so, Mackie is wrong. Even type (3) claims are inductively supportable. But in neither of these cases does our evidence include the past instances of the constant conjunction. Davidson's formulation is vindicated.

How could it be a coincidence, someone might ask, that all events of type *F* are followed by events of type *G*? If we have reasons for projecting the constant conjunction, it cannot really *be* a coincidence that all events of the one kind are followed by events of the other. In one way, this thought is perfectly correct, but it shifts between the wider and narrower sense of 'It is coincidental that . . .', that I mentioned earlier. If we have reasons for thinking that there will be no more events of the antecedent type, or that the coincidence will continue in the future, it follows that it cannot be coincidental that the generalisation is true, in the wide sense. But in the narrow sense, it is perfectly coincidental that the unrestricted constant conjunction generalisation is true, because the reasons have nothing whatever to do with any connection between the events of the two types, however indirect. In neither of the two cases do events of one type cause events of the other, or are they joint effects of any single previous cause, however remote. Any intuitions to the contrary about its not being coincidental I diagnose as arising from a shift between the wide and narrow sense of 'It is a coincidence that . . .'.

How does any of this relate to what I have called 'causal scepticism'? I do not, of course, deny that experience has something to do with justified causal belief. A necessary, if insufficient, condition for having justified beliefs about causal generalisations is having observed constant conjunctions of events of two types. But the argument has been that such experience, on its own, cannot yield probable general causal conclusions. What more is needed to make, as it were, experience support or justify general causal beliefs always includes other causal beliefs, which in their turn are supported by experience only with the help of other causal beliefs. There is no way in which experience on its own, or with *causalitätfrei* principles, can support the causal structure. And this means,

too, that the sort of argument form we sketched, (B), is incapable of yielding an argument which could serve as an answer to this question: Why do we think that there are *any* causal constant conjunctions whatever? Do we have any reasons for believing that there are not *just* cosmic coincidences?

The analogy between this view of causal belief and scepticism concerning justified belief in a physical world is indeed striking. Even a physical world sceptic will accept that it is possible to 'justify' some physical object claims on the basis of others. Similarly, one can 'justify' causal beliefs on the basis of others. Some will see in this, as they would in the case of physical objects, a scepticism about the real possibility of our having any genuine causal knowledge at all. What sort of 'justification' is it, they may ask, if some causal beliefs have to be used to support others? Others will applaud the result as showing the bankruptcy of the classical paradigm of having justified belief. If justification involves this sort of circularity, so be it, they will say. For my part, I do not mind which lesson is drawn, because I cannot see that there is any real difference between those two ways of looking at the matter. If there is some argument which concludes, on the basis of our experience alone, that there are some causal constant conjunctions rather than none at all, or that there is a physical world, rather than only sets of sense data, I certainly have not sketched any such argument here.

*Department of Social Sciences and Humanities*
*The City University*
*Northampton Square*
*London EC1V 0HB*

*Ratio*, XXIV, 2, 1982

# MACHINES AND MISTAKES

*Stuart Goldkind*

## 1. INTRODUCTION

The central claim of Mechanism is that human beings are nothing more than (very complicated) machines. This claim has been found quite disturbing by some people, and over the years there have been numerous attempts to remove the annoyance by providing a convincing refutation. The most familiar form that such attempts at refutation take is that of an argument purporting to show that humans are possessed of some quality, feature, characteristic, or ability, X, which machines cannot possibly have. In order to show that machines cannot have the trait X, some sort of appeal is generally made to the 'nature' of machines. (For example, in arguments based on some famous results of Kurt Gödel, the claim is made that machines by their very nature cannot produce certain theorems). The evidence for humans possessing X, on the other hand, is usually empirical: we have merely to look in order to see that humans have X.

In this paper we will be looking at the concept of error to see if it provides a plausible candidate for X; that is, we will be examining the claim that machines, unlike humans, lack the ability to make mistakes. We will be considering, as possible support for this claim, the view that there is some sort of logical or linguistic contradiction in the very notion of a completely programmed entity which is capable of error. The 'nature' of machines that is appealed to in arguing for this view is their 'rigid' 'programmed' constitution; this, it might be held, is what makes it impossible for a machine to make the same sort of mistakes that humans make.

In considering this view it will be found that: (1) The by now fairly widely recognized dissolution of the hardware-software distinction supports a perspective on machines which makes it more plausible to regard them as beings which can perform actions (rather than as ones which can 'merely follow their programs'); (2) Even present day machines can make something which is like a human mistake in at least the respects that (i) the machine makes the 'mistake' when functioning normally and (ii) the 'mistake' is such that under other circumstances the machine could have avoided making it.

In his celebrated article, 'Computing Machinery and Intelligence,'