

Perceiving causation and causal singularism

Victor Gijbbers

Abstract

Elizabeth Anscombe's classic paper *Causality and Determination* claims that causation can be perceived. It also defends causal singularism, the idea that the causal relation is fundamentally between the particular cause and effect and does not depend on regularities holding elsewhere in the universe. But does the former furnish an argument for the latter? The present paper analyses a special type of causal experience involving emotional reactions to present stimuli; for instance, being frightened by a spider. It argues that such experiences are *strongly local*, in the sense that they justify belief in a local causal relation independently of our knowledge about events elsewhere in the universe. If this analysis is correct, Humean regularity theories of causation are false; and all other non-singularist theories of causation face a difficult explanatory challenge. This means that the case for causal singularism comes out considerably strengthened.

1. Introduction

Elizabeth Anscombe's seminal paper *Causality and Determination* (1971) is famous for claiming that (a) universality is not part of the concept of causation, and (b) causation can be observed in the individual case. While these claims appear to be logically independent, Anscombe indicates that a tradition stemming from David Hume connected their denials in the following way:

not (b) \rightarrow not (a).

The idea would be that if causation can *not* be observed in the single case, our knowledge of causal relations can only come from the observation of regularities, and so universality must be an element of the concept of causation. The weaknesses of this argument are palpable and need not detain us. But one may wonder whether there is not a more promising argument from the perceptibility of causation to causal singularism:

(b) \rightarrow (a).

For if causation *can* be observed in the individual case, does this not prove that causation must in some sense be present *here*, in this single observed cause-effect pair? And does that not in turn show that the truth of the individual causal statement must be independent of the obtaining of any universal regularities? Anscombe does not explicitly commit herself to this argument but has been read as implicitly defending it, for instance by Ernest Sosa and Michael Tooley:

One reason for thinking [that Anscombe wishes to defend causal singularism] is that there would seem to be a close relation between a singularist conception of causation and the thesis that causation can be directly perceived – a thesis which she certainly is defending. For if two events could be causally related only if that relation were an instance of some law, to observe that two events were causally related would be to observe that there was some relevant law, and it is not easy to see how a single observation could serve to establish such a conclusion. So it is hard to see how causation could be directly observable if a singularist conception of causation were not true. (Sosa & Tooley 1993, p. 13.)

Whether or not Anscombe wanted to embrace it, the availability of such an argument for singularism is well worth investigating. And indeed it has been investigated, perhaps in most detail by Helen Beebe in two

excellent papers (2003, 2009).¹ Now Beebee does *not* accept singularism, and argues that while Anscombe is correct in claiming (b), the singularist claim (a) does not follow from it. In other words, Beebee maintains that a non-singularist like the Humean regularity theorist is able to account for what she calls ‘thick’ causal experiences, perceptual episodes that present causal content.

In the present article, I will argue that there is a type of thick causal experience that the Humean regularity theorist cannot account for. Given the popularity of the regularity theory, this is already a significant result, although it falls short of establishing causal singularism. However, this type of thick causal experience also creates an explanatory burden for any theory of causation, a burden that the causal singularist can easily discharge but that non-regularity non-singularist theories may have more trouble with. The causal singularist position is therefore strengthened in two ways: by defeating one of its main rivals, and by putting an extra explanatory burden on its remaining rivals.

Clearly any argument starting from a particular type of experience of causation will have to spell out in detail where this type of experience is to be found. This will be the burden of sections 3, 4 and 5, which also serve as a (constructive) criticism of the existing literature. But before one starts a search for something, one needs a good idea of what one is searching for – one must at the very least be able to recognise it when found. We need to have a good idea of *how* causal experience might spell trouble for a regularity theory in order to know *what kind of* experience the singularist is looking for. For this reason we will first develop the logic of our argument before analysing any particular experiences of causation.

2. The argument

Before we talk about experiences of causation, we should first clarify an ambiguity in the term ‘experience’. Suppose that I walk into a room and I seem to see a lion there; in fact I unhesitatingly² come to believe that there is a lion there. However, it soon turns out to be a friend of mine prancing about in a rather realistic lion costume. There are several ways one could describe this situation. One could say that I experienced a lion, but that the experience was not veridical. One could also say that I merely thought I experienced a lion, but in fact I experienced a guy in a lion suit. Finally, one could say that what I experienced was neither a lion nor a guy in a lion-suit, but a lion-like visual impression. And of course one’s choice will be bound up with other philosophical commitments – e.g., the third option fits well with a traditional empiricist foundationalism that needs experience to be immediately accessible and not in need of justification.

