Meeting the hare in her doubles: Causal belief and general belief

R. M. Sainsbury

I

> A CAUSE is an object precedent and contiguous to another, and so united with it, that the idea of the one determines the mind to form the idea of the other, and the impression of the one to form a more lively idea of the other.

This is standardly called Hume's second definition of causation. So construed, it gives rise to familiar problems. As a definition of causation it is circular, for 'determines' is a synonym for 'causes'. Given that Hume himself criticizes using synonyms in attempts to define causation,[1] it would be odd if he were to engage in the practice himself. It is also visibly not equivalent to the first definition, which makes no mention of minds.

These worries disappear if we treat it not as a definition of causation, but rather as a definition of what it is to believe in causation: for short, as a definition of causal belief. One way to make this explicit, guided by Ramsey's phrase 'a habit of singular belief',[2] is as follows:

> S implicitly believes that Fs cause Gs iff S is in the grip of an F–G regularity, i.e.
>
> (i) coming to believe an F has occurred causes S to believe that a G will occur nearby and later;

(ii) coming to believe that a G has occurred causes S to believe that an F has occurred nearby and earlier.

Although this is not equivalent to the first definition, equivalence is no longer to be expected, for one would not suppose that causation and causal belief are the same thing; there can be unsuspected causes and false causal beliefs. Moreover, relative to the ambition of defining causal belief, the definition is not circular, for the concept of causal belief does not occur in it: although the concept of causation is used, it is not used within the belief context. The definition could therefore be useful to someone who understands causation but needs an account of the kind of psychological state a person is in if they believe some causal generalization.

The definition is not the same as Hume's second definition, for I have made no attempt to do justice to Hume's first phrase ('an object precedent and contiguous to another'), which may perhaps be a condensed repetition of the first definition. Perhaps, in this context, 'an object' means 'a kind of object'; and precedency and contiguity between kinds of objects is naturally understood as entailing the constant conjunction of all objects falling under the kinds.

On this interpretation of the second definition, the first and second definitions would not only be definitions of different things, they would also differ in another way. Whereas the first definition, in my opinion, defines singular causation (in terms of the particular cause and its effect being subsumed under a regularity) the second definition addresses not belief in singular causation but belief in causal generalization. I will not defend my exegetical claims (which are far from idiosyncratic), for the main theme of the present paper is a question which is of interest even if the exegesis is defective: can a definition of belief in a causal generalization differentiate it from belief in a non-causal generalization? Before approaching this directly, the next section sets some background, including the notion of implicit as opposed to explicit belief.

II

The beauty of the second definition, under this construal, is that it explains causal belief as something not requiring belief in a causal content, and this makes room for various reductive or eliminative options for Humeans (see §III below). It also makes causal belief something that could be possessed by creatures lacking the concept of causation. Let us call a belief concerning the unexperienced 'inductive'. It is beyond doubt that Hume thought it important to point out that non-human animals have inductive beliefs (despite not being capable of argumentation) and that such beliefs normally arise only through the experience of constant conjunction. According to the present account of the second definition, it would follow that any creature with a normally produced inductive belief has a causal belief, that is, implicitly believes something of the form Fs cause Gs. This is indeed what we find Hume claiming:

> they [sc. non-human animals] become acquainted with the more obvious properties of external objects, and gradually, from their birth, treasure up a knowledge of the nature of fire, water, earth, stones, heights, depths, &c., and of the effects, which result from their operation.[3]

The final phrase, about 'the effects, which result from their operation', shows that the knowledge of fire, water and so on includes causal knowledge. In the light of this context, we can read Hume as attributing causal knowledge to the old and canny greyhound who 'meets the hare in her doubles': the dog's 'conjectures' are that being chased one way by the younger dogs will, given the lie of the land, drive the hare back this way.

Even those sympathetic with an account of causal belief which would make it correctly attributable to creatures lacking the concept of causation may feel that we should distinguish between that way of having a causal belief and a way of having such a belief available only to a creature who does employ the concept of causation. For example, it might be held that we need to employ the concept if we are going to make use of Hume's 'rules to judge of causes and effects', and so entertain causal generalizations merely hypothetically. This requires an account of a causal generalization being before the mind without being believed. If evidence then leads to belief, this is explicit belief, apparently involving a relation to a causal content. It was in deference to this view that I took the definiendum of the second definition to be implicit causal belief. But what is implicit belief?

An explicit belief involves, at least superficially, a relation between a believer and a content. Being in a state of believing a content generally makes a difference to dispositions to act. This enables us to indicate the schema of a definition of implicit belief:

S implicitly believes that p iff S has dispositions characteristic of explicit belief that p.

