The British Journal for the Philosophy of Science Advance Access published February 13, 2007

Notions of Cause: Russell's thesis revisited*

Don Ross and David Spurrett

ABSTRACT

We discuss Russell's 1913 essay arguing for the irrelevance of the idea of causation to science and its elimination from metaphysics as a precursor to contemporary philosophical naturalism. We show how Russell's application raises issues now receiving much attention in debates about the adequacy of such naturalism, in particular, problems related to the relationship between folk and scientific conceptual influences on metaphysics, and to the unification of a scientifically inspired worldview. In showing how to recover an approximation to Russell's conclusion while explaining scientists' continuing appeal to causal ideas (without violating naturalism by philosophically correcting scientists) we illustrate a general naturalist strategy for handling problems around the unification of sciences that assume different levels of naïveté with respect to folk conceptual frameworks. We do this despite rejecting one of the premises of Russell's argument, a version of reductionism that was scientifically plausible in 1913 but is not so now.

1 Russell's Naturalistic Rejection of Causation

Russell ([1913])[1] characterized what he called 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 reasons include the descriptive claim that practitioners of 'advanced' sciences, particularly physicists, do not seek causes, and the normative claim that it is improper for philosophers to

* We thank John Collier, Harold Kincaid, James Ladyman, Adriano Palma, David Papineau, and an anonymous referee for this journal for comments and criticism of earlier versions of this article.

[1] All references to 'Russell' in this article are to 'the Russell of 1913' unless otherwise stated. Citations are to 'On the Notion of Cause' as reprinted in *Mysticism and Logic* (1917).

doi:10.1093/bjps/axl027

legislate on whether they should. His claim about advanced sciences construes cause-seeking as the quest for, or assumption of the existence of, 'invariable uniformities of sequence' ([1913], p. 178) or laws of constant succession. According to Russell, belief in the 'law of causality' leads us erroneously to expect and to prize these uniformities.

Our aim is to re-evaluate Russell's thesis in the light of subsequent developments in science and philosophy. Our motivation is not Russell scholarship but interest in philosophical naturalism. By 'naturalism' we refer to any thesis according to which metaphysical hypotheses should be derived, as far as possible, from scientific discoveries and arguments, rather than extra-scientific sources like ontological intuitions, 'common sense' or 'first philosophy'.[2] Naturalism is, then, an epistemological doctrine that will lead its proponents to favour some ontological views over others. The adequacy of such naturalism has lately become the focus of lively philosophical debate. (See McDowell [1994], Almeder [1998], Rea [2001] for the tip of the iceberg). Though we promote naturalism, this article is not a polemic in its favour. Rather, it is an exploration, focussed by way of the topic of causation, of a perennial naturalist issue: the conflict between what Sellars famously referred to as the manifest and scientific images. The naturalist is acutely aware that there is a difference between the world as it appears to so-called 'common sense observation' and the modally structured, perspective-independent world that science aims in the limit to describe, and holds that insofar as there can be any legitimate role for metaphysics in the description of reality, it must develop its concepts and arguments by reference to the scientific account. This interest led us to Russell, whose 1913 essay on causation is a particularly illuminating (partly because of its occurrence in a different philosophical era) expression of this contemporary naturalist problematic.

Let us first indicate what makes Russell's 1913 argument representative of the sort of naturalism we have demarcated. His contention that causation is not a significant concept in metaphysics is partly based on his claim that it is an artefact of an anthropocentric perspective that science supersedes. He also blames philosophers for supposing, against the evidence in science, that there exists a 'law of causality' identifying causal relations with 'invariable uniformities of sequence'. Here Russell anticipates Cartwright's (e.g., [1980]) contention that laws understood as generalizing or quantifying

[2] An anonymous referee wondered whether we fully intend the force of the idea that metaphysical hypotheses should be *derived from* science as opposed to something weaker such as 'informed by' science or the yet weaker 'compatible with' science. 'I don't see,' writes the referee, 'why a naturalistic metaphysician cannot think that philosophical modes of thinking do provide distinctive evidence of their own that bear on metaphysical hypotheses.' Ladyman and Ross ([forthcoming]) give arguments for doubting that these weaker forms of naturalism are viable. Therefore, it is the strong form that interests us here, and we intend the formulation seriously.

event regularities have numerous counterexamples. Russell agrees with Cartwright that such laws do not describe how bodies actually behave. Instead of concluding from this that the laws of nature are false, he maintains that they are not about event regularities per se, defending this through an account of the features of scientific practice that both generate the counterexamples and explain the inappropriateness of an event-regularity view of laws. First, he argues, the Humean interpretation of laws as regularities depends on a qualitative basis for individuating events that cannot be rendered in objectively quantitative terms; yet a mature science is one that aims at strictly quantitative description in the limit ([1913], p. 183f). Second, the idea that uniformities of *sequence* might be 'laws' makes asymmetry between the past and the future fundamental, whereas maturing sciences aim always at laws that avoid reference to time in this sense (though they often refer to intervals, some of which may be rightly called 'temporal') ([1913], p. 185f). Russell concedes that in the 'infancy' of a science ([1913], p. 178) the principle of 'same cause, same effect' may prove useful for discovery, but maintains that advance in science consists in moving away from such simple formulae, and towards understanding phenomena in terms of mathematically specified relations of functional interdependence in which 'past' and 'future' are not systematically different,[3] which the 'law of causality' is too crude to embrace, and which Russell therefore thinks should not be called 'causal'.

In calling Russell's argumentation here 'naturalistic', we do not mean to imply that Russell consistently observed naturalism as his first allegiance. He was first and foremost a Platonist. Though Platonism has sometimes been included among the foils of naturalism, there are in fact versions of Platonism compatible with naturalism (Brown [1999]; Ladyman and Ross [forthcoming], Chapter 4); and Russell's Platonism was motivated by facts about mathematics and its relationship to science, not by *a priori* logic. In any case, Russell's main basis for argument in the essay on causes is description of the general content of science. Such appeals have no force without an accompanying normative thesis to the effect that science should have authority over philosophy. Thus Russell cites James Ward as complaining about physics on the grounds 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 which Russell objects by saying that 'philosophy ought not to assume such legislative functions' ([1913], p. 171). Later, having found Bergson attributing an event regularity principle of causation to scientists, Russell scolds that 'philosophers . . . are too apt to take their views on science from each other, not from science' ([1913], p. 176). Towards the end of his essay, Russell diagnoses commitment to a law of causation as stemming from projections into the metaphysical

[3] That is, as Field ([2005]) notes, being '*deterministic in both directions*'.

interpretation of science of 'anthropomorphic superstitions' based on the practical human predicament with respect to the asymmetry of past and future in our memories and capacities for control. Finally, at one point ([1913], p. 184) he suggests a very radical naturalist response to the conflict between everyday and scientific conceptions: the idea that sound metaphysical insights cannot be stated in natural language, but require expression in a formalism from which anthropocentric distortions have been purged.

Histories of naturalism in philosophy typically describe Hume as the original radical naturalist, and then, perhaps after a nod at Peirce, leap straight to Quine and the subsequent rise of naturalized epistemology inspired by cognitive science. This pattern provides part of the interest in finding such forthrightly naturalistic argument coming from Russell in 1913. Further, Russell's argument, like that for the version of naturalism most often debated by philosophers of science today, derives its conclusion mainly from putative facts about what scientific theory implies there to be, rather than from taking human psychological dispositions to construe experience in particular ways as evidence that the construals in question have some content that *prima facie* merits preserving in a more sophisticated account. (Certain kinds of reflections on typical human psychology *are* relevant to any version of naturalism, including Russell's. We return to this in the next section.) We think we can learn important things about naturalism in the philosophy of science by examining its antecedent expression in the quite different context of 1913.

Let us say something about the context in question. Condensing breathlessly, what mainly preoccupied Western philosophy of science between Hume and Quine were two bodies of ideas: Kantianism and positivism. Indeed, as Friedman ([1999a]) has persuasively argued, positivism was initially a variant of Kantianism. Quinean naturalism, anticipated and directly inspired by the later Carnap (Creath [1991]), emerged from the gradual morphing of Kantian positivism into Humean empiricism during the 1930s and 40s. Kantians, including Kantian positivists, privilege science over philosophy in *one* sense: they agree that science authoritatively informs us about empirical matters. However, they reject naturalism as we understand it, in supposing that philosophy retains a distinctive, extra-empirical, task of elucidating the forms of understanding under which the empirical world is 'grasped'. Positivistic Kantians differ from non-positivists in denying a transcendental basis for this elucidation, reducing the forms of understanding to conventions. Thus the non-naturalistic commitment of Kantianism is its belief that philosophers of science have a distinctive task—either transcendental inquiry or logical analysis—that can and should be pursued *independently* from empirical science (Friedman [1999b]). Russell, writing before positivism took hold, represents one last expression of a version of naturalism descending directly from Locke and not passing through Kant. It is an interesting question, one we will not directly

pursue here, whether some contemporary naturalists might be rediscovering this abandoned thread.

