Three ways of worrying about 'causation'[1]

For 'Causal Republicanism' Conference, Sydney, July 2003

David Spurrett (University of Natal, Durban: spurrett@nu.ac.za) and Don Ross (University of Alabama at Birmingham, and University of Cape Town: dross@commerce.uct.ac.za).[2]

1. Introduction

Russell (1913)[3] famously characterised the "law of causality" as a harmful "relic of a bygone age", and urged the "complete extrusion" of the word 'cause' from the philosophical vocabulary. His justification for this proposal is, in large part, the descriptive claim that physicists, in particular, but more generally the practitioners of any 'advanced' science, do not in fact seek causes, conjoined with the view that it is no part of the proper business of philosophy to legislate to the effect that they should. The descriptive claim that 'advanced' sciences do not seek causes depends partly on a specific construal of cause seeking, as the quest for, or assumption of the existence of, "invariable uniformities of sequence" (1918: 188). It is these uniformities, according to Russell, that belief in the 'law of causality' leads us, erroneously, to expect and also to prize. It is these aspects of Russell's paper that provide our point of departure.

We distinguish two questions regarding Russell's position. First, is it indeed the case that scientists, at least when 'advanced', do not seek 'invariable uniformities of sequence'? Second, if they *don't*, does it follow that there is no defensible way of saying that they can, or do, seek causes? Our answers are, respectively, a qualified yes and a no. Our route to these conclusions depends largely on our distinguishing between three general sorts of question about 'causation' (the 'worries' of our title) and defending some definite views about the relationships between the worries, and between the answers to the questions particular to each area. We begin by introducing the three types of worry: cognitive, scientific and metaphysical.

2. Three Worries

2.1 Worries 'cognitive'

It is relatively uncontroversial, even among fierce sceptics about causes in science and metaphysics, that we *think* in causal terms.[4] Most serious philosophical dispute about causation, that is, centres on whether we should think causally in science or metaphysics, and if so how we should do so. There are, of course, a variety of views on what is going on when we think causally, including Descartes' a priori approach, Hume's associationism, Kant's synthetic a priori. Let us characterise the phenomenon here

---

1 This is a working draft version. Please do not cite without permission.

2 Earlier versions of this paper were presented by David Spurrett at the annual Dubrovnik Philosophy of Science Conference (April 2003), to the Moral Sciences Society of St John's College, Cambridge (June 2003) and at the Causal Republicanism Conference in Sydney (July 2003). We are grateful to the audiences on all occasions for criticism and feedback.

3 All references to 'Russell' in this paper should be read as references to 'the Russell of 1913' unless otherwise stated. Specific citations are to the version of 'On the Notion of Cause' reprinted in Russell's (1918) *Mysticism and Logic*.

4 Hume: "The appearance of a cause always conveys the mind, by a customary transition, to the idea of the effect. Of this also we have experience. We may, therefore, suitably to this experience, form another definition of cause, and call it, an object followed by another, and whose appearance always conveys the thought to that other." (*Enquiry Concerning Human Understanding* section 7, part 2).

quite broadly in terms of *anticipation*: some systems in the world (including people, other animals, plants, other life-forms, and sub-systems within some of them[5]) allocate their limited resources in ways that make sense if we attribute beliefs about the relative likelihood of various possibilities (that is, beliefs about conditional probabilities based on projected[6] patterns and regularities), and what difference this or that behaviour might make.

We regard anticipation in this sense as a cognitive phenomenon, and expect that cognitive science (including some philosophy) will be able to tell us all manner of things about how (and how well) it works, how it relates to other cognitive phenomena, how it can malfunction, and in what ways it may differ across species, learning histories and environments. If we agree that the study of anticipation can reasonably be called (as it often is) the study of 'causal cognition' or 'causal learning', etc., then one general set of worries about causation is just the various ways of wanting to know what is going on with anticipation, in general or in specific cases.

2.2 Worries 'scientific'

Russell, as noted, maintains that 'advanced' sciences have no use for the 'law of causation' and don't seek causes understood in terms of that law, going so far as to claim that in such sciences the *word* 'cause' "never occurs". We consider the question of what to count as an advanced science in more detail below. For now, whether or not Russell is correct about the relevance of the 'law of causality' to science, advanced or otherwise, we take the following two observations as primary data: First, there are a very large number of sciences. Second, many of them, perhaps all, explicitly claim to pursue and provide answers to causal questions. By this latter claim we mean only (for now) the weak claim that the output of communities of scientists of a wide range of sorts includes papers and other documents that refer to the 'causes' of various phenomena, or the 'mechanisms' that 'produce' them, or ways in which various effects can be 'inhibited' or 'prevented' or how this or that intervention can make a difference to the likelihood of some outcome. This is supposed to be an empirical claim, although it's not obvious just what would count as good evidence for it. At a level somewhat above the anecdotal, though, consider the following.

We consulted a single issue of *Science* (volume 300 – June 2003)[7] and noted that things referred to as being 'caused', 'determined' or being the 'result' of something within the first few pages of the issue (excluding letters, pieces on policy matters and editorials) include 'broadening of lines' in spectroscopy, 'global extinctions', 'a gradual drop in IOM [Indian Ocean Monsoon] intensity over 5000 years', 'an electrical barrier', 'how genes are expressed in two closely related species of fruit fly', decrease in 'the reactivity of some state-to-state channels', 'efficient, spontaneous cell-cell fusion' (p1617), 'the common cold' (p1619) 'oceanic primary production' (p1625), and 'pear blister canker and cadang-cadang disease of coconuts' (p1633). More significantly, perhaps, in the same issue there are thirteen full-length 'reports'. Of them eight explicitly refer to cause, effect, or some cognate term in their abstract, although in one case the 'consequence' is admittedly of a formal property of a theory. While the remainder don't use explicitly causal language in their abstracts, the content of each at least sometimes explicitly refers to causes, or is (sometimes) amenable to being described in causal terms without serious hermeneutical violence.

A simple search for articles in which the word 'cause' appeared in the on-line archives of *Science* between October 1995 and June 2003 returned a list of results containing 8288 documents, which

---

5 Immune systems, for example (Hershberg and Efroni 2001).

6 We say 'projected' rather than projectible, because we often project in the absence of projectibility, for example attributing current, rather than merely counterfactual, beliefs and desires to the dead, or projecting patterns that are in fact artefacts of experimental control as though they were not. (See section 3.1 below.)

7 CONFERENCE NOTE: Compare with Chris Hitchcock's survey if that appears in his final paper.

indicates an average of just under 90 documents per month in which the word 'cause' occurs. 'Effect' is more popular – 10456 documents for the same period, just over 112 per month. 3646 documents include the word 'influence', 8805 the word 'response', 2683 the word 'consequence'.[8] (In case anyone was wondering, 'philosophy' turns up in 553 documents, just under 6 times a month on average, or approximately half as often as 'soil' with 1007.)

So scientists, it seems, talk about causes, and do so a great deal, and (unless intolerable strain is to be placed on the notion of an 'advanced' science) Russell is simply wrong that the *word* cause 'never occurs'. We take it that what is typically being provided, or attempted, when scientists talk about causes is *causal explanation*. A well-established set of philosophical worries about explanation in general, and causal explanation in particular, arise here, and these are our 'scientific worries'. We regard these worries as being focussed on the problem of what makes for a good explanation in a particular science, in general, or for some group of sciences.

For the purposes of this paper we adopt, without argument, a broadly Kitcherian approach to these questions. As Kitcher has emphasised, the explanatory relations accepted in any special science are reciprocal functions of the explanatory schemata accepted in that science. For now, we understand by 'cause' in the 'scientific' sense, some subset of the parochial explanatory relations peculiar to a given special science. We also note that some of the sciences are concerned with functional explanation (in a normative/teleological sense of function) and teleology, and that this shows in characteristic reference to certain types of causes in their explanations. The sciences in question are the behavioural sciences, broadly understood. For our current purposes we also take this fact as a primary datum.[9]

Note here, that the trend manifest in the efforts of functionalist philosophers to turn back behaviourist and positivist campaigns against reference to agency run in the opposite direction to broadly 'Humean'[10] moves to drain it out in other parts of science and philosophy. (Indeed behaviourism was inspired by this Humean spirit.) As we will see this provides one source of semantic strain on talk of 'causation'.

