

THE VISUAL EXPERIENCE OF CAUSATION

BY SUSANNA SIEGEL

Abstract
Abstract
Abstract
Abstract
Abstract
Abstract

How is causation represented in the mind? We often *believe* that one event has caused another. But can we *visually experience* two things as causally related? If so, then experiences represent causation. A different question in the vicinity is whether we can ever see that something is causing (or has just caused) something else to happen. In the relevant sense of ‘seeing’ here, seeing is factive: you can see that p only if p . By contrast, experiential representation of properties or relations is not factive, so you can represent that p even if p is not true.

These questions are distinct. There are several ways in which visual experiences might represent causation, even if causation cannot be perceived. Most simply, there might be no such thing as causation – it might be a fabrication of the mind. In that case visual experiences could in principle still represent that causal relations obtain, though no such experiences would be veridical. (Some theories of intentionality for experience rule out this possibility, saying that an experience represents a property F only if experiences of that type are normally caused by instances of F .) Alternatively, experience could represent (even correctly represent) that causal relations obtain, but these experiences might fail to count as perception, if their relation to what they represent is never sufficiently direct.

In this paper I shall argue that causation is represented in visual experience. The first part of the paper makes a positive case for this claim. In §I, I introduce and clarify the central thesis. In §II, I discuss the bearing on this claim of the psychologist Albert Michotte’s data, which are often said to

show that adults ‘perceive causality’. In §III, I introduce a strategy for defending the view that some experiences represent causal relations. In §IV, I present the central consideration which favours this view. The defence continues in §§V–VI. The last section of the paper is devoted to defusing arguments that we cannot see or experientially represent causation. So while most of the paper addresses the question about visual experience, the discussion at the end bears on the question about perception as well.

I. THE CAUSAL THESIS

The central thesis defended in this paper is this:

Causal thesis. Some visual experiences represent causal relations.

What is a causal relation? The relations of pushing, pulling, lifting, stopping, moving, supporting, hanging from, and preventing something from happening might naturally be considered modes of causation: they are specific ways of causing something else to happen. There also seems to be a more general relation which these relations exemplify – causation itself. There may be relations which are less specific than the modes I have listed, but more specific than causation itself, such as the varieties of ‘launching’ discussed by Michotte.¹ In plain launching, an object *a* approaches a stationary object *b*, which upon contact starts to move in the same direction as *a* was moving in. In launching-by-entraining, *a* pushes *b* along. In launching-by-expulsion, *b* moves as if thrown by *a*. The causal thesis should be interpreted as saying that visual experiences represent either the more general relation, or one of its specific modes.

What it is for a visual experience to represent a causal relation? The notion of representing a property in experience derives from the notion of representing that such and such is the case in experience. If experience represents that such and such is the case, then if such and such did not obtain, the experience would be inaccurate. For instance, if your visual experience represents that the square before you is red, then it will be inaccurate if there is a square before you that is not red. None of this is a definition of what it is for experience to represent that such and such is the case, but it is a start. To finish it, what is needed is an account specifying which accuracy-conditions are the ones at issue. Here the main constraint is that the accuracy-conditions are supposed to reflect the phenomenal character of the experience, so that a statement of the accuracy-conditions of the experience would characterize the subject *S*’s experience, from *S*’s point of view.

¹ A. Michotte, *The Perception of Causality* (London: Methuen, 1963).

One reason to think that visual experiences are assessable for accuracy is that we regularly classify some of them as ‘veridical’ or ‘illusory’. A straightforward account of this is that they are *correct* when they are veridical, and *incorrect* when they are illusory. Some philosophers think this account is not correct, and they deny that such experiences are assessable for accuracy.² Here I am just going to assume, contrary to these claims, that visual experiences are assessable for accuracy.³

Once it is granted that visual experiences have accuracy-conditions, it is a short step to defining what it is for experiences to represent a property. In the visual experiences of the sort at issue, the properties that are experientially represented are the properties that the objects *S* sees appear to *S* to have. For instance, ordinary objects appear to have colour and shape properties, and the causal thesis says that they can appear to stand in causal relations as well. If the causal thesis is true, causal relations are represented in experience in the same way as colour and shape properties are.

To maximize the interest of the causal thesis, it will be useful to make two stipulations about the kind of mental state that can count as an experiential representation of a property. First, since no one doubts that we form beliefs about the causal relations between things, beliefs will not count as experiential representations (if they did, then the causal thesis would be true just in virtue of the fact that people believe that one thing caused something else to happen). Secondly, some philosophers think of experiences as a conjunction of judgements and sensations. In the case of the experiences to be discussed here, this view categorizes them as conjunctions of judgements that one thing has caused another thing to happen, where the judgement itself is not systematically or intrinsically related to any phenomenal character; and a sensation which is intrinsically phenomenal, but is not systematically or intrinsically related to any representation of causation, judgemental or otherwise. On such a view, experiential representations of causation are structured much as Thomas Reid supposed perceptions to be: the sensational and the representational aspects can vary independently. I shall interpret the notion of an experiential representation in such a way that the causal thesis would not be made true by such a conjunction of states.

² Further discussion of this point can be found in E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/spr2005/entries/perception-contents>, §2.1.

³ Even if experiences are not assessable for accuracy, something very much like the question of whether causation is represented in experience would still arise. For instance, if some experiences consist in a subject’s perceiving an object as *F*, as some naïve realists about perception hold, then one can ask whether the values of *F* can include causation. Likewise, if experiences consist in adverbial modifications of the subject whereby one is ‘appeared-to *F*-ly’, one can ask in which ways one can be appeared to, and specifically whether one may be appeared-to by two (or more) relata in a causal manner.

