Living Causes[1]

John Dupré, University of Exeter

Abstract

This paper considers the applicability of standard accounts of causation to living systems. In particular it examines critically the increasing tendency to equate causal explanation with the identification of a mechanism. A range of differences between living systems and paradigm mechanisms are identified and discussed. While in principle it might be possible to accommodate an account of mechanism to these features, the attempt to do so risks reducing the idea of a mechanism to vacuity. It is proposed that the solution to this problem requires the development of a philosophical account of process adequate to apply to living systems.

I. *Introduction.* There has been a considerable amount of significant work in the philosophical study of causation in recent years. However, it has yet to provide an account of causality adequate to comprehend the complexities that are emerging in recent work in the biological sciences. Or so I shall argue. In this paper I shall very briefly survey some of the more recent philosophical ideas in the area, with special attention to recent work on mechanism, which has been proposed as grounding the causal features of complex systems including living beings. I shall suggest that the assimilation of living systems or parts of living systems to mechanisms is problematic in several respects. I shall propose that an adequate account both of the nature of living systems and of the causality that they exhibit will require more serious attention to the thoroughly processual character of living systems, and ultimately a better understanding of processes themselves.

II. *Recent theories of causation.* Theories of causation are generally divided into two broad categories. On the one hand there are theories that emphasise the mode of production of an effect. At a fairly abstract level, these would include accounts based on the idea of causal powers or capacities (Cartwright 1989; Mumford and Anjum 2011)—accounts famously castigated by Hume—and the concept of causal process developed by Salmon (1984) and refined by Dowe (2000). More recently the most widely discussed accounts of causation as production have been grounded in the concept of mechanism (Machamer, Darden and Craver 2000)[2]. Whereas these accounts allude to something—a power, a mechanism—that can be located in the individual cause, the alternative approach respects Hume's strictures, and restricts itself to general relations between causes and effects. A recently influential version of the second approach has been grounded in the idea of a cause as potentially making a difference, and more specifically, James Woodward's (2003) idea of a cause as a possible intervention (though not necessarily a human intervention). But also notable are counterfactual-based accounts, most notably in the work of David Lewis (1968), and statistical analyses by authors such as Eells (1991) and Hitchcock (2001). These general approaches might be referred to as, respectively, productivist and difference-making.[3]

[1] The paper has benefitted from comments on an earlier draft by James Woodward, Alex Powell, and Federica Russo.

[2] This has been an extremely active programme in the last decade or so, and a large philosophical literature has been generated. Some of this addresses many of the points emphasised in the present essay, and reaches a view similar in many ways to that defended here (e.g. Bechtel 2011). Though it is doubtless possible in principle to stretch the idea of mechanism to apply to whatever turns out to be the way living systems are organised, it seems to me more important to stress the differences than the similarities with standard conceptions of a mechanism.

[3] Though all such accounts share Hume's scepticism about causal powers, not all envisage a reduction of causation to relations between events. Woodward's, in particular, is not reductive in this sense.

I do not propose to examine all of these approaches to the topic in detail. Indeed, I believe that causality is a complex and diverse set of phenomena, and most or all of these accounts provide valuable and complementary perspectives on the topic. Such a pluralistic view is quite a common one among contemporary philosophers; however, there are significant differences in the form that such pluralisms can take. Nancy Cartwright (2004) defends what might be called brute pluralism: there are just different kinds of causes. Peter Godfrey-Smith (2010) supports an epistemic pluralism about causality, but is sceptical about the ontological status of causes. Federica Russo and Jon Williamson (2007), finally, concur with the epistemic pluralism, but think there is an underlying and monistic causal reality that these various methods seek to explore.