2.3 Worries 'metaphysical'

There is a set of standard problems about the metaphysics of causation that are not our primary interest here. They include questions about whether the causal relata are events, facts, features, property instantiations, tropes, states of affairs, situations, aspects or anything else, and questions about how many causal relata there are (see Schaffer 2003 for a survey). These problems are all, in various ways, problems concerning the regimentation of intuitions about the *concept* of causation; and they will seem pressing problems only to philosophers who believe that conceptual regimentation is a fruitful way of arriving at metaphysical conclusions. Such belief lives uneasily with Russell's project, and ours, of wondering whether intuitions about a concept of causation identify *anything* metaphysically significant in the first place. On the other hand,there are some paradigmatically metaphysical questions that do not, in being posed, take a general concept of causation for granted, but that can be seen, and have been seen, as having something to do with causation, and they are the focus of the worries that are of interest to us.

We take it that the variety of sciences noted in the preceding section is important, and that questions arise about whether, or in what ways, the findings of various sciences 'hang together'. No single

---

8 The on-line archive does not indicate how many documents are searched. Note that the archive includes letters, news, articles on policy, editorials, etc., as well as reports and research articles. Note also that the search engine uses 'stemming' so that a search for 'cause' will also return as results articles including words such as 'causal' and 'causality'.

9 Although see Ross & Spurrett (forthcoming) for a defence of such explanations in ways complementary to the present work.

10 CONFERENCE NOTE: Perhaps relate to Helen Beebee's discussion of Hume.

science, and perhaps not any combination of sciences *as* sciences, can answer those questions. Three aspects of the worry about how things hang together that can readily be regarded as having something to do with causation are the tasks of:

(a) Identifying a unifying ontological structure that justifies the argument patterns accepted across all of the sciences.

(b) Saying something genuinely enlightening about the ontologies of all sciences by reference to general structural relations of some kind.

(c) Identifying the 'glue' that holds all objective relations in place.[11]

We don't claim that there are the *only* metaphysical worries related to causation. Note, though, that from a realist perspective (and we're assuming realism throughout here)[12] these are at least among the crucial metaphysical questions relating to causation. The last of the three is perhaps the one most clearly associated with causation in the relevant philosophical tradition – what we call glue is, roughly, what Hume called the 'cement of the universe', or has more recently been called 'Humean bondage' (Hitchcock 2003). Much metaphysical inquiry seeks to justify the idea that unification of a scientific worldview could consist in providing grounds for thinking that there is some *one,* general, sort of glue. Such glue, if it exists, must obviously be a generic relation of some sort. If something worthy of the name `causation' were this relation, then it would follow that the metaphysician would be positing an overarching kind of causation distinct from (even if subsuming as instances, though this would itself require further investigation), the parochial causal relations of the special sciences. It is then a proper metaphysical question as to whether any such relation actually structures the world.

3. Arguing for Separation

The quick survey of three ways of worrying above is only intended to put some preliminary flesh onto our claim that there are (at least) three broad and distinct areas in which questions about causation can be seen to arise. One question that naturally pops to the fore at this point is: *what is the relationship between the three?*

Attention to this question will occupy the concluding section of our paper. First, however, we must spend more time with the differences amongst concepts of causation, and by doing so isolate the parameters and constraints around any attempt to identify their possible basis for family resemblance.

3.1 Cognitive, Scientific and Metaphysical

'Cognitive' questions about causation, recall, arise given the apparent fact that some sorts of systems *anticipate*. Around this phenomenon cluster a collection of questions: *How* do we anticipate? *What* exactly is it that we do when we anticipate? Is the cognitive process like an inference, and if so what *sort* of inference? How different is what *we* do when we anticipate to what nonhuman anticipating systems of various sorts do? How much of what is going on here is learned, and how much innate? Is any of it unlearnable? Cognitive scientists of various sorts have shed some light on these questions,

Contrary to what some kinds of behaviourist expected, for example, some associations seem easier than others to learn. Rats, for example, find it easier to learn (and show appropriate anxiety when

---

11 We used the same set of questions to a related purpose in Ross & Spurrett (forthcoming).

12 Our preferred brand of realism, the 'ontic structural' flavour, is articulated by James Ladyman (e.g. 1998, 2000; French 1998; French & Ladyman 2003a, b; see also Worrall 1989).

presented with the stimulus) an association between eating a certain food and nausea, and a loud flash and an electric shock, than the complementary pairings.[13] Developmental psychology, for its part, already tells us a great deal about the causal (in this cognitive sense of causal) expectations of children around 3 months old, who are, *inter alia*, surprised when apparently cohesive objects seem spontaneously to fragment, to pass through one another, or fail to exhibit a single trajectory through space and time (Spelke et al 1995). Other work (Baillargeon et al 1995) indicates that at the same age infants are disposed to find apparently unsupported objects that do *not* fall surprising. Some headway is being made with the difficult question of what features of an object's behaviour and/or structure lead children to regard it as capable of self-motion, either at all or of various particular sorts (Gelman et al 1995).

Before starting to argue for separation, one crucial point about our causal thinking (in the sense where it is an object of study for cognitive science) must be noted. It is not recognizably 'Humean' on any strict construal of 'Humean'. We take it that a strictly Humean notion, whatever, exactly, *Hume* might have thought, of causal thinking would require that effects followed causes infallibly – and Hume did write of an effect being the 'infallible consequence' of a cause.[14] He also offered that 'we may define a cause to be an object, followed by another, and where *all* the objects similar to the first are followed by objects similar to the second'[15] (emphasis added). This is the 'constant conjunction' view of our cognitive processes, which forms a historical ancestor for similar views about cause in science, and the nature of scientific law. We can also find strong statements of it in, for example, Kant[16] and Mill.[17]

There is ample evidence, though, that we do not think in terms of constant conjunctions at all, but rather in terms of multiple chains of influence similar to directed causal graphs, or Bayes nets. Children between the ages of 3 and 5 show the ability, *inter alia*, to reason to unobserved causes, to plan suitable novel interventions in response to requests to prevent a process they had only observed in operation (i.e. one that they had never seen prevented), and to make inferences about the 'direction' of causal dependence in cases of simultaneous change (see, e.g., Gopnik et al, in press). The data are inconsistent with any simple conditioning or associationist view, effectively demanding recognition of a human[18] capacity to construct causal models at least broadly equivalent to Bayes nets. The evidence noted above that some associations are more readily learned than others also undermines Hume's

---

13 See, e.g. Garcia and Koelling (1966) and the papers in Seligman and Hager (1972), as well as Palmerino *et al* (1980).

14 *Enquiry Concerning Human Understanding*, Section 7, part 1.

15 *Enquiry Concerning Human Understanding*, Section 7, part 2. Also: 'similar objects are always conjoined with similar. Of this we have experience' (Enquiry: 60).

16 Kant (1933, p.218) "All alterations take place in conformity with the law of the connection of cause and effect"; "Everything that happens, that is, begins to be, presupposes something upon which it follows according to a rule." Kant also asserted that "experience continually presents examples of such regularity among appearances" (KdRV B123).

17 Mill (1872, Bk. III, Ch. V, §2): "The law of causation, the recognition of which is the main pillar of inductive science, is but the familiar truth that invariability of succession is found by observation to obtain between every fact in nature and some other fact which has preceded it, independently of all considerations respecting the ultimate mode of production of phenomena and of every other question regarding the nature of 'things in themselves'." At times Mill (for example when writing of the composition of causes) made clear that his view was probably not, at least not consistently, strictly Humean: 'there are such things in nature as parallel cases; that what happens once will *under a sufficient degree of similarity of circumstance*, happen again' (System of Logic III, 3, emphasis added). More strikingly distinguishing laws from constant conjunctions, while discussing composition of causes, Mill wrote: 'The law, in short, of each of the concurrent causes remains the same, however their collocations may vary; but the law of their joint effect varies with every difference in the collocations' (Mill 1972: 469).

18 It is not clear to what extent non-human animals are capable of this sort of causal learning. [TASK: Add references to Povinelli, Tomasello and Call].

confidence that experienced events are 'entirely loose and separate'[19] or that a person without experience could form no causal beliefs at all.[20]

If our actual psychology is *not* strictly Humean, then arguments pointing out that science isn't strictly Humean may fail to establish the sort of interesting or important contrast that might otherwise be expected. Similarly the temptation to make talk of causes in other areas, including science and metaphysics, be strictly Humean may lose some of its urgency. How, though, does our psychology, however it works, relate to science and metaphysics?