We have identified the naturalist as someone who holds that metaphysical hypotheses should be derived from science. Naturalists sharing this commitment may yet differ sharply on what science in fact tells us, and how, if at all, metaphysical conclusions can be derived from it. Some with good claim to be called naturalists (at least with respect to the philosophy of science) such as van Fraassen ([1980]) and Fine ([1986]) oppose metaphysics entirely. van Fraassen's constructive empiricism retains the positivist's suspicion of metaphysics, seeing no role for philosophy as a *unifier* of science and expressing this through denial that we have any sound reason for believing that there is an objective modal structure to the world (van Fraassen [2002]): there are simply observable[4] events, and models by which we simplify and summarize parts of the observational record that are salient to us and relatively stable in our experience. This position suggests the retirement of the philosopher. Within the ranks of those naturalists who think there *is* work for philosophers, we note a significant division between those who think that science suggests a unified account of reality and those who think it does not. Cartwright's (e.g., [1999]) view that the world is 'dappled' is the leading contemporary instance of the latter mood of naturalism. In the present article we do not argue directly against the view that the world is dappled, instead following a line of argument internal to the concerns of 'unificationist' naturalists.

Claims about the way in which the world might be unified necessarily amplify our empirical data. Unificationist naturalists hold that this amplification can be legitimated in the way that ampliative inferences are justified within scientific disciplines, using only data endorsed by science. (Disagreements over the acceptability of these ampliative inferences are, of course, part of what is at issue with the anti-unificationists.) What would make some facts about the world—the unifying ones—more general than those they unify would be that they are projectible from data drawn from domains of greater generality. So, perhaps, only fundamental physical facts, true everywhere in the universe at least since soon after the Big Bang, can unify branches of non-fundamental physics that describe special aspects of our galaxy. In this sense—the sense of their being highly general facts to which less general would-be facts are 'answerable'—the world for the unificationist must be modally structured, even if the existence of the universe itself might be ultimately contingent.[5] This is the position that Ladyman ([1998], [2000]; French and Ladyman

[4] See Ladyman ([2000]) for argument that van Fraassen, as a modal anti-realist, is inconsistent in concerning himself with whether science is explanatorily adequate to observable (as opposed to observed) events.

[5] That is to say: a physicist (for example) might, in working up a theory, treat some relations as holding necessarily *given that* there is any universe, or any universe containing matter, or any

[2003]; Ladyman and Ross [forthcoming], Chapters 2 and 3) has baptised and defended as 'ontic structural realism'. The naturalist metaphysician differs from the special scientist in combining sets of data that individual sciences leave separate. Needless to say, a scientifically justified hypothesis about the way in which the world is unified is fully as fallible as any other generalization inducted from data.[6]

The Russell of 1913 is a unificationist naturalist in this sense. He denies ([1913], p. 185f) that scientists make inferences from the hypothesis of the 'uniformity of nature,' because such a procedure would treat this hypothesis as at least implicitly *a priori*. However, he then goes on directly to argue ([1913], p. 186f) that science advances generalizations about 'relatively isolated systems' only for practical purposes: its guiding aim is to achieve generalizations of maximum scope. Thus 'it should be observed that isolated systems are only important as providing a possibility of *discovering* scientific laws; they have no theoretical importance in the finished structure of a science' ([1913], p. 187). Despite anticipating Cartwright's criticism of the event regularity interpretation of laws, then, Russell is not led from this toward the view that the world is dappled. As sciences widen the scope of the generalizations they produce, he says, they abandon simple regularities among 'causes' and 'effects' which they used as ladders to discovery in their immature phases, and identify progressively 'greater differentiation of antecedent and consequent' and a 'continually wider circle of antecedents recognized as relevant' ([1913], p. 178).

Where more recent work in naturalistic metaphysics has been preoccupied with questions about unification, this has usually been in the specific context of trying to understand the relationship between causation and explanation (Friedman [1974]; Kitcher [1981], [1989]). Whether or not all science reduces to physics (in which case responsibility for providing unity would fall to a single discipline) the naturalistic metaphysician concerned with unity addresses questions about whether, or in what ways, the findings of the various sciences 'hang together'. In this she supposes that the distinctions amongst the sciences are ontologically significant—that is, reflect real distinctions in the structure of the world, rather than just epistemic or institutional conveniences for us. In recent debates, a flagship instance of which is Kim ([1998]; see also Walter and Heckmann ([2003])) a crucial aspect of this problem arises from the multiplicity of causal claims, and *types* of causal claims, made by the various sciences. Where there are multiple (for example mental and neural) causal

universe in which the value of Planck's constant is what it is, or etc., while allowing that the relevant condition on which the necessity in question is premised is an accident.

[6] Naturalists claim only that science pursues a unified account of the world as its aim; they do not assert that we can know, at least in advance of the completion of science, that the world in fact has a fully general unifying structure.

claims about what happens in some part of the world, it seems to many philosophers that some of them must be false, unless we accept reduction or overdetermination. This is known in recent analytic metaphysics as 'the causal exclusion argument' (CEA).

Naturalistic anti-reductionists must, if they deem it valid, regard the conclusion of the CEA as a reductio against the conjunction of its premises. For example, Kincaid ([1997]) takes as his central problem the challenge to unification of the sciences raised by the failure of reductionism, and by the fact that (he argues) different sciences display a proliferating range of explanatory strategies distinguished by reference to special, parochial causal patterns—patterns which indeed provide much of the basis for the locations of the boundaries between the scopes of the sciences in the first place. Because of the priority Kincaid attaches to the empirical evidence against reductionism, he regards the CEA as a reductio against the premise that where there are multiple (for example, both mental and neural) causal claims about what happens in some part of the world, all but one of these must be false unless we accept reduction or overdetermination. In later work, Kincaid ([2004]) makes realism about causes the basis for an argument that special sciences (including social sciences) discover laws of nature, despite the fact that the laws in question must usually be stated with ceteris paribus clauses. Note the point of agreement here: we find a reductionist champion of the manifest image (Kim) agreeing with an anti-reductionist naturalist (Kincaid) that the scientific worldview obtains such unity as can be had by reference to causation.

In the context of this current dialectic, it is interesting that Russell combines the commitments to naturalism and unification shared by us and Kincaid with the latter's reductionistic physicalism, and to the mereological atomism embraced by the CEA's champions. In our view, atomism and naturalism were compatible when Russell wrote his essay on causes but were about to be rendered incompatible by the progress of fundamental physics. 1913 was the very year in which the old quantum theory was consolidated with Bohr's quantization of Rutherford's theory of atomic orbits. The old quantum theory, in turn, made the late nineteenth-century impasses that had inspired Mill's emergentism about chemistry (and, by implication, the life and behavioural sciences) look much less likely to remain pressing problems. Specifically, the key breakthrough of old quantum theory, Einstein's introduction of photons, made it seem likely that wave-particle duality with respect to light was about to be dissolved in a way favourable to atomism. That the duality would instead be generalized to all matter by de Broglie—let alone that quantum entanglement would be discovered—was something Russell could hardly anticipate. Russell kept a close eye on the frontier of science, however. By the time of his *ABC of Relativity* a few years later (Russell [1925]), he knew that de Broglie's work spelled trouble for metaphysical atomism. As Ladyman

and Ross ([forthcoming]) demonstrate, many analytic, metaphysicians writing now still do not seem to know that entanglement kills classical atomism in fundamental physics stone dead, and the relevance to actual science of at least Kim's version of the CEA crucially depends on this ignorance (see also Glymour [1999], Ross and Spurrett [2004]).

These considerations might suggest at first glance that Russell's 1913 essay is unlikely to be of anything but purely historical interest. Furthermore, the fact that the one point of agreement we just identified between the naturalist Kincaid and his (perhaps unwittingly) anti-naturalist opponent Kim, that causation is ontologically fundamental, constitutes *prima facie* grounds for supposing that Russell is simply incorrect about the irrelevance of causes to science and metaphysics. Those who now look to causation to provide the unifying structure of the world have in mind a notion of cause as in some sense the 'cement of the universe' (see Hume [1978], p. 662 for the leading source of this metaphor) or what we elsewhere (Ross and Spurrett [2004]) call the 'glue' holding all objective relations in place.[7] The causation as glue tradition contains a substantial body of work, developing what are collectively called 'causal process theories'. Such glue would be or would provide the necessity that Hume could not find in the impressions and their regular relations. Causal process theories, that is, can be understood as attempts to answer 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. They purport to show that we can at least observe something that is precisely *diagnostic* of what we have traditionally meant by 'causation'. Contra Russell's complaints about the law of causality, causality, or the glue, is not itself an event regularity; it is supposed to be that in virtue of which there are projectible regularities in the first place.[8]

Salmon's ([1984]) causal process theory followed Reichenbach ([1957]) in describing real processes in terms of the transmission of *marks*.[9] According to this view a genuine process is one that can be modified (or 'marked') at some stage, and observed to carry the same modification at subsequent stages. Some processes, that is, transmit information about their antecedent stages, while others do not; those which do are genuine processes, the others pseudo-processes. A crucial scientific motivation for Salmon's project of distinguishing

[7] An anonymous referee observed that cement and glue might be regarded as inapt metaphors to the extent that causation is supposed to be dynamic, while the metaphors suggest stasis. This is grist to our mill, insofar as not all naturalistically respectable lines of thinking about 'causation' are primarily dynamic. See the final two sections of the present article.