Whenever causal relations are represented in an experience, the experience has accuracy-conditions, and so represents that such and such is the case, where a causal relation R is part of such and such. This raises the question of what the relata of the causal relation are. The causal thesis leaves this open. The contents of an experience illustrating the causal thesis could be naturally (albeit approximately) expressed by a sentence of the form x causes y to happen, in which case y would be an event. Or the contents could be naturally expressed by a sentence of the form $x R y$ which posits a relation between two things – sometimes objects ('Bob lifted Jack', 'The laundry hangs on the line'), sometimes events ('The scream caused the shudder'). Still another option would be contents roughly expressed by a sentence of the form x causes y to ϕ ('Bob caused Jack to fall', 'The scream caused me to shudder'), where y and ϕ are pairs of objects and properties that bear some close relation to events. If you applied a systematic semantic theory to each of these sentences, it would posit different truth-conditions for each of them, involving different ontological commitments.

The causal thesis is neutral on which of these truth-conditions (if any) experiences have when they represent causal relations. Experiences may not always be sensitive to these differences in truth-conditions, so that there is sometimes some indeterminateness in how exactly the world must be, if the experience is accurate. In other cases experiences may have more determinate truth-conditions. All that matters for the causal thesis, however, is that the truth-conditions include a causal relation.

II. MICHOTTE'S RESULTS

In the 1960s Michotte published the results of about 150 experiments which showed that adults are inclined to describe scenes of launching and entraining in causal terms. In launching, object a moves toward a stationary object b , makes contact with it, and b then without delay begins to move in the same direction as a while a stops. Entraining is the same except that a moves along with b . It will be useful to keep these vividly in mind for the subsequent discussion.⁴

Michotte's experiments were attempts to isolate the exact parameters of motion that elicited such descriptions from adults. In some of the experiments the 'objects' were lights or shadows projected on a screen, in others they were ordinary hefty objects such as wooden balls, and some experiments combined both. Subjects knew that no actual causal relations were

⁴ Readers are encouraged to view these at <http://research.yale.edu/perception/causality/capture-launchAlone.mov> and <http://research.yale.edu/perception/causality/entraining.mov>.

operative in many of the situations they saw (such as where a shadow 'launches' a ball or *vice versa*), but that did not influence their descriptions. This led Michotte to posit a representation of causality which was something other than a belief about what was happening in the environment. He called it an 'impression of causation'.⁵ Michotte's results establish that adults regularly describe launching and entraining in causal terms, saying things like 'The red square moved the blue one along', or 'The ball pushed the shadow'.

What, if anything, do Michotte's results establish about the contents of the visual experiences their subjects report? In general, how people report what they see is a poor guide to the contents of their visual experience. People may use different expressions to describe a scene that looks the same – for instance, two people might try to guess how large or how far away something is, and differ in how good they are at making such estimations. Taking their reports as a guide would force us to conclude that tables (say) look as if they are of different sizes and at different distances from each subject, when in fact there might be no such differences at all. So if Michotte's results tell us anything about the contents of his subjects' visual experiences, this will not be because we can simply read the contents off their reports. In general, one cannot do that.

Michotte's results do not by themselves entail the causal thesis. At best, they entail it when combined with further theoretical assumptions about the relation between experience, beliefs and reports. Suppose in a case of launching the 'launching' object (sometimes called the 'motor object') is a ball, and the 'launched' object is a patch of light. Adults know that hefty objects cannot cause patches of light to move. Suppose also that we take the following substantive assumption as a rough guide to the contents of experience: these are the contents of what one would believe if one relied on what one knows about the relata. Then we should probably say that the contents of experience are roughly that there is a certain pattern of continuous motion between the ball and the 'moving' patch of light. If we changed the substantive assumption so that the guide to the contents of experience is instead what one would believe if one bracketed what one knows about the nature of the relata, then we might get a different result – that the ball causes the patch of light to move. Either way, we get a verdict that bears on the causal thesis only when Michotte's results are combined with substantive theoretical assumptions linking reports to background beliefs and contents of experience.

⁵ Psychologists today use 'perceptual causation' to refer to the kind of mental representation that Michotte posited to explain this description, though sometimes this term is used to label the kinds of scenes that adults are prone to describe as causal – launching, entraining or expulsion.

Furthermore, Michotte's results are compatible with the claim that the visual experiences reported in his experiments (and others like them) merely represent the input-conditions that elicit the causal reports, as opposed to representing causation (or a kinetic mode of causation) itself. The input-conditions are the specific spatiotemporal pattern of continuous motion that elicits the reports. These conditions are features of the world that obtain (we may assume) independently of the subjects' having any experiences. When we ask what contents experiences have, we are asking how experiences present things to the subject. We are asking, roughly, how a specific pattern of motion looks. Michotte's results are compatible with the claim that the subjects' experiences have contents which exactly match this pattern of motion. It is this pattern that elicits the causal descriptions, even when people know that there is no actual causation. Moreover, there are also cases of actual causation, such as heat causing water to boil, or pressure on a button making a light go on, that do not elicit the causal descriptions, as well as cases without actual causation that do elicit it. So an opponent of the causal thesis might say that the natural conclusion to draw is that Michotte's subjects represent the spatiotemporal pattern of continuous motion, rather than causation.

The central question that Michotte's results pose for the causal thesis, then, is whether there is a defensible principle relating report-contents and experience-contents which can be added to Michotte's results to support the causal thesis. This would be one way to try to defend the causal thesis. But instead of trying to develop such a principle here, I am going to pursue a different strategy.