This is only a schema, for it does not specify the relevant 'characteristic' dispositions. Not all can be involved, as some of these are dispositions to exercise the concept of causation. The schema ensures that explicitly believing that p entails implicitly believing that p, but not vice versa. Those unpersuaded by the explicit-implicit distinction can just set it to one side.

III

If one has a definition of causation, why would one need an additional definition of causal belief? If p is defined as q, then surely the belief that p just is the belief that q?

This assumes that definition induces a relation which permits substitution salva veritate even in hyperintensional contexts. It is reasonable to suppose, however, that this is not how

Hume thought of his first definition, that is, his definition of causation, for he thought that this deflated the ordinary conception. In causal belief we aspire to believe not just that things are constantly conjoined, but that they are also necessarily connected. Since we tend not to have the right view even of what we are aspiring to in exercising the concept of causation, it would be incorrect merely to equate our causal beliefs with beliefs in constant conjunctions.

The considerations which follow are supposed to apply, with appropriate adjustments, regardless of the precise version of Humeanism adopted. However, the presentation requires a definite version, one which settles, for example, whether causation is to be reduced to constant conjunction, so that it is something real and potentially knowable, or whether it is to be eliminated in favour of constant conjunction, so that there is no intelligible concept of causation, constant conjunction being merely the closest intelligible approximation. The version I will discuss in detail is a form of reductionism: the content of a general causal proposition is the same as that of a proposition affirming constant conjunction among kinds of events, along with priority and contiguity in the pairings of tokens of the kinds. Causation is thus a perfectly genuine, real and knowable phenomenon, and is so even for Hume, provided that we do not read him as sceptically denying that we possess inductive knowledge. (I think the non-sceptical reading is correct, and it is supported by what I have already quoted of Hume's view of the inductive knowledge which non-human animals possess.) When a generalization involves a causal idiom, however, it can be believed in a different way from the way in which the same proposition, or a proposition of the same kind, can be believed when expressed in causal-free idiom. This difference in mode of belief is easily confused with a difference in content believed, thus giving rise to the widespread myth that causation is something more than constant conjunction. We generally fail to recognize that a general causal fact is just a fact of constant conjunction (plus priority and contiguity). This is why the second definition is important: it

describes the mode of belief, being in the grip of a regularity, which constitutes a special way of believing in a constant conjunction, a way which makes us wrongly think we are believing in some kind of necessity in the world. The projectivist confusion could be expressed as the move from

believing something to be an F caused the belief that it is a G

to

its being an F made it a G (by a necessary connection).

The issues on which we are about to embark have an interest beyond the exegesis of Hume. For example, some versions of functionalism, the thesis that each mental state is individuated by its location in a causal network of other states, could use something along the lines of the second definition. We cannot expect to be able to define what it is implicitly to believe arbitrary kinds of content in the way proposed for causation. For example, the following are dubious:

i) S implicitly believes that not-p iff it is not the case that S believes that p.

ii) S implicitly believes that a is F iff, concerning a, S believes that it is F.

iii) S implicitly believes that something is F iff there is something which S believes to be F.

iv) S implicitly believes that (p or q) iff (S believes that p) or (S believes that q).

v) S implicitly believes that (p because q) iff (S believes that p) because (S believes that q).

Since we can be agnostic, we can fail to believe p without believing not-p. Non-conceptual creatures might exemplify the equivalence in (ii), but according to Fregeans a conceptual creature could believe of Hesperus that it is visible in the morning without believing that Hesperus is visible in the morning. Unless intuitionist norms are invariably adhered to, there are cases in which one believes that something is F without believing of anything in particular that it is F; the point carries over to disjunction. The final case, (v), is close to our official target. The most obvious counterexamples are ones which involve causal connections between things neither of which is the cause of the other. For example, I may believe that the toast is done because I believe it is emitting a certain smell, but the proposition 'the toast is done because it is emitting a certain smell', understood as an identification of a cause, is not one which I accept: smell-emission does not cause doneness. In such cases there is a causal connection (the heat that caused the doneness caused the smell), even if it is not one which sustains the application of 'because'. As I will elaborate at the start of §IV, I will not press Hume's account to adapt to the distinction between causal connection and causation.

A causal belief is a belief with causal content; but on the version of Humeanism adopted here, a belief with causal content is a belief with merely a certain kind of constant conjunction as its content. What makes some such beliefs special is the mode of believing, not the content believed. We can also have beliefs about beliefs: we believe we have beliefs with a content involving not merely causation (that is, on the present formulation, constant conjunction) but also necessary connection, which is unintelligible. These second order beliefs do have a distinctive content, namely, the special mode of causally believing. The distinction can easily be slid over, and this is just what we find in Hume: sometimes the 'idea of necessary connection' is explained in terms of the determination mentioned in the second definition, and sometimes in

terms of awareness of that determination (an awareness not mentioned in the second definition). The present formulation equates what is distinctive in causal belief with the determination, and this may become a distinctive content for a second-order belief, one involving awareness of the determination.