Where we say 'cognitive' Russell and his contemporaries might have said 'psychological'. Furthermore, Russell (1918: 184) objects to one of the definitions of 'cause'[21] in Baldwin's *Dictionary* (1901) that it is psychological, focusing on 'thought or perception' of a process where what is required is a definition of 'the process itself'. Since the full title of Baldwin's (1901) *Dictionary* indicates that it is a *Dictionary Of Philosophy And Psychology*, Russell's complaint that such a work contains some psychology is a little bizarre. The entry in question also takes pains to distinguish the 'psychological treatment of causation' from the 'logical', a detail Russell opts to ignore. In fact the authors of the 'Cause (notion of)' entry are quite clear that questions about how science proceeds, or how it *should*, are separate from questions about how we, in point of fact, happen to think – one of the very claims that Russell goes on to defend.

We would like to offer a complementary argument to Russell's for the conclusion that the psychological notion of causation should not be allowed to legislate for physics. Russell's claim is similar to one advanced by Redhead (1990) who regards the common sense notion of causation as an unwelcome, anthropocentric intruder in physics. In this Redhead echoes some of Russell's remarks (e.g. 1918: 206) about how our expectation that science should discover causes results from our projecting into it a model of volitional influence he thinks we discover in introspection. Russell, and Redhead, both argue, in effect, that nothing in physics conforms to the psychological notion of cause, that physics is fine the way it is, and hence that any difference between physics and psychology does not mean that physics has to change in any way. We agree, but not simply *because* some view about causality is psychological, but because there are good independent grounds for thinking that our psychology is both pretty impressive, *and* of limited legislative value in science and metaphysics.

On the first claim, our anticipating systems are probably highly efficient, by any standard, at extracting and integrating information at the scales and paces typical of our normal embodied lives. We wouldn't be here if they weren't. Not only that, it could be that some kind of folk notion of agent causation (that is, of making things happen by pushing and pulling, and of there being the kinds of thing that can push and pull and have reasons for doing so) is part of this makeup. It may even be that having some version of the folk concept, and applying it to some features of the world, is biologically necessary for some types of agency, including our own. Even if that were so, it would not follow without further argument that the folk concept would be a good guide to what should happen in science and metaphysics. That's not what it is *for*.

Furthermore, the 'folk conception' of causation manifest in the things we *say* about what causes what, and what we anticipate in any given case, may be a *less* accurate instrument than our non-linguistic mechanisms for anticipation. It *would* be so if the production of digitally coded signals required Procrustean manipulation of the output of the non-digital anticipating systems, and that would in turn be the case if the requirements for getting an effective digital signalling system going included drawing fewer distinctions, expressing a smaller number of relationships, etc. (This may be bad news for conceptual analysis.)

---

19 [TASK: Fix reference. The claim about the person without experience ("*Adam"*) is in the abstract of the *Treatise*.]

20 [CONFERENCE NOTE: Also the site of a possible connection with Helen Beebee's paper.]

21 The authors of the entry in question are J. M. Baldwin himself, and Dr. G. F. Stout.

We take it that if we indeed have such things as pre-theoretical 'intuitions' about cause at all, then some of them, at least, are facts about the folk conception, and what we say (including to ourselves) about causes. But then our 'intuitions', if we have such things, are evolved mechanisms for getting around in the world under various resource (including time) pressures, and navigating social games, not cognitive gadgets designed to produce systematically worthwhile guidance in either science or metaphysics. That is to say, no matter *how* things are with science or metaphysics, we contend that our intuitions are best regarded with cautious suspicion. Here are a few quick reasons for being suspicious of our intuitions:

First, the distinguishing feature of an intuition is that it be immediate, and particularly that it not involve inference or reasoning.[22] But (and Descartes, for example, knew this perfectly well[23]) what counts as immediate depends partly on how much practice we have devoted to something – some of our 'intuitions', that is, will be developmental achievements, produced by long training. (There is, in this light, something at least faintly absurd about philosophers of all people, after years of very specialised training, and with highly developed skills in ritualised forms of debate, claiming to *have* such things as 'pre-theoretical' intuitions at all.)

Second, there is reason to think that what counts as anyone's intuitive (in the sense of immediate, unforced) response to a question, including a 'causal' question, is subject to cultural variation. Although data in this area is often subject to controversy, there is a growing body of evidence that adults in different cultures engage in strikingly different causal reasoning (See Morris et al 1995 for a survey). These often manifest in different 'intuitions' – Chinese and American newspaper reports of very similar local multiple murders differed strikingly in the extent to which they referred to dispositional properties of the perpetrator (more common with Americans) and situational ones (more common in Chinese reports). Further, Chinese and American subjects in an experimental task involving counterfactual judgements about the murders had differing views about what sorts of changes (dispositional or situational) would have been likely to *prevent* the murder.

A possible response to here would suggest that some sort of 'core' human notion of causation could perhaps be recovered, even allowing for the fact that environment, including cultural environment, can play a role in constructing distinctive local manifestations of causal thinking. That seems reasonable, but it does not damage our key point. That point is that naked appeals to 'intuitions' about causation are not innocent moves, and that difficult work, not typically undertaken by philosophers who often seem very confident of the universality of their own intuitions, would be needed to establish what counted as the 'human' intuitions on such matters. And even then, a justification would be needed for thinking that intuitions were worth appealing to at all.

Third, data that might plausibly be taken to indicate what our intuitions are in various cases can, and sometimes does, show that they are depressingly wrong by the lights of our scientific and mathematical theories of the relevant domain. Again, interpretive controversies about some of this data rage on, but it seems reasonably clear that people are, when relying on their intuitions, relatively poor at judgements about conditional dependence (as measured in the Wason selection task), probability (where many people will assert that the conjunction of two possibilities is more likely than one of the conjuncts alone) and in other areas, compared to communities of scientists. The very infants

---

22 Locke (*Concerning Human Understanding* Book IV, Chapter 2): "For if we will reflect on our own ways of thinking, we will find, that sometimes the mind perceives the agreement or disagreement of two ideas immediately by themselves, without the intervention of any other: and this I think we may call intuitive knowledge."

23 Rule 11 in Descartes' *Rules for the Direction of the Mind* suggests that we should run through deductions 'in a completely uninterrupted train of thought' to help us 'form a distinct and, as far as possible, *simultaneous* conception of several of them' (1985 volume 1: 37, emphasis added). In Descartes' case, of course, some of his other commitments concerning the question of his own nature and origin gave him great confidence that what seemed utterly obvious to him just *had* to be true, so practising arguments until they became 'simultaneous' or intuitive was epistemologically advisable.

who seem to have pretty solid expectations about what an object can do, turn out to be rather bad judges of how shadows cast by objects should behave, for example being surprised when a shadow moves at the same time as the object apparently casting it (Spelke *et al*, 1995). Our response in such cases is not to take the mathematical theory of probabilities, or logic, or the laws of optics to have been impugned, and nor *should* that be our response. One way of thinking of this last point is as follows. Just as the cognitive scientific study of vision has enabled us to explain some visual illusions, so the study of causal cognition in a scientifically informed setting might be expected to show us that we were prone to some 'causal illusions' (Gopnik *et al*, in press). In a Dennettian spirit we propose that the proper task of philosophy and science in cases of causal illusion is to explain the answers people give to some causal questions, rather than to try to make the answers turn out to be correct.[24]

Having said all that, something remains to be added in defence of our capacities for causal cognition. They are, as noted, probably highly efficient. They can plausibly be seen as providing a good first approximation to a science of 'goings on that an ape like that needs to pay attention to and make many successful predictions about'. We can also allow that folk-physical and folk-psychological explanations could *both* be first approximations of sciences. But being that isn't the same as being serious guides to science or metaphysics.

A possible objection to our conclusion would be to point out, or argue, that cause talk plays an important pedagogical or pragmatic role in science.[25] We'd be willing to grant such a role, but maintain that doing so would not significantly undermine our main point, but rather reinforce it: even if some idiom or way of expressing scientific work was cognitively useful to us, it would not follow that it was more than useful. We also warn against making the supposed cognitive properties of the single individual the bottleneck through which scientific discovery and justification must pass. Through the division of intellectual labour, and augmentation of our cognitive abilities by sensory and other prostheses, communities of scientists come to have cognitive powers quite different from those of single individuals (See Hutchins 1997 for an account of collaborative cognition in a navigational setting).