III. HOW TO DEFEND THE CAUSAL THESIS

If the causal thesis is true, then some visual experiences represent causal relations. One candidate for illustrating the causal thesis is experiences of certain static scenes. Suppose there is a cat sitting on a loose mesh hammock, forming a sitting-cat-shaped dip in the hammock's surface. The causal thesis allows that your experience of seeing the cat represents that the cat is pulling the hammock downwards. If the causal thesis is false, then your experience can represent that the part of the net directly under the cat is closer to the ground than the rest of it, but not that the cat is pressing on it. The opponent of the causal thesis holds that experience simply remains neutral on whether any force is being continuously exerted in such a case.

The descriptions of these static experiences championed by an opponent of the causal thesis do not strike me as phenomenologically apt. But the

phenomenological aptitude of a proposed content for visual experience can be hard to judge when we are given only a single case. Given only a single case, it can be difficult to discern what the phenomenal character of the visual experience is, hence hard to discern which contents most adequately reflect it. How, then, might we proceed if we want to know whether the causal thesis is true?

One way to proceed is to focus on a pair of cases that fairly obviously differ phenomenologically. At this stage, what is important is that people who may ultimately disagree about what, if anything, is represented in experience would agree that these are cases that differ phenomenally. All that introspection is relied upon to do is to recognize a phenomenal contrast, without taking any stand on whether it is also a difference in what is represented, let alone on what the representational difference, if there is one, may be. Once we have such a phenomenal contrast in hand, the following method is available for reaching a reasoned verdict about whether the contrasting experiences represent causal relations. The method is to reason about the best explanation of the phenomenal difference, ruling out some hypotheses and making the case for others. This method is general: it can in principle be used with respect to any property or relation that is putatively represented in experience. How successful it is may vary from case to case, depending on the properties or relations in question, and on the sort of considerations that are available for assessing alternatives.⁶

So far I have spoken of representational differences that 'accompany' the phenomenal ones. One might wonder what 'accompaniment' is, and more specifically whether representational features are explanatorily more fundamental than phenomenal features, or the converse, or neither. All these options can be respected by the causal thesis. If the fact that visual experience has certain contents – say, causal ones – explains why it has the phenomenal character it does, then it will be natural to say that the fact that experienced causal contents in one of the cases but not the other account for the phenomenal contrast. On this way of thinking about things, the opposing sides in a debate about whether experience has causal contents may agree about what phenomenal character the experience has, but disagree about its contents. In contrast, if the fact that experience has a certain phenomenal character explains why it has causal contents, then it will be more natural to say that the view that the experience has causal contents offers a further characterization of the phenomenal character. On this way of thinking about things, what the opposing sides disagree about is in part what the phenomenal character of the experiences in question is in the first

⁶ I discuss and defend this method in 'How Can We Discover the Contents of Experience?', *Southern Journal of Philosophy*, 45 (2007), pp. 127–42.

place. I shall talk as if the causal thesis is one way to explain the phenomenal difference, but this is just a manner of speaking. My defence of the causal thesis is indifferent to which of these is the way in which it is developed.

IV. CAUSAL UNITY IN EXPERIENCE

In defending the causal thesis, I shall discuss pairs of experiences which differ phenomenally. These pairs will bring into focus a specific and familiar kind of unity in experience. Suppose you are playing catch indoors. A throw falls short and the ball lands in a potted plant, with its momentum absorbed all at once by the soil. You see it land, and just after that, the lights go out. The ball's landing in the plant does not cause the lights to go out, and I shall suppose that you do not believe that it does. Nevertheless it may seem to you that the ball's landing somehow caused the lights to go out. This is the first case. In the second case, you likewise see the ball land and the lights go out. But this case is unlike the first: you do not have any feeling that the ball's landing *caused* the lights to go out. Your visual experience represents the ball's trajectory and landing, and it also represents the lights going out, but so far as your visual experience is concerned, these events merely occur in quick succession.

It seems plain that there can be a phenomenal difference between two such experiences. This provides a starting-point for the method described above. But in this case, a bit more can be said at the outset about the phenomenal contrast to be explained. This contrast seems to have something to do with the connection between events: the difference stems from how the lights' going off seems to you to be related to the landing of the ball.⁷ The successive events seem to be unified in experience in a way that is not merely temporal.

There are also cases of a similar sort of experienced unity involving simultaneous events. Suppose you draw open the curtain from a window and let in some light. Here there is one event, the uncovering of the window, that occurs simultaneously with another event, the increased illumination of the room. These events occur simultaneously. They can also, let us assume, be experienced as occurring simultaneously. Contrast this with a case in which you draw open a curtain that does not block out any light in the first

⁷ I have focused on how the observed events seem to be related: the light's going off seems related to the trajectory of the ball. Arguably, the associated experiences (or parts of experience) also seem to be related: the experience of seeing the lights go off seems related to the experience of seeing the ball land in the plant. I am setting aside the question of how these two kinds of unity may be related.

place (perhaps because it is translucent). Just as you uncover the window, the sun comes out from behind a dark cloud, causing the room to lighten progressively as the curtain is drawn open. Here we have a very similar event of uncovering the window, coupled with an event of the room gradually lightening, and (we can suppose) these events have and are experienced as having the same duration. But it seems that there could none the less be a phenomenal difference between the two experiences, with respect to how the window's uncovering and the room's illumination are experienced as related to one another. I am taking this to be an intuition. In this application of the method for investigating whether a certain content is a content of visual experience, I am taking the *explanandum* to be not merely a phenomenal contrast, but more specifically a contrast with respect to whether the events in question seem to be unified – leaving open, at the outset, just what kind of unity those events seem to have.

In what way might the events seem to be unified? A natural suggestion is that in the one case there is an experiential representation of causation, whereas in the other case there is not. If so, then in the ball case one of the experiences represents the ball-landing and the lights' going off as causally unified, and in the curtain case one of the experiences represents that the uncovering of the window lets in the light. The main merit of the causal thesis is that it provides a plausible account of the phenomenal difference between each pair.