IV

The main questions of the paper can now be expressed as follows: is there a distinctive way of believing a generalization which has this feature: when we believe a generalization in this way, we (wrongly, in Hume's view) take ourselves to be believing something which involves necessity? If so, does being in the grip of a regularity constitute that mode of belief? We do not need to raise more familiar questions, like whether there are counterexamples to the reduction of causation to constant conjunction. Counterexamples are likely to be allegedly non-causal constant conjunctions. But it might be that if they are believed in the special way, then they are, wrongly even by the lights of those who allow necessity, treated as necessity-involving. So our present question has some independence from questions about the correctness of Hume's first definition.

In the present formulation, the Humean is committed to the possibility of at least two ways of believing a constant conjunction: a way distinctive of causal belief, and some other way. This is because, as I assume, there are generalisations which are merely 'accidental' and are not believed in the causal way and are not supposed to correspond to any causal fact, for example, 'All the coins in my pocket are silver'. This is consistent with the reduction of causation to constant conjunction, for the reduction identifies only a specific kind of generalisation with causation: generalisations quantifying over specified kinds of events and affirming priority and

contiguity of the pairings of the tokens. However, there are complications in saying in more detail just what a Humean could reasonably take as her aim.

The second definition does not speak directly to the mode of belief distinctive of belief in a singular causal statement. According to the first definition, a singular causal statement is true only if some regularity subsumes it, but because no regularity is entailed, belief in the singular causal statement intuitively requires only belief in the existence of a regularity and does not require believing one; hence it cannot require being in the grip of one. Perhaps belief in the existence of a subsuming regularity can in turn be indirectly explained in terms of grip, possibly as the disposition to fall into the grip of any regularity which appears to subsume any singular case which is believed. But I leave on one side this problem for the Humean, and focus only on belief in causal generalisations. Of these, there are at least two kinds. There is a kind which simply summates a number of singular causes, for example:

All the players on the team make their spouse happy.

These are 'accidental' (or non-projectible) general causal truths. There is nothing (I am assuming) about being a player on the team which promotes making a spouse happy. Given the relation of such generalizations to singular causal statements, it seems no more appropriate to say that belief in them requires being in the grip of a regularity than to say the same of any of the singular causal beliefs they summate. By contrast, there are causal generalizations which we think of as non-accidental, for example, 'Dropping wine glasses makes them break'. In possibly non-Humean terms, the contrast is that in the accidental cases no causal regularity is specified, not even in a vague or incomplete way, whereas in the non-accidental cases there is at least a sketch of a relevant regularity. We can put matters this way only if we have some adequate conception of a kind of regularity not exemplified by the regularity that all the players on the team make their spouses happy. The idea behind the second definition is that Hume can mark the

contrast between these kinds of regularity not in terms of difference of content or difference of kind of fact affirmed but merely in terms of difference in mode of belief: to believe in the non-accidental way is to be in the grip of a regularity, but this is not so for belief in accidents, even when these accidental cases contain a causal content. (The accidental causal content's relation to belief can be addressed only after an account of belief in singular causal statements is in place.) The grip mode of belief needs to be what is distinctive of non-accidental generalizations, the other cases involving one or more kinds of non-grip modes of belief. On this Humean position, the contrast between accidental and non-accidental generalizations, which at first appeared to be a contrast of content, is dissolved into a contrast of mode of belief. It is something subjective, in that, with no cognitive error, the same thing may be believed in the grip way by some but not all subjects, even if they coincide in informational state; in this sense the contrast is, as Hume might have put it, 'in the mind'.

The first phase of the argument is to try to discover a mode in which generalizations can be believed which does not amount to grip: a non-grip mode. One could just stipulate some primitive mode, but this would be unenlightening. Since believing that all F*s* are G*s* does involve a commitment to believing to be G whatever one believes to be F, we need to have as much detail as possible about how being in the grip of a regularity is something more than this. The theme of this section is that the account of the non-grip mode of belief requires successive enrichments to avoid counterexample, and we end up with the second definition. This is not what was wanted, for it means that the second definition does not, after all, identify a distinctive mode of belief, the grip mode.