[8] To suppose that unification might consist in the identification of some privileged 'glue' is, then, to suppose that there is a master *relation* (in this case causation) that is implicated in all *projectible* generalizations.

[9] Process theories differ in what they take real processes to consist in. For Dowe (e.g., [1992]) they are possessors of conserved quantities, and for Collier ([1999]) they are transfers of information.

causal from pseudo-processes, as with Reichenbach, is the alleged relativistic requirement not to count a faster than light process as causal (Salmon [1984], p. 141; Reichenbach [1957], pp. 147–9). A commonly used illustration of the problem concerns a spot of light moving on some surface because the light source is rotating. For large enough distances between spotlight and surface, the spot cast by the rotating light will move across the surface at a superluminal velocity (Salmon [1984], p. 143). Even so, it would not be a causal process for Salmon, because it cannot be marked. A filter that turned the beam red, placed just before one part of the surface, would make the light that struck that bit of surface red (diagnosing the process between the filter and surface as causal) but not change the spot elsewhere, showing that the successive stages of the spot were not causally connected.

The metaphysic suggested by process views is effectively one in which the universe is a graph of real processes, where the edges are uninterrupted processes, and the vertices the interactions between them. Thus process views, if correct, would make putatively causal claims by scientists subject to a critical test. Those that pick out real processes could be causal, those that do not, cannot. Then, when scientists find causes in their domains of enquiry, whether they put the matter this way or not, they will find ways of picking out real information-carrying relations. That is to say, if macroeconomics, for example, has its own notion of 'macroeconomic causation,' this notion will turn out to correspond to some information-carrying relations really instantiated on the general graph of real processes.

Russell's 1913 claims about the irrelevance of cause to science and the tradition of causal process theories are, then, in conflict. It is bound to seem at first glance that unless Russell is completely wrong, the process theory tradition must be a mistake. There are at least two central reasons why most contemporary philosophers *are* apt to regard Russell as wrong. First, Russell, on the other side to us of decades of reflection and debate over forms of reductionism and the autonomy of the special sciences, takes for granted a now quaint-seeming view (among philosophers of science though not, unfortunately, analytic metaphysicians—see Ladyman and Ross [forthcoming], Chapter 1) of the capacity and right of *physics*, rather than science more generally, to populate our ontology. As we have explained, attention to the proliferation of non-reducing patterns of causal explanation in special sciences largely defines the contemporary naturalistic (unificationist) metaphysician's mission. For this, process theories are one possible piece of useful ordinance. Second, Russell relies on a Humean conception of causation in pronouncing for eliminativism. Many, perhaps most, contemporary philosophers of science would agree with him that scientists do not seek causes in *that* sense, but would then go on to say that that is not the sense in which the idea of causation is either scientifically or

metaphysically interesting. For these reasons taken together, Russell's thesis might be written off as of historical interest only.

Yet Russell's thesis is not dead. Some contemporary philosophers of physics continue to endorse and argue for it. Redhead ([1990]) explicitly acknowledges it when developing his own argument that talk of causes has no place in physics. He considers an explanation of the fall of a body using Galileo's law, and asks what in the explanation might be regarded as the cause. As he notes, closely following Russell's logic (Russell [1913], p. 183f) the position at an earlier time 'can hardly be cited as the cause', that the acceleration is merely 'defined by the kinematic relationship' expressed in the law, that the law itself cannot be regarded as a cause because it is not an event, and finally that citing the force as a cause is 'a very anthropocentric notion' (Redhead [1990], p. 146). Developing his charge of anthropocentrism, Redhead maintains that '[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], p. 147). Batterman ([2002], p. 127) makes similar remarks. Norton ([2003], pp. 3–4) endorses Russell's thesis for Russell's reasons, and says 'mature sciences . . . are adequate to account for their realms without need of supplement by causal notions and principles. The latter belong to earlier efforts to understand our natural world or to simplified reformulations of our mature theories, intended to trade precision for intelligibility'. Glymour ([1999], p. 463) also denies that there are causal laws in physics.

The thought that, somehow, *both* the process tradition and Russell of 1913 could be correct raises perplexing possibilities. Could process theories unify the special sciences but *not* unify them with physics? That would be surprising and confusing, not least because one of the best known and most carefully articulated criticisms of Salmon's causal process theory is, roughly, that the opposite is the case, i.e. that it works better for physics than for the special sciences (Kitcher [1989], see also Ross and Spurrett [2004]). Were the thought that process theories worked for everything except physics instead taken as grounds for doubt about process theories, then perhaps the rebound view might be that Kim's causal exclusion problem and the attractions of Russell's thesis in physics mutually support one another: causal overdetermination is evidence for reductionism—in particular, for doubt that parochial causal patterns in special sciences are metaphysically significant—and then, once we go reductive, causation, absent in the reduction base, drops out altogether. Two overdetermining causes are not *one* cause too many, but *two*.

We agree with Russell—and with most current philosophy of science—that scientists do not generally seek 'invariable uniformities of sequence'. But it does not follow from this that they seek nothing worth calling causes. We argue that scientists do seek causes, in the everyday sense we describe, and that they find

them. This might look like a straightforward confutation of Russell's thesis. We go on to argue that it is not, since Russell's argument is most interestingly read as attacking the idea that there is a 'master' idea of causation, independent of anthropocentric bias, that should be expected to feature in the metaphysical unification of the scientific worldview. This might be a basis for dissolution of Kim's problem, but also seems to imply rejection of the point of process theories. However, we will skirt that conclusion too. Instead, we will defend the following consilience between Russell (and philosophers of physics who agree with him) and Salmon: the naturalist metaphysical project of seeking universal glue is well motivated; process theories are on the right track in this search; but 'causation' is a semantically unfortunate name for the glue to which they lead our attention. Part of our defence of the conclusion that it is semantically unfortunate that the cement of the universe gets called 'causation' involves identifying what we call conflicting centres of semantic gravity for 'causation'.[10] Roughly, there is a folk centre, a philosophical centre, and a centre that has a useful place in special sciences. It is the existence of this third centre that stops us from following Russell all the way to outright eliminativism about causation.

2 Psychology, Folk Notions and Intuitions

To decide whether there are causes, or whether anyone seeks them, one needs to decide (i) what to count as a cause, and (ii) how to tell whether such things are real, or are sought. Russell partly settles (i) by reference to psychology. He also, as noted, regards psychological considerations as irrelevant to philosophy,[11] maintaining that the content of some cause-talk is anthropomorphic projection, failing to correspond to what is discovered by science. Russell does not claim, though, that the 'law of causality' itself arises from psychological considerations—he blames philosophers for the law. Philosophers might reply that in taking 'invariable succession' as causation they were following a long tradition, claiming to find the law of causation through psychological observation. This is particularly true of the history of the account Russell criticizes, namely, the Humean view that causes are event regularities or 'constant conjunctions'. Both Hume and his successors, including Kant ([1933], p. 218) and Mill ([1974], pp. 326–7), derive the philosophical significance of

[10] In previous versions of this article we spoke sometimes of different *concepts* or *notions* of causation. We agree with those, including Crane ([1995]), an anonymous referee, and David Papineau (personal correspondence), who argued that this was sometimes confusing and unhelpful.

[11] Russell ([1913], p. 174) objects to one of the definitions of 'cause' in Baldwin's ([1901]) *Dictionary* that it is psychological, focusing on 'thought or perception' of a process where what is required is a definition of 'the process itself'.

constant conjunctions from purported observation of their importance to both everyday and scientific cognition. Naturalists, trying to shun appeals to intuitions and conceptual analysis, are particularly inclined to look for evidence in scientific psychology.

Some of what we know from the study of 'causal cognition' or 'causal learning' undermines the idea that our causal cognition is Humean. Contrary to early behaviourist dogma, associations are not all equally learnable. Rats find it much easier to learn (and show appropriate avoidance behaviour when cued for) an association between eating a certain food and nausea, and a loud flash and an electric shock, than they do with either complementary pairing (Garcia and Koelling [1966]). We might say that natural selection has projected a regularity in the world of rats, so that any given rat does not experience some types of phenomena as 'entirely loose and separate' (Hume [1975], p. 74). The projected regularities here, referred to by Garcia and Koelling as 'genetically encoded hypothes[es]' are not, furthermore, 'invariable uniformities of sequence'. Learning in the nausea case can take place even when the delay between stimulus and response is considerable (over an hour) and avoidance or aversion can also be learned following a *single* exposure.

The work of developmental psychologists already tells us a lot about the causal expectations of, for example, children around 3 months old, who are surprised when apparently cohesive objects seem spontaneously to fragment, to pass through one another, or not to exhibit a single trajectory through space and time (Spelke *et al*. [1995]), or when apparently unsupported objects fail to fall (Baillargeon *et al*. [1995]). 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 (Gelman *et al*. [1995]). Here too some events are not experienced as 'entirely loose and separate.'