For our purposes, however, we do not need a fully worked out theory of experience as long as we are clear about how we use the term in this paper. As a matter of terminology, then, let us stipulate that the phrases ‘an experience of causation’ and (equivalently) ‘a thick causal experience’ indicate an experiential episode of a subject that *prima facie* justifies belief in a statement of the form ‘*c* caused *e*’ without relying on belief in a causal regularity linking events of types *C* and *E* (where *c* is an event of type *C* and *e* an event of type *E*).

Having clarified what we mean when we talk about an experience of causation, we go on to ask: why might an experience of causation be troublesome for non-singularists? The crux must be that such an experience is a *local affair*, and nevertheless an experience of *causation*. As Sosa and Tooley say in the quote given in section 1, if causation is a matter of law, then an experience of causation would have to be an experience that justifies belief in a law. And that may seem to be a bit much for any single experience to do. After all, we take experiences to be determined by what is going on here and now; that is, we take them to be independent of what happens in distant places or at distant times. So when we say that an experience of causation is a

1 Beebee tells us that while she hasn’t found the singularist argument in print, she heard it “in conversation many times”; and she points at Armstrong 1997 (213-4) and Cartwright 2000 (47) as coming close to defending it. One can perhaps add Ducasse (1965) as such an ‘almost’ defender.

2 It is not uncommon to say that in such an experience one ‘non-inferentially’ comes to believe that there is a lion there. I don’t think it is possible to make a distinction between beliefs one arrives at inferentially and beliefs one arrives at non-inferentially. But that is not the topic of this paper, and so I will simply avoid the phrase.

local affair, we are at the very least claiming that the subject could have had the same experience even if other parts of the universe had been wildly different from how they actually are.

But if this is a merely phenomenological claim, it cannot spell trouble for the non-singularist. For we can hold that distant events make no difference to the phenomenology of our experience; and hold, at the same time, that those distant events *do* make a difference when it comes to determining whether the experience we have reveals a *real* relation of causation or instead offers no more than the *illusion* of causation. Suppose we have had a 'thick causal experience'. And now we learn that there is no global regularity covering the local cause-effect relation that we believe we have experienced. What happens then? Does the experience still justify belief in the local cause-effect relation? Or will its justificatory role be negated by our knowledge that there is no global regularity? The singularist will no doubt claim the former and the regularity theorist will claim the latter, but the experience itself does not seem to favour either explanation.

So in order to create trouble for the regularity theorist, a thick causal experience must be local in a stronger sense than the merely phenomenological one; it must be local in an epistemic sense. The experience must be such that knowledge about distant events cannot negatively affect its ability to justify the local causal claim. Let us therefore define a *strongly local causal experience* as a causal experience that can perform its justificatory role no matter what we learn about distant events.³ If *that* type of experience exists, then the regularity theorist might be in trouble.

Let us consider in more detail how strongly local causal experience creates trouble for the Humean regularity theorist. As a regularity theorist, she holds that

(A) Singular causal statements imply global regularities.

As a Humean, she holds that the world is a Humean mosaic. This involves the following reductionist claim:

(B) The truth of global regularities is grounded in the global pattern of local matters of fact.

where grounding is a relation of ontological dependence. Together (A) and (B) entail that

(*) The truth of singular causal statements is at least partly grounded in the global pattern of local matters of fact.

Belief in the Humean mosaic also involves the following independence claim:

(C) Local matters of fact place no constraints on events elsewhere.

The existence of strongly local causal experiences is incompatible with the conjunction of (*) and (C). If the truth of singular causal statements is at least partly grounded in the global pattern, then knowledge about that global pattern must be relevant to our singular causal beliefs. So whether or not we should believe in the existence of some local causal relation must be sensitive to what we know about distant events. Now a causal experience is strongly local just in case perform its justificatory role no matter what we learn about distant events. This is logically compatible with (*) only if the local matters of fact revealed by such an experience constrain what the distant events can be, and constrain them in such a way that those distant events must exhibit a global regularity. But according to (C) there are no such constraints. Hence: *if there are strongly local causal experiences, then the Humean regularity theory is false.*

Assuming that strongly local causal experiences exist, this conclusion would have a big impact on current debates in the philosophy of causation, since Humean regularity theories are still quite dominant. But it falls short of proving causal singularism. Given the above argument, one can avoid contradiction by denying (A), which is to embrace causal singularism. But one can also deny (B) or (C) or both. To deny (B) is to claim that global regularities are not grounded at all or are grounded in something that is not the global pattern of local matters of fact; e.g., in relations between universals or in dispositional essences. To deny (C) is to claim

³ This does not make a strongly local causal experience incorrigible; the possibility is left open that it would lose its justificatory role when we learn more about *local* events. We will briefly return to this at the end of section 5.

that given certain local matters of fact, not all patterns of events elsewhere remain possible. Usually, (C) would be denied on exactly the same grounds that (B) would be denied – because local matters of fact imply truths about universals, perhaps, or because dispositional essences must always act in the same way, or for some such metaphysical reason.