A first proposal for a non-grip mode of believing that all F*s* are G might be:

1) For all *x*, S believes that if *x* is F then it is G.[4]

Although universal quantification occurs in both definiens and definition, there is no circularity relative to the aim of defining general belief, for in the definition the quantifier lies outside the belief context. The definition looks quite different from the second definition, so there is a good prospect that the definitions pick out quite different psychological states: the grip mode needs to guarantee the presence of non-grip mode, but not conversely. However, (1) is not correct, for it is not the case that anyone who believes any generalization is thereby belief-related to every object in the universe and believes a singular conditional concerning it. Even if the domain of the quantifier in the definition is somehow restricted, for example to just the objects S encounters, it is still highly implausible to suppose that general belief requires having so many singular conditional beliefs about wholly irrelevant objects (non-Fs).

This problem could be addressed by getting the conditional out of the belief context:

2) For all x, if S believes that x is F then S believes that x is G.

This might be without counterexample (i.e. there might be no object S believes to be F without also believing it to be G) even when S does not believe that all Fs are Gs, for S believes that there are Fs which are not G. I believe there are white geraniums, and, being consistent in at least this respect, I do not believe that all geraniums are red. But all the geraniums I have come across have been red, so it is true of each thing that if I believe of it that it is a geranium I believe of it that it is red. (2) can be true even when S does not believe the generalization.[5]

It seems that we need to get some kind of actual or potential transition into the picture. I can think of two (potentially combinable) ways to achieve this: by relying on counterfactuals or by relying on causation. One might improve (2) by modifying it so that the belief that x is F

causes the belief that x is G. In the state envisaged in the previous paragraph, it is arguably false that my believing something to be a geranium causes me to believe it to be red; so that counterexample would be avoided. But the price is high: the 'improvement' would lead to something barely distinguishable from the second definition, and so would not be an advisable route to take in the service of finding a distinct, non-grip, mode of belief in generalizations.

This is a route we will, in the end, be forced to take; but let us first try instead to bring counterfactuals to bear:

3) For all x, if S were to come to believe that x is F then S would come to believe that x is G.[6]

This is not right as it stands, for some ways of coming to believe of an object that it is F are ways of coming to believe of it that it is not G, and these would typically lead S to abandon the generalization rather than to believe that the object in question is G (as well as not-G). (3) would wrongly represent S as not in fact believing that all Fs are Gs. For example, it might be true that I believe that all the people in the room speak English, but also true that were a reliable informant to assert something I did not previously believe, namely that Pierre is a monolingual Frenchman and is in the room, I would come to believe that Pierre is in the room but would not go on to believe that Pierre speaks English. I would change my mind about whether everyone in the room speaks English. But as things are, I do believe that everyone in the room speaks English, even though I am not in the state specified in (3).

Peacocke gestures towards a solution in terms of modes of presentation: (3) should hold for any object presented under a minimal mode of presentation like 'the next one to be encountered'.[7] Relating this to the example just given, the idea is that if Pierre were presented

simply as the next person in the room to be encountered, I would go on to believe that he speaks English. Any single suggestion of this kind will be inadequate, since if, for example, one believes that

> I'll never encounter an honest politician (where this equals: all honest politicians are things I'll never encounter)

it will not be counterfactually true that were an honest politician presented under the mode 'the next one to be encountered' one would believe that one will never encounter it.

Fixing up (3) will no doubt involve introducing caveats about what other information is acquired in the course of coming to believe of something that it is F: ideally, this information should not itself carry reason to believe that the object is not G. Such caveats are in any case required by the present formulation of the second definition. All beliefs are defeasible, even those in the obtaining of generalizations one takes to be non-accidental. One way in which they can be defeated is by finding a counter instance, an F that is not a G. That one would abandon belief in a generalization were one to encounter a counter instance does not show that one does not in fact believe it. So both the second definition and (3) will need to be supplemented by some proviso allowing the possibility of coming to disbelieve a generalization one in fact believes. One could add something like the following: provided that S does not, in coming to believe of x that it is F, come to believe anything else of x which causes (or would cause) S to believe that x is not G.

We are now in danger of obliterating any substantive difference between the grip mode of belief, as characterized in the second definition, and any other. If we set aside those features of the second definition specially adapted to events, the difference between it and (3) is just that

shown in the following skeletons (omitting the envisaged provisos), where p and p′ range over singular instances of the generalization in question:

coming to believe that p causes one to come to believe that p′;

if one came to believe that p, one would come to believe that p′.

I will suggest that each skeleton needs to add the feature of the other which it lacks: the first skeleton needs to add counterfactual dependence and the second needs to add causation. Then the skeletons will converge, and the difference will well and truly have been obliterated, and thus Hume's project will have run into the ground.