In our discussion thus far we've allowed an ambiguity between at least two sense of science to pass unresolved, and we hope without serious consequence. The ambiguity is between science as a practice, and science understood in a realist mood as a body of claims about the mind-independent objects of scientific enquiry, whatever they may be. Bringing that distinction into the foreground, though, is to raise metaphysical questions explicitly, and it is to the connections and differences between science and metaphysics that we now turn.

3.2 Scientific and Metaphysical

To begin with, we need a little more detail regarding science. From one perspective the survey of *Science* magazine described above shows very little. Russell could indeed by wrong about the *word* 'cause' being used in science, yet either be correct that it shouldn't, or be correct that when it is, the 'law of causality' is not what is being invoked. We've said, and maintain, that when scientists refer to causes they are attempting causal explanations. To defend this claim we need to consider the substance of Russell's description of science.

Russell's more serious claim regarding scientific practice, when one leaves aside the question of what words are used, is that scientists don't seek causal laws, when that is understood to mean seeking laws

---

[24] By way of a brief aside, this is why, unlike e.g. Hitchcock (2003), we aren't persuaded that the failure of our intuitions to give decisive answers, or unanimous ones across several people, to some causal questions tells us anything of scientific or *metaphysical* significance, although it may tell us something interesting about our intuitions. [CONFERENCE NOTE: Possible opportunity to connect with Chris Hitchcock's paper.]

[25] Norton (2003) may be an example of such a view.

of invariable succession. More specifically Russell says that this holds of 'advanced' sciences. (Of course if an advanced science just is one that makes no reference to causes, or at least has no use for the 'law of causality' then Russell's claim is true by definition.)

Russell's argument depends, then, on taking a scientific (causal) law to be a law of invariable succession – that is holding scientific practice up against what Cartwright later (1980) called the 'facticity' view of laws. Russell in effect anticipates Cartwright by pointing out that laws thus understood have numerous counterexamples, although instead of emphasising their resulting falsehood, focuses on the features of scientific practice that both produce many of the counterexamples, and show the inadequacy of the facticity view of laws. So, Russell concedes that in the 'infancy' of a science (1918: 188) the principle of same cause same effect may prove useful (he does not say how), but maintains that every advance in science is a move away from such simple formulae, and towards understanding phenomena in terms of more complicated, mathematically specific and time-symmetric,[26] forms of interdependence, for which the law of causality is unsuitable, and which Russell therefore thinks should not be called 'causal' at all.

Russell's position does not involve commitment to anything like Cartwright's more recent view that the world is 'dappled' and that fundamental physical laws are to be regarded with suspicion (1999). In fact Russell takes (to put it in current terminology) the failure of the facticity view of laws to be worse news for the facticity view than for the laws themselves. He contends that while at most needing the 'law of causality' as a prop in its early stages, any science assumes the 'uniformity of nature' (1918: 194), and grants that controlled conditions under which constant conjunctions are observed can be useful in the discovery of the relations of interdependence that he takes mature science to seek. Mature science, on his view, tends to produce 'a greater differentiation of antecedent and consequent' and a 'continually wider circle of antecedents recognised as relevant', leading to a view of the relationships between measurement values at times and places, where simple generalities about regularities in the actual flow of events are at best of limited practical value and minimal theoretical interest.

We agree that this picture is roughly correct, and observe that when contemporary scientists talk about causes what they are typically doing is documenting that kind of relationship. Although Russell's descriptive claim about what scientists say may be incorrect, he can be read as making the prescriptive proposal that doing *that* isn't seeking causes, and that the fruits of such work *shouldn't* be understood, or referred to, as finding 'causes'. We, though, don't see that anyone, philosopher or otherwise, is holding a trump the playing of which enables them to tell scientists that when they're seeking the sorts of relations of interdependence they do, they aren't seeking 'causes', or shouldn't say that they are.

One objection here might be that the work of scientists does not involve seeking 'Humean' uniformities, and that since *that* is what 'cause' refers to, scientists are mis-using the term. But if we're correct that our own psychology isn't Humean in that sense (3.1 above) then one key temptation to make event regularity definitive of causation disappears. Furthermore, suppose that, in the spirit of Russell, we sought in science for something rather like our own causal cognition to check whether what goes on in science should be called causal. Then, we suggest, we'd find analogies as well as differences – neither we nor science rely on constant conjunctions, instead we both consider co-determination by multiple factors. On the other hand we seem committed to the unidirectional determination of the future by the past, a notion that finds no comfortable home in many scientific theories. So if our own psychology is to exert any regulative pressure on what scientists can call causes, the question seems open; and furthermore we *don't* think it should exert such pressure.

---

[26] That is, being time-reversible, and not having any theoretical home for the notion of the direction of causation from past to future that seems central to (at least) the folk notion of causation. We agree with Field (forthcoming) that the lack of past-future asymmetrical dependence in fundamental physics is not Russell's main argument against the relevance of causes to science.

A more difficult objection would insist that the problem of distinguishing genuinely causal from non-causal processes is, at least partly, a *metaphysical* problem, and hence that metaphysicians should have something to say about whether scientists are behaving appropriately in referring to causes.

Recall the three tasks we identified as aspects of the overarching metaphysical worry, in realist mood, about how things hang together, each of which can be understood as being related to causation:

(a) Identifying a unifying ontological structure that justifies the argument patterns accepted across all of the sciences.

(b) Saying something genuinely enlightening about the ontologies of all sciences by reference to general structural relations of some kind.

(c) Identifying the 'glue' that holds all objective relations in place.

The first question is most strongly associated in the recent philosophy of science with the work of Philip Kitcher (e.g. 1976, 1981, 1982, 1989), the latter two with that of Wesley Salmon (e.g. 1984, 1994). We begin by briefly glossing each programme without, for reasons to be discussed shortly, making direct reference to causation.

Kitcher has argued that *ontological unification,* within one special science or across two or more of them, consists in the justification of common *argument patterns* that hold within or across the science(s) in question. (An argument pattern is a set of ontological and structural primitives featuring recurrently in the explanans of accounts given in the unified domain.) Kitcher (1989) characterizes his work as analyzing `top-down' explanation, where phenomena are explained by having their roles in wider ensembles of regularities fixed. He contrasts this with `bottom-up' explanation, the sort analyzed by Salmon.

According to Salmon, science aims to discover the structure of the world. In his (1984) he attempted to say something enlightening about the ontologies of *all* sciences by reference to general structural relations holding, at a micro level, between all real objects and processes. He argued that some processes transmit information about their antecedent stages, while others do not, and that those which *do* are genuine processes, the others mere pseudo-processes.

Tasks (a) and (b) are most naturally seen as complementary, and also the two most readily understood as having no essential connection to causation. Task (c), though, is deeply associated in the philosophical tradition. What we call glue is what Hume called the 'cement of the universe' and, furthermore, Salmon's project was in part an attempt to wrestle seriously with Hume's epistemological challenge, to say how anyone could know, by any amount of observation, which links between processes are causal and which are not. Salmon's process account is supposed to show that we can, at least, observe something that is precisely *diagnostic* of causation.

With Salmon and Reichenbach (1957)[27] we can describe this in terms of the transmission of *marks*. In the absence of specific structure-preserving (and constituting) activity, entropy will eliminate marks on objects that carry information about their histories. More generally we can describe this in terms of negentropic *information*.[28] A structure is, by definition, something that resists entropy, at least briefly.

---

[27] It is worth noting that Russell himself provides a crucial inspiration for process accounts of causality in his (1948). Some of the connections are described in Dowe

[28] As Dowe (1992) pointed out, and Salmon (1994, 1997) accepted mark transmission cannot be fundamental since some real processes cannot be marked. Dowe's conserved quantity version of the process view does not have this limitation. Our preferred approach, equivalent to Dowe's with respect to the mark-transmission problem, is Collier's (1999) which makes information, but not marks, fundamental. Neither Collier's nor Dowe's approach rely on counterfactuals in the formulation of their

Therefore, wherever marks are preserved we have structure. The goal of science is to discover the structures in nature.

`Structure', furthermore, admits of rigorous, and realist, account. Wherever local negentropy is high (Schrödinger) we have an information channel. Some such channels are exploited for convenience by local limited information-processors (such as ourselves), but are ontologically redundant, in that the information they carry is recoverable by reference to structures of greater computational depth. Information channels that would *have* to be used, by any physically possible information processor, to describe all events and regularities are *real patterns* (Dennett 1991). These are the structures that science seeks to discover.

