

The Effective Indexical

Huw Price

Centre for Time, University of Sydney

[Draft at November 23, 2007]

Abstract

In a famous paper in *Noûs* in 1979, John Perry points out that action depends on indexical beliefs. In addition to “third-person” information about her environment, an agent need “first-person” information about where, when and who she is. This conclusion is widely interpreted as a reason for thinking that tensed claims cannot be translated without loss into untensed language; but not as a reason for realism about tensed facts. In another famous paper in the same volume of *Noûs*, Nancy Cartwright argues that action requires that agents represent their world in causal terms, rather than merely probabilistic terms: for, Cartwright argues, there’s a distinction between effective and ineffective strategies, that otherwise goes missing. This is widely taken as a reason for thinking that causal claims cannot be translated without loss into merely probabilistic claims; and also – in contrast to Perry’s case – widely regarded as a reason for realism about causation. In this paper I ask whether the latter conclusion is compulsory, or whether, as in Perry’s case, the need for causal beliefs might merely reflect some “situated” aspect of a decision-maker’s perspective.

I Two bits of *Noûs* from 1979

In his well-known paper ‘The Problem of the Essential Indexical’ – the first article in Vol. 13 of *Noûs* in 1979 – John Perry points out that action depends on indexical beliefs. In addition to “impersonal” or “third-personal” information about her environment, an agent need “situated” or “first-personal” information – information about the “I, here and now” of her situation.¹ If she wants to get to a meeting, and believes that it starts at 1:00pm, it takes the belief that it is *now* 1:00pm to get her moving.

¹Perry, J. 1979. ‘The Problem of the Essential Indexical’, *Noûs* 13, 3–21.

Note that Perry is not offering this as a reason for thinking that the world contains additional indexical facts or ontology, over and above the facts properly recorded on non-indexical maps. Some people may believe that there are additional facts of this kind – at least in the temporal case, some people certainly claim to believe something like this.² But this isn't Perry's claim. Rather, Perry is revealing to us a fact about the cognitive foundations of action: the fact that it requires first-personal, situated thoughts, as well as third-personal, unsituated thoughts. This is a psychological discovery, not a metaphysical discovery. It is a reason for thinking – contrary to what was widely believed at the time, apparently – that indexical talk cannot be translated without loss into non-indexical talk. But it isn't a reason for populating the world with indexical ontology.

In another famous paper in the same volume of *Noûs*, Perry's (then) Stanford colleague, Nancy Cartwright, argues that action requires that agents represent their world in causal terms, rather than merely probabilistic terms; for, Cartwright argues, there's a distinction between effective and ineffective strategies, that otherwise goes missing.³

Cartwright begins with the familiar distinction between laws of association and causal laws. Noting Russell's argument that, as she puts it, 'laws of association are all the laws there are, and that causal principles cannot be derived from the causally symmetric laws of association', she goes on to argue 'in support of Russell's second claim, but against the first.' (1979: 419) She agrees with Russell that '[c]ausal laws cannot be reduced to laws of association', but maintains that 'they cannot be done away with'. (1979: 419) Cartwright's argument for the latter conclusion is that causal laws are needed to ground an 'objective' distinction between effective and ineffective strategies. She illustrates this distinction with some 'uncontroversial examples'. The most famous concerns a letter she received from an insurance company (Figure 1).

It simply wouldn't be true to say,
"Nancy L. D. Cartwright. . . if you own a TIAA life
insurance policy you'll live longer."
But it is a fact, nonetheless, that persons insured by
TIAA do enjoy longer lifetimes, on the average, than
persons insured by commercial insurance companies
that serve the general public.

Figure 1: Cartwright's example – the letter from TIAA

Cartwright notes that the objective fact reported in this letter – viz., that buying life insurance from this company would be an *ineffective* strategy for living longer –

²Though it isn't easy to characterise exactly what they believe, in my view.

³Cartwright, N. 1979. 'Causal Laws and Effective Strategies', *Noûs*, 13, 419–437.

depends on *causal* rather than merely *probabilistic* facts about the world. But, she argues, the ‘objectivity of **strategies** requires the objectivity of **causal laws**’. In other words,

causal laws cannot be done away with, for they are needed to ground the distinction between effective strategies and ineffective ones. ... [T]he difference between the two depends on the causal laws of our universe, and on nothing weaker. (1979: 420)

In this case, then, the conclusion is supposed to be metaphysical, rather than psychological. This is intended to be an argument for *realism* about causal powers, not merely a point about the cognitive architecture of deliberating agents.