Although strongly local causal experiences thus do not force one to embrace singularism, the singularist can nevertheless issue a challenge to the non-Humean non-singularist. The singularist has an obvious story about how strongly local experiences are possible: causal relations simply *are* purely local. But the non-singularist will have to explain how it is not only possible for humans to experience something that has non-local causal implications, but how it is also possible that the justificatory role of such experiences cannot be undermined by learning more about non-local facts.

Take the universals approach as an example. If there are strongly local causal experiences, then the defender of the universals approach would not only have to construe those causal experiences as clues to relations between universals, but also as clues that *could not be undercut by experiences of distant events*. It is far from obvious how to do this, since it would seem that if a local experience here-and-now can be positive evidence for a claim about universals, then a local experience somewhere else could also be negative evidence for that same claim. This doesn't depend on details about the universals approach but is a general problem for any non-singularist theory of causality. So although strongly local causal experiences cannot prove the truth of singularism by themselves, they do raise the burden for non-singularist approaches. (Whether this burden can be discharged cannot be investigated within the confines of this paper.)

Of course, all of this becomes relevant only if strongly local causal experiences exist! Whether they do is the issue that will concern us in the bulk of the paper. First, in section 3, we will discuss the ideas of Albert Michotte and Susanna Siegel, who are interested in showing that we truly *perceive* causation. Beebe argues that these types of perception are compatible with a regularity theory; and indeed, we will see these experiences are not necessarily *local*.

This motivates section 4, in which we look at the way in which C. J. Ducasse, Mary Shepherd and Nancy Cartwright bring in the notion of the 'single experiment' when discussing the possibility of perceiving causation. Their discussions show that local observations can give us causal knowledge. However, I will argue that the local causal conclusions justified by such observations are still liable to revision in the light of evidence about distant events, and so are not *strongly* local.

In section 5 I develop an idea of the Gestalt psychologist Wolfgang Köhler, according to whom the relation of causation becomes especially clear in emotional reactions such as being frightened. When an evil clown suddenly jumps at you from behind a bush, what you experience is *being frightened by the clown*; the connection between the cause and the effect seems particularly intimate. I analyse such experiences and argue that they are indeed *strongly local*: their justificatory role can be undermined by new information about the local situation, but not by new information about non-local events. Thus, the conclusions drawn above follow – the Humean regularity theory is false, and other non-singularist theories of causation have an added explanatory burden.

3. Perceiving causation

The psychologist Albert Michotte initiated a famous research programme investigating causal perception, culminating in his 1946 book *La perception de la causalité*. A typical experiment done by Michotte might involve the test subject watching a screen on which the following happens: a little circle moves rightwards towards a little square; the circle stops when it touches the square; and at the same moment, the square starts moving to the right. The test subject will then be asked to describe their experience, and might typically say something like: "I saw the circle bump into the square and propel it forward." This description is in causal terms. Michotte found that whether test subjects describe the experience causally is highly sensitive to the

exact parameters of the experience; e.g., the time delay between the circle touching the square and the square starting to move.

A curious feature of Michotte's experiments is that while the test subjects in some sense may be said to 'perceive causation', their experiences are in that respect not veridical: the circle, being only an imagine on a screen, does not *really* cause the square to move. It merely *seems* to cause the movement. But of course there are analogous cases, such as that in which a ball bumps into a milk carton and pushes it forward, where we *do* have veridical experiences. And this, Michotte might say, means that we can also perceive causation in a veridical sense.

Are Michotte's experiments of the right kind to spell trouble for the regularity theorist? Are they strongly local? Beebe (2009) helpfully points out that there is certainly something special about the experiences Michotte investigated: they are not sensitive to background information. Suppose that *I see a witch*, in the sense that I see an old woman with a weird hat and a large wart on her nose, and this visual impression is accompanied by the judgement: 'this is a witch'. Clearly, this judgement is sensitive to new information. When I learn that the woman before me is in fact a saintly abbess dressed up for a Halloween party, I revise my judgement, and I will no longer be tempted to claim that *I see a witch*.