I shall next turn to a consideration of Russo and Williamson's thesis, and from there lead into a discussion of the very influential mechanistic position that they partially endorse. While recognising real insights that this new mechanism offers, I shall argue that it is ultimately inadequate for understanding causality within biological systems. Rather, I shall suggest that a more radically processual account of first living things, and then causation at least in so far as it applies to living things, is needed for this purpose. Unfortunately, extant process-based accounts of causality are wholly grounded in consideration of physical causes, and the features of causal relations on which they are based are difficult or impossible to translate into biological contexts. Also unfortunately, the concept of a process lacks a widely agreed and satisfactory explication. So this paper will be more the statement of an agenda for future work than the proposal of a detailed account of the nature of living causes.

III. *Epistemic causal pluralism.* Russo and Williamson (2007) assign complementary roles to the two basic concepts of cause just distinguished. There are two problems with difference-making analyses, they argue. The first problem is that ultimately such approaches cannot, according to Russo and Williamson, successfully resolve the classic problem of distinguishing causation from; the second is that they cannot solve the problem of external validity. The latter problem is that there is no way of knowing whether results reached by population level studies can be extended beyond the specific and perhaps idiosyncratic conditions of the original investigation. But both these problems can be solved by discovering the mechanism responsible for the population level relationship: the mechanism discloses the causal connection between the correlated items; and it shows what it is about a population that makes the correlation hold, and hence what must be true about further populations to which it is extended.

But mechanisms alone, they note, are insufficient to generate the population level causal results we often require. They cannot tell us whether or to what extent the conditions required for the operation of mechanism actually occur in the population. If, to take a rather well-worn example, we want to know whether smoking causes lung cancer we need both to understand the mechanism that connects the behaviour to the disease, and to know whether the conditions for this mechanism to work exist in actual human populations. Of course, this should not be taken to deny the evidential weight of, for example, Richard Doll's (1950) famous epidemiological investigation of the relation between smoking and lung cancer. But evidential weight may still be distinct from proof; and the ability of the tobacco industry to keep the debate going to some degree for a further half century may be ironic support for Russo and Williamson's position.

Russo and Williamson are anyhow surely right to insist that both these perspectives on causality are important, and are taken to be important in scientific practice. Consider the US Institute of Medicine Report (2011), "Adverse Effects of Vaccines: Evidence and Causality". The report states:

> 'Two streams of evidence support the committee's causality conclusions: epidemiologic evidence derived from studies of populations (most often based on observational designs but

> randomized trials when available), and mechanistic evidence derived primarily from biological and clinical studies' (p. 2)

However, the committee did not perceive these approaches as being on a quite equal footing. Discussing different degrees to which they took causal hypotheses to be supported, they write:

> 'The framework allows for a causality conclusion of "convincingly supports" based on an epidemiologic weight-of-evidence assessment of high in the direction of increased risk (which requires at least two well-conducted epidemiologic studies). Strong mechanistic evidence, which requires at least one case report in which compelling evidence exists that the vaccine indeed did cause the adverse event, always carries sufficient weight for the committee to conclude the evidence convincingly supports a causal relationship.' (p. 7)

Rightly or wrongly, I think it is clear that the committee believed, first, that to say C causes E is to say that C has the power to produce E, where this is understood as requiring that there be a mechanism leading from C to E; and, second, that decisive evidence that C has the power to cause E requires both understanding the mechanism by which this result is produced and demonstrating at least once the operation of this mechanism.

The connections between on the one hand the statistical relations between a cause and its putative effect, and on the other the mechanism(s) that connect the former to the latter can be complex. Consider genetics. Mendelian genetics is a difference-making science par excellence. In classical Mendelian experiments the available alleles are variables (in Woodward's (2003) sense) and the phenotypic outcomes are the effects that can be manipulated through this variable[4]. The basic idea lives on in a more complex polygenic world. Richard Dawkins once wrote: 'when a geneticist speaks of a gene "for" red eyes in Drosophila ... [h]e is implicitly saying: there is variation in eye colour in this population; other things being equal, a fly with this gene is more likely to have red eyes than [is] a fly without this gene' (1982, p. 21). On the face of it, this characterization of the gene is consistent with many different alleles with minor contributions to a polygenic effect. There are, however, many problems with Dawkins' proposal, not least that it does indeed invite a conflation of causation with correlation.