Comparing these two pieces of *Noûs* from 1979, it is hard not to be struck by a strange mix of analogy and disanalogy. On the analogical side, both Perry and Cartwright argue that beliefs of one sort are not reducible to beliefs of another sort. In both cases, moreover, the argument turns on the role that the former (irreducible) beliefs play in action. But on the disanalogical side, as I’ve noted, the intended morals of the two stories are quite different: in one case, an argument for realism about causation, in the other, a discovery about the cognitive foundations of agency.

This paper is motivated by the question whether the “metaphysical” interpretation of Cartwright’s data is compulsory – whether there isn’t a “Perry option” in this case, too. In other words, couldn’t the causal representations required by agents be in some sense “merely indexical” – features of the way agents need to conceptualise their world, rather than features of the world itself? Couldn’t Cartwright’s observation rest on an “effective indexical”?

The rest of the paper goes like this. In the next section I say more about Perry’s “essential indexicals”, with which I’m aiming to compare causation: about why they’re essential, and especially about why we shouldn’t interpret that fact as supporting an ontological view of indexicality. I’ll also introduce a further species of temporal indexicality, to which I’ll make some appeal in considering the causal case.

I’ll then turn to the causal case, and try to show that analogous considerations support a non-ontological view of causal judgements (and in particular of the distinction on which Cartwright’s observation depends). Finally, to situate my view, I compare it to a characterisation Jim Woodward offers of his own manipulability view of causation, in a section of his recent book⁴ in which he aims to contrast it with the apparently more

⁴Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*, New York: Oxford University Press.

subjective view that Peter Menzies and I once defended.⁵ I'll argue that the intuitions about the objectivity of causal judgements to which Woodward appeals have analogues in the Perry indexical case, too: a fact which suggests that they are not inconsistent with an indexical interpretation of causality.

2 Perry's indexicals: why essential, why not ontological?

2.1 Why "essential"?

Why are indexicals "essential"? In brief, because we are "located" creatures: located in temporal and "personal" space, as well as "spatial space". It is worth noting that while Perry stresses the importance of our location (in all these respects) as *agents* – he focusses on the role of indexical beliefs in the explanation of *action* – the same point also applies on the input side. Located *observers* need indexical beliefs, too, in order properly to record the information they gather on "third-personal" maps.⁶

For contemporary audiences, a helpful intuition pump is provided by "virtual worlds", such as Second Life. In these cases, it is easy to see how access to all the third-person information about the game isn't sufficient to enable one actually to play. One also needs locating information about one's own avatar: which one it is, and where and presumably when it is, within the space (or spacetime) provided by the game.

2.2 Why not "ontological"?

There are three main kinds of argument that count in favour of the non-ontological diagnosis of Perry's observations – i.e., in favour of *not* concluding that these observations call for indexical facts or states of affairs. I'll call the combination of these factors the Three-fold Path to indexicality:

1. **Variability.** Confronted with what looks initially like a label for an objective feature of the environment, we discover that differently located creatures use the "same" indexical term to pick out different places, times or people. The cases in question are too systematic to be dismissed as accidental ambiguities, and there's

⁵Menzies, P. & Price, H., 1993. 'Causation as a Secondary Quality', *British Journal for the Philosophy of Science* 44, 187–203.

⁶In both cases, the crucial point is to identify a distinctive functional role for indexical beliefs, of a kind that couldn't be replicated by some suitable application of an account of the functions of belief *in general*. It is this conclusion that implies that we need something new in our cognitive architecture at this point – something tailor-made for the functional role in question.

an apparent symmetry among these different uses of the same term – no one speaker appears privileged. Breaking this symmetry is not out of the question. It is the option embraced by presentists, in the temporal case (and presumably by solipsists in the personal case), who claim priority for just one of the standpoints concerned. Whatever the merits of this option in the temporal case, it is clearly unattractive in the other cases. Hence the appeal of the non-ontological view, which – by locating the variability at the conceptual or “mode of presentation” level – avoids the need to break the symmetry.⁷

It might be thought that there is an intermediate possibility, that of restoring the symmetry at the ontological level by reading indexical claims in relational terms. But Perry’s essential point is that this option won’t do. If we try to cash the indexical term “I” as a relational term “I-*for*-X”, then either “X” itself is indexical (as in “I-for-me”), or it isn’t (as in “I-for-HP”). In the former case we’ve made no progress, while in the latter case we’ve lost the essential indexical – what Perry’s examples show is precisely that relational beliefs cast in such non-indexical terms will miss a crucial element in the explanation of actions.⁸