In the first skeleton, corresponding to the second definition, 'causes' matches 'determines', a semantically somewhat mysterious 'continuous present'. The present tense in the definiendum can be taken as 'Believes (timeless) at t'; but one can believe some generalization at a time without the causal transaction the second definition envisages occurring at that time, so we cannot treat the present tense in 'determines' in the same way. If we interpret the definiens as a quantification over all actual acts (by S) of coming to believe that p, past, present and future, to the effect that each causes an act of coming to believe that p′, we fail to allow for the fact that one can believe a generalization today and not tomorrow, in which case one would (today) believe the generalization while not being in the (transtemporal) state the second definition would require. If we interpret the continuous present (in a way not generally justifiable – but we are trying our best!) as saying that belief in p has caused belief in p′ up to now (the time of the belief in the generalization), the definition is insufficient, since the believer may have abandoned the belief after the last occasion for the causal transaction. The second skeleton overcomes this difficulty, in that the time associated with belief in the generalization can be linked with the time

of the believing in the definiens. Expanding the skeleton a bit: for all times t, one believes a certain generalization at t iff if one came to believe at t that p, one would come to believe at t + ε that p′. This solution is available because a modal fact can be true at a time when nothing relevant is actually happening. It is hard to find an interpretation of Hume's 'determines', other than the counterfactual one, which gives the second definition any promise of correctness. So let us amend it accordingly, taking the new skeleton of the second definition to be:

(*) if S came to believe that p this would cause her to come to believe that p′.

The attempt to describe the non-grip mode of belief in a generalization, using the skeleton 'if one came to believe that p, one would come to believe that p′', is inadequate. Believing even an accidental generalization, like 'All the coins in my pocket are silver', requires it to be the case that coming to believe an instance of the antecedent should potentially produce belief in the corresponding instance of the consequent. In some actual or counterfactual circumstance in which I come to believe that this is a coin in my pocket, my coming to believe that it is silver does not manifest my belief in the generalization if it arises independently (e.g. by perception); to manifest my belief in the generalization it should be caused by the belief that it is a coin in my pocket. If we add this causal feature to the earlier counterfactual skeleton, we get:

(**) if S came to believe that p this would cause her to come to believe that p′.

This is the same as (*), the skeleton we arrived at when trying to improve the second definition towards something acceptable. We have converged from different directions on essentially the same condition. This does not bode well for the attempt to find substantively different accounts

of belief in grip mode and belief in non-grip mode. However, the modal notions which have been brought to bear are rich, and it is worthwhile considering other applications of them, without trying to achieve a simple definition.

V

The hypothesis now to be considered is that to believe in grip mode is to believe with a high degree of resilience, and to believe in non-grip mode is to believe with a low degree of resilience. Resilience is a modal notion, and reflects the likelihood of the subject revising his belief (or revising downwards his degree of confidence in it) in the light of potential conflicting information. Resilience is not the same as degree of confidence or credence, which is a this-world matter of strength of belief, and could be measured by betting quotients. High confidence might combine with low resilience. Someone in the motor trade assures me that Ford does not make a magenta car. I believe this with complete confidence (I have no independent information and no reason to doubt the word of someone in a position to know). But I might believe it with low resilience, in that I would abandon it with no struggle were I to see a Ford-shaped magenta car with the Ford logo. A more resilient believer would discount the new evidence, supposing that the Ford-shaped magenta car had been resprayed by an enthusiast or was a Ford look-alike produced by another manufacturer.

Examples lend some initial plausibility to the identification of grip with resilience. Consider two responses available to someone who comes to learn that Mary was at the party, given that the person previously believed both that everyone at the party got drunk, and that Mary is a model of sobriety. One option is to abandon the generalization, holding to the belief in Mary's sobriety. Another is to abandon the belief in Mary as a model of sobriety in favour of the generalization about the party. What factors might influence the choice between these revisions?

One possible factor is this: if you supposed it was the kind of party which of its nature ensured that the guests became drunk, you would be more likely to think that Mary's sober disposition had been over stretched, and to retain the opinion that everyone at the party got drunk; whereas if you supposed it just an accident that those who got drunk did so, you are more open to the thought that Mary stayed sober, and that the generalization should be qualified or abandoned. In short, if you take the generalization to reflect something non-accidental, you would be less likely to qualify or abandon it were you to be faced by some potentially defeating evidence than if you take it not to reflect anything non-accidental.

A Humean distinguishes 'laws of nature' from 'accidental generalizations' without supposing that there is a difference of content between them. What we call a statement of a law of nature, like 'All animals are mortal', contains no more 'nomic content' than an accidental generalization like 'Everyone in the room speaks English'. The difference is not a difference in content, an implausible view given the invisibility of a relevant semantic difference, but a widespread difference in mode of belief: laws are believed in grip mode, that is, resilient mode; accidental generalizations in non-grip mode, that is, non-resilient mode. If we learn of something we had supposed to be immortal that it is an animal, we will typically revise our belief in its immortality. If we learn of someone we believe to be in the room that she does not speak English, we will typically revise our belief in the generalization that everyone in the room speaks English.