It was once supposed that molecular genetics would be the mechanistic science that would ground the general relations asserted by Mendelian genetics. But a number of philosophers and biologists have shown this expectation to be highly problematic (e.g., Hull 1974; Kitcher 1984; Dupré 1993). A central difficulty, reflecting the complexity of biological systems, is that there is no reason to expect a one-to-one mapping of Mendelian relations onto molecular mechanisms[5]. Statistical relations of the kind envisaged by Dawkins' definition, at any rate, may well be underpinned by a wide variety of molecular mechanisms. And while a gene 'for' x will presumably initiate some causal pathways that lead to the trait x, it may very well also be involved in mechanisms that tend to prevent the appearance of trait x.

The more general point emerging from the preceding remarks is a relatively familiar or even banal one: biological systems are extremely complex, and parts of them frequently engage in multiple distinct processes. Instances of biological kinds typically differ in causally significant respects. Hence,

[4] Manipulation or intervention, for Woodward, does not require a manipulating agent.

[5] The occasional classically Mendelian relations, in which the famous numerical relations between phenotypes are found, have turned out to be very much the exception rather than the rule. Such relations reflect an unusually simple relation between genotype and phenotype.

further, the presence of multiple mechanisms mediating the interaction between various aspects of variable systems is likely to be typical. For these reasons we should be genuinely—not merely in principle—suspicious of the legitimacy of exporting causal results between contexts.

This should begin to explain our worries about the ambitious monism in the project of Williamson and his collaborators. In an earlier paper Williamson (2006) writes: "The epistemic theory of causality is an example of a general strategy for developing a determinate, monistic metaphysics that is true to the epistemology of a concept" (2006, p. 80); and in the same paper: "the causal relation is just the causal belief graph of an omniscient rational agent (an agent whose evidence is exhaustive)". Not believing in omniscient agents, I have serious doubts about the feasibility of tasks for which their assistance is solicited. I doubt whether, in the real world of finite agents, there is such a thing as exhaustive evidence. I focus on Williamson here, however, not to make cheap shots at fairly harmless metaphors, but because not only does he well represent tendencies that are characteristic of the new mechanism, but the embedding of mechanism in theories of general causality provides a plausible way of using mechanism to understand causality. Now, however, I shall turn to some concerns about mechanisms that are intended to motivate a somewhat different metaphysical picture.

IV *Finding "the" mechanism.* Let me start with an example. Consider the question whether a proclivity for violent behaviour is inherited. In other words, does having a violent parent make one more likely to be violent? Let us assume, plausibly enough, that there is a statistical correlation: Violent behaviour by parents increases the probability of violent behaviour by children by x%, let us say. An intervention that reduced the violence of parents would very likely reduce the violence of children. Twin studies lead to the conclusion that having the same genotype as a violent person increases the probability that one will be violent. Physiologists tell us that exposure to violence raises cortisol levels, and that higher cortisol levels reduce the threshold for violent behaviour. And psychologists tell us that children tend to imitate the strategies of their parents for solving problems. We conclude that there are several causal pathways underlying the correlation between generations in tendency to violence. On the other hand there is no doubt that there is a great deal of variability among violent parents and among their children, and a great deal of diversity in forms of violent behaviour. And there are very likely opposing connections. Probably some children are so appalled by the violent behaviour of their parents that they are fully resolved never to imitate them. Perhaps others become extremely fearful through exposure to violence and learn to avoid confrontation at all costs. And so on.

My suggestion is that in such a case there is no general truth about the causal relation between having violent parents and becoming violent. There are various correlations with more or less consistency across different populations. It is not that these correlations are non-causal: they are the consequences of many causal connections. But the prevalence of the circumstances that give rise to these connections varies from one population to the next. I argued some years ago that the central general causal relation between two factors was that one of them (the cause) raised the average probability of the other (the effect) in a specific population (Dupré 1984); the problem of external validity is in general insoluble, though certainly we may have reasons to think populations are sufficiently similar that it is likely that the same causal relations will apply. I still endorse this position. There may be other relations, for example relations of lawlike connection, but in common with many philosophers of biology, I am somewhat sceptical about these in biological contexts. At any rate, relations of average effect seem to me the important ones in biological systems. If, on the other hand, we ask about the possible causal relations between individual events or other appropriate entities, I agree with Hitchcock (2003) that there are many such.