2. **Parsimony.** The previous argument appealed to symmetry at the level of (actual or possible) users of indexical terms, arguing that where there is the relevant kind of Variability, no one user (or equivalence class of users) has any claim to priority. This argument appeals to a corresponding symmetry in the world – i.e., to the attractions of not having to accept that one time, place or person is ontologically privileged. The non-ontological interpretation thus claims an attractive economy, perhaps backed up by physics.
3. **Genealogy.** To avail ourselves of the attractions advertised by two previous arguments, however, we need to show that we can explain the relevant linguistic and psychological phenomena, without the postulate that the function of indexical terms and beliefs is to represent indexical ontology. This is the genealogical

⁷Is there another ontological option, preserving the symmetry by allowing the ontology itself to be indexical? Perhaps, but unless the ontology in question is the kind of deflationary ontology we get for free on the back of language games, it is going to have trouble with the next two steps in the Three-fold Path.

⁸More precisely, Perry’s argument shows that there’s a functional role played by indexical beliefs which cannot be simulated by in any generic way by non-indexical beliefs – i.e., which cannot be engineered simply by plugging some suitable content into a model of the cognitive functions of beliefs in general. The functional difference in question is not therefore merely a product of a difference in *content*, but needs to be regarded as a primitive difference in *use*. Indexical concepts hook into our functional architecture in some distinctive way, which can’t be simulated by wholly non-indexical concepts.

aspect of Perry's project – the task of explaining the use of indexical vocabulary, without assuming indexical ontology. And it is an essential ingredient of an argument for a non-ontological interpretation of the data.⁹

2.3 Some working assumptions about time

The non-ontological interpretation of Perry's data treats *now* as merely indexical. For the remainder of the paper I want to take for granted what I'll call *Boltzmann's hypothesis*, viz., that the *direction* of time is also indexical – also 'in the eye of the beholder', as it were. There are objective temporal asymmetries, of course, at least in our region of the universe – the second law of thermodynamics describes the most prominent of these – but I'm going to assume that these don't establish any objective sense of temporal priority; in particular, they don't provide any objective sense in which (what we call) the future arises out of (what we call) the past, rather than *vice versa*.¹⁰ In so far as our ordinary temporal conception embodies a stronger sense of directionality, it is a feature of our representation of the temporal world, rather than a feature of the world as it is in itself. In other words, it is non-ontological, in much the same sense as the idea of the present moment itself is non-ontological.

I'm aware that Boltzmann's hypothesis will strike some people as surprising and controversial. In my view the case for accepting it is very strong, however, and has been so for more than a century. As Boltzmann saw (at least darkly), it consists in an appeal to what I'm calling Variability – the fact that physics seems to allow that there could be creatures elsewhere with the opposite temporal orientation (and to provide no sense in which their viewpoint is any worse than ours). The reason the hypothesis *seems* surprising is that the features of our temporal outlook it treats as merely perspectival are very deeply ingrained. We don't find it easy to think of them as anything less than fully objective, and it isn't obvious where else the contagion will spread, when we attempt to do so – and what else will get caught in the indexical net.

This last issue is my present concern. I think that much of ordinary modal and causal thinking gets caught in the net, when the direction of time is rendered indexical in this way. I don't have space or time enough to argue the temporal case here, so I'll simply assume Boltzmann's hypothesis in considering the causal case.

⁹Again, the point turns on the *pragmatic* character of the distinctive role of indexical terms and beliefs. Sometimes we need to explain the functional role of beliefs in terms of an ontological account of what we take them to be beliefs about. But this isn't a case of that kind.

¹⁰In other words, it isn't an objective matter whether entropy is actually increasing or decreasing in our region of the universe – there isn't a distinct possible world which differs simply in the fact that everything is happening in the opposite temporal sense.

3 Causation and the Three-fold Path

I'm interested in the project of leading causation down the Three-fold Path – in other words, of showing that considerations closely analogous to those that support the non-ontological interpretation of Perry's essential indexicals also support a non-ontological view of causal judgement. In a slogan, I want to show that causal judgments are indexical, too.