It is also clear that we do not *spontaneously* think in terms of constant conjunctions, but rather in terms of networks of influence similar to directed causal graphs, or Bayes nets. Young children (between the ages of 3 and 5) can, *inter alia*, reason to unobserved causes, plan appropriate novel interventions to prevent a process they had only observed in operation (but never seen prevented), and make inferences about the 'direction' of causal dependence in cases of simultaneous change (see, e.g., Gopnik *et al*. [2004]). The evidence is inconsistent with a simple conditioning or associationist view where 'invariable uniformities of sequence' are what is learned.

The relevance of this to a naturalist should be clear: if 'our' concept of causation is an image of the kind of causal cognition that we actually engage in, then any feat of conceptual analysis starting from a model of 'Humean' constant conjunctions, or an appeal to alleged intuitions about such constant conjunctions, seems unlikely to be relevant to us, to science, or to metaphysics. To the extent that we have culturally universal intuitions about causation,

this is a fact about our ethology and cognitive dispositions, rather than a fact about the general structure of the world. For naturalists, intuitions are not evidence for their own content, in either science or metaphysics.

Empirical evidence in the specific case of intuitions about causation provides a particular exemplification of this general scepticism. Intuitive (in the sense of immediate, unforced) responses to causal questions are culturally varied. In particular it appears that adults in different cultures reason differently about causes. 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 in America) and environmental ones (more common in China). 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 (Morris *et al*. [1995]). Nor are our causal intuitions a reliable basis for basic scientific inference. It seems reasonably clear that people are relatively poor at judgements about conditional dependence, as measured in the Wason selection task (see Cosmides [1989]); probability, where many will, for example, assert that the conjunction of two possibilities is more likely than one of the conjuncts alone (Tversky and Kahneman [1983]), and in other areas, compared to communities of scientists. Our response in such cases is of course not to take our scientific practice as thereby impugned. Just as the scientific study of vision enables us to explain some visual illusions, so the study of causal cognition might be expected to show us that we were prone to 'causal illusions' (Gopnik *et al*. [2004]). 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.

The upshot of this section so far is twofold. First, human causal cognition is not Humean. One temptation to take regularity as definitive of causation is thereby removed. Second, our intuitions about causation are neither detailed enough nor stable enough for the entire folk concept as exemplified by this or that culture to be used a basis for constructing the concept as it is appealed to by special sciences. To the limited extent that the folk concept of causation has a common core across cultural elaborations, its content is as follows. First, it construes causal relations as centred on some agent of change (animate or otherwise), thereby distinguishing agents and patients. Second, it postulates various transformative principles (often conceived as 'forces') proceeding *out from* an agent *to* the recipient of causal influence. Kim's ([2005]) endorsement of a 'generative' account of causation may illustrate the attractions of this feature of the core folk centre of gravity, as we suspect does Armstrong's ([2004]) contention that the philosophical theory of causation begins with our experience of 'biff' in acting in the world. Third, the core folk notion of

causation incorporates assumptions about time-asymmetry: causal influences flow from the past into the future. These three features seem empirically to exhaust the core folk concept.[12] The artificial intelligence researcher Patrick Hayes ([1979]) spent years trying to work out deductive models of the everyday concept of causation as the basis for designing a simulation of a person. These models began with the three features just identified, then added knowledge people have of the mechanics of their own bodies to produce programs for distinctive manipulations of solid objects and liquids. Folk psychology would be conjoined with the core folk causal model to produce programs for causally influencing other agents.

It is open to a naturalist to allow a folk notion to fix what is to count as a cause for purposes of inquiry into anthropically constructed or centred domains and practices. However, the naturalist cannot allow anything but science to say whether there *are* such things as causes. With a watchful eye open for illicit appeals to folk judgements, then, let us turn to consider science.

3 Causes in Science

Russell maintains that 'advanced' sciences have no use for the 'law of causation' and do not seek instantiations of that law, even claiming that in such sciences the *word* 'cause' never occurs. Whether or not Russell is correct about the relevance of the 'law of causality' to science, many of the sciences, at least claim, explicitly, to pursue and answer causal questions. By this we mean (for now) the weak claim that the output of many communities of scientists 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'.

Consulting a single issue of *Science* (volume 300—June 2003) we 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', 'the common cold', 'oceanic primary production', and 'pear blister canker and cadang–cadang disease of

[12] Norton ([2003], pp. 18–1), in trying to articulate the structure of the folk notion of causation, suggests that it has somewhat more structure than this. However, it seems to us that he reads too much philosophical sophistication and restrictiveness into the folk concept. Norton himself inadvertently provides some evidence for this when he goes on to show how the various folk properties of causal relations that go beyond the ones we have identified are all subject to being waived in special applications.

coconuts'. A 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, averaging around 90 documents per month, in which the word 'cause' occurred. 'Effect' was more popular—10456 documents for the same period, around 112 per month. 3646 documents included the word 'influence', 8805 the word 'response', 2683 the word 'consequence'. (In case the reader is wondering, 'philosophy' turned up in 553 documents, much less often than cause, and approximately half as often as 'soil' with 1007).

So scientists talk about causes a great deal and, unless intolerable strain is to be placed on the notion of an 'advanced' science, or 'advanced science' as just a misleading way of denoting fundamental physics, Russell is simply wrong that the word cause 'never occurs'. This is a result of very modest significance. As already noted, most philosophers of science today agree with Russell that scientists do not seek laws of constant succession. What, then, *are* the things scientists call causes? We take it that what is being provided, or attempted, when scientists talk about causes is causal explanation. A well established set of philosophical questions about explanation in general, and causal explanation in particular, arise here, focussed on the problem of what makes for a good (causal) explanation in a particular science, in some group of sciences, or in general. For present purposes we adopt, without argument, a broadly Kitcherian approach to these questions. As Kitcher has emphasized (e.g., [1981]), the explanatory relations accepted in any science are determined by the 'general argument patterns'[13] accepted in that science. For reasons noted above, and returned to below, we are much more concerned than Russell to take seriously the fact of the number and variety of sciences. We agree with Kincaid ([1997]) that scientists typically rely on special causal relations that are parochial to their individual disciplinary traditions, and that do not, at least on the face of things, reduce to or welcome analysis by reference to a single overarching conception of the kind that philosophers typically seek.

It thus turns out that when contemporary scientists talk about causes what they are typically doing is documenting the very kind of relationship of multiple dependence that Russell takes to be distinct from putative 'causal' relationships. Russell's descriptive claim about what scientists say is incorrect, and he gives no compelling naturalist reason for saying that explanatory interdependencies should not be called causes. Further, if we are correct that cognitive science tells us that our own cognition is not Humean, and if supposed intuitive judgements that it is Humean are not intuitions but

[13] Where a general argument pattern is 'a triple consisting of a schematic argument, a set of sets of filling instructions containing one set of filling instructions for each term of the schematic argument, and a classification for the schematic argument' (Kitcher [1981], p. 516).

culturally developed philosophical theories, then a key temptation (for a naturalist) to make event regularity central to causation disappears. Since scientists also do not seem to make event regularity central to causation, the naturalist is left with no standing motivation for the Humean analysis. Might our verdict then be that scientists instead join the folk in using causal ideas to indicate hypotheses about 'biff' and 'flow'?

Can we seriously attribute to scientists such hazy and immature metaphysical generalizations as refer to 'flow' and 'biff'? If the resolution we are canvassing at this point were intended as anything more than a way station in the argument, then as naturalists we would have to give this question careful attention. Since we do not ultimately attribute the hazy metaphysical conviction to scientists however, we can just indicate suggestively how someone might most plausibly try to do so. First, one would maintain that scientists have no distinctive general idea of causation of their own; they just inherit folk custom here, and then tighten application in any given instance (so as to make causal claims suitable for testing and measurement) when they hypothesize or aim to discover specific *mechanisms.* It is a regularly encountered thesis in the methodological literatures of special sciences that explanations have not been provided until a mechanism, of a kind distinctive to the special-science in question, is specified. (See, for example, Elster [1989] and Kincaid [2004]). 'Mechanism', as this idea is used outside of literal mechanics, typically seems to mean nothing more definite than 'stable, quantitatively parameterized and measurable transmitter of influence'—which in turn looks as if it could be rendered without loss as 'biff-flow channel in a specified model'.[14] The substantive point here is that 'influence', as restricted to the sort of thing one needs a mechanism to transmit, is plausibly equivalent to the folk concept that Armstrong calls 'biff'.