Some philosophers have found the idea of a causal relation that applied only within a specific population unsatisfying, "a sorry excuse for a causal concept" (Eells and Sober 1983, p. 54). The

alternative is to restrict populations in ways that ensure 'contextual unanimity' (Dupré 1984), ensuring that the cause has a constant effect throughout the range of a generalisation (Cartwright 1983; Eells 1991). I don't want to discuss this debate here, but rather look at a somewhat parallel response to the problem just raised about multiple mechanisms. Just as we can try to filter populations to remove some of the disorderliness of multiple often conflicting causal influences, so we can filter mechanisms to restrict our causal analysis to those that fit it. According to many exponents of the new mechanism, "mechanisms are functionally individuated by their phenomena" (Illari and Williamson 2012), an idea attributed to Stuart Glennan and sometimes referred to as Glennan's Law (Glennan 1996). Since the phenomenon of becoming violent has no unique mechanism, presumably it is not explicable. Perhaps this is right: there is no unique explanation for becoming violent. What about Johnny's becoming violent? This might be an instance of one of the several possible mechanisms, perhaps imitation of his violent parents, or it might be a combination of several such mechanisms.

But this is not a promising path to follow in, say, cell biology. We want an explanation, for example, of mitosis, the process of cell division, not an explanation of why one particular cell undergoes mitosis. So consider one crucial step in mitosis, the formation of the microtubules that constitute the mitotic spindle, which, in turn separates the two strands of DNA in the chromosomes of the dividing cell. This is, surely, a biological function. Yet it turns out that it does not individuate a mechanism. The reason for this is just that there are a number of different ways in which microtubule growth (or degradation) is directed; if one of these is prevented by some kind of intervention, then other mechanisms take up the slack. Spindle formulation continues, if more slowly (Duncan and Wakefield 2011). Moreover, and this seems very typical of constructive biological processes, the development of a spindle involves the upshot both of processes that extend microtubules, and of processes that decompose them. This combination of net effect and redundancy of process seems typical of many molecular processes in biological systems. And though it may be annoying to those who would like to trace *the* pathway or mechanism to an outcome, the resources that this redundancy of process provides first, for overcoming obstacles to any particular route to a particular state, and second, for responding rapidly to specific and variable surrounding conditions, seem highly adaptive. A similar and crucial example is found in studies of the expression of large numbers of genes in a cell progressing towards the specific expression pattern characteristic of a particular cell-type. These studies have shown that quite diverse trajectories of gene expression can nevertheless arrive at the same final state (Huang et al. 2005). The researchers who obtained this result describe cell fates as high-dimensional attractor states.

This suggests something teleological in the functioning of the cell, though not—I hasten to add—teleological in any mysterious or spooky sense. The problem with this kind of teleological system, or system with a robust tendency to end up in particular preferred states, is that the relation between functions and mechanisms is more complex than that supposed by standard mechanistic theories. A variety of mechanisms exist that contribute to the serving of a biological function, but none of them is necessary, and in many cases none is sufficient, to bring about the performance of the function. Of course we could at this point, in strict accordance with Glennan's law, move up to some higher level and say that this was where the relevant mechanism must be found. Whatever it is that assures the production of mitotic spindles by some route or other is the mechanism that causes spindle production. Now the trouble is that even if (which I'm quite unsure about) there must be such a higher level mechanism, this just isn't what scientists typically mean when they say they are investigating the mechanism(s) of spindle production: they are referring to the various processes that can, at different times in different contexts, contribute parts of spindles.