3.1 Variability

Here there is certainly a contrast between the case of causation and that of the familiar indexicals. With the latter, the possibility of inter-speaker variability is forced upon us (though for slightly different reasons in the spatial, personal and temporal cases). In the case of causation, however, it can easily appear that there is no such variation – that we are all in the same boat, so to speak. But I think there are three reasons for regarding causal judgments as similarly though less obviously variable:

1. **Causation and time's arrow.** As Russell already noted in his famous (1913) paper, our causal judgements seem to depend on accidental features of our temporal constitution – e.g., as Russell puts it, on the fact that we remember the past but not the future. While Russell didn't have all the details worked out, I think his intuitions were spot-on. Interestingly, much recent work on this topic consists in reactions to a third seminal paper in that lucky thirteenth volume of *Noûs* from 1979, Lewis's 'Counterfactual dependence and time's arrow' (*Noûs*, 13, 455–476). Discussion of the issue raised by Lewis has confirmed how hard it is to give an "ontological" account of the apparent asymmetry and temporal orientation of causation. At best, moreover, we seem to have to reduce it to something that physics gives us reason to regard as an entirely contingent and possibly local feature of our region of the universe, namely, the thermodynamic asymmetry. It is doubtful whether this is basic enough, but in any case it allows for variation, once we realise that physics allows for the possibility of creatures in regions in which this asymmetry points the other way. Thus in the causation case it is arguable that we get support for Variability from physics, if not in ordinary life.¹¹

¹¹This is a big topic, of course, and even in physics our thinking about experimentation, counterfactual reasoning, and the like, is still deeply conditioned by the causal view. (This is what makes the case so interesting, of course . . .)

2. **Causation and orientation in epistemic space.** There is a deeper point lurking beneath Russell's comment about the relevance of the fact that memory works backwards and not forwards. Roughly, the causal arrow points from what we know to what we don't know, reflecting not *ontological priority* but *epistemic priority*. (Actually, it points to a special subset of what we don't know, namely, what we take it we might fix by fiat – more on this in a moment.) Thus its direction is associated with the orientation of a certain sort of boundary in what we might call “epistemic space” – and if we imagine creatures for whom this boundary runs differently, we are imagining creatures whose causal arrow differs from ours.¹²

Here's an example of what I have in mind. Consider the perspective of someone writing the prequels to the original Star Wars movies. The prequels are required to be consistent with the original episodes, and hence choices about the later times in the originals determine facts about earlier times in the prequels. If we intervene to replace Luke Skywalker by Harry Potter in Episode 4, the ramifications will spread backwards as well as forwards in the temporal dimension of the series. This is an illustration of how non-standard epistemic access leads to non-standard causal arrows. With a little imagination, it is easy to support the idea that our conventional standpoint is contingent and parochial, precisely as the argument from Variability requires.

3. **Causation and contextuality.** I think it is arguable that the familiar contextuality of causation is best treated in these terms, as a kind of implicit epistemic variability. I don't have space to try to develop this idea here, and I think the proposal is initially less telling than the above considerations – largely because we are familiar with other ways of accommodating contextuality – but this is not to say that the familiar ways will still seem the best ways, in the light of a deeper understanding of the indexical character of causal judgements. But I leave this on the table for future work.

3.2 Parsimony

The parsimonious attractions of a non-ontological reading of causation begin to come into view in Hume. Famously, Hume regarded the necessity associated with our ordinary idea of causation as much better explained in tidy psychological terms, than added as a mysterious ingredient in reality. Russell adds a new element to the argument, in

¹²For more on this idea see my ‘Causal Perspectivalism’, in Huw Price and Richard Corry, eds., *Causation, Physics and the Constitution of Reality: Russell's Republic Revisited* (OUP, 2007), 250–292.

the form of the argument that physics has no use for the notion of causation – though arguably this was already apparent earlier, e.g. to positivists such as Mach. Perhaps Russell’s main contribution was to call attention to the issue of time-asymmetry: the lack of time-asymmetry in fundamental physics provides a powerful new element to the case for parsimony.

However, it may be that the strongest ingredient in the case for parsimony has only been coming into view in the past fifteen years or so (from a direction telling related to Cartwright’s argument), in the form of what I’ll call *interventionism* – a new understanding of the importance of notions of intervention and manipulation to our understanding of causality. If we accept first that causation needs to be understood in terms of intervention, and second that “natural” interventions are hard to find in the physical world (at least in the required numbers), then we have a new motive for seeking a non-ontological understanding of causation. Here are some remarks from Judea Pearl on the latter point:

If you wish to include the entire universe in the model, causality disappears because interventions disappear – the manipulator and the manipulated loose [sic] their distinction. However, scientists rarely consider the entirety of the universe as an object of investigation. In most cases the scientist carves a piece from the universe and proclaims that piece *in* – namely, the *focus* of investigation. The rest of the universe is then considered *out* or *background* and is summarized by what we call *boundary conditions*. This choice of *ins* and *outs* creates asymmetry in the way we look at things and it is this asymmetry that permits us to talk about “outside intervention” and hence about causality and cause–effect directionality.¹³

3.3 Genealogy

Interventionism is also an important step towards the kind of genealogical proposal that would support a non-ontological view of causation (paralleling Perry’s understanding of indexicals). Interventionism helps to bring the target into view. The genealogical project is to try to show that, and how, creatures might benefit from representing their world in terms of interventions, even if there are no corresponding “facts” (again, think of the indexical parallel, if you are unsure what this would mean).