But this is not the case in Michotte's experiments. For suppose we find out that that there is no causation going on on the screen; then we would *still* be tempted to describe our experience in causal terms, and we would *still* experience a phenomenological difference between the 'causal' (e.g., short time delay) and 'non-causal' (e.g., long time delay) visual sequences. (If one doubts this, simply recall that Michotte's own test subjects must have *known* that there was no actual causation going on.) Like well-known visual illusions, the experience of causation persists in the face of information that shows it to be non-veridical. This would also mean that they will not be influenced by information about distant events.

Susanna Siegel (2009) attempts to establish the point that we experience causation in a somewhat different way. She suggests that we can experience the same sequence of events in two ways that are identical *except for* a causal component. I throw a ball; it falls into a potted plant; and at the same moments the lights in the room go out. According to Siegel, one can experience this in a phenomenologically unified or a non-unified manner. She then argues that the best explanation of this phenomenological difference is not some difference in our background knowledge – in part because our having the unified experience is possible even when we do not *believe* the ball to have caused the darkness – but through the fact that in the one case the situation is represented as causal in our experience, and in the other it is not. Again the experience would be insensitive to information about distant events.

Beebe (2009) argues that while Michotte and Siegel show that causal experience may be independent of *some* background knowledge – knowledge that there is in fact no causation in this particular case – they have not even attempted to show that causal experience can be had independent of *all* background knowledge. (Nor was this their purpose.) And it is hard to see how anyone *could* show that, since *all* of us have a large amount of background knowledge. Again and again have we seen the sequence of one thing bumping into something else and the other thing starting to move. Again and again have we seen some obvious change in the environment (often the flicking of a light switch, but also sometimes just the waving of an arm) being followed by a light turning on. And so we may without contradiction hypothesise that causal perception as described by Michotte and Siegel is possible because of a prior acquaintance with regularities; thus, that perception of causation in the individual case is parasitic on a prior grasp of causal regularities. And if this is a coherent hypothesis, then this type of experience cannot create trouble for the Humean.

More fundamentally, though, even *if* the causal experiences described by Michotte and Siegel could be had independently of all background knowledge, they still would not be strongly local. For it would only be the *phenomenology* of the experiences that is independent of what we know about distant events; but this in no way shows that these experiences can play their justificatory role independently of such knowledge. We need

to move beyond causal phenomenology to causal justification. With that in mind, let us investigate three authors who have written about gaining knowledge of causation from a single experiment.

4. Causal experiments

Anscombe's examples of perceiving causation are very different from those of Michotte and Siegel. She mentions the causal verbs "*scrape, push, wet, carry, eat, burn, knock over, keep off, squash, make* (e.g. noises, paper boats), *hurt*" as being crucial parts of our observation language; and goes on to suggest that even substantives like 'wind' and 'fire' are causal terms (Anscombe 1971; p. 93 as reprinted in Sosa & Tooley 1993). In this she is followed, or rather anticipated, by C. J. Ducasse, whose 1965 (in part based on his 1926) tells us:

The plain fact to which I refer is that every person has *perceived* – and I say *perceived*, not *inferred* – that, for example, a particular tree branch was *being caused to bend* by a particular bird's alighting on it; that a particular bottle was *being caused to break* by the fall on it of a particular rock; that a particular billiard ball was *being caused to move* by a particular other billiard ball's rolling against it; that a particular match was *being caused to ignite* by friction of it on a particular rough surface; that a particular footprint was *being created* – that is, being caused to exist – by a particular man's foot walking on a beach, etc. (Ducasse 1965, p. 177.)

As indicated earlier, Anscombe doesn't say much more about the nature of this causal experience, or how exactly the causation shows up in the experience; but Ducasse does go on to develop an argument for his claim that causation is perceived rather than inferred. This argument is based on his somewhat idiosyncratic definition of causation. For Ducasse, the crucial notion here is that of an *experiment*, which is a local state of affairs (not necessarily created by human beings) in which only two changes occur, one just before the other. For example: the tree is motionless. Only one change occurs: a bird alights on the branch. A second change begins to occur very soon after the first change has begun to occur: the branch bends. This is an experiment in Ducasse's sense.