To defend the causal thesis properly, what is needed is reason to think that this is the best account of the phenomenal contrast. There seem to be two main alternative types of hypotheses. Either the difference is due to a cognitive element, such as a disposition to form a causal belief, or else it is due to a difference in exclusively non-causal contents.⁸ In the next section, I discuss the latter alternative – agreeing with the causal thesis that experiences have contents, but holding that in the pairs of cases in question these contents are exclusively non-causal.

V. NON-CAUSAL CONTENTS

One might think that the phenomenal difference between the two cases in each pair (the ball and the curtain) is a difference exclusively in non-causal contents. The central question about this proposal is whether the contents of the experiences in question could be exhausted by non-causal contents.

⁸ The further alternative that the experiences in question are simply 'raw feels', with no content at all, is at odds with the starting assumption that the visual experiences in question have contents. Some responses to the two-component view in §VI address this view as well.

I shall focus for a moment on motion, illumination and temporal contiguity in the curtain case. An opponent of the causal thesis might say that if the two events – the uncovering and the lightening – unfold in such a way that the increases in light are correlated with the movement of the curtain in just the right way, then there will be an impression of unity between these events. But this need not be an impression of causal unity, the opponent might say; it is merely an impression of a special correlation between the rates at which two things happen.

Could it be the rate at which the window is uncovered and at which light streams in that makes it seem as if uncovering the window has let in the light? If so, then while these rates could be the same in each of the curtain cases, they could not be visually represented as the same – since if they were, then (on this proposal) there would be an impression of causation in both cases, whereas by hypothesis there is such an impression only in one.

This opponent of the causal thesis thus has to hold that if there were a pair of curtain experiences with the phenomenal difference described, one of the experiences would have to be falsidical. *Mutatis mutandis* for the ball experiences.

This prediction seems implausible. Why must either experience involve any error in how fast the uncovering or the lightening is experienced as occurring? Taken separately, each case in the pair can perfectly well be imagined to be a completely veridical experience. So this alternative requires illusion where intuitively there does not seem to have to be any. If so, then both the actual rates of the events and the rates at which they are experienced as occurring can be the same, on both occasions. It seems that opponents of the causal thesis will have to look elsewhere to account for the phenomenal difference.

Where else might they look? The curtain cases were described in such a way that the colour and thickness of the curtain differs in the scenes, and there were corresponding differences in the contents of experience. The difference in perceived colour, however, seems irrelevant to experienced unity between the uncovering of the window and the lightening of the room. Changing the perceived colour of the curtain does not seem to affect whether or not the visual experience represents the light as being let in as a result of the curtain's moving, so long as the curtain is thick. (Moreover, if one could have such an experience at all, it seems that one could have it without even registering the colour of the curtain.) What about the curtain? Suppose, for the sake of argument, that the curtain's luminous properties are unlike its colour in that at least one of the experiences has to represent that the weave of the curtain is either opaque or transparent. (This leaves open whether or not the other experience is neutral on the curtain's

luminous properties, either by failing to represent anything about them at all, or by representing a more general luminous property of which opacity and translucence are determinates.) One might claim that the phenomenal difference between the two curtain-openings is due to a difference in representation of luminous properties. In effect, this proposes to explain an impression of unity between the window-uncovering and the room's getting lighter by focusing exclusively on the illumination properties of the curtain.

If the impression of unity has any content at all, however, then it would concern the relation between the lightening and the uncovering of the window, since these are the events that seem to be unified. This proposal looks for the key representational element in a different relational property – the relation between the curtain and light in general. So it seems to be looking in the wrong place.

Another option is that visual experience represents the events in question as unified, without representing them as causally unified. What kind of unity might this be? Perhaps there is a kind of unity analogous to the kind discussed by Husserl in connection with the experience of the passage of time. Suppose you hear a series of five sounds: the clink of a cup against a saucer, the groan of an accelerating bus, a creak from a chair, a snippet of a loud voice, and the honk of a car's horn. Compare this auditory experience with hearing five notes of a melody. We experience the notes of the melody as unified in a way in which we need not experience the five sounds as unified, even if at each moment we remember the sounds from the previous moments. Husserl observed that such remembering does not suffice to make us experience any remembered series of motley sounds as integrated, as we hear the notes of a melody to be integrated. Rather, when we hear the sounds as a melody, the previous ones remain present to us in a distinctive way even after they have ceased. Husserl had a special term, 'retention', for this way of remaining present to us. In the case of the melody, the previous notes are 'retained', whereas in the motley series they are not, and that is what makes it the case that we hear the notes as a melody, but not the motley sounds.⁹

There seems to be a phenomenal difference between these two experiences, one which stems not just from the different qualities of the individual sounds, but from the way in which the sounds seem to be related to one another in each experience. One might think that analogously, the ball's landing in the plant seems in the first case to be unified with the lights' going

⁹ Husserl discusses the notion of retention, in connection with hearing melodies, in lectures given between 1893 and 1917, in *On the Phenomenology of the Consciousness of Internal Time*, tr. J.B. Brough (Dordrecht: Kluwer, 1980). I have discussed the contrast between melody and motley sounds in Siegel, 'The Phenomenology of Efficacy', *Philosophical Topics*, 33 (2005), pp. 265–84.

out, while it does not seem like that in the second case. The general idea seems to be that the events somehow belong together as a unit – perhaps in something like the way in which the *Gestalt* psychologists thought that the geese in a flock appear to be a unit.

As stated, this proposal is not a very specific one. Moreover, it does not make clear what accuracy-conditions it is positing. Likewise, in the case of retention, it is difficult to say how the world has to be in order for experiences of retention to be accurate. It is clearly not enough that the notes of the melody occur in succession. The opponent of the causal thesis I have been considering draws an analogy between retention, on the one hand, and a kind of unity that can be visually represented in the case of the ball or the curtain, on the other, where the unity experiences have exclusively non-causal contents. This raises the question of which contents these are. A proposal which we can assess needs to say, at a minimum, how things in the world have to be in order for the experience to be accurate.