¹³Pearl, J. 2000. *Causality*, New York: Cambridge University Press, New York, at pp. 349–350. Compare this to what physics has taught us about say the temporal indexical: as Hermann Weyl puts it in a famous passage in *Philosophy of Mathematics and Natural Science* (1949), “the world *is*, it does not *happen*. Only in the gaze of my consciousness, crawling up the life-line of my body, does the world fleetingly come to life.”

Nevertheless, the kind of functional story required for causation is apparently considerably more complex than for indexicals. In my view, it is likely to involve an intersection between two major elements:

The predictor's perspective.

Inter alia, I think, causal notions are notions irreducibly associated with the perspective of non-omniscient creatures, and inherit certain features from that perspective: their conditional/hypothetical structure, and (in our case) a typical temporal orientation, associated with the fact that we acquire knowledge from the past but not the future. We have no non-inferential knowledge of the future, but our welfare depends on what happens there – hence our interest in making the best guesses we can, in the circumstances. In my view, the realisation that this is our situation and our fate is the beginnings of a very fruitful approach to the topics of natural law, probability and chance: in each case, roughly, we can see why creatures in this situation should need (or at least profit by inventing) such notions, and this is the beginnings of a genealogy. (Once again, the argument for regarding this as a non-ontological account follows the three-fold path.)

The agent's perspective.

Important as it is, however, the predictor's perspective cannot be the full story about the peculiar indexicality of causation. There is a new ingredient which gets added to the picture only when the predictor is also an actor, whose own actions lie in the epistemic field in question.

(The relation between these two is that any agent is necessarily non-omniscient, but not every non-omniscient creature is necessarily an agent. We can imagine non-omniscient creatures who are merely *observers*, at least in some specified realm.)

In my view, one of the complexities of thinking about causation by analogy with indexicals is that some analogue of indexicality arises at both stages. The first stage is closely related to the temporal perspective itself, and to what we might call the merely *predictive* modalities of law and probability. Causation depends on this, and takes over its epistemic asymmetry, but depends on a new element, that of *agency*. Thus causation turns out to be doubly indexical, in my view.

However, it's the second bit that matters in the present context. Accordingly, I want to try to focus on the issue as to why probability *as assessed from an agent's point of view* is properly different from what we might call third-person probability. I'll try to draw a veil over the more ambitious claim that even "third-person" probability is necessarily indexical in a deeper sense.

4 The indexicality of agent probabilities

4.1 Evidential autonomy

“[F]reedom consists, in part, in being free to believe what one wants about one’s own actions.” – Joyce

I want to begin with some suggestive remarks by Jim Joyce.

[M]any decision theorists (both evidential and causal) have suggested that free agents can legitimately ignore evidence about their own acts. Judea Pearl (a causalist) has written that while “evidential decision theory preaches that one should never ignore genuine statistical evidence . . . [but] actions – by their very definition – render such evidence irrelevant to the decision at hand, for actions change the probabilities that acts normally obey.” (2000, p. 109) Pearl took this point to be so important that he rendered it in verse:

Whatever evidence an act might provide
On facts that precede the act,
Should never be used to help one decide
On whether to choose that same act. (2000, p. 109)

Huw Price (an evidentialist) has expressed similar sentiments: “From the agent’s point of view contemplated actions are always considered to be *sui generis*, uncaused by external factors . . . This amounts to the view that free actions are treated as probabilistically independent of everything except their effects.” (1993, p. 261) A view somewhat similar to Price’s can be found in Hitchcock (1996).

These claims are basically right: a rational agent, while in the midst of her deliberations, is in a position to legitimately ignore any evidence she might possess about what she is likely to do. . . . A deliberating agent who regards herself as free need not proportion her beliefs about her own acts to the antecedent evidence that she has for thinking that she will perform them. Let’s call this the *evidential autonomy thesis*.¹⁴

It seems to me that Pearl’s verse makes a different point from the one that Joyce is looking for – effectively, it is just the prescription that Causalists think we need, to avoid Evidentialist mistakes. It doesn’t sufficiently stress the indexical character of the deliberator’s viewpoint. It suggests that there is evidence but that actors should ignore it, whereas the point should be that there isn’t evidence in the first place, *from the actor’s own perspective*. Thus I think the point that Joyce is after is better put like this:

¹⁴Jim Joyce, ‘Are Newcomb Problems Really Decisions?’, ms.