It has been held that the content of a law-statement somehow entails, sustains or supports a counterfactual, whereas this is not so for the content of an accidental generalization. A difficulty with such views is that it is hard to say in any way at all, let alone a systematic way, what expressions encode the relevant difference of content. A Humean can mark the difference not as a difference in content but as a difference in mode of belief, itself reflected in what

resilience-counterfactuals are true of typical believers. Strictly, there is no such thing as being a law or non-accidental generalization, but we call something a law if it is believed in grip mode by us, or by us and our friends, or by the experts, or by the majority (the relativization is normally not determinate).

There is some initial plausibility in the identification of grip with resilience; but there are also difficulties. Here are some possible counterexamples. Suppose that I believe that every guest at the party got drunk, but I regard this as pure accident: I know of nothing about the party which made it specially conducive to drunkenness. The generalization should be believed in non-grip mode, and let us suppose that it is in fact thus believed. But suppose I came to this accidental belief by having been present (say as a spy rather than a guest, so that I can soberly observe without being a counter instance of the generalization), and having observed every guest get drunk. Were I to learn that Mary, whom I had supposed a model of sobriety, was among the guests, I would revise my belief in Mary's sobriety rather than abandon the generalization. So a belief held in what ought, for a Humean, to be non-grip mode might meet a condition for being resilient, that is, for being held in grip mode. A belief in an accidental generalization may be resilient because of the way in which it is acquired, and may not reflect a view of the generalization as non-accidental.[8] Other examples suggest the converse kind of failure: one may have a non-resilient belief in something one treats as non-accidental; that is, one may regard as non-accidental something one believes in non-resilient mode. For example, a firm has been working on alternative unbreakable drinking vessels for wine. I am convinced of their success, know that they are about to drop their prototype onto concrete, and believe that this occurrence will not be followed by a breaking. I also believe that dropping wine glasses (vessels made of glass) is followed by their breaking. This is a causal generalization, and being, arguably, more than just a summation of singular causal facts it is the kind of generalization that I should, and

we will suppose do, believe in grip mode, reflecting my perception of what I believe as something non-accidental. But were I to come to learn that the dropping of the prototype is in fact the dropping of a wine glass (the innovation was to develop unbreakable glass), I would maintain my confidence in the firm and abandon the belief that all droppings of wine glasses are followed by breakings. On the present account, this lack of resilience means that I am not in the grip of a regularity; but this is a case of non-accidental belief, and so should be a case of grip. Resilience is almost never total, for almost all beliefs can be defeated by conflicting evidence, so we cannot demand that total resilience is what reflects the grip mode of belief.

Confidence can mimic resilience. In the first counterexample, I am totally confident in the accidental generalization that every guest at the party was drunk because I have observed each guest to be drunk, and am well placed to know that none had escaped my observation. The confidence in this case held in place the counterfactual that I would incline to discount, if not the new evidence offered (Mary was at the party), at least its impact on the hypothesis given previous beliefs. The mechanisms are complex and depend upon both kind and weight of evidence. Hume was concerned only with inductive beliefs, ones believed on grounds which do not include observation of every instance, and inserting this as a restriction of the subject matter would help protect his approach. But confidence must still be separated from resilience, as shown by the example of the magenta Ford. In the party case, it makes a difference whether the apparent counter evidence takes the envisaged form, or whether it simultaneously provides an explanation of my mistake. For example, if I were to be told that Mary had been hiding in a cupboard, so I hadn't seen her, it is more likely that I would revise my belief in the generalization.

There are a morass of factors at work in the notions now under discussion (degrees of confidence, evidence and its weight, and the probabilities of revision in the light of new