Returning to the case of the ball landing in the plant and appearing to turn off the lights, I shall suppose that the ball's landing does not really turn off the light. So the proponent of the causal thesis (if in agreement that this experience illustrates the thesis) predicts that the experience is falsidical, because it represents two events as causally related when they are not. What does the opponent predict about the experience? Is it veridical because the events really are unified in the way in which experience presents them as being (and there is no other illusion), or is it falsidical because the events only appear to be 'unified'? Similar questions apply to the experience of seeing the light coming on at 6 a.m. and then hearing the neighbour's whirring coffee grinder. There does not seem to be any obvious relation that events in the world can have when such experiences are veridical but lack when they are falsidical. This suggests to me that the analogy with Husserlian retention should not be taken to suggest that there is non-causal unity content in the curtain and ball cases. The lesson of the analogy should just be that there is a special kind of temporal relation that one can experience events as having, and that sometimes such events also seem to be causally related, but at other times they do not.

How else might one try to account for the feeling of unity in the ball and curtain cases (and others like them), without accepting the causal thesis? A more specific proposal is that in the ball case, experience represents that if the ball had not landed in the plant, then the lights would not have gone out. According to this proposal, what is represented is a relation of counterfactual dependence: specifically, the lights' going out counterfactually depends on the ball's landing in the plant.

Some philosophers think that causation is a kind of counterfactual dependence, and they may hear this proposal as a proposal in favour of the causal thesis, rather than a proposed alternative to it. But even those who favour counterfactual analyses of causation do not identify this relation with causation itself. For example, according to Lewis, if c and e are two events such that e would not have occurred without c , then c is a cause of e – so counterfactual dependence is sufficient for causation. But he denies that it is necessary. Others who have developed counterfactual theories of causation take counterfactual dependence as a starting-point rather than a stopping-point of analysis. Because it has struck many philosophers as such a reasonable starting-point for analyses of causation, one might think that it is the perfect candidate for being a non-causal relation which is represented in experience: it is (relatively) simple, but close enough (by their lights) to causation itself to account for unity phenomenology in the cases at issue.

More generally, the proposal says that for two events x, y that occur, if an experience represents that y would not have occurred had x not occurred, then it will seem to the subject as if x and y are unified in the way in which the ball's landing and the lights' going off seem to be unified. This proposal predicts that if visual experience represents that an event y counterfactually depends on an event x then those events will be experienced as unified in the way exemplified by the ball and the curtain cases.

Some philosophers think that counterfactuals cannot be represented in visual experience. Colin McGinn seems to voice this claim:

You do not see what *would* obtain in certain counterfactual situations; you see only what actually obtains. When you see something as red you do not see the counterfactual possibilities that constitute its having a disposition to appear red. Your eyes do not respond to *woulds* and *might have beens*.¹⁰

There are several different things McGinn might have in mind. One is that we do not see events other than the actual ones that occur at the time of the experience. This is true. Another claim is that if you see that an object has a property F , then F has to be instantiated by the object you see. This claim is also true. But it does not follow from either of these claims that it is impossible to see that an object has a dispositional property. If a thing has a disposition, then it is part of 'what actually obtains' that the thing has the disposition. If a counterfactual is true of an object you see, then the counterfactual is part of what actually obtains too.

A different thing McGinn might have in mind is simply that counterfactuals cannot be represented in visual experience – not because they are not part of what actually obtains, but because their antecedents and

¹⁰ C. McGinn, 'Another Look at Color', *Journal of Philosophy*, 93 (1996), pp. 537–53, at p. 540.

consequents describe occurrences that are not occurring in the very situation being seen (or if it is a case of hallucination, in the situation that is experientially represented). If this claim is true, then it is incorrect to propose that instead of representing causal relations between the ball's landing and the lights' going off, the subject's experience represents that the lights' going off counterfactually depends on the ball's landing.

I think this proposal is false, but not because counterfactuals in general cannot be represented in visual experience. The Gibsons suggested that we sometimes perceive 'affordances' of things, where these are possibilities for interaction with them – such as the rollability of a ball, or that a flat solid surface would support us, or that we shall not find support beyond a 'cliff' where the floor suddenly drops off.¹¹ If we can perceive affordances, then there will be cases where we can perceive possible – and sometimes merely possible – continuations of what we actually see. Perhaps one of the simplest cases of such an affordance would be a ball's continuing along the path on which it is thrown. If someone catches it, the catcher interrupts its path, and arguably there are cases (curve balls aside) in which we can perceive and experientially represent the direction in which it would have continued had it not been caught. More exactly, we can experientially represent that the ball would have continued in the direction it was moving, had its path not been interrupted. Some other examples: you see Joe trying to grab hold of a cup sitting on a high shelf. He stands on his toes, stretching one arm up high, straining to reach a bit higher. But the cup is just out of reach. In the counterfactual *If he reached just a bit farther, he could grab the cup* the antecedent describes a continuation of a motion that one can see – the motion of reaching for the cup. Similarly, suppose you are watching a powerful wind-storm, with wind blowing so hard that a cloth stuck to a branch of the tree is waving furiously, and remains attached to the branch only by a small corner. The antecedent of the counterfactual *If the wind blew harder, the cloth would fly away* describes a strengthening of the wind. One can imagine a rock balancing precariously on the pointed tip of another rock. Could your visual experience represent that the rock would tip over if pushed, Joe could grab the cup if he reached farther, or that the cloth would fly away if the wind blew harder? As proposals about the contents of the visual experience, these are not obviously incorrect. If these counterfactuals can be represented in visual experience, then McGinn goes too far if he suggests that no counterfactuals can be so represented.