The evidence my choice to *you* would provide
On earlier matters of fact,
Is irrelevant to *me* as *I* decide
On whether to perform an act.

However, the main point I want to extract from these remarks by Joyce is that they are indicative of what seems to be an emerging consensus that an agent's epistemic maps are necessarily "gappy" maps – they necessarily encode less information than impersonal maps about the actions of the agent in question, *as she deliberates about those actions*. Thus an agent's subjective probabilities have to be "indexical" in a new way – a way associated specifically with the fact that she is an agent. To repeat, this isn't just the general 'subjectivity' of any epistemic viewpoint – it is a new dimension of indexicality, associated specifically with agency.

So far, however, this consensus falls short of what we need. We need to distinguish two distinct modes of "gappiness", associated with an agent's conditional and unconditional probabilities, respectively. Some of the considerations in the above remarks bear on the latter – on an agent's *unconditional* probabilities concerning her own contemplated actions. But it is the distinctive character of the conditional probabilities that matters to the genealogy of the causal viewpoint.

4.2 Ramsey contingency: the deliberative blind-spot

Consider first the **unconditional probability of a contemplated action**, as considered from the agent's point of view. The crucial point is that an agent can't meaningfully assign credences to events which are the subject of a current deliberation. (She could assign numbers, of course, it is just that they don't count as credences in the normal sense – they don't guide betting, for example.) An agent puts a belief about what she is going to do into her belief box – onto her epistemic map – *as she decides to act* ... but that only makes sense if it wasn't there beforehand, from her point of view (or, perhaps better, that whatever was there before is irrelevant – agentive freedom trumps epistemic constraint, as it were). So an agent's epistemic map must show *terra incognita*, in regions of current deliberation.¹⁵ I take this to be a fundamental fact about the cognitive phenomenology of agency, which I'll call *Ramsey contingency*, because I think it is what Ramsey had in mind when he said, "My present action is an ultimate and the only ultimate contingency."

¹⁵We could put this in terms of the Principle Principal: an agent is bound to have "inadmissible" information about her own actions, *as she acts*.

This fact about an agent's probability space – the fact that it is undefined at contemplated actions – shows already that agent probabilities differ from third-person probabilities, in a distinctively “first-personal” kind of way. (It is precisely *our own* actions that occupy the blind-spot, after all.) However, this isn't yet the difference crucial to the causal perspective. For that, we need to turn to *conditional* probabilities. Recall that the basis of Cartwright's argument was that a merely evidential decision strategy would err in the case of spurious correlations. To associate effective strategies with an agent's indexical probabilities, we need to show that an agent's conditional probabilities behave in the right way: in effect, that they ignore the spurious correlations present in the third-person probabilities.

4.3 From epistemic blind-spot to sphere of influence

The task of an agent is to find out what else wiggles, in response to her possible present actions. This is a matter of accumulating judgements about conditional probabilities – probabilities of outcomes, conditional on actions – *but in the special probability space associated with the agent's perspective*. In the typical case,¹⁶ these conditional probabilities define a cone in spacetime, an expanding “sphere of influence”, defining all the possible futures the agent takes to be under her control. (Events in this cone have no meaningful unconditional probabilities, from the agent's point of view, prior to her choice.)

Now the challenge. Aren't spurious correlations a problem for this picture, as in Cartwright's insurance example and the so-called medical Newcomb problems? No, in my view. Explaining exactly why not is a matter of some delicacy, but my own view¹⁷ is that a crucial ingredient is to note that an agent who takes statistical data to be applicable to her own case in such a situation is guilty of an error of statistical reasoning: she is irrational, but not because she follows the wrong decision procedure. Rather, the problem is that having made a judgement of conditional dependency – e.g., of the probability of her having the cancer gene, conditional on deciding to smoke – she fails to take into account new relevant information about her own case, namely, that she now holds that very same conditional credence. So if she makes the wrong decision on this basis, the error is one of failing to assess probabilities correctly – in effect, a failure of the principle of total evidence – not a problem with her decision procedure.

In any case, as I also explain in Price (1991), this view is compatible with the claim that giving causal information is a way of short-circuiting the possibility of such error,

¹⁶Ignoring, e.g., backward causation, and the products of other strange epistemic constraints.