These reflections might seem at first glance to furnish a basis for resolving all questions about what a contemporary naturalist ought to think about Russell's thesis. Neither the folk nor scientists conceptualize causation in a Humean fashion, in part because both consider more complicated interdependencies than bare regularities, and both can (more arguably) be taken as generally supposing, when attributing causal influence from some x to some y, a 'flow of biff' from causal agent x to patient y. If we were to accept all this, then we could

[14] Consider, for example, how little additional structure the following examples of causal mechanisms given by Kincaid ([1997]) have in common: strategies of organizations produced by the mechanism of differential birth and survival in different environments (sociology) (p. 19); rational expectations as mechanisms explaining monetary policy failures (economics) (p. 140); signal sequences attached to proteins as protein-sorting mechanisms in cells (cellular biology) (p. 53); protein structures as mechanisms determinant of binding strengths in antibodies (biochemistry/physical chemistry) (p. 60); distributed rule followed by networked individual people as the mechanism that steers an aircraft carrier (cognitive anthropology) (p. 124); homeotic genes as determinants of development paths from imaginal disks to specific organs in organisms (developmental biology/genetics) (p. 61).

say that we let folk custom fix what general sort of relation counts as causal, and then assign to science the job of telling us in detail what instantiates the set or network of actual causal relations.[15] Then our verdict on Russell's thesis would be that Russell conceded too much semantic priority to Humeans after all, and was thereby blinded to the compatibility of everyday and scientific ways of thinking.

This sort of resolution would be better than either simple eliminativism about causation, or an attempt to save causation by a defence of Humeanism. It is very close to a proposal of Norton ([2003]), whose general view is in turn the closest one to ours that we find in the literature. However, we think it is not quite adequate. For one thing, it is silent on any connection between causal relations and ontological unity. Norton explicitly denies such a connection, for reasons similar to Kincaid. Certainly, it is a plausible view to hold that causation has nothing to do with the glue the unificationist seeks, and that Kim and others are likewise confused in regarding causal multiplicity as being a problem for folk metaphysics. But unless one thinks that fundamental physics is the only metaphysically serious science—as Norton's failure in his article to consider any cases from special sciences as possible counterexamples to Russell's thesis suggests that he might—then leaving the problem at this point will be fully satisfying only to the advocate of metaphysical disunity.

As we have urged, Russellian eliminativism about causes is a plausible position in *fundamental physics*. No resolution of the problem of unity in the face of multidisciplinarity, with respect to causation or any other putative glue, can by itself resolve the tension between Russell's thesis *as applied to fundamental physics* and causal process theories. An advocate of metaphysical unification could just deny that causal process theories have got anything right at all, and set off in search of alternative glue. But this would imply an error theory of the underlying metaphysics of special sciences—Kincaid ([2004]), for example, would come out as flatly wrong—and that in turn leads to a position suggested by Dennett ([1991]) according to which fundamental physics studies the real *illata* and special sciences study ontologically second-rate *abstracta*. This would be reductionism (though not of Kim's kind), and we fear it is where Norton's position leads.

Let us recall our main motivation for revisiting Russell's discussion of causation. One goal of a naturalistic unifying metaphysics is to say genuinely enlightening things about the ontologies of all sciences by reference to general structural relations of some kind ('glue'). This (somewhat vague) formulation raises questions about the appropriate conceptual source for the 'general

[15] Note that the fit between the folk notion and the output of science need not be perfect in this scenario—the folk can perfectly well be surprised about what causes what, and also be called upon to make various modest revisions to their view of what causes are without being subject to the wholesale violence of Russell's proposal that there are no causes at all.

structural relations' that might do the job. We read Russell as depending on the idea that when anthropocentric intuitions and science fail to accord, the naturalist philosopher who is trying to describe the objective world (not the phenomenological data) must let science hold the trumps. Now, our evolved natural cognitive dispositions, as discussed above, tend to package reality for pursuit of anthropically centred goals; the selection that produced us could not have 'cared' about anything else. To the extent that informing our metaphysics with science involves replacing the anthropocentric picture with a more objectively oriented one, then, there must be conceptual revision. However, as we just saw, unresolved issues surrounding the metaphysical status of special sciences leave different revisionary possibilities open. Suppose we did conclude, like Dennett, that special sciences—but not fundamental physics—are pragmatic pursuits made necessary by the conjunction of human epistemic limitations and parochial concerns. In that case it might not be surprising if it turned out that minimal conceptual revision was necessary for understanding the role of causal generalization and explanation in special sciences; we should expect anthropocentric concepts to serve us well in anthropocentric pursuits. But it then would be surprising if the homely model of causation were vindicated in application to physics, since this would amount to the discovery that our natural conceptual apparatus had been well evolved for domains of no consequence to our fitness. The more ambitious unificationist hope, that special sciences and physics can justifiably be interpreted as pursuing an objective and coherent description of reality in tandem, should lead us to look for comprehensive revisionary principles first, then ask how these can accommodate the practices we find in special sciences, and what relations these have to patterns of folk usage.

We can make the above agenda more explicit by framing it in terms of the theses with which we have been contrasting Russell's, viz., causal process theories and Kim-style reductionism. The point of process theories, as discussed earlier, is to distinguish real causal processes from merely apparent ones. Salmon, like Russell and many philosophers of science, holds that, on questions about objective reality, not only should science trump anthropocentric perspectives (that is, we should be naturalists), but that within science physics should trump the special sciences whenever they conflict. Thus for Salmon the paradigm instances of real causal processes are to be sought in physics, particularly microphysics. Such commitments fuel the recent controversies based on worries about causal exclusion.

Kitcher's ([1989]) early critique of Salmon centred on just this point, arguing that our identification of causal interactions in most, probably all, macro-processes depends more on our knowledge of the general causal structure of the macro-world (as uncovered by sciences other than fundamental physics) than it does on the sorts of micro-process paramount for Salmon. However,

for Kitcher such considerations only give reason for not attempting to extract direct methodological guidelines[16] for conducting the business of the macro-sciences from process accounts, and do not seriously 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. That is, in Kitcher's view, a sound point about explanatory practice in science need not by itself establish anything about fundamental ontology. This is why Kitcher is able to describe his approach and Salmon's as complementary, characterising his own work as analysing 'top-down' explanation, where phenomena are explained by having their roles in wider ensembles of regularities fixed. He contrasts this with the 'bottom-up' explanation, the sort he says is analysed by Salmon—who in turn explicitly endorses Kitcher's suggested complementarity (Salmon [1990]). The problem with this rapprochement is that it sacrifices the ambition for unification of fundamental physics and special sciences, if it is agreed that special sciences explain by reference to causal processes but fundamental physics does not. This is slightly ironic in the light of the fact that Kitcher's overarching aim is to defend the idea that the scientific value of explanation consists in the unifications it furnishes. The Kitcherian rapprochement, at this point in our dialectic, supports explanation as unification within disciplines, but frustrates general unification—which is the only justifiable point of metaphysics, according to naturalists.

This ecumenicism tries to leave all the folk intuitions intact at once. If we accept it, we could refine the agent-centred idea of causation into a general process theory and make that central to microphysics, *and* hold on to the familiar ontology of the macro-world on grounds that it paid its way in science through delivery of top-down explanations. Kitcher's top-down explanations are precisely not supposed to be causal. The recent emergence of interest in causal exclusion can be understood as recognition that ecumenicism—criticized by Kim ([1998]) as hope for a 'free lunch'—dodges some questions about causation instead of answering them.

We have now argued our way into a tight corner. Getting out requires us to seek an account that (i) explains why special scientists can carry on about causal processes without embarrassment in the company of physicists (or, therefore, metaphysicians) but (ii) follows Russell in denying intuitions (folk or Humean) about causation *any* role in informing the metaphysical foundations of physics, or science more generally. The view we outline joins Russell, Redhead, and Norton in leaving the core folk concept of causation (as explicated above)

[16] Except, at least, for the imperative not to traffic in parochial relations of interdependence that have no counterparts at all on the graph of real processes, i.e. which are spooky according to fundamental physics.

intact. We appeal to two considerations, neither of them derived from natural or folk intuitions, to explain why the core folk concept of causation need not conflict with naturalistic philosophy of science or naturalistic metaphysics. One consideration is the very thin-ness of the core folk conception discussed in Section 2. The folk are committed only to the existence of processes, but not to the metaphysical idea that causal processes are glue. The second consideration will be a set of reasons for thinking that if Russell, Redhead, and Norton are right about fundamental physics, then the exact sense in which they might be right—given the actual state of contemporary physics—suggests physical and metaphysical explanations for the existence of processes, while implying that causal processes are not metaphysical glue. Our view grants Salmon and other process theorists a good deal—in particular, it accords with their belief in objective relations isomorphic to those that special sciences identify as 'causal', and with their idea that those relations should be explicated in terms of processes. In requiring of physical theory that it explain why special sciences are needed to study such processes it serves unification. However, in elaborating on how causation might play no role in fundamental physics, it accommodates Kitcher's objection to Salmon and locates what is right in Russell's thesis once his primitive reductionism is set aside.