When we have an experiment, we can say that the first change caused the second change. So in our example we can say that the alighting of the bird causes the branch to bend. In fact, Ducasse claims that this is what causation *is*. The relation between the events and surrounding circumstances in an experiment "is *not a sign* that causation, in some mysterious sense, is occurring, but is *causation itself*, and is perceived by the performer or observer of a well-conducted experiment." (1965, p. 178.) Since we can perceive – even if not infallibly, as Ducasse stresses in his 1967 – changes or absences of changes, we can perceive causation. And it also follows that "*one single experiment* is sufficient in principle to make evident that the particular change C in S did cause the particular change E in S. For repetition of the experiment is called for, if at all, only to make sure that the experiment was conducted strictly, and its outcome observed accurately." (1965, p. 178.)

We can compare Ducasse's reasoning here to that of a much earlier critic of Hume: Mary Shepherd in her brilliant 1824 book *An Essay upon the Relation of Cause and Effect* (partially reprinted in Shepherd 2018). Shepherd also points to the possibility of establishing causation in a single experiment: "One trial is enough, in such circumstances, as will bring the mind to the following reasoning. Here is a new quality, which appears to my senses: But it could not arise of itself; nor could any surrounding objects, but one (or more) affect it; therefore that one, (or more) have occasioned it, for there is nothing else to make a difference". (Shepherd 2018, p. 37.) In a footnote on the same page she also echoes Ducasse's point about repetition of experiments: "[w]hen more trials are needed than ONE, it is order to *detect* the circumstances, not to lay a *foundation for the general principle*, that a LIKE Cause repeated, a LIKE effect will take place."

There is, however, a crucial difference between Shepherd's approach and Ducasse's, for Shepherd does not *define* causation in terms of experiments. Nor does she need to, since her argument targets Hume's claim that only custom – and thus repetition – can convince us of a causal claim. Ducasse, on the other hand, wants to

prove singularism. And it is his definition of causation in terms of experiments that allows him to do so. For it is clear that *given* that definition, ‘being an instance of causation’ is a purely local property of a purely local situation; and so Ducasse can write that an “occurrence therefore could be unique in the history of the universe, and yet be, and be known to be, a case of causation.” (p. 178.)

But this should give us pause, for observation has dropped out of Ducasse’s argument here. Whether or not causation can be *observed* in the single instance is irrelevant; it is his *definition* of causation that establishes it as a local relation. But this means that despite initial appearances, the observability of causation does not strengthen Ducasse’s argument one whit. All the weight is on his definition of cause. And it is clear that this definition simply begs the question against the non-singularist.

Still – there remains the interesting fact that we often *do* reach causal conclusions on the basis of a single trial or observation, as Shepherd and Ducasse point out. Does this mean that such experimental observations are strongly local causal experiences? It is helpful to look at Nancy Cartwright’s (2000) discussion of how causal regularities and singular causal claims interact in our scientific reasoning. Any experimental test of a causal claim, she suggests, will presuppose the truth of some singular causal claims (e.g., “setting this button to 5 causes the magnet to produce a field of 5 Tesla”) as well as of some causal regularities (“magnetic fields accelerate electrically charged particles”). Once such presuppositions are in place, one-shot experiments *can* decide singular causal claims. But without the presuppositions, no experiment is possible.

Clearly Cartwright’s one-shot experiments are *not* examples of strongly local causal experiences. For they explicitly rely on non-local knowledge; and finding out that we were wrong about the presupposed causal regularities would undermine our confidence in the local causal claim. If it turns out that magnetic fields do not accelerate electrically charged particles, then the causal conclusions drawn from our experiment will lose their justification. So what we learn from Ducasse, Shepherd and Cartwright is that, yes, we can draw causal conclusion from one-shot experiments in science; but, no, this ability does not involve strongly local causal experiences. The search must continue.

5. Strongly local causal experience

On the first nice day of spring, I walk to the back of my garden, open the shed door, and immediately scream and jump back, having been scared by a big hairy spider. Clearly I am in no doubt that it was the spider that frightened me. Doubt about the origin of the fear seems impossible; as if I might have stood there panting in the doorway of the garden shed, wondering *which* of the objects in the visual field caused my fright; and by taking a good look first at the rake, then at the bicycle, and then at the spider, I slowly formed a hypothesis about the cause of my fear. Of course not. The fear I experienced is *fear of the spider*; not even fear of spiders in general, but a fear caused by that *particular* spider over there, *that* one, the big black hairy beast from which I instinctively flee. And indeed it is not just the fear, but also the fleeing that I experience *as caused by the spider*. This example of experiencing causation is significantly different from those we discussed in sections 3 and 4, and, I will suggest, much more fruitful for the singularist.