But what about causation? Notice that, as intended, the process view enables putatively causal claims by special scientists to be subjected to a critical test. Those which pick out real processes could be causal, those that don't, can't, because they don't pick out real processes of *any* kind. So when special scientists seek causes in their domains of enquiry and get it right, whether they know it or not, they find ways of picking out real informational relations. That is to say, if macroeconomics, for example, has its own notion of 'macroeconomic causation,' and process views are correct, then this notion will turn out to correspond to some specific informational relation, one that really holds on the general topology of the universe. Saying this does not entail saying anything specific about how an isolated macroeconomist might reason her way to a generalization.

That process views have this payoff when it comes to distinguishing real from pseudo processes counts in their favour. And if we're allowing that when scientists seek causes they may indeed be seeking true causal explanations even if they're not hunting Humean regularities, then process views turn out to play the right regulative role with respect to science, without undermining the explanatory autonomy of specific sciences. Does that mean that process views are best thought of as analyses of causation? We contend that it does not, and that instead it is an open question, depending on a semantic decision, whether to construe them as analyses of causation. There are considerations weighing in on both sides.

One argument *against* regarding process views as analyses of causation can be dispensed with relatively quickly. Kitcher (1989) criticizes Salmon (1984), by pointing out that our identification of causal processes and causal interactions in most, probably all, macro-processes depends more on our knowledge of the general causal structure of the macro-world than it does on the sorts of micro-process that are paramount for Salmon. However, while such considerations give good reason for not attempting to extract methodological guidelines[29] for conducting the business of the macro-sciences from process accounts, they do little to undermine the value of Salmon's approach and its descendents when they are understood as saying something about the kinds of real structures in the world that science aims, in the limit, to discover. Or, more directly, a fair point about *explanatory* priority in science need not by itself establish anything about fundamental ontology.

As hinted at already, there is an argument in favour of regarding process views as being precisely about causation, and that is that they are the heirs of a philosophical project, descending from Hume, through Russell (1948), Reichenbach and others seeking to purge folk elements from the concept of causation in favour of an analysis giving priority to considerations arising from fundamental physics. The tension between how things are with fundamental physics and how they are with the folk notion, that is, can be interpreted in at least two ways. Redhead (1990) explicitly acknowledges Russell (1918) when developing his own argument for the thesis that talk of causes has no place in physics.

---

basic concepts, thereby lacking a recognized limitation of Salmon's (1984) approach. While important, these refinements within the process programme are not crucial for our purposes here.

[29] Except, at least, for the imperative not to traffic in parochial information-transmission relations that are spooky according to fundamental physics.

> In what sense does the event consisting in the object having fallen a distance s_1 at t_1 cause the object to reach s_2 at t_2? … The position at t_1 can hardly be cited as the cause of the position at t_2. Surely the cause must be something that is introduced to account for why the position of the object has changed. Well, is not the answer to cite the acceleration? But the acceleration is just defined by the kinematic relationship expressed in Galileo's law. So can the cause be Galileo's law? But Galileo's law is not another event which happens at t_1, `causing' the particle to move with constant acceleration and so to reach $s_2$ at $t_2$. It is just an expression of that acceleration. So should we not retreat to citing the force, such as gravity, which `causes' the acceleration? But the idea that forces `cause' bodies to move is a very anthropocentric notion (Redhead 1990: 146).

Redhead here is granting physics priority in telling us how the world is, but allowing the folk notion priority when it comes to fixing what is to count as a cause. As a result, like Russell, he argues that nothing in the world discovered by science corresponds to the folk notion,[30] and urges its elimination.

We've already indicated some reasons for rejecting the view that the folk notion is a regularity notion. For now note that, *whatever* kind of notion it is, a possible response to any difference between the practice or content of science, including fundamental physics, and the folk notion, could be to give priority not only with respect to how the world is, but also to what causes are, to science. In that case one would say that the folk are *wrong* about what causes are, and that even if what goes on in physics is strange or unfamiliar, it is where we'll find out what causes really are.

In opposition to this argument giving semantic priority to process analyses of causation, understood as heirs to Hume's project of removing folk elements from the concept, though, are heirs whose ancestry goes back to the very folk notions that process analyses undermine. Most striking here are functionalist approaches to mind, and in the behavioral sciences more widely. We noted their existence, and their position of tension with respect to Humean approaches including early behaviorism, in section (2.2) above. So semantic ancestry does not give unequivocal support either for or against regarding process views as analyses of causation, and we are still faced with a choice between taking Salmon's 'real causal processes' to be a subset of the real patterns, implying that only microphysics identifies metaphysically 'genuine' causes, or, following Kitcher in regarding any non-redundant explanation that picks out a real pattern as identifying a genuine cause, thereby analysing all scientific causation in terms of real patterns.

The choice is if anything more tricky for the metaphysician working in a Quinean spirit, and considering how to identify the 'glue'. Appropriating a familiar word that is also used in science, such as 'causation', could seem like an exercise in telling scientists what to think. Invoking a folk concept, on the other hand, would involve making folk intuitions epistemically relevant to metaphysics, which we've argued is a bad idea and is also un-Quinean. Our preference happens is to call the glue 'information',[31] and allow scientific practice to trump both folk and philosophers' usages with respect to what gets called 'causation'. There are at least two complementary ways of defending this preference. In other, ongoing, work with others, we motivate it by a survey of high-level theories across several sciences and conclude on the basis of this that an information-theoretic metaphysic can directly unify an impressive range of otherwise ontologically disconnected explanatory frameworks. Here, we will take another tack, one more traditional amongst analytic philosophers. That is, we explicate Russell's own response to the problem raised and diagnosed in his essay, and see where it leads us with respect to the semantic impasse we have just identified.

---

[30] Developing his charge of anthropocentrism, Redhead continues: 'There is no such thing as a non-natural motion. To most physicists the old-fashioned idea of cause arises from the idea of our interfering in the natural course of events, pushing and pulling objects to make them move and so on. In modern physics there are just regularities of one sort or another' (1990: 147).

[31] Recall that, consistently with the point being made about 'causation', we use 'information' in a way that gives scientific semantics precedence over folk semantics.

4. Letting science hold trumps

At the end of the previous section, we found that methodological percepts we called `Quinean' left us conflicted over what to do about the failure of fit amongst folk, scientific, and metaphysical concepts of `causation.' One option, of course, would simply be to declare and define analytic distinctions to fix three notions of `causation$_1$,' `causation$_2$' and `causation$_3$'. This is the kind of `solution,' however, that invites satire of analytic philosophy. As a more substantial problem, this approach achieves such resolution as it can through simply bypassing questions about the *relationships* among the concepts. If we have succeeded in persuading the reader that there is not just one monolithic concept of causation, then going on to inquire into interrelationship amongst the resulting family of concepts seems the truly interesting philosophical topic to pursue.

We noted that the problem is made interestingly *hard* under the discipline of a `Quinean' attitude to philosophical method. This attitude is, of course, the view that it is science, not a priori philosophical reflection, that should be entrusted with the primary role in populating our ontology, and that it is reflection upon the history of science that should guide epistemology. The stance gets named after Quine because he gave it vivid and repeated stress. However, both Carnap and Reichenbach defended the same perspective earlier – and ahead of them was Russell, in the very essay that inspires the present one. "Dr. James Ward," Russell complains, " … makes [it] a ground of complaint against physics [that] the business of those who wish to ascertain the ultimate truth about the world … should be the discovery of causes, yet physics never even seeks them. To me, it seems that philosophy ought not to assume such legislative functions …" (1918: 173). On the basis of this meta-principle, says Russell, the conclusion we should derive from the absence of reference to causes in mature physical theory is that "there are no such things" (*Ibid*).

On questions about conflicting grounds for semantic usage as between science, metaphysics and folk practice, then, Russell's preference is clearly to let science hold trumps. As we have seen, however, it greatly overstates the facts to say that scientific papers don't advance many of their explanations in terms of causation, even if the concept indeed drops out of the formal statements of physical theories. Given this fact, it seems that if we can find no over-arching concept of causation supplied by any of the historically plausible sources – physics, philosophy, or everyday expression – then we must suppose that each special science invokes its own parochial causal notion, so that we will have to provide analysis not just of `causation$_1$,' `causation$_2$,' and `causation$_3$,' as above, but of `causation$_1$' ,…, `causation$_n$'! Undertaking *that* task would *surely* justify satire.