4 Letting Science Hold Trumps

As we pointed out above, one premise of Russell's, deriving from his naturalism in application to the state of empirical science when he wrote, is a direct and simple form of reductionism about the relations among the sciences. When he champions letting science hold the trumps on conceptual revision, he supposes that it is specifically fundamental physics that has the cards. He argues that because the functional generalizations of physical theory involve no temporal asymmetries, whereas causation 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 anticipate by constructing causal relations. He thus assumes what current worriers about causal exclusion do, that whatever general patterns of relations are identified by any sciences must ultimately be (at least in principle) identifiable and explicable by physics.[17]

[17] Neither Russell nor the causal exclusionists expect physics to step in and take over the explanatory projects of the special sciences. Kim ([1998]), unlike Russell, has a sophisticated account of why this will not happen that is not just pragmatic. That is, Kim argues that special sciences study the distributions and effects of properties that are 'micro-based' but not themselves microproperties; and these latter are the direct targets of physical generalizations. However, because for Kim special-science properties supervene without cross-classification on physical ones, physics must still be able to identify and explain, at least in principle, all general relations that hold among groups of these micro-based properties.

As we pointed out in Section 1 above, when Russell wrote his essay this was still a reasonable bet. However, such reductive physicalism is not at all well supported by contemporary science. Most types of entities and relations that feature in interesting scientific generalizations in special sciences do not decompose into types and relations featuring in generalizations of fundamental physics. All that the state of current science justifies in this philosophical neighbourhood is the weaker principle that physics provides constraints on other sciences. In particular, it is a consistently enforced methodological principle in recent 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 [2004]; Ladyman and Ross [forthcoming]), we have discussed this as 'the primacy of physics constraint' (PPC). It might be metaphysically explained by a combination of affirming global supervenience of the non-physical on the physical while denying local supervenience. If this is the strongest commitment to physicalism that scientific practice expresses, then it is the version of physicalism with which a consistent naturalist should stop. Then: if physical theory abjures reference to causal relations but biology or economics invokes them, this generates no contradiction, and so no inference from this possible fact to Russell's conclusion is warranted by the primacy of physics as it features in scientific practice.

Once we abandon reductionism, we can no longer identify existence with mereological composition out of basic constituents. On the other hand, as Kim's worries about 'free lunches' remind us, if we aim at a serious metaphysical account then we need some alternative account of what it is to exist that makes existence an objective matter, rather than a function of the pragmatic constructive efforts of an arbitrary group of epistemically limited beings. Furthermore, if this metaphysic is to be naturalistic, it must both leave science as the arbiter of which particular things and types of things exist, and it should cohere with the PPC.

We 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]) argues that Dennett's own formulation of this 'existence principle' is not fully successful in avoiding instrumentalism, and argues for replacing it with the following formulation:

To be is to be a real pattern, and a pattern is real if

(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.

While avoiding instrumentalism, this formulation suffers from some imprecisions that emerge when one tries to unpack all the information-theoretic concepts it invokes in a vocabulary derived strictly from fundamental physics, as a naturalistic metaphysics of existence requires. Ladyman and Ross ([forthcoming], Chapter 4) therefore construct a precisification in physical language. Introducing this here, however, would carry us into lengthy technical issues that would unduly distract from the issues at hand, so the reader is referred offline while offering an intuitive explanation of the imprecise formulation. The Dennett/Ross theory's reliance on physically *possible* perspectives, and on what is project*ible* rather than on what is actually project*ed*, makes it realistic rather than anthropocentrically instrumentalistic, since it abstracts away from actual, contingent observers and appeals ultimately to physical theory in the limit. Its reliance on *physically* possible observers makes it non-idealistic, and respects both the PPC and our naturalistic suspicion of any conceptual framework in metaphysics that claims wider motivational scope than what can be had from fundamental physics. Simultaneously, it avoids the implication of reductionism. Many real patterns will be robust only given restrictions in special regions of the universe in which they occur, thus requiring knowledge of various contingent initial conditions not deducible from generalizations of fundamental physics (though constrained by it). Echoing Russell ([1913], p. 187) 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.[18] 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

[18] 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 with reference to the distribution of traces relative to our backward light cone, rather than by the examiner of the whole universe, which might be a block (Lockwood [2005])) will cease to make any informational difference to anything. At some point, that is, he is 'forgotten' not just by people but also by the universe, in the sense that no traces of him *qua* the distinctive properties that made him unique are detectible by any physically possible information detector.

be able to pick out macroeconomic or Napoleonic patterns (except insofar as they're undistinguished instances of general ones), though it explains why and where they are possible. Furthermore, as detection of these patterns may require cross-classification of other patterns typically thought of as more basic by reductionists, *no* reductionism is implied by the view at all.

Special sciences are, then, not responsible for finding generalizations that support counterfactuals arising outside of their scope boundaries. By contrast, fundamental physical theory[19] is responsible for supporting generalizations that hold across, at minimum,[20] the entire actual universe. Commitment to the idea that a measurement from anywhere is a potentially relevant counterfactual to any thesis of fundamental physics is commitment to precisely the asymmetry that is the basis for the PPC. The PPC is not guaranteed to block only false beliefs, partly because physics can not be infallible, but more importantly (practically speaking) because fundamental physics is not finished, and some parts of the universe are, at any given time, inaccessible to our measurement and hence the testing of our generalizations. It is still rational to project the structural character of the generalizations in the limit from the sample we currently have, and this is what our fundamental physical theories do. (See Ladyman and Ross [forthcoming] for details.)

We want now to see where reasoning that resembles Russell's leads if we replace his reductionist version of the priority of physics with the weaker PPC. Consider the set of fundamental physical theories currently being taken seriously by physicists—that is, those used or investigated in funded research programmes that aim to unify physics and so shade into metaphysics. The theories in question are the serious candidates for accounts of quantum gravity (Greene [2000]; Smolin [2001]). Suppose, hypothetically, that none of the generalizations in these theories invoke the kinds of causal asymmetries characteristic of the flow (or process) understanding of causation, while some or all special-science causal principles do presuppose such asymmetries (as Kincaid [2004] persuasively argues). Then three possibilities are open with respect to the status of these asymmetries in metaphysics.

One possibility is that fundamental physics ultimately endorses a 'background-free' theory of the sort promoted by Smolin ([2001]). 'Background

[19] This is *not* coextensive with all of institutional physics, i.e. some branches of physics are special sciences. Scientists do not apply the PPC to everything physicists believe true. We can identify the parts of physics that are taken to be fundamental by identifying the topics of the generalizations that asymmetrically constrain the rest of science in its working institutional life.

[20] We say that physics is *at minimum* responsible for describing the entire actual universe—rather than just saying that that is what it describes, period—because we think that physics makes some claims with modal force, and so is answerable to counterfactuals that can be interpreted as being about at least mathematical possibility. Our formulation avoids embroiling us, at least here, in argument with van Fraassen ([2002]), who denies that there are modal physical truths.

freedom' is the property that a theory has no, or almost no, structure, that it is independent of dynamics. If this is the future of physics, then the search for metaphysical glue is forlorn, and the substance of radical disunity views might be satisfied. Smolin favours Loop Quantum Gravity (LQG) as the candidate unifying theory for General Relativity and Quantum Mechanics in part because he thinks it is (largely) background- free.[21] Notice, however, that even here both Russell and the causal process theorists could feel substantially vindicated. The eliminativist would not be able to rely on the claim that causation is based on asymmetry and fundamental physics insists on universal symmetries. However, the core folk causal notion is committed to more than just asymmetry; it is anchored in the anthropomorphic experience of flow of influence. In the limiting case of background freedom, nothing fundamental will be available to stand proxy for the universal direction of causal influence, and so the anthropocentric experience fails to find any analogical grip there at all. At the same time, process theories might be necessary for explaining and predicting the real patterns in our part of the universe, in which the Second Law holds sway. It would be the responsibility of LQG to explain the local asymmetries, which would in turn explain the utility of specifically *time*-asymmetric processes as book-keeping devices for observers taking measurements around our region of the universe.

However, a *completely* background-free physics is in any case just a limiting possibility, which most physicists (including Smolin) regard as unlikely. We can sort other perspectives in quantum gravity by reference to the distance they stand away from this limit. Shifted just over from it are possible theories that give only geometrical properties of the universe, which neither constrain nor are consequences of any local processes. Here talk of 'glue' might still seem inappropriate, and the following would all look like motivated speculations for the metaphysician: the universe is minimally unified, there is no causation at the level of fundamental physics, process theories might still describe every real pattern we can measure except some features of the large-scale structures of spacetime, but if they do this cannot be explained directly by fundamental physics, though they must be compatible with it.

A more interesting possibility (by reference just to the issue of the dialectic between Russell and process theorists) compatible with both non-reductive naturalism and the Russell–Redhead thesis about physics is that fundamental physical theory helps us to say something about what all of the parochial special-science models of causal processes have in common, though these processes make no first-order appearance in fundamental physical theory

[21] We worry that this may be an instance of a physicist commiting a sin more typical of philosophers: letting metaphysics drive speculations in physics, instead of the other way around.

itself. If this possibility obtains, the metaphysician is justified in studying physics in search of universal glue, but causation cannot *itself* be that glue. The triumph of superstring/M theory would support this metaphysical picture. Investment of energy in superstring/M theory is usually justified by the fact that it takes seriously the alleged induction over the history of physics that respect for symmetries is the key principle of any would-be 'fundamental' physical theory (Greene [2000]). As we have argued, the folk understanding of causation may not be committed to very much, but it insists on asymmetry between causes and their effects. Thus if M theory turned out to yield the content of fundamental physics, then causation could not be universal glue. ('Stringhood' might be the universal glue.)