We find this type of example in the writings of the psychologist Wolfgang Köhler, who discusses a whole array of them in the final chapter (‘Insight’) of his 1947 book *Gestalt Psychology*. To give a few: I admire the beautiful singing voice of the opera diva; I am charmed by the first smile of my child; I enjoy the coolness of the beer I’m drinking; I am annoyed by the mess someone made in my lab. These are all cases where we experience some external state of affairs *as causing* an internal reaction; and where this reaction is *directed at* that external state of affairs. The voice is *experienced as* causing my admiration, and my admiration is *directed at* (is admiration of) the voice.

With these examples, Köhler’s is explicitly taking aim at Humeans. But his arguments against their theory of causation are not very convincing. For example:

Suppose that [a Humean] is riding in a bus in which he has to stand, because it is crowded. At a certain moment a heavy man chooses the feet of our Humean as a good place to stand on. As a matter of principle, the Humean would have no way of deciding whether the [...] face of a girl in the bus or this man's conduct is the cause of the anger which he, the Humean, now feels. He might find out by experimentation and induction, or he may have learned to connect such things correctly in the past. But [...] he could not possibly tell without such indirect evidence. (p. 343)

We have learned from Beebe that the Humean can just accept that, yes, a lot of indirect evidence may have been incorporated into how we experience the world. So what? If the singularist is to make any headway, we need experiences that are strongly local.

I want to argue that the spider experience *is* a strongly local causal experience, as are the experiences in Köhler's own examples. I was frightened by the spider. I was experientially aware of the spider as that which frightened me. And my certainty that it was the spider that frightened me cannot be undercut by information about non-local events. I may learn that other people have not been frightened by spiders. I may find out, by experiment, that I myself am not frightened by the spider when I return to the shed. I may even find that the next spring, having forgotten all about the incident, and being just as surprised by the spider as I was the first time, it nevertheless does not frighten me. None of that seems to cast the slightest doubt on the veracity of my original causal experience.

Let us put some more pressure on this claim.⁴ Suppose that I try very hard to find some regularity under which my experience can be classified. I systematically bring myself into contact with the spider, even going as far as touching it. No fear. I systematically expose myself to other – even bigger and hairier – spiders. Again, I never feel the slightest tinge of panic. I ask my parents whether I had any traumatic experiences with spiders when I was young. Nope. I have brain scans made while I'm watching spiders. No fear-related signals are seen, not even weak ones. Would this not undermine the justificatory role of my experience? Would it not make me wonder whether perhaps something else – something glimpsed only from the corner of my eye perhaps, or some random event in my brain – caused the original fear?

If we were totally committed to the truth of the regularity theory, then we might indeed start to have doubts about the experience. As Quine (1951) suggested, we *can* hold on to any belief if we are willing to make adjustments elsewhere. If I have decided to hold on to the regularity theory no matter what, then beliefs about distant events will *always* be relevant to the justifications I have for local causal claims, because I will always give up local causal claims if they fail to instantiate regularities; and so there would, for me, *be* no strongly local causal experiences. But for the current discussion we should presuppose a dialectical situation in which we have not yet settled on the true theory of causation. It is important to see that everyone, no matter which of the current theories of causation they accept, will agree that a Michotte-style sequence on a screen is not really an example of causation; and that similarly everyone will agree with Cartwright that experimental conclusions about causal relations can be undercut by new information about other events. But this is not true in the spider case. If the *only* reason to reject the causal justification offered by the spider experience is that it doesn't fit one's favourite theory of causation, one is walking a far more dubious path.

It might be suggested that we often reject experiences that do not fit the regularities of nature. If one (seems to) see an elephant flying outside one's window, or if one hears a report about a wooden statue weeping blood, one might be justifiably sceptical. But such events – purported miracles – fail to fit the regularities of nature in the *strong* sense that they are inconsistent with what we believe about those regularities. We have good reasons to believe that elephants cannot fly and that wood cannot produce blood. The spider case fails to fit the regularities of nature only in the *weak* sense that it does not fall under any known regularity. But in this it is no different from many of the most mundane events. As Cartwright (1994) points out, we cannot apply the laws of mechanics to predict where a dollar bill swept away by the wind will land. Nor can we

4 I am indebted to an anonymous referee for some of the following examples.

explain why the glass shattered in exactly eighty-three pieces when it hit the floor; or why, of all the people we met, it was this one person we fell in love with. So there is in fact no salient analogy with miracles.