information) and it would be little help to a Humean to be told that he has a clear thesis only once the issues have been sorted out. I think, however, that one can sketch a promising direction by trying to hold as much constant as possible, and by transforming the essentially quantitative notions into qualitative ones. We need to make explicit the implicit relativity of resilience: just about every belief is resilient to some but not other potential new information. To call one belief more resilient than another is a way of saying that one is resilient to "more" information than another. If the beliefs are different, the comparison will not be straightforward, since new information relevant to one may be irrelevant to the other. Every belief is resilient to irrelevant information, so there is a danger of incorrect measurement of relative degrees of resilience of distinct beliefs. In the promising approach, we consider a single belief, and a single set of background beliefs (or a single distribution of subjective probabilities) and, for a series of pieces of potential new information, we ask two questions: is there a way of holding the belief such that, given this background, it would be resilient to this information? And is there a way of holding the belief such that, given this background, it would not be resilient to this information? Here '(believing that p) is resilient to (q)' means something like: the probability of the believer revising her belief that p on encountering the information that q is low. Those pieces of information for which the answer to both questions is 'Yes' are the 'sensitive' ones: one might or might not revise in the light of them. Given a psychological state of believing, we can rank its resilience in terms of how many of the sensitive pieces of potential information it is resilient to: the number of pieces of sensitive information, those which might cause a revision in belief, which would not do so. Finally, we can identify grip mode for that belief with having high resilience and non-grip mode with having low resilience. (Or we could rework grip as a matter of degree, and so claim yet another illusion in the supposition that there is a difference of kind between the accidental and the non-accidental.)

To show the promise of the approach, I briefly apply it to two of the problematic examples. There are two ways to believe that every guest at the party got drunk: one may think of it as an accident, or one may regard the party as the kind of party which makes drunkenness likely or even inevitable. The first way of believing is less resilient: there is a piece of information, that Mary was at the party, coming to learn which would probably lead one to revise the belief about the party, given one's background confidence in Mary's sobriety. The second way of believing is more resilient, for coming to learn this would probably cause no revision, but would instead cause a revision in the background belief about Mary's sobriety, even though this was held with the same degree of confidence as in the previous case. There are two ways to believe that glasses break when dropped: I may believe that this just happens to be so, or I may believe that it is held in place by causation. In both cases, the belief is not resilient with respect to the potential information that the firm's unbreakable drinking vessel is made of glass. But there is information to which the latter way of believing is resilient and the former not, for example, seeing a magician drop what appears to be a glass and it not break. In one case there is a fair chance that the evidence will be taken at face value and the generalization revised, in the other a fair chance that it will be discredited and the generalization retained (what appears to be glass is not always glass).

As the examples reveal, resilience will always be hard definitively to assess, for it involves consideration of potential responses to every piece of relevant information. But there is nothing circular about an application of induction here: the fact that this belief is believed with resilience to <u>these</u> pieces of sensitive information can give reason to believe that it is also resilient to <u>other</u>, or <u>many more</u>, or even <u>all</u> such pieces of information.

VI

We have come some way from Hume's second definition, and I close by estimating this distance. There were two initial Humean hopes: that causal belief could be distinguished as a mode of belief rather than as belief with a distinctive content, and that this could be achieved by the second definition. The second hope was disappointed, because even belief in an accidental generalization required the relevant kind of causal and counterfactual connection between singular beliefs. Since resilience is defined for beliefs in generalizations, beliefs some of which the second definition was supposed to define, it is not likely that resilience can be used to produce a new definition which would look anything like what Hume offered. With this, many bright consequential hopes have to be abandoned: we can no longer expect the kind of definition of causal belief which will help a functionalist (or not one who takes it that there is a substantial difference between grip and non-grip modes), nor can we expect the kind of definition which will make it straightforward that non-human animals have such beliefs.[9]

Despite this, the first hope remains appealing. Recognizing non-accidental non-causal beliefs and accidental causal ones, we modified it into the hope that we could single out distinct modes of belief in terms of which to mark the distinction between accidental and non-accidental. Thus far, it seems, we remained more or less within a Humean perspective. But the attempt to mark this distinction in terms of the modal notion of resilience may seem both antithetical to Hume's philosophy and hopelessly circular. It may seem antithetical because Hume apparently had no truck with this kind of modality, and it may seem circular because the relevant counterfactual facts themselves involve the concept of the non-accidental.

Hume saw nothing wrong with a kind of necessity and possibility arising from a comparison of ideas, and not considered to hold 'between the objects themselves'; this may be approximately equated with apriori necessity and possibility. On this view, the possible is the

coherently thinkable, and this notion played a pivotal role in his philosophy. The range of worlds involved in the account of resilience does not exceed the coherently thinkable ones, but we must restrict our attention to a proper subset of them, those in which the subject assigns the same subjective probability to all save the target belief and the newly acquired evidence as to her actual beliefs. Restrictions like this do not involve a new notion of possibility: we can stipulate that these are the worlds to be considered, and a 'comparison of ideas' can check whether, for an arbitrary alternative to the actual world, this restriction has been honoured. The claim that the approach is positively antithetical to Hume (as opposed merely to being not something one could attribute to Hume) would need to be based on something more than that it exploits a modal notion.