Let me return more specifically to the question of causation. What is the cause of mitosis? Unsurprisingly, this question is somewhat ambiguous. On one natural reading it asks, what initiates the sequence of cell changes that results in cell division? The answer to this remains obscure, though

we know many things that enhance or inhibit the tendency of cells to divide. For example, we know that cigarette-smoking has a tendency to cause a lot of unwanted mitosis in certain lung cells. To supplement this well-established statistical result in the manner of Russo and Williamson (2007), we would like to know how some constituent of cigarette smoke initiates or triggers the tendency of cells to divide excessively. In another sense of the question we might want to know the sequence of molecular events underlying the series of changes that constitute this part of the cell-cycle. Here we encounter the "mechanisms" such as those of spindle formation but, as just explained, there is no unique and definitive sequence of molecular events that fills this role. Finally we might be looking for the higher level control systems that generate the reliable production of the sequence of stages, including spindle production, that constitute this vital biological process.

When I have suggested to prominent new mechanists that "mechanism" is, for various reasons that I shall amplify further in a moment, a misleading term to apply to biological systems, they tend to reply that this is the term biologists always use, and we would do well to conform to their usage. So the following quote from researchers on the mitotic spindle is instructive:

> The traditional view of the mitotic spindle apparatus as a molecular machine which is built through a defined irreversible set of instructions is gradually being replaced. It can instead be envisaged as a self-regulating dynamic structure where multiple pathways of MT [microtubule] generation are spatially and temporally controlled and integrated, constantly "talking" to one another and modifying the behaviour of their MTs in order to maintain a flexible yet robust steady-state spindle. Through taking a holistic view, methodologically and conceptually, we can continue to learn more about how this fundamental biological process takes place. (Duncan and Wakefield 2011, p. 330).

The "pathways of MT generation" are the sorts of things that are more naturally described as mechanisms; the kind of process referred to in the penultimate sentence of the quote, even if we agree to call it a mechanism, is not the mechanism of spindle formation. Crucially, this latter process may very well be part of a much wider control system that regulates much more than the proper production of mitotic spindles; in which case once again we have failed to find a mechanism that satisfies Glennan's law.

V *From mechanism to process.* I suppose that it may well be possible to shoehorn descriptions of biological systems into talk of mechanisms if one is sufficiently determined. Indeed there is a serious danger of vacuity in some treatments of the topic, in which it seems that mechanisms just are whatever explains whatever happens. If the concept of a mechanism is to do any work we must surely have some sense of what isn't a mechanism or at least what might constitute an explanation that wasn't a mechanism. It seems to me that there are good reasons to think that biological systems—organism, cells, pathways, etc.—are in many ways quite misleadingly thought of as mechanisms. Paradigmatic machines—cars, dishwashers, computers—consist of a number of parts, typically more or less rigidly connected. The constituent parts gradually wear out, and the machine lasts as long as they are replaced piecemeal. Certainly we can relax the assumption of rigid connection, and extend the concept perhaps to the reactions in a vat in a chemical works. Here very large numbers of parts move around freely, and undergo predictable kinds of interactions when they encounter appropriate partners. We have here fairly clear representatives of the entities and activities demanded by new mechanists.

One thing that is added when we move to biological systems is that these, organisms for instance, constantly rebuild and replace their worn parts. Contrary to some versions of mechanism, there is no a priori reason why the process should end, and hence no terminal condition. Lineages of organisms have no mechanistically inbuilt tendency to terminate, though in the end no doubt most will do so. I

suppose that this could be accommodated with a relatively easy modification to standard accounts. However, the more such modifications are needed, the stronger the suspicion that the concept of mechanism may be heading for vacuity. And other necessary modifications strike me as increasingly difficult to accommodate.