None the less there is something correct in McGinn's general doubt. He seems correct about cases in which the antecedent of the counterfactual does

¹¹ E.J. Gibson and R.D. Walk, 'The "Visual Cliff"', *Scientific American*, 202 (1960), pp. 67–71.

not specify any natural continuation of an event that you see. In the ball case, it is the event of the ball's landing and the event of the lights' going off that seem to be unified. So if the proposal challenging the causal thesis by appealing to counterfactual dependence were correct, then we should expect that the visual experience in this case represents that if the ball had not landed in the plant, then the lights would not have gone off. In the curtain case, it is the event of moving the curtain and the lightening of the room that seem to be unified. So if the proposal were correct, we should expect that the subject's visual experience represents that if the curtain had not moved, the room would not have lightened.

In both these counterfactuals, the antecedents do not describe natural continuations or movements of scenes that the subject sees. Instead, the antecedents merely negate a description of something that happens, without any indication of what would have happened instead. The reason to think that this could not be represented in visual experience derives from the intuition McGinn voices when he says that our eyes do not respond to woulds and might have beens. The antecedent of the counterfactuals in the ball and curtain cases are not closely enough connected to what the subject actually sees to be represented in visual experience. The possibilities they describe are visually unconnected with the actual events we see. The hypothesis that we represent counterfactual dependence in the ball case (and others like them) demands much more powerful insight into modal space than is demanded by the proposal that we represent counterfactuals about the direction in which the ball will continue once it is thrown. Perception is not that powerful.

VI. TWO-COMPONENT VIEWS

So far, I have criticized one alternative to the account of the phenomenal differences provided by the causal thesis, the alternative that the phenomenal differences discussed can be accounted for by a difference between the non-causal contents of the experiences in each pair. I now turn to another alternative to the causal thesis, which chalks up the unity phenomenology to a cognitive element, such a disposition to form a causal belief.

According to the simplest version of this view, experience proper consists of a state without causal content – perhaps a state with exclusively non-causal contents, or perhaps a raw feel with no contents at all. But, this position says, when two events seem to be unified in a visual experience, in addition *S* is disposed to form a causal belief (as it might be, a belief that uncovering the window let in the light). According to a different version of

the view, experiences are identified with a sensory–cognitive complex, rather than with the sensory component alone. Either way, there are two components in the vicinity of the experiences in the ball and the curtain cases: the sensory component, and a disposition to form a causal belief.

What, according to these views, explains how it is that one has a disposition to form a causal belief? I shall first consider the version that identifies experiences with the sensory–cognitive complex. It is not open to proponents of this view to say that *S* has the disposition because *S* has an experience, since by their lights the disposition is part of the experience. The reasonable alternative is that it is the sensory component alone that gives the subject this disposition. The simple version of the view, which identifies experience with this sensory component alone, can agree that the sensory component makes *S* disposed to form this belief – and proponents of this view can thereby simultaneously embrace the natural idea that the experience is responsible for *S*'s being so disposed.

Whichever state or pair of states is identified with experiences, the crucial question is what exactly makes *S* disposed to form a causal belief. Here is a crucial difference between the Michotte cases and the case of the ball and the light. In Michotte's cases, an impression of causation is inevitable. It is easy to imagine a version of the ball case, however, where one did not feel that the ball landing in the plant and light going off were unified. One might just experience this as a succession of events that are not unified as they are in the initial example. Indeed, the two-component view is committed to allowing such cases, since each component could in principle be had without the other. But if the sensory component can be had without *S*'s being disposed to form a causal belief, then it is doubtful that the sensory component is what gives *S* that disposition. The two-component view, then, seems to lack a good account of why in the ball case *S* is disposed to form a causal belief.

Another challenge to two-component views concerns the relation between the putatively separate components of experience. Since it is possible to have the sensory component without the cognitive component, the relation must fall short of lawlike regularity. One option is that the sensory component causes the causal representation, without there being any lawlike connection between the property of being that kind of sensory state and the property of being a belief or disposition to believe with the specific causal content posited. (This is compatible with there being some other sort of nomic connection between the two events, connecting some other properties they each have.) But such a causal link does not seem to be able to account for how closely the cognitive component would be *felt* to be associated with the sensory component. A pang of sadness could cause you to remember

that you need to buy light bulbs, but this would not be enough to make it seem to you that the sad feeling and the memory content have anything to do with each other. In contrast, even by the lights of the two-component views, the putative sensory component and causal contents seem to *S* to be associated, since the sensory component presents the relata that seem to be causally related. So the relation between these components would have to be something other (or more) than mere causation.

The proponent of the two-component view might posit a higher-order state which represents that the sensory component and the component representing causation are appropriately connected. This would be one way to ensure that the two components seem to *S* to go together. But it is pretty clearly not necessary to have such a higher-order state in order to have the kind of experience of unity of the sort illustrated by the ball and curtain cases. Such higher-order states about relations between other mental states are comparatively rare, since they require elaborate self-reflection.

One might wonder where in the scheme to fit Hume's claim that the idea of causation arises when the mind makes a transition between ideas. There are two ways in which this might be done. First, the transition between ideas could be part of the content of the experience, in which case the result would be a higher-order content of the sort just discussed. Secondly, the claim that there is such a transition could just be a description of some psychological structure underlying experience. If so, it is neutral on the causal thesis.¹²

A different two-component view, perhaps loosely related to Hume's claim, is that one component is non-causal contents (with whatever phenomenal character goes along with having those), while the second component is a sensation, or raw feel, which is associated with causation. According to this view, when these two components are put together, we get the distinctive feeling of causal unity.