Perhaps we can find the something more in the explanation of the alleged circularity. The allegation is that the concept of counterfactuality involved in the concept of resilience itself involves the concept of the non-accidental. But how does this enter? Not, if the point of the previous paragraph is accepted, with the notion of a non-actual possibility as such. Perhaps it enters with the assessment of the likelihood of belief revision in the light of the envisaged new information. To assess the likelihood we must bring to bear knowledge of non-accidental matters, laws relating to the subject's dispositions and psychology. Does this not show that the approach is both circular and anti-Humean?

Resilience is just a pairing of worlds with a number, a number which measures the likelihood that the target belief will be abandoned at the paired world. As such, it contains nothing about psychological laws. But even if it did, this would not amount to circularity: what is to be defined is belief in the non-accidental, not the non-accidental as such (though given the character of the former, there is room for a Humean deflation of the latter). The structure here is that of the second definition: causation is used in the definition of causal belief, not in the

definition of causation (though given the character of the former, there is room for a Humean deflation of the element of necessity supposedly contained in the latter).

Perhaps there is circularity in that we need beliefs about psychological laws in order to have evidence for which number measures resilience at a world. This is the right content for an accusation of circularity, but it falls in the wrong place. It would be damaging only if this content, belief in the non-accidental, belonged to the concept of resilience. There is nothing wrong if it belongs merely to what is involved in possessing evidence for its application. Similarly, an analysis of perception which uses causation is not rendered circular by the fact that much, or even all, causal knowledge is based upon perception.

Hume's second definition is extraordinary not only for being a brilliant answer, but also for being an answer to a brilliant question, one which Hume thought, correctly as far as I am aware, had never been asked by any of his predecessors. If the arguments of this paper are accepted, his second definition is not a correct answer, but the underlying idea that we can distinguish a distinctive way of believing a generalization, a way which in some sense puts us in its grip, remains promising, and highly relevant to many currently debated issues in metaphysics and philosophy of mind.

References

David Hume, Enquiry concerning Human Understanding, , ed. L A Selby-Bigge, 3rd ed. revised P H Nidditch (Oxford: Clarendon Press, 1975)

C. Peacocke, 'Causal modalities and realism', in Reference, Truth and Reality, ed. M. Platts (Henley: Routledge and Kegan Paul, 1980), pp. 41–68

F. Ramsey, 'General propositions and causality', in Foundations: Essays in Logic, Mathematics and Economics, ed. D. H. Mellor (London: Routledge and Kegan Paul, 1978), pp. 133–51

---

[1] 'If a cause be defined, that which produces anything; it is easy to observe, that producing is synonymous to causing', in David Hume, Enquiry concerning Human Understanding, , ed. L A Selby-Bigge, revised P H Nidditch (Oxford: Clarendon Press, 1975, 3rd ed.), Section 8, Part 2, fn, p. 96. The context makes it plain that this is regarded as unsatisfactory.

[2] F. Ramsey, 'General propositions and causality', in Foundations: Essays in Logic, Mathematics and Economics, ed. D.H. Mellor (London: Routledge and Kegan Paul, 1978), pp. 133–51, p. 136. Ramsey applies the phrase to 'variable hypotheticals', sentences like 'All men are mortal'. He argues that laws are special cases of such hypotheticals.

[3] Hume, Enquiry, Section 9, p. 105.

[4] 'To believe that all men are mortal—what is it? Partly to say so, partly to believe in regard to any x that turns up that if he is a man he is mortal.' (Ramsey, 'General propositions and causality', p. 136). Ramsey would not approve of the use to which his idea is put here, for he takes 'All men are mortal' to be an example of a non-accidentally true sentence, whereas I use his account in the search of a non-grip mode of belief. Another difference is that Ramsey will conclude that 'variable hypotheticals' are not really propositions at all. This is a good route for an eliminativist: if beliefs apparently in a complex content can be reduced to relations between beliefs in a simpler content, one is free to say that the supposed complex content is unintelligible nonsense, and that the question of something satisfying it does not so much as arise.

[5] The example can be made transtemporal by supposing that I never have the good fortune to encounter a white geranium.

[6] Cf. C. Peacocke, 'Causal modalities and realism', in Reference, Truth and Reality, ed. M. Platts (Henley: Routledge and Kegan Paul, 1980), pp. 41–68.

[7] Peacocke, 'Causal modalities and realism, p. 55; cf. Ramsey, 'General propositions and causality', p. 136: for 'any x that turns up' (cited above in fn 4).

[8] Hume would not regard generalizations all of whose instances had been observed as even candidates for being believed with grip. For him, the relevant cases involve inductive projection, which is what grip sustains.

[9] Though it is not nothing that a development of the second definition will ground an ascription of general beliefs to creatures lacking the concept of universal quantification.