If brute conceptual splitting is, then, a route to analytical bankruptcy, it seems that our only hope lies in trying to find a suitably abstract conceptual foothold from which we might try to state something that various concepts of causation have in common. We have already given reasons for denying that folk intuitions could provide such a foothold; though from whatever foothold we find, if we can find one, we ought to be able to say something about the relationship between those intuitions and the other understandings of causation. Our Quinean commitments preclude a search for the foothold from a priori principles. And we have argued that if we simply induct directly over scientific practice, we get the very disconnected medley of concepts we hope to organize. We seem to be stuck.

Some recent philosophers, such as Dupré (1993), would welcome this impasse as evidence for fundamental disunity of science. We suggest, however, that Russell, despite his general empiricism that makes the very idea of metaphysical organization suspicious, can help us see how to relax our apparent bind. As a reductionist about the relations among the sciences, Russell can let physics trump quite directly. As we have already noted, on this basis he inclines to the view that our cleanest semantic reform should be to dispense with causation as a serious scientific idea. However, Russell's empiricism is not relentlessly consistent. Thus we also find, running throughout his essay, a less metaphysically cautious route to the same conclusion. He argues that because the functional generalizations of physical theory involve no temporal asymmetries, whereas causation as understood in the anthropomorphic sense *depends* on such asymmetry, the universe itself cannot have causal relations as a feature, except in the sense that the universe includes cognitive systems and cognitive

systems achieve anticipation by constructing causal relations.. The premise about physics here is rooted in classical theory, but this is unimportant in the present context. The important point is that Russell's argument vacillates between a purely epistemological one to the effect that causal concepts are unjustified in science, and one that draws an ontological conclusion about the non-existence of causes from the assumption that a certain sort of physical theory is descriptively true.

Now, Russell's lifelong inconsistency as between empiricism and realism, reminiscent of Locke's, is famous. Our Quinean stance allows us to be realists about objects of scientific reference if, on independent grounds, we wish to be. Thus if we were also prepared to follow Russell in his assumption that (unless dualism is true) physics directly fixes the scientific ontology, and if we thought that contemporary physical theory cites no causal relations, then we could embrace Russell's eliminativist conclusion straightforwardly. This, however, would finesse our original problem quite illegitimately. That problem is generated by the prolixity of parochial causal relations in the special sciences. Letting physics determine ontology directly simply denies the ontological significance of this fact. However, we don't think the actual practice of science justifies such reductionism. All it does justify is a much weaker principle to the effect that physics provides epistemic constraints on other sciences. In particular, it is a consistent methodological principle across the whole history of modern science that no special science may offer explanations at time *t* that contradict the generalizations of whatever is roughly or nearly consensual amongst physicists at *t*; and there is no symmetrical constraint in the other direction. Elsewhere (Ross 2000, Ross and Spurrett forthcoming), we have discussed this as `the primacy of physics constraint.' If physical theory abjured reference to causal relations but biology or economics invoked them, this would generate no contradiction, so the constraint would not warrant inference from the hypothetical fact to Russell's conclusion.

We will now argue, however, that the constraint *does* pack enough punch to allow us to take Russell's reasoning as a template, requiring us only to complicate it a bit. Better still, the complications in question will yield a conclusion much less radical in its revisionary implications for the special sciences than Russell's own. Put in terms we introduced earlier, the constraint will provide the foothold for conceptual organization we were looking for. The constraint is strong enough to tell us to look to physics for *general* ontological structures – so we don't have to violate our self-imposed Quinean discipline by resorting to a priori analysis. If there is universal glue, and if the nature of that glue can help us see what the various concepts of causation have in common, then perhaps we can find it in physics after all.

Is there something that can explain the force of the primacy of physics constraint other than the sort of mereological reductionism presupposed by Russell? We have argued elsewhere (Ross 2000, Ross and Spurrett forthcoming) that the answer is `yes,' and that the something in question is just a pragmatic fact about the organization of inquiry, rather than an a priori metaphysical hunch. We have many disciplines that study temporally and/or spatially bounded parts of spacetime. They are, by fiat, not responsible for finding generalizations that support counterfactuals outside of those boundaries. Then we have one discipline, fundamental physical theory,[32] that is responsible for supporting counterfactuals across the entire physical universe, whether this is taken to be infinite, as in classical physics, or finite and limned by space-like and time-like singularities, as in relativistic physics.[33]

---

[32] This is *not* coextensive with all of physics. Some branches of physics can be special sciences in the sense just given. We are here *defining* `fundamental' by reference to what we take to be a real working convention in science.

[33] The infinite classical universe raises special epistemological (inductive) problems for the principle of full counterfactual support that do not arise in the finite relativistic one. In the second case, we can invoke a sensible degree of verificationism and absolve science of responsibility for information lying across singularities that can't, in principle, be retrieved. We can thereby avoid some issues that tie Russell in deep knots towards the end of his essay. Einstein's universe is a far less perplexing place in which to be a philosopher than Newton's.

Saying that a measurement from *anywhere* is a potentially relevant counterfactual to fundamental physics *is* just a fancy philosopher's way of stating the asymmetric primacy constraint.

Let us now take a few argumentative steps which mirror Russell's reasoning exactly, and none of which depend on the mereological reductionism we won't allow ourselves. We of course never actually have the full set of generalizations, good across the whole universe, at which we aim in the limit, because some parts of that universe are, at any given time, inaccessible to our measurement. This includes all points outside of our collective lightcone, particularly, at any given time *t*, all those points later than *t* along the time-like dimension. But it is rational to project the *logical* character of the generalizations in the limit from the sample we have at *t*. Now suppose, hypothetically, that none of the generalizations in the sample invoke the kinds of causal asymmetries characteristic of the folk concept of causation. Suppose that some or all of the parochial special-science causal principles *do* presuppose such asymmetries. Then: *because* special sciences are scope-restricted and metaphysics, by definition, is not*;* and *because* our Quinean rule prohibits our seeking metaphysical premises from anywhere more general in scope than fundamental physics, it *follows that* whatever concept we use to organize the special-science causal concepts, if there is any such concept to be found, cannot itself be a causal concept.

Is there such a non-causal organizing concept for causal concepts to be found *in fact*? We are going to propose `information-transmission' as such a concept. Before we do so, however, we must point out some subtleties without which our argument here will be misunderstood. First, while we think it possible to argue that information is the best current candidate for the role of the famous `glue' (Ross *et al* forthcoming), we will not be pressing this direct metaphysical conclusion now. Here we are staying close to Russell's project and restricting ourselves to analysis. Second, in the full metaphysical voice we are keeping quiet here, we in fact think, following Collier (1999), that information-transfer actually *does* reduce all causal concepts as a matter of ontology. This is because, in that voice, we deny Russell's premise that the current sample of physical generalizations – our *t* being of course importantly different from Russell's – respects his principle of functional symmetry. But our subject here is just causal concepts as these feature in the relations between the three ways of worrying, and this subject is more abstract than any specific metaphysical thesis that must, in Quinean spirit, rely on some contingent facts about the actual physical universe as we find it.

From this abstract perspective, Russell's hypothetical universe shares at least one key feature with any actual one, whether or not our special theory of the latter is right. In the limit from which the whole general structure of that universe would be surveyed, it must be a Paremenidean `block universe' (Capek 1976). That is, all general relations amongst all spacetime points must there appear on one static topology. It follows from this that if causal concepts depend on real processes, then causal concepts cannot be required to specify the topology. Put metaphorically, if `causation' is for identifying `goings on,' then a conception of the ideal of physical theory like ours and Russell's can allow for no non-anthropocentric goings-on in the limit. Russell stresses this point at length, leading him to conclude his essay with reflections on its implications for the free will problem. These reflections of his are conditioned on special inductive difficulties raised by the idea of a (classical) *infinite* block universe, and so need not detain us here. What interests us are the implications of a limit-conception of a block universe, as the target subject of physical generalizations, for the nature of the concept that might be used to in turn organize causal concepts.