In at least the paradigm cases of machines, we can very precisely identify the causal role of parts and the effects of manipulating or removing them, and their interactions with other entities are generally specific and limited. These characteristics are again inscribed in the details of recent accounts of mechanism: it is the specification of parts and their behaviours (activities) that provides the mechanistic explanation. However, in biological systems parts often have multiple roles and multiple causal interactions. The term 'moonlighting protein', introduced when it was first discovered that proteins might have more than one function, both names this phenomenon, and tellingly indicates the strength of the antecedent assumption that a biological entity should have exactly one function. Nowadays it is thought that a substantial proportion of proteins moonlight. Arguably this was an insight that was obscured by traditional mechanistic thinking. It is fascinating how this expansion of our picture of protein behaviour is tracking the earlier history of understandings of gene behaviour, during which phenomena such as alternative splicing gradually made it clear that particular genes might end up serving a wide range of cellular functions.

And of course it is obvious that organisms interact with other organisms in multiple ways: aggression, cooperation, communication (accurate and inaccurate) by chemical, visual, etc. means. Such interactions are not merely the activities of the largest biological entities, but they are also involved in the constitution of larger entities. For many microbial cells proper functioning requires integration in complex communities, such as biofilms, comprising a range of different kinds of cells and an elaborate division of labour. As cells join and leave these communities the boundaries between organisms shift or blur. And I have argued that the deep interdependence between multicellular organisms such as ourselves and the trillions of symbiotic microbes with which they coexist justifies us in considering the latter genuine parts of the whole multicellular entities in which they live (Dupré and O'Malley 2007, 2009).

This, I think, leads to the deepest problems with new mechanist accounts of biological systems and causality. All these accounts start with an inventory of entities, and though these may not be called 'things', they are pretty clearly conceived as a fairly stable inventory of fairly stable things. But the entities that form the hierarchy of biological ontology are not stable. They are, rather, stabilised over a very wide variety of timescales, and the processes of stabilisation are a fundamental part of the explanation of the activities of living systems. Living things are the explanandum in biological sciences at least as much as they are the explanans. What are stable and robust in biology are not things, but processes.

The idea that biology is fundamentally processual is hardly a new one. Evolution, whether or not it is required to make sense of everything else in biology, is surely a process. Development is a process, and those who have emphasised the importance of development to understanding evolution have argued persuasively that what evolves are not particular stages in development, but the whole life cycle. At the cellular and molecular levels it is even clearer that nothing stands still: a static cell is a dead cell. The maintenance of life processes at the higher levels (organisms) is sustained by countless lower level processes. It would in principle be possible to treat these processes merely as changing things. But in fact the entities we treat as things are typically very specifically stable, which is to say stabilised, stages in processes. The constant flow of material through a cell or an organism, a flow that includes many other elements that we are inclined to think of themselves as things—notably often very complex macromolecules or, for organisms, whole cells in the form of symbiotic microbes—seems much better understood in general as a hierarchy of processes than of things.

An obvious concern with the above line of argument is that we have a very poor philosophical grasp on what processes are. Worse still, it is widely feared that taking processes seriously as a basic ontological category will require reading Whitehead, who is rumoured to be unreadable and possibly unintelligible. (And there is no smoke without fire.) I agree that the proper characterisation of processes is a major challenge, and I don't propose to offer here any kind of detailed account here of what processes are. My main goal will rather be to show that a focus on processes will make much clearer the distinctions we need to think about biological causes and, therefore, that we have a real motivation to think seriously about process—even if it does mean reading Whitehead[6]. I will make just three preliminary comments on the general topic of process.

First, a plausible prima facie distinction between processes and things is that the former, but not the latter, require some kind of internal change to continue to be the processes that they are. A process, as mentioned above for the case of cells, needs activity to sustain it. One can imagine a rock undergoing no changes at all for, say, a minute without thereby ceasing to be a rock; a mouse in the same state of stasis is an ex-mouse. This point might also be thought of in relation to more classical analytic metaphysics. Those who think of objects as four-dimensional space-time worms often think of these as a sequence of stages, each of potentially infinitesimal temporal thickness (Lewis 1983). My suggestion is that while it is possible to have such a temporally thin time-slice of a rock, the same cannot be said of the mouse. Some degree of time-occupying process is necessary to make something a mouse. Similarly, if one prefers an ontology in which temporal stages are basic (e.g. Sider 2000), then if these are to be stages of organisms, they had better have some temporal thickness.