¹⁷Defended in Price, H. 1991, 'Agency and Probabilistic Causality', *British Journal for the Philosophy of Science*, 42, 15–76.

by handing the agent the information she needs about her agent probabilities. This no more implies that causal information needs to be ontologically basic than the fact that I can hand you information about what time it is now shows that “now” must pick out something ontologically basic.¹⁸

Obviously, this needs to be spelled out at greater length. The point I want to emphasise here is that in the conditional case, too, there is a reason why probabilities assessed from an agent’s point of view are different from those assessed from a third-person point of view, when the probabilities in question concern the agent’s own actions – i.e., in this case, when they are probabilities of outcomes conditional on actions.¹⁹ Once again, then, we have the tell-tale signs of indexicality: a reason why an agent’s way of thinking about her own situation (and its relevance to other matters) is properly different from a third-person way of thinking about the same matters.

4.4 Summary: an indexical genealogy for causal judgments?

This sketch suggests that the key elements of an indexical genealogy of causal judgments are in place, or at least have some realistic prospect of being put into place. Roughly, we have the outline of an understanding of why agents see matters in terms of interventions – why their distinctive epistemic maps take that form. And hence, I think, we have the beginnings of all three steps in the Three-fold Path, for the case of causation – the beginnings of a case for thinking of interventionist causation as a conceptual species of the same genus as Perry’s essential indexicals.

5 Situating the view: comparison with Woodward

In a chapter of his recent book in which Woodward aims to contrast his own manipulability account of causation with the (as he sees it, more subjective) view proposed by Menzies and myself, he characterises his view as follows:

[O]n the view I am advocating, our notion of causality developed in response to the fact that there are situations in which we could manipulate X , and by so doing manipulate Y . This fact led us (3.3.1) to form the notion of a relationship between

¹⁸There’s a separate issue about the classic Newcomb problem. There, one can have a divergence between causal and evidential prescriptions, at least on a sufficiently realistic understanding of causation. I argue: so much the worse for realistic understanding of causation; let’s use the analogy with the Principal Principle (and the status of ‘inadmissible evidence’) to conclude that one-boxing should be uncontroversial.

¹⁹The reason, to repeat, is that judgments about such probabilities can feed into the very actions such judgments are about, in a way which isn’t true of judgments made from the third-person perspective.

X and Y that would support such manipulations and to contrast this with the notion of a mere correlation that would not support such manipulations. However, it is built into the notion of a relationship that will support manipulations in this way that (3.3.2) such a relationship would continue to hold even if we do not or cannot manipulate X , or if our beliefs and attitudes were different, or even if we did not exist at all. If it is asked why (3.3.2) is built into our notion of causation, my response is that any other view of the matter would involve a bizarre and magical way of thinking, according to which our ability to manipulate X or our practical interest in manipulating X or our beliefs about the results of manipulating X somehow make it the case that a means–end connection comes into existence between X and Y where this connection would not exist if we did not have the ability or interest or beliefs in question. Taken literally, such a view, if intelligible at all, would require human beings to have god-like powers that they plainly do not possess. (Woodward 2003: 120)

I want to make two responses to this challenge. First, I have argued that the indexical term “we” in Woodward’s assertion that “our notion of causality developed in response to the fact that there are situations in which *we* could manipulate” is ineliminable. The judgement at the heart of claims about what we could manipulate is irreducibly indexical, due the first-person character of the agent probabilities that it involves. This is reflected in the fact that creatures with different epistemic “situations” would make different judgements about what could be manipulated by manipulating what, and there’s no objective sense in which we are right and they are wrong – to think otherwise is to accord our viewpoint a god-like priority that as Woodward says, it plainly does not possess. (Here, as in many other cases in the history of science and philosophy, it is the modest, ‘subjectivist’, Copernican view that does the better job of recognising the contingencies and limitations of the human standpoint; and the objectivist view that confuses us with gods.)

Second, the analogy with more familiar indexicals reveals that something like Woodward’s (3.3.2) can hold, even in cases in which which the non-ontological status of the facts in question is not in doubt: if this is Sydney and today is Sunday, then my thinking or desiring otherwise won’t make any difference; nor would my non-existence, or yours. So these kinds of facts have the kind of ‘solidity’ and objectivity that Woodward thinks that subjectivists must deny to causal facts . . . and yet the *here* moves, as Galileo might put it. We all agree (I take it) that the distinction between indexical and non-indexical “facts” is not a distinction between two elements of reality. It is a distinction between two ways of representing reality, the former being essential to us in virtue of certain features of our own situation. If this is right, then Woodward’s intuitions about the observer-independence of causal facts would seem to be no bar to

an indexical reading in that case, too: the kind of reading I've argued to be pressed on us by the recognition of the potential variability of the manipulator's perspective.