Once notions of `process' are deleted from the concept of causation, as in a block universe they must be, what remains? We can see, looking from a godlike perspective at the overall structure, that from any more limited perspective *within* that structure only restricted information is accessible. A billiard ball, let us imagine, is at spacetime coordinates ($s_{x, \ldots, z}$, t_1) and also at ($s_{x', \ldots, z'}$, t_2). The field equations that determine the topology of the universe will have as implications that at t_1 the ball could not have received information from any points $t_{x>1}$, or any information from spacelike separated points *s*

sufficient to distinguish $s_{x', \ldots, z'}$ from $s_{x'', \ldots, z''}$.[34] No observer may explain the ball's trajectory by reference to informational relations amongst its locations that can't physically obtain. Perspectivally restricted observers with cognitive capacities, who can process more information than the ball can, may of course be able to predict its location at t_2 from t_1, though they can't simulate the godlike observer in the limit and locate its entire trajectory if this is sensitive to complex microdynamics that they can't compute, or that are in fact irreducibly statistical. If the object is a cat, or a grain of pollen in the sea, or an electron, instead of a ball, then the observers at t_1 will have correspondingly less information about its location at t_2 than the godlike observer, depending on the information-gathering sophistication of their physical theory at t_1. How much information about its location at t_2 the electron itself can have at t_1 is a contingent consequence of the structure of the universe, but if Copenhagen quantum theory is both complete and true then our current physics already supplies the answer to this question. The physical theory in the limit answers the question by definition.

The point of the foregoing is to illustrate our contention that what remains of causation as a physical concept once it is bleached of process notions in a block universe is a *physical* concept of `information.' We say `a' physical concept rather than `the' physical concept because there is more than one exact candidate. All candidates are conceptually descended from that of Shannon and Weaver (1949), but developments in computation theory since that time have introduced various possible refinements, and important philosophical issues are raised by efforts to choose among these for application in metaphysics. Elsewhere, we ourselves use the refinement promoted by Collier (1999). However, at the abstract, purely analytical, level of the present discussion we need not go into these details. Our claim here is just that in the block universe of physics-in-the-limit, information is the appropriate organizing concept for the various special-science and folk concepts of causation. We will now go on to say, in general, how this organizing can work. Given the rules we have set for ourselves, our proposal accomplishes something if and only if the organizing does not imply reductionism of special-science kinds and relations to physical ones.

We here make use of earlier work by Ross (2000), building on Dennett (1991). Pressed to defend the possible reality of intentionally characterized states in a way that implies none of reductionism, eliminativism or anthropocentric instrumentalism, Dennett introduced the idea of a `real pattern' in terms of informational compressibility. Ross (2000) works up a more precise formulation of Dennett's idea as a definite theory of existence as follows:

> To be is to be a real pattern, and a pattern is real iff
>
> (i) it is projectible under at least one physically possible perspective
>
> and
>
> (ii) it encodes information about at least one structure of events or entities *S* where that encoding is more efficient, in information-theoretic terms, than the bit-map encoding of *S*, and where for at least one of the physically possible perspectives under which the pattern is projectible, there exists an aspect of *S* which cannot be tracked unless the encoding is recovered from the perspective in question.

This theory's reliance on physically *possible* perspectives makes it realistic rather than antropocentrically instrumentalistic, since it abstracts away from actual, contingent observers and appeals ultimately to the sort of physical theory in the limit discussed above. Its reliance on *physically* possible observers makes it non-idealistic, and captures the primacy-of-physics constraint in a way that respects our Quinean refusal to acknowledge the possibility of any justifiable conceptual

---

[34] These are of course contingent facts about the topology, not a priori principles. In our examples, we just assume absence of paths for backwards time-travel by macroscopic objects, and, when we discuss a case involving electrons below, absence of quantum tunnels. If the actual universe does include these things, they will be reflected on the topology of its block representation.

framework that has wider empirical scope than that of fundamental physics. At the same time, however, it avoids the implication of reductionism. Many real patterns will be detectable only within small bounded regions of spacetime and so will be invisible to generalizations captured in terms of the kinds and relations of physical theory. Echoing Russell (1918: 189) in our terms, perhaps one must restrict one's attention to the surface of the earth after the dawn of life to encode and project macroeconomic patterns, and to a much smaller span of centuries and locations along that surface to encode and project the pattern constituted by the person of Napoleon.[35] Since we *define* fundamental physics as the science that takes patterns detectable from anywhere in the universe as its subject matter, it follows that physics will not be constructed so as to be able to pick out macroeconomic or Napoleonic patterns. Furthermore, as detection of these patterns may require cross-classification of physical or other patterns, *no* element of reductionism is implied by the theory at all. Meyering (2000, see also Menzies 1988) defends anti-reductionism of this sort in terms of `multiple' and bi-directional supervenience, and we are pleased to note that Russell (1918: 191-192) anticipates this point very closely in his discussion of possible causal interpretations of the relationship between mind and matter in the block universe. Ross and Spurrett (forthcoming) dissolve Jaegwon Kim's (1998) `causal exclusion' argument for reductionism about mental states in just this way.[36]

Now, to the organization of causal concepts.[37] Our analysis explains why parochial causal concepts are more robust in special sciences than in physics, and also why they are ineliminably *parochial*. First, special sciences, by definition, are not cast from the vantage point of the limit. By this we mean not only that they aren't cast from the *physical* limit; they also aren't cast from the limits achievable at their *own* scope boundaries. Thus, although their causal claims can be translated into information-theoretic relations on a block representation as analytic exercises, this has no empirical point: special-science accounts are accounts of what are, for their retailers, ongoing processes. Descriptions of information-transferring processes in terms of process and `flow' just *are* full-blown causal accounts (Barwise and Seligman 1997; Collier 1999). Of course, it is true that *actual* physical theories, unlike the ideal physical theory, are not cast from the limit either, which is why causal concepts play an ineliminable role in *doing* physics, as opposed to writing it up for the ages. We eliminate causal concepts from very formal *statements* of physical generalizations, where and if we do so, for the sake of making their logical commitments as regards counterfactual support more transparent. Second, it is *because* special sciences are concerned with relatively isolable regions of spacetime, which involve cross-classification from the perspective of the limit, that the causal relationships on which they variously focus will appeal to different aspects of their information-carrying potential. This emerges as Kitcher's point that explanatory argument patterns distinguish special sciences, and are reflected in parochial restrictions on what kinds of relations count as causal for each of them.

Finally, we can make sense of the relationship of folk causal concepts to scientific ones. Ross's theory of existence makes it a necessary condition on a pattern's reality that it not be informationally redundant in the limit. When actually formulating scientific theories short of the limit, we cannot *know* that our posited patterns would satisfy this criterion, but we *care* whether they do, and so we make our existence claims at least implicitly provisional. Furthermore, as science progresses we adjust our ontology in accordance with our concern for ontological parsimony. However, practical folk have no systematic reason to be interested in this constraint. Nor could natural selection attend to it when it designed the native anticipatory apparatus used by the practical folk. In coping with problems of scarcity, process is almost everything. Attention to wider informational dynamics in which processes are embedded typically deliver few if any additional payoffs, or even gets in the way of payoff

---

35 That is, no temporal coordinates prior to those coincident with Napoleon's childhood are relevant, and at some point in the future his existence *as* the distinctive pattern he was (`was' here being indexed by us now, rather than by the examiner of the block universe) will cease to make any informational difference to anything. At some point, that is, he is `forgotten' not just by people but by the universe!

36 We regret that when we wrote that paper we had not yet noticed Russell's prescience on the issue.

37 [CONFERENCE NOTE: From here on there are various points of contact (and conflict) with Hue Price's projectivist pragmatism, and opportunities for direct reference to his position.]

maximization because of computational costs.[38] Therefore modeling causal relations as sequences of collisions of objects in time is a sensible heuristic. Furthermore, for social animals like people, or even mobile asocial animals that prey on and/or are preyed on by other mobile animals, it might even make sense to model all causal relations on volition, with rocks then being treated as limiting cases of agents with minimal sophistication of purpose and adaptability. Such models of causation of course won't generalize when one tries to do science, that is, when one's scope of concern widens beyond what nature could build us to worry about. New worries, we conclude a bit tritely, motivate new concepts.

References (incomplete):

Baillargeon, R., Kotovsky, L. & Needham, A. (1995) 'The acquisition of physical knowledge in infancy' in Dan Sperber, David Premack and Ann James Premack (eds.) *Causal Cognition*, pp. 79-116. Oxford: Clarendon.

Baldwin, J. M. ed. (1901) *Dictionary of Philosophy and Psychology*, London: Macmillan.