It should be noted that M theory and LQG are not strictly exclusive of one another. There are possible extensions of M theory—albeit not yet mathematically formulated—which would explain why LQG holds (Greene [2000]; Smolin [2001]). This would be a particularly inviting theoretical environment for reconciling Russell and the process theorists. Though LQG has at least one consistent model in the context of a block universe, it is usually presented in tensed terms. Thus it has resources for potentially providing a correct description of the historical development of the region of the universe in which we live, or even of the whole of what physically exists, that would be given in terms of processes. M theory might then offer the unifying structural account at one higher level of abstraction. At this level Russell's thesis would be confirmed.

We claim that on all these possible futures, there might be real processes—chemical ones and evolutionary ones and macroeconomic ones, etc.—but that there is nothing they would have in common beyond satisfying the thin folk conception of time-asymmetric influence-flow described earlier. We are thus in basic agreement with the claim of Kincaid's, cited earlier, that many of the mechanisms for transmitting influence found across the sciences have nothing more general in common than that they transmit influence. On some scenarios now seriously entertained (background freedom), fundamental physics would be merely compatible with asymmetric influence-flow in some regions of reality. On another account—M theory—there would be no fundamental asymmetries, hence nothing corresponding to causal processes at the fundamental level; yet fundamental physics would explain why such asymmetries hold with respect to real patterns measurable in special-case regions where values of some of the fundamental 11 dimensions of strings take vanishingly small values.

It should now be clear in what sense our view accords with Russell's. Causation, we claim, is an epistemic book-keeping device used by agents whose world of experience is time-asymmetric and dynamic. It is a useful device, at least for us, for locating some real patterns, and fundamental physics might

play a role in explaining why it is. Or, if fundamental physics does not directly do this, then the PPC at least requires that the time-asymmetric measurements characteristic of our region of the universe be compatible with fundamental physics; and the ubiquity of these measurements would in turn explain the value of the 'causation' book-keeper in special sciences. However, the current state of fundamental physics supports Redhead's and Russell's expectations that it will not directly generalize over causal processes.

How much does this allow us to concede to causal process theories? Salmon initially sought to provide an account of causation as universal glue. Taken strictly, this is incompatible with Russell's thesis. However, faced with the problem that quantum theory requires uncaused parameter values, Salmon ([1990]) ultimately expressed stark pessimism about the idea of causation as universal glue. Note that this reason, while different from those we have pursued here, is still a sound one for a naturalist. What Salmon considered more important to the enterprise of process theorists was the attempt to distinguish cases of 'real' causation from pseudo-processes; and process theorists still think that this enterprise is getting somewhere. Now, this may be thought to also be incompatible with Russell's thesis. After all, if there is not 'really' any such thing as causation—as a metaphysical kind—then how can there be a metaphysical distinction to be drawn between real and pseudo-processes?

The answer to this question rests on reference to Russell's mistaken belief in reductive physicalism. For a non-reductionist, the failure of causal relations to appear in generalizations of fundamental physics does not imply that these relations simply denote fictions.[22] All our current evidence tells us that we live in a region of the universe—which might or might not be coincident with the whole universe[23]—in which the degrees of freedom of every system we approximately isolate for measurement and projection are restricted by various asymmetries. This region is highly isotropic, and supports robust (that is, projectible across all counterfactual spaces within the boundaries of the region) distinctions between sources and recipients of information. That is sufficient to establish the scientific utility of directional flow. Thus it is sufficient to account for the *scientific* utility of the *folk* idea of causation. We said earlier that we do not attribute to scientists a metaphysically serious

[22] Norton ([2003], p. 3) says something that sounds similar: 'I will characterize causal notions as belonging to a kind of folk science, a crude and poorly grounded imitation of more developed sciences. While this folk science is something less than our best science, I by no means intend to portray it as a pure fiction.' As we will go on to do, Norton characterizes this not-altogether-fictional status in terms of 'heuristics'. However, in our view, special sciences, where causal relations are ineliminably quantified over, are not 'crude and poorly grounded imitations' of physics; they track real patterns that fundamental physics does not and cannot. Therefore, the sense in which we regard causal relations as 'heuristics' separates them more sharply from fictions than does Norton's sense of the term.

[23] If BT turns out to be fundamental, then this coextensivity must hold, and, as noted, Russell's thesis will be of merely historical interest.

commitment to hazy folk intuitions about biff and flow. We can now say what we meant by that. Taking directed, irreversible transmission of influence, plus conservation principles, for granted is rational when you know you're taking all your measurements from regions governed by a particular asymmetry, the Second Law of Thermodynamics. This practice conforms so closely to what is enjoined by a folk metaphysical endorsement of causation that no serious risk of practical (as opposed to philosophical) misunderstanding arises if the scientist helps herself to the culturally inherited folk idea of causation. But she need not thereby endorse folk metaphysics, any more than she must endorse folk physics. It is either to fundamental physics directly, or to some domain less general than fundamental physics but more general than our domain of immediate measurement (e.g., LQG, in the scenario where that is explained by M theory) to which one must turn for an explanation of how regions governed by the Second Law can come to be and remain stable (at least temporarily).

Is this enough to show that there are real processes, and so also cases where we might misattribute process, that is, pseudo-processes? Or does it merely show that it is useful for us limited beings, stuck in a restricted region of reality governed by special asymmetries, to book-keep our experience in terms of the idea of causal process? This question amounts to asking once again how it is possible to be a naturalistic realist (in denying anthropocentric instrumentalism) and a non-reductionist (in denying that fundamental physics exhausts or strictly implies everything that is true) at the same time.

The PPC and the theory of real patterns explain why parochial causal relations are robust in special sciences, but also why nothing stronger than the thin notion of flow or process unifies them. First, special sciences, by definition, are not constructed from the vantage point of the limit. By this we mean not only that they are not constructed from the physical limit; they also are not cast from the limits achievable at their own scope boundaries. They are developed in, and are accounts of parts of, a restricted world of informational asymmetries. We do not know, until fundamental physics tells us, whether in this respect the world as it appears when we restrict its boundaries to conform to the edges of our experience agrees with the correct description of the actual world available from the limit. So the importance of causation in special sciences cannot license an inference to the *metaphysical* significance of causation. It is because special sciences are concerned with relatively isolable regions of space-time, which involve cross-classification from the perspective of less-than-fundamental but nevertheless limiting domains for them, that the causal relationships on which they variously focus will appeal to different aspects of their information-carrying potential. This emerges as Kitcher's and Kincaid'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.

This perspective helps us make sense of the relationship of the folk causal concepts to scientific applications. The Dennett/Ross existence principle 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.[24] 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, tracking the trajectories through local space and time of bundles of rewards is almost everything. Attention to wider informational dynamics in which processes are embedded typically delivers few if any additional payoffs, and may get in the way of payoff maximization because of computational costs.[25] Therefore modelling causal relations as sequences of collisions of objects in time is a sensible heuristic. Furthermore, for social animals like people, or 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 amounting to limiting cases of agents with minimal sophistication of purpose and adaptability. Such models of causation will not 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. However, special extensions of the thin concept of flow causation, such as animism, can then be dropped, leaving us with the kind of stripped down notion of process sought by Salmon and other process theorists. What these theorists give us insofar as their analytical efforts are successful is not *universal* glue, but it is indeed a binding agent that holds special sciences together in parts of reality where they have a point.

We close by returning to the two questions we identified at the outset. First, we asked, is it indeed the case that scientists, at least when 'advanced', do not seek 'invariable uniformities of sequence'? We agree with Russell that they generally do not. Because we are persuaded that cognitive science shows that everyday people do not anticipate in terms of such uniformities, we are not convinced that this fact is important. Second, we asked, if they do

[24] As naturalists, we do not claim that greater relative parsimony provides any reason to think a scientific theory relatively more likely to be true. The point is simply that it does not help to unify science, and thus makes no contribution to metaphysics, to say that objects generalized over only by redundant theories exist.

[25] 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 is wise to book-keep your structures and relations more like a real scientist.

not, does it follow that they do not seek what are legitimately called causes? We worked up some concern over a seeming tension lying within the basis for Russell's answer. On the one hand, his crucial premise for justifying the answer 'no' is that physicists do not mention causes in their fundamental generalizations. But since special scientists carry on about causes all the time, Russell's denial entails the conclusion, terrible for a naturalist, that most scientists are ontologically confused in ways philosophers can diagnose. We eventually found the solution to this tension just where a naturalist would want it, in correcting empirical assumptions about physics that were not borne out by the subsequent development of science. Because reductive physicalism has (to this point) failed, there is no justification for the projection of causation all the way down to fundamental physics and metaphysics. Nor does science presently motivate us to conclude that special sciences do not measure real patterns that must be modelled as processes. While thus not then exactly reaching Russell's original conclusion, we have got to a slightly more nuanced one by following naturalistic rules of which he would have approved.