The passage I quoted above from Woodward's book continues as follows:

This conclusion is reinforced by the naturalistic, evolutionary perspective endorsed in chapter 2. According to subjectivist accounts, causal relationships have their source in facts about us – facts about our expectations, attitudes, and so on – which we “project” on to the world. If we think about this claim in the light of the argument of section 2.1, the subjectivist picture looks rather peculiar. To begin with, what is the evolutionary story about the benefits we derive from this projective activity? After all, our projectivist tendencies systematically lead to beliefs that, by the subjectivist's own account, are mistaken or ungrounded – mistaken in the sense that they ascribe a false objectivity to causal claims or involve thinking of the distinction between causal and correlational claims as having an objective basis in nature rather than in facts about us. Why should we and other animals go to the trouble of distinguishing between causal and correlational relationships if all that is “really out there” in the world are correlations? All that projecting seems wasteful and gratuitous.²⁰ (Woodward 2003, 120–121)

Once again, I want reply by calling attention to the analogy with familiar indexicals. In this case, Perry offers us the beginnings of an understanding “about the benefits we derive from this projective activity”, as Woodward puts it. To paraphrase Woodward some more, why should we go to the trouble of distinguishing between here and there, now and then, self and other, if all that is “really out there” in the world are the bare non-indexical facts? All that projecting seems wasteful and gratuitous. Well, Perry and others have shown us why it isn't wasteful and gratuitous, in the indexical case. On the contrary, as Perry puts it, the indexical is essential, for creatures in our circumstances: creatures who need to coordinate their own actions and observations with third-person maps of their environment.

²⁰This passage continues with the following challenge:

Moreover, why do our projective activities take the particular form they do? In some cases, a single co-occurrence is sufficient for belief in a causal connection, whereas in other cases, repeated co-occurrences have no such effect. Why is this? Saying that what distinguishes causal relationships from mere correlations are facts about our attitudes and projective activities gives us no insight into why our attitudes are such that they lead us to interpret a single episode of mushroom consumption followed by nausea as indication that the former causes the latter, but do not lead us to make a similar inference when confronted with evidence of an extensive correlation between purchase of life insurance and increased longevity. (Woodward 2003, 121)

I have argued that this issue is taken care of by the reflexive character of agent probabilities. In the usual cases of spurious correlation, such as Cartwright's TIAA example, the fact that a judgement of conditional dependency would screen off the very correlation on which it should supposedly be based ensures that there can be no such dependency, from the agent's distinctive (reflexive) point of view.

I've suggested that Perry's project be seen as a beacon for the causal case, too. I've argued that Variability and Parsimony recommend an indexical reading of causal beliefs – less obviously, certainly, than in the case of the spatial, temporal and personal indexicals, but the basic motivations are the same. This leaves the task of completing the genealogy, but here, I think, the work of Woodward, Pearl and others provides a very large part of what we need. Above all, it begins to show us why natural creatures in our situation – creatures with limited information, capable of acting in their own environment – should benefit by representing that environment in causal terms. What the indexical analogy provides is a framework to assess the metaphysical significance of these results. And the biggest element in this assessment – the element that speaks loudest for the modest, non-ontological interpretation – is the recognition of the importance of the sheer contingencies of our temporal and epistemic situations.

William James characterised pragmatism as a view that looked for “the trail of the human serpent” in our concepts and conceptual schemes. In the case of the temporal, epistemic and agentic aspects of conceptual schemes, in my view, James's metaphor is inappropriately delicate: the human contribution is more elephant than serpent, and it is only the familiarity of our temporal viewpoint that prevents us from seeing this substantial creature, sitting squarely in the middle of our field of view. We mistake the shape of the elephant for the shape of reality, as it were, taking for granted that our peculiar epistemic, temporal and agentic perspective is a view of the world in itself. But what a true commitment to naturalism demands is a willingness always to challenge the appearances, always to ask whether we can explain why things should look this way to creatures of our particular nature, in our particular circumstances. This is the genuinely progressive, genuinely naturalistic project, in my view, and nowhere is it more interesting and more important than in the case of causation. What appeals to me about the analogy with Perry's treatment of the essential indexical is the way it seems to provide new traction for this crucial project – a new lever to convince us that the elephant is a moveable beast, so to speak, and not part of the furniture after all.²¹

²¹This piece is greatly indebted to discussions with Jenann Ismael over several years. I am also grateful to audiences in College Park, Armidale, Tilburg and Venice for comments on recent versions. My research is supported by the Australian Research Council and the University of Sydney.