Barwise, J., and Seligman, J. (1997). *Information Flow: The Logic of Distributed Systems.* Cambridge: Cambridge University Press.

Capek, M. (1976) 'The Inclusion of Becoming in the Physical World', in M. Capek (ed.) The *Concepts of Space and Time*. Boston Studies in the Philosophy of Science. Vol. XXII. Dordrecht, Reidel, pp.501-524.

Cartwright, N. (1980) 'Do the Laws of Physics State the Facts?', *Pacific Philosophical Quarterly* 61, pp. 75-84.

Cartwright, N. (1983) *How the Laws of Physics Lie*. Oxford: Oxford University Press.

Cartwright, N. (1999) *The dappled world.* Cambridge University Press.

Collier, J. (1999) 'Causation is the Transfer of Information'. In H. Sankey (ed.) *Causation, Natural Laws and Explanation*. Dordrecht: Kluwer, 279-331.

Creary, L. G. (1981) 'Causal Explanation and the Reality of Natural Component Forces', *Pacific Philosophical Quarterly,* 62, 148-157.

Dennett, D.C. (1991) 'Real Patterns', *Journal of Philosophy*, 88(1), 27-51.

Descartes, R. (1985) *The Philosophical Writings of Descartes*, 2 vols., (translated and edited by J. G. Cottingham, R. Stoothoff, and D. Murdoch), Cambridge: Cambridge University Press.

---

[38] This statement needs some qualification. Some folk engage in practical activities of a kind that evolution could not anticipate, like buying and selling assets in futures markets. When you do that sort of thing, it's wise to book-keep your structures and relations like a real scientist.

Dowe, P. (1992) 'Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory', *Philosophy of Science*, 59(2): 195-216.

Dowe, P. (1997) 'Causal Processes', *The Stanford Encyclopedia of Philosophy (Fall 1997 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall1997/entries/causation-process/>.

Dupré, J. (1993). *The Disorder of Things.* Cambridge, MA: Harvard University Press.

Field, H. (forthcoming) 'Causation in a Physical World', to appear in M. Loux and D. Zimmerman, eds., *Oxford Handbook of Metaphysics*.

French, S. & Ladyman, J. (2003a) 'Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure' *Synthese*, 136 (1), 31-56.

French, S. & Ladyman, J. (2003b) 'The Dissolution of Objects: Between Platonism and Phenomenalism' *Synthese*, 136 (1), 73-77.

French, S. (1998), 'On the Withering Away of Physical Objects', in E. Castellani (ed.), *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*, Princeton University Press, pp. 93-113.

Garcia, J. and Koelling, R. A. (1966) 'Relation of Cue to Consequence in Avoidance Learning', *Psychosomatic Science*, 4, 123-4. Reprinted in Seligman and Hager (1972) pp 10-14.

Gelman, R., Durgin, F. & Kaufman, L. (1995) 'Distinguishing between animates and inanimates: not by motion alone' in Dan Sperber, David Premack and Ann James Premack (eds.) *Causal Cognition*, pp. 150-184. Oxford: Clarendon.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., and Danks, D. (in press). 'A theory of causal learning in children: Causal maps and Bayes nets', *Psychological Review.*

Hershber, U. and Efroni, S. (2001) 'The Immune System and Other Cognitive Systems', *Complexity*, 6 (5), 14-21.

Hitchcock, C. (2003) 'Of Humean Bondage', *British Journal for the Philosophy of Science*, 54, 1-25.

Hogan, C. J. (2002) 'The Beginning of Time', *Science* 2002 March 22; 295: 2223-2225.

Hume, D. (1975) *Enquiries concerning Human Understanding and concerning the Principles of Morals*, ed. by L.A.Selby-Bigge. Oxford: Oxford University Press.

Hume, D. (1978) *A Treatise of Human Nature*, ed. by L.A.Selby-Bigge. Oxford: Oxford University Press.

Hutchins, E. (1995) *Cognition in the Wild*, Cambridge, Mass.: MIT Press.

Kant, I. (1933) *Critique of Pure Reason*. Trans. N. Kemp Smith. London: MacMillan.

Kitcher, P. (1976) 'Explanation, Conjunction and Unification', *Journal of Philosophy* 73, 207-212.

Kitcher, P. (1981) 'Explanatory Unification', *Philosophy of Science*, 48: 507-531.

Kitcher, P. (1982) 'Two Approaches to Explanation', *Journal of Philosophy*, 82 (11), 632-639.

Kitcher, P. (1989) 'Explanatory Unification and the Causal Structure of the World', In P. Kitcher and W. Salmon, eds., *Scientific Explanation*. Minneapolis: University of Minnesota Press, 410-505.

Ladyman, J. (2000), 'What's Really Wrong with Constructive Empiricism?: van Fraassen and the Metaphysics of Modality', *The British Journal for the Philosophy of Science*, 51, pp. 837-856.

Ladyman, J. (1998) What is Structural Realism? *Studies in History and Philosophy of Science* 29: 409-424.

Locke, J. (1975) *An Essay Concerning Human Understanding*, (Edited and with a foreword by P.H. Nidditch), Oxford: Clarendon.

Menzies, P. (1988) 'Against causal reductionism', *Mind* 98: 551-574.

Mill, J. S. (1872) *A System of Logic: Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation* (8th edition). London: Longman, Green, and Co., 1916.

Mill, J. S. (1972) *A System of Logic Ratiocinative and Deductive, Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*, (2 vols.) J. M. ROBSON (Ed.), Volumes VII and VIII of J. M. ROBSON (General editor) *The Collected Works of John Stuart Mill* (Toronto: University of Toronto Press).

Morris, M. W., Nisbett, R. E. & Peng, K. (1995) 'Causal attribution across domains and cultures' in Dan Sperber, David Premack and Ann James Premack (eds.) *Causal Cognition*, pp. 577-612. Oxford: Clarendon.

Norton, J. D. (2003) 'Causation as Folk Science' http://philsci-archive.pitt.edu/archive/00001214/

Palmerino, C. C., Rusiniak, K. W., and Garcia, J. (1980) 'Flavor-illness aversions: The peculiar roles of odor and taste in memory for poison', *Science* 208: 753-755.

Redhead, M. (1990) 'Explanation'. In D. Knowles, ed., *Explanation and Its Limits*. Cambridge: Cambridge University Press, 135-154.

Reichenbach, H. (1957) *The Philosophy of Space and Time*. New York: Dover.

Ross, D. & Spurrett, D. (forthcoming) 'What to say to a skeptical metaphysician: A defense manual for cognitive and behavioral scientists', *Behavioral and Brain Sciences*.

Ross, D. (2000) 'Rainforest Realism: A Dennettian Theory of Existence,' in D. Ross, A. Brook and D. Thompson (eds.) *Dennett's Philosophy: A Comprehensive Assessment*, pp. 147-168. Cambridge, Mass.: MIT Press.

Russell, B. (1948). *Human Knowledge.* New York: Simon and Schuster.

Russell, B. (1913) 'On the Notion of Cause' *Proceedings of the Aristotelean Society*, reprinted in (1918) *Mysticism and Logic* (London: Allen and Unwin), 180-208.

Salmon, W. (1984) *Scientific Explanation and the Causal Structure of the World.* Princeton: Princeton University Press.

Salmon, W. (1994) 'Causality Without Counterfactuals', *Philosophy of Science*, 61(2): 297-312.

Salmon, W. (1997) 'Causality and Explanation: A Reply to Two Critiques', *Philosophy of Science*, 64(3): 461-77.

Schaffer, J. (2003) 'The Metaphysics of Causation', *The Stanford Encyclopedia of Philosophy (Spring 2003 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2003/entries/causation-metaphysics/>.

Seligman, M.E.P. & Hager, J.L. (1972) *Biological Boundaries of Learning*. New York: Appleton Century Crofts.

Shannon, C., and Weaver, W. (1949). *Mathematical Theory of Communication.* Urbana: University of Illinois Press.

Spelke, E. S., Phillips, A. & Woodward, A. L. (1995) 'Infants' knowledge of object motion and human action' in Dan Sperber, David Premack and Ann James Premack (eds.) *Causal Cognition*, pp. 44-78. Oxford: Clarendon.

Spurrett, D. (2001) 'Cartwright on laws and composition', *International Studies in the Philosophy of Science* 15(3): 253-68.

Worrall, J. (1989), 'Structural Realism: The Best of Both Worlds?', *Dialectica*, 43, pp. 99-124