The case for the strong locality of the spider experience can also be strengthened in a more positive way, namely by analysing what makes it and experiences like it special, and how these special features enable them to play their justificatory role independently of information about events elsewhere. I want to follow Köhler himself by drawing attention, first, to the *directedness* of the reaction to the stimulus. The fear is fear *of* the spider; the admiration is admiration *of* the voice. This is a feature that examples from mechanics clearly lack. When one of Hume's billiard balls hits another and sets it into motion, the motion of the latter object is just *motion*; it is no way 'directed at' the motion of the former object. We can think away the former object, and the motion remains the same; whereas we cannot think away the voice and still have a subject admiring that voice. And this inseparability of cause and effect has not been achieved by means of a question-begging verbal definition, in the way that an object can only become 'hammer-propelled' by being caused to move by a hammer. The inseparability of the admiration and the voice is right there in the experience.

Again following Köhler, we may notice that we experience our reaction as *apt*, that is, we experience it as the natural or appropriate reaction to the stimulus. The voice is experienced as *admirable*: it is part of our experience of the voice that admiration is the appropriate reaction to hearing it. In the same way, when we are frightened by the spider, we experience our fright as the natural and appropriate reaction to the eight-legged monstrosity. To be sure, we may *believe* or even *know* spiders are not really fearsome. Afterwards we may laugh about our reaction and we will perhaps call it foolish. But in the moment, we experienced fright as the appropriate reaction to the spider. And just as with directedness, mechanical examples do not exhibit this feature of aptness. It makes no sense to ask whether it is appropriate for a mechanical sequence to take place. It takes place or it does not take place, but it is not a subject for normative judgement. Köhler draws a strong conclusion from this difference, arguing that in his examples we have direct experience of the 'functional relation' between cause and effect, which no inductive method can ever give us:

I did directly experience, first, that my admiration was related to the singing rather than to any other thing, and secondly, that admiration was the natural answer to such a way of singing. As a consequence, I needed no indirect criteria, no scientific investigations, no coefficients of correlation, in order to know about the connection between the singing and the admiration. As a matter of fact, my experience told me more than any scientific induction could. For induction is silent as to the nature of the functional relation which it predicates, while in the present example a particular fact of psychological causation was directly experienced as an understandable relationship. (p. 324.)

Can this be made more precise? Here I turn to Zwart 1967, the only philosophical discussion of Köhler's theory of causal experience that I am aware of. Zwart considers the Humean idea that we may have simply learned through long experience to associate, say, feeling anger with someone hitting us. We were hit, we felt anger. We were hit again, we felt anger again. After a while, we came to associate the anger with the hitting, and *now* we experience the anger as being directed at the hitting. Does this make sense? According to Zwart it does not:

[N]ow we must assume that when it comes to the young child, being hit (A) can give rise to anger (B) without the child at first being aware of any relation between these two events, and that only after some time does the child learn to see that B is an effect of A. But this is [...] absurd, for a child too can only feel anger after being hit, when it has first learned to understand being hit as a reason for being angry. (pp. 25-26; my translation from the original Dutch)

I read Zwart as arguing that we do not just experience our anger as apt when we are hit, but that the nature of the case is such that we could not be angry after being hit *unless* we grasped this aptness. Anger is of such a nature that it requires not just directedness at a stimulus, but it requires us to understand the stimulus as being the kind of thing that ought to induce anger – and as in fact inducing *this* anger, that we are feeling right now.

Thus, Zwart's argument would be that the type of experiences that Köhler is talking about can only appear to us whole and complete, *with* the intrinsic grasp of causality, because this particular causal relation cannot exist without being grasped as at the same time a rational relation. I will call this feature of the experience its *rational transparency*.

These three properties of directedness, aptness and rational transparency form an intertwined trio. It is only because the reaction is directed towards the stimulus that the question of its aptness comes up. Rational transparency then is the property that the aptness of the reaction – and hence also the directedness of the experience – must be grasped for the experience to be possible.

Now it is clear, and has been emphasised above, that these properties are not to be found in mechanical sequences, or indeed in any sequences among external objects. But more importantly, these properties perfectly explain why intuitions gained from Humean example of billiard balls will not carry over to psychological reactions. For the crucial feature for Hume is that it is always conceivable that the purported cause and effect could exist independently; that in experience we never find any event that is not 'loose and separate'. I'm not taking a stance here on whether that is really true about external events. But it is obviously not true for experiences that essentially involve directedness, aptness and rational transparency. It is here that we find *in the single experience itself* the kind of 'necessary connexion' that Hume vainly sought among external experiences. And that is of course precisely what is needed for the experience to be strongly local, that is, for it to be able to play its justificatory role independently of what we learn about global patterns.⁵

In the end, one will have to insist on the phenomenology – on the fact, if it is a fact, that we do indeed have experiences characterised by directedness, aptness and rational transparency. But such experiences seem to exist. And this means that strongly local causal experiences exist.