In order for these components really to be separate, it must be coherent to suppose that the alleged raw feel of causation could occur without the non-causal contents – since if these could not occur separately, then there would not be two components after all. To focus on the supposed raw feel, the proponent of this view is committed to the possibility that someone could have an experience with this very raw feel, but without the overall feeling of causation. This possibility is not obviously coherent. Our initial grip on what phenomenal character this is comes from the sorts of examples I have discussed, in all of which there is uncontroversially some feeling of

¹² I am focusing just on Hume's claim about a transition between mental items, ignoring other commitments about the nature of ideas that Hume may have had that might rule out the causal thesis.

causation (just *which* kind of feeling is what is in question). Do we really have any grip on what the phenomenal character is, apart from those experiences of causation? Proponents of this two-component view are committed to saying that we have, since for them the relation between phenomenal character and the feeling of causation is contingent. They are committed to there being something recognizable as the same feeling as is had in my examples, but without being a feeling of causation.

I have argued that the causal thesis is a good explanation of the unity experienced in the ball and the curtain cases. There are of course alternative explanations, and I have not addressed every possibility. One alternative I have set aside is that there might be no contents involved in the experiences in question at all. But if there are, then the options seem to fall exhaustively into two broad categories: there are non-causal contents, and there are combinations of cognitive causal contents with sensory component lacking causal contents. I have tried to address the main considerations surrounding each of these alternatives.

In the rest of the paper, I consider an argument against the causal thesis which also opposes the thesis that causation can be perceived. My attempts to defuse it thus bear on both of the questions with which the paper began – the question about perception, and the question about visual experience.

VII. ARE ARCAINE FEATURES OF CAUSAL RELATIONS REPRESENTED IN EXPERIENCE?

Someone might doubt the causal thesis on the ground that causal relations are the wrong *kind* of thing to be represented in visual experience, given the nature of our perceptual apparatus. Causation, the thought goes, is a kind of counterfactual dependence, or a kind of energy transfer, or a kind of lawful necessitation; none of these things can be represented in visual experience. I shall criticize two specific instances of this doubt, and then respond to the doubt in its general form.

This doubt about the causal thesis has a two-part structure. The first part identifies a factor **X** such that to represent experientially that *a* is causing *b* would be to represent factor **X**. Factor **X** could be a relation which (some true theory in metaphysics tells us) is the causal relation, or it could be a related property such that any instance of causation is also an instance of **X**. The second part provides what is supposed to be a principled reason to deny that factor **X** can be so represented.

One instance of the doubt focuses on counterfactual dependence. According to this version of the doubt, to represent experientially that *a* and

b are causally related would be to represent that certain counterfactuals hold. For instance, if your visual experience represents that the knife is cutting the bread, then it would have to represent that if the knife stopped moving and no other force were applied to the bread, then the bread would not continue to be sliced; if your visual experience represented that uncovering the window lets in the light, then it would have to represent that if the window were covered again, then the light would be blocked, and so on.

The second part of this instance of the general doubt is the claim that counterfactuals are not represented in visual experience. This claim contradicts the central idea in Gibson's work on perception, that we perceive 'affordances' of things, such as the rollability of a ball. To oppose the claim, however, one need not think that we perceive nothing but affordances. It is enough to refute McGinn (on this interpretation of the remark I have quoted) if there are some cases in which we see that *p*, where *p* would be naturally expressed by using a counterfactual. I have examined earlier some cases in which it is not obviously incorrect to say that experiences represent counterfactuals.

Another instance of the doubt focuses on what is useful to call *nomic generalism* about causation, where this is the view that two (or more) relations are causally related only if the relation they stand in instantiates an (exceptionless) law. In contrast, *singularism* about causation holds that causal relations need not instantiate such laws.¹³ It has been suggested that nomic generalism about causation is incompatible with the thesis that causation can be perceived. I shall call this 'the incompatibility claim'. In a summary of Anscombe's paper 'Causality and Determination', Sosa and Tooley endorse this incompatibility claim, and suggest it as an interpretation of Anscombe's main point.

... 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 [Anscombe] 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 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.¹⁴

Suppose nomic generalism about causation is true. Taking '*S* observes that *p*' to be non-factive, and substituting '*S*'s visual experience represents' for '*S* sees that' (plus minor adjustments), yields the following argument:

¹³ Sometimes 'singular causation' is used differently, to pick out causal relations holding between particulars as opposed to properties.

¹⁴ E. Sosa and M. Tooley (eds), *Causation* (Oxford UP, 1993), introduction, p. 13.

1. If S 's visual experience represents that a causes b , then that visual experience represents that (a, b) instantiates a law
2. For any events (a, b) , one's visual experience cannot represent that (a, b) instantiates a law
3. Therefore one cannot represent in visual experience that a caused b .¹⁵

There are many ways to elaborate what it would be for two causal relata to instantiate a law, and how this is elaborated will determine exactly what (1) and (2) come to. For now, I set this complication aside. On virtually any account of what nomic subsumption is, (1) seems open to question. One may grant that if a causes b , then (a, b) instantiates a law. It does not follow from this that if one's experience represents that p , and p entails q , then one's experience represents that q . That there is an ice cube on the table entails that there is a chemical compound on the table. But one can see that there is an ice cube on the table without seeing that there is a chemical compound on the table. An analogous point holds for representations in visual experience. So the argument seems to start off badly.