Department of Philosophy
and Department of Finance,
Economics and Quantitative Methods,
University of Alabama
at Birmingham, School of Economics,
University of Cape Town,
900 13th St. So.,
Birmingham, AL,
USA 35294
dross@commerce.uct.ac.za
dross1@uab.edu

School of Philosophy and Ethics,
University of KwaZulu-Natal,
Durban, 4041,
South Africa,
spurrett@ukzn.ac.za

References

Almeder, R. [1998]: *Harmless Naturalism*., La Salle: Open Court.

Armstrong, D. M. [2004]: 'Going through an open door again: counterfactual versus singularist theories of causation' in J. Collins, N. Hall, L. A. Paul (*eds.*), *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 445–58.

Baillargeon, R., Kotovsky, L. and Needham, A. [1995]: 'The Acquisition of Physical Knowledge in Infancy' in D. Sperber, D. Premack, A. Premack (*eds.*), *Causal Cognition*, Oxford: Oxford University Press, pp. 79–116.

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

Batterman, R. [2002]: *The Devil in the Details*, Oxford: Oxford University Press.

Brown, J. [1999]: *Philosophy of Mathematics*, London: Routledge.

Cartwright, N. [1980]: 'Do the Laws of Physics State the Facts?' *Pacific Philosophical Quarterly*, **61**, pp. 75–84.

Cartwright, N. [1999]: *The Dappled World*, Cambridge: Cambridge University Press.

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

Cosmides, L. [1989]: 'The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task', *Cognitive Psychology*, **31**, pp. 187–276.

Crane, T. [1995]: 'Mental Causation', *Proceedings of the Aristotelian Society*, **69**(Suppl. 1), pp. 211–36.

Creath, R. (*ed.*). [1991]: *Dear Carnap, Dear Van: The Quine-Carnap Correspondence and Related Work*, Berkeley: University of California Press.

Dennett, D. [1991]: 'Real Patterns', *Journal of Philosophy*, **88**, pp. 27–51.

Dowe, P. [1992]: 'Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory', *Philosophy of Science*, **59**, pp. 195–216.

Elster, J. [1989]: *Nuts and Bolts for the Social Sciences*, Cambridge: Cambridge University Press.

Field, H. [2005]: 'Causation in a Physical World', in M. Loux, D. Zimmerman (*eds.*), *Oxford Handbook of Metaphysics*, Oxford: Oxford University Press, pp. 435–60.

Fine, A. [1986]: *The Shaky Game*, Chicago, IL: University of Chicago Press.

French, S. and Ladyman, J. [2003]: 'Remodelling Structural Realism', *Synthese*, **136**, pp. 31–56.

Friedman, M. [1974]: 'Explanation and Scientific Understanding', *Journal of Philosophy*, **71**, pp. 5–19.

Friedman, M. [1999a]: *Reconsidering Logical Positivism*, Cambridge: Cambridge University Press.

Friedman, M. [1999b]: *Dynamics of Reason*, Stanford: CSLI.

Garcia, J. and Koelling, R. [1966]: 'Relation of Cue to Consequence in Avoidance Learning', *Psychosomatic Science*, **4**, pp. 123–4.

Gelman, R., Durgin, F. and Kaufman, L. [1995]: 'Distinguishing Between Animates and Inanimates: Not by Motion Alone', in D. Sperber, D. Premack, A. Premack (*eds.*), *Causal Cognition*, Oxford: Oxford University Press, pp. 150–84.

Glymour, C. [1999]: 'A Mind is a Terrible Thing to Waste–Critical Notice: Jaegwon Kim, Mind in a Physical World', *Philosophy of Science*, **66**, pp. 455–71.

Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., and Danks, D. [2004]: 'A Theory of Causal Learning in Children: Causal Maps and Bayes Nets', *Psychological Review*, **111**, pp. 1–31.

Greene, B. [2000]: *The Elegant Universe*, London: Jonathan Cape.

Hayes, P. [1979]: 'The Naïve Physics Manifesto', in D. Michie (*ed.*), *Expert Systems in the Micro-electronic Age*, Edinburgh: Edinburgh University Press, pp. 242–70.

Hume, D. [1975]: L. A. Selby-Bigge (*ed.*), *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, (3rd edn), Oxford: Oxford University Press, Revised by P. H. Nidditch, Original publication 1777.

Hume, D. [1978]: L. A. Selby-Bigge (*ed.*), *A Treatise of Human Nature*, Oxford: Oxford University Press, Orignial publication 1740.

Kant, I. [1933]: N. Kemp Smith (*ed.*), *Critique of Pure Reason*, London: Macmillan.

Kim, J. [1998]: *Mind in a Physical World*, Cambridge, MA: MIT Press.

Kim, J. [2005]: *Physicalism, or Something Near Enough*, Princeton: Princeton University Press.

Kincaid, H. [1997]: *Individualism and the Unity of Science*, Lanham, MD: Rowman and Littlefield.

Kincaid, H. [2004]: 'There Are Laws in the Social Sciences' in C. Hitchcock (*ed.*), *Contemporary Debates in Philosophy of Science*, Oxford: Blackwell, pp. 168–85.

Kitcher, P. [1981]: 'Explanatory Unification', *Philosophy of Science*, **48**, pp. 507–31.

Kitcher, P. [1989]: 'Explanatory Unification and the Causal Structure of the World' in P. Kitcher, W. Salmon (*eds.*), *Scientific Explanation*, Minneapolis, MN: University of Minnesota Press, pp. 410–505.

Ladyman, J. [1998]: 'What is Structural Realism?' *Studies in History and Philosophy of Science*, **29**, pp. 409–24.

Ladyman, J. [2000]: 'What's Really Wrong With Constructive Empiricism? Van Fraassen and the Metaphysics of Modality', *British Journal for the Philosophy of Science*, **51**, pp. 837–56.

Ladyman, J. and Ross, D. [forthcoming]: *Every Thing Must Go: Metaphysics Naturalised*, Oxford: Oxford University Press.

Lockwood, M. [2005]: *The Labyrinth of Time*, Oxford: Oxford University Press.

McDowell, J. [1994]: *Mind and World*, Cambridge, MA: Harvard University Press.

Mill, J. S. [1974]: A System of Logic Ratiocinative and Deductive, Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation, (volume 1), in J. M. Robson (*ed.*), *The Collected Works of John Stuart Mill*, Vol. VII, Toronto: University of Toronto Press.

Morris, M., Nisbett, R. and Peng, K. [1995]: 'Causal Attribution Across Domains and Cultures', in D. Sperber, D. Premack, A. Premack (*eds.*), *Causal Cognition*, Oxford: Oxford University Press, pp. 577–612.

Norton, J. [2003]: 'Causation as folk science', *Philosophers' Imprint*, **3**(4), <www.philosophersimprint.org/003004/>

Rea, M. [2001]: *World Without Design: The Ontological Consequences of Naturalism*, Oxford: Oxford University Press.

Redhead, M. [1990]: 'Explanation', in D. Knowles, (*ed.*), *Explanation and Its Limits*, Cambridge: Cambridge University Press, pp. 135–54.

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

Ross, D. [2000]: 'Rainforest Realism: A Dennettian Theory of Existence', in D. Ross, A. Brook, D. Thompson (*eds.*), *Dennett's Philosophy: A Comprehensive Assessment*, Cambridge, MA: MIT Press, pp. 147–68.

Ross, D. and Spurrett, D. [2004]: 'What to Say to a Skeptical Metaphysician: A Defense Manual for Cognitive and Behavioral Scientists', *Behavioral and Brain Sciences*, **27**, pp. 603–27.

Russell, B. [1913]: 'On the Notion of Cause', in *Proceedings of the Aristotelean Society*. Reprinted in Russell, B. 1917, *Mysticism and Logic*, London: Unwin, pp. 171–96.

Russell, B. [1925]: *ABC of Relativity*, London: Unwin.

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

Salmon, W. [1990]: 'Scientific Explanation: Causation and Unification'. *Critica Revista Hispanoamericana de Filosofia*, **22**, pp. 3–21.

Smolin, L. [2001]: *Three Roads to Quantum Gravity*, New York: Perseus.

Spelke, E. S., Phillips, A. and Woodward, A. [1995]: 'Infants' Knowledge of Object Motion and Human Action', in D. Sperber, D. Premack, A. Premack (*eds.*), *Causal Cognition*, Oxford: Oxford University Press, pp. pp. 44–78.

Tversky, A. and Kahneman, D. [1983]: 'Extensional Versus Intuitive Reasoning: the Conjunction Fallacy in Probability Judgement', *Psychological Review*, **90**, pp. 293–315.

Van Fraassen, B. [1980]: *The Scientific Image*, New York: Oxford University Press.

Van Fraassen, B. [2002]: *The Empirical Stance*, New Haven: Yale University Press.

Walter, S., Heckmann, H.-D. (*eds.*), [2003]: *Physicalism and Mental Causation*, Bowling Green, KY: Academic Imprint.