6. Conclusion

Where does that leave us? Beebe is right that many examples of experiencing causation do not spell trouble for the regularity theorist, since they leave open the possibility that either the causal experience itself or its veracity depends on other events happening at other places and times. But I have argued that the type of emotional reaction described by Köhler is different. Here our grasp of the causal relation between the stimulus and the reaction is so tight that it cannot be epistemically undermined by information about what happened in other situations. This type of experience is strongly local.

In section 2, we set out the consequences of finding strongly local causal experiences. First, it follows that the Humean regularity theory is false. Second, it follows that any other non-singularist theory of causation has an added burden: for it must show how its rejection of singularism is compatible with the existence of experiences like the spider example, experiences which essentially involve awareness of their own causal structure, and the causal implications of which cannot be undermined by learning more about non-local facts. This *may* be possible, but it is one dialectical challenge that singularism itself simply avoids. The question how much trouble this challenge makes for various forms of non-singularism must be left for future research.

One final note. The present paper was focused on scenarios of emotional reaction to stimuli. But equally fascinating questions arise about the experience of agency (where perhaps we grasp, in a strongly local way, a causal relation between our own decision and our physical movements) and even perception (where one might suggest that experiencing one's perceptions *as such* – as being *of* an external world – involves

⁵ It might be tempting to say that the justificatory role of such experiences cannot be undermined at all. But I do not want to go that far. If there is one thing that analytic philosophers are good at, it is coming up with strange sceptical scenarios. Perhaps it is possible that there is a neuroscientist who manages to give me the exact experience of being frightened by a spider even though no spider, and no spider-like thing, is there at all. I don't want to take a stance on the possibility of this scenario. But note that it undermines the causal experience in a purely local manner: in this scenario the local events themselves are different from what I supposed them to be: they involve evil neuro-machines rather than spiders. So this in no way undermines the strong locality of the experience.

precisely the triad of directedness, aptness and rational transparency). But these too are questions we cannot address here.

References

- Anscombe, Elizabeth. (1971). *Causality and Determination: An Inaugural Lecture*. Cambridge University Press.
- Armstrong, David. (1997). *A world of states of affairs*. Cambridge University Press.
- Beebe, Helen. (2003). Seeing Causing. *Proceedings of the Aristotelian Society* 103: 257-280.
- Beebe, Helen. (2009). Causation and Observation. In Helen Beebe, Christopher Hitchcock & Peter Menzies (eds.), *The Oxford Handbook of Causation* (pp. 471-497). Oxford University Press.
- Cartwright, Nancy. (1994). Fundamentalism vs. the Patchwork of Laws. *Proceedings of the Aristotelian Society* 94: 279–292
- Cartwright, Nancy. (2000). An empiricist defence of singular causes. *Royal Institute of Philosophy Supplement* 46: 47-58.
- Ducasse, C. J. (1926). On the Nature and the Observability of the Causal Relation. *The Journal of Philosophy* 23: 57-68.
- Ducasse, C. J. (1965). Causation: Perceivable? Or Only Inferred? *Philosophy and Phenomenological Research* 26: 173-179.
- Ducasse, C. J. (1967). How Literally Causation is Perceivable. *Philosophy and Phenomenological Research* 28: 271-273.
- Köhler, Wolfgang. (1947). *Gestalt Psychology: An Introduction to the New Concepts in Modern Psychology*. Liveright Publishing Corporation.
- Michotte, Albert. (1946). *La perception de la causalité*. Editions de l'Institut Supérieur de Philosophie.
- Ott, Walter. (2009). *Causation & Laws of Nature in Early Modern Philosophy*. Oxford University Press.
- Quine, Willard V. O. (1951). Two Dogmas of Empiricism. *Philosophical Review* 60: 20–43.
- Shepherd, Mary. (2018). *Selected Writings*. Edited by Deborah Boyle. Imprint Academic.
- Siegel, Susanna. (2009). The Visual Experience of Causation. *The Philosophical Quarterly* 59: 519-540.
- Sosa, Ernest & Michael Tooley (eds.). (1993). *Causation*. Oxford University Press.
- Zwart, P. J. (1967). *Causaliteit*. Van Gorcum.