In its general form, the doubt considers a feature X closely tied to causation itself – either by being a feature had by any relata of causal relations, or by being the *analysans* of an analysis of causation. I shall call X a 'distinctive feature of causation'. The doubt then claims that the distinctive feature X cannot be represented in visual experience, and concludes that causal relations cannot be either. So far, I have considered two possible values for X – relations of counterfactual dependence, and lawfulness. But something more general can be said. From the fact that X is a feature distinctive of causal relations, it does not follow that an experience representing a causal relation would have to represent X . An experience could correctly represent that a succession (or a pair of simultaneous relata) is a case of causation, without representing any of the features that (putatively) make it a case of causation. Even if feature X is the property that some theory of causation tells us *is* causation itself, and even if that theory is correct, one could still represent in experience that one thing causes something else to happen without representing X in experience. Visual experience might just be neutral on whether anything, or any relation between things, has feature X .

Holes provide an analogy. It is hard to say what makes something a hole. But it is clearly possible to see holes, and to represent them experientially as

¹⁵ Sosa and Tooley's remarks could also be developed into an argument about seeing, taking ' S observes that p ' to be factive, and substituting 'see' for 'observe': if S sees that a causes b , then S sees that (a, b) instantiate a law; we cannot see that a law obtains; so we cannot see that a causes b . The criticism raised against the argument presented in the text applies to this argument as well.

holes. Here we seem to have a case in which there is a property – the property of being a hole – that is represented in experience, even though we do not represent in experience any more specific feature that distinguishes holes from non-holes.

Suppose you see some cheese with a hole in it, and your visual experience is veridical. According to one theory of the metaphysics of holes, they are immaterial particulars. According to an opposing theory, they are material but negative parts of material particulars.¹⁶ When the cheese appears to have a hole in it, does it appear to host an immaterial particular, or does it appear to have a material part – or neither, or both? If the cheese appeared to have both of these more specific properties, then the experience could not be veridical after all, since the hole cannot be both material and immaterial. The best answer seems to be 'neither': visual experience just seems neutral on whether the hole in the cheese is a material but negative part or an immaterial part hosted by a material particular. Visual phenomenology does not seem to count for or against either of these theories, even if there are faults to be found with the theories on other grounds. If so, it seems likely that visual experiences do not take a stand on the metaphysics of holes.

Suppose some metaphysical theory of holes turns out to be correct (one of the two I have mentioned, or some other). If experience is neutral in the way I have described, then we have a case in which an object has a distinctive feature H if it has a hole, and experience represents the property of having a hole, but does not represent distinctive feature H . So – returning to the general doubt – it will not in general be true that if X is a distinctive feature of causal relations, an experience represents that a causal relation holds between two things only if it represents X . To support the doubt, then, some argument is needed that the case of causation is not like the case of holes.

An opponent of the causal thesis might say that the case of causation is indeed disanalogous to the case of holes or ice cubes: whereas it is not analytic that ice is H_2O or that holes are, say, immaterial particulars, it is supposed to be analytic that causation is a kind of counterfactual dependence. Or one might try to make a similar point using the notion of concepts instead of analyticity: it is supposedly a conceptual truth, if it is true at all, that causation is a kind of counterfactual dependence, but supposedly not a conceptual truth that ice is H_2O , or that holes are immaterial particulars – assuming that they are. But while this point shows a disanalogy between the ice case and the case of causation, there does not seem to be any disanalogy

¹⁶ The first theory is defended by R. Casati and A.C. Varzi, *Holes and Other Superficialities* (Cambridge: Bradford Books, 1994), the second by D.D. Hoffman and W.A. Richards, 'Parts of Recognition', *Cognition*, 18 (1985), pp. 65–96; see also D.K. and S.R. Lewis, 'Holes', *Australasian Journal of Philosophy*, 48 (1970), pp. 206–12, repr. in D.K. Lewis, *Philosophical Papers*, Vol. 1 (Oxford UP, 1983), pp. 3–9.

between the case of causation and the case of holes. Whatever relation theorizing about causation has to the concept of causation or the meaning of 'causation' seems to be in place in the case of theorizing about holes and the meaning of 'hole' as well.

I have argued that causation may be like holes, in that we can represent causation in visual experience without representing its distinctive features. This rebuttal emphasizes what you do *not* have to represent in visual experience in order for your experience to represent causation. One might then wonder whether there is any way causation *must* look, if it is represented in visual experience. At one extreme is the position I have just re-argued, which says that causation could be represented only if its distinctive nature were revealed in experience as well. At the other extreme is the position that things could look in *any* way to *S* compatible with *S*'s visual experience representing causation (e.g., an experience could count as representing causation, even if all that *S* seemed to see was an undifferentiated expanse of yellow). This extreme seems implausible as well. If both positions are false, then there are only some ways in which events (or other relata) can appear when they appear to stand in a causal relation. But which ways are these? To put it differently, what are the phenomenal constraints on the representation of causation in visual experience? Are there any?

One proposal for a minimal constraint is that things look as if they stand in a causal relation only if they look contiguous spatially or temporally. (The events in the ball case are not and do not appear to be spatially contiguous, but they do appear to occur in quick succession.) Michotte's results seem to provide some support for this minimal constraint. If you show adults (and, as it happens, chimpanzees) a scene of launching with a temporal gap (the second ball does not start moving until after the first one has collided with it and stopped), they do not report what they see as a case of causation.

Although this constraint cannot differentiate representation of causation from the representation of mere contiguity, it may be a necessary phenomenal condition for experiential representation of causation. The deeper question in the vicinity, however, is why there should be any phenomenal constraints at all on how causation may look.¹⁷

Harvard University

¹⁷ Versions of this paper were presented at the Arizona Ontology Conference, Bates College, the On-Line Philosophy Conference and the University of Warwick. Thanks to audiences for their responses. For further discussion, thanks to Stephen Butterfill, Elizabeth Camp, Andy Egan, John Hawthorne, M.G.F. Martin, Casey O'Callaghan, Laurie Paul, Brian Scholl, Jessica Wilson, and especially David Chalmers.