Second, the designation of an entity as a thing or a process is often best seen as relative to a time scale. A mountain is paradigmatically a thing when one is walking up its side; but for the geologist it may be part of a process of, for example, tectonic upheaval. In biological systems entities are continuously being both stabilised and disrupted. An entity stabilised over a particular time scale may, relative to that time scale, be considered a thing. A cell, for example, is maintained for a period of time by virtue of a number of highly energetic chemical processes. At longer time scales it can be seen as part of the processes that maintain an organ, and in which cell division and cell death (apoptosis) are regular features. Hence in different contexts it may be treated as a thing or a process. However, for the reasons indicated in the preceding paragraph, I hold that the latter is always a more fundamental perspective.[7]

Third, the reification of stages of processes can have serious, sometimes epistemically harmful, consequences. The paradigmatic instance must be the neo-Darwinian model of evolution so effectively promoted by Richard Dawkins, which effectively treats evolution as a sequence of genomes, with the relation between those genomes and their subsequent ontological trajectories as something that can safely be ignored from the point of view of evolutionary models. This model has, of course, been the subject of critique by evolutionary developmental biologists ("evo-devo") for many years. The correct upshot of this critique, as also argued from a slightly different perspective by developmental systems theorists, is that the only adequate conception of an evolutionary lineage is as a sequence of life cycles, which is to say processes. Of course the root of this erroneous reduction of evolutionary history to a sequence of genomes is in the false assumption that in all important respects the developmental trajectory is determined by the genome; and if in fact development was fully determined at some point in the life cycle it would in theory be possible to

[6] Whitehead (1929/1979) is the classic twentieth century work on process philosophy. But there are, of course, other sources. Nicholas Rescher's work (e.g. Rescher 1996) is clear and accessible, and a number of authors have offered interpretations of Whitehead. Daniel Nicholson has interestingly suggested to me that one does better, in this context, to attend to 20th century biologists, such as Waddington and Woodger, who were strongly influenced by Waddington, and interpreted his key idea for biological purposes.

[7] The points sketched in this paragraph are treated in more detail in Bapteste and Dupré (2012).

use this stage as a proxy for the whole cycle. But this would be a very special and fortuitous circumstance that would enable us to model evolution as if it were merely a sequence of static stages; and of course we now know that development is not determined so simply, but emerges from a sequence of interactions between biology and environment that continues throughout the life cycle.

So how does a processual perspective help us to understand biological causality? If biological systems consist of a hierarchy of processes, then we might plausibly start with the idea that causation involves the intervention in a process. As indicated above, this immediately points to some important distinctions we might make between kinds of biological causal intervention.

First we may distinguish interventions that disrupt a process. These would include the shootings, poisonings, stabbings and so on that have been the mainstay of most traditional philosophical discussions of causality: these are interventions that drastically disrupt the process that constitutes a living organism. Here there is a hierarchy of disruptions. A bullet to the heart prevents the circulation of blood, and thereby the provision of oxygen to cells throughout the body which, finally, suffers the disruption of multiple processes that require energy for their continuation. This is the kind of case where the provision of something like a mechanism does, in the way supposed by the committee on the safety of vaccines, obviate the necessity for some highly unethical randomised controlled trials. Investigations of process disruption are, obviously enough, central to the study of disease.

The idea of an intervention begins to limp a little when we move from the study of catastrophic injuries to the study of health, or of the normal maintenance of living processes. Given the temporal perspective on an organism I have urged, these might be called processes of development; though, as I shall discuss in a moment, at appropriate relatively short time scales development amounts to the maintenance of stable states. Consider again, what causes the formation of the mitotic spindle. Most cells undergo division at some point in their normal life history, which will require the production of the spindle, and we can ask what provides the signal to do this. We can also look at cases in which the normal timing of cell division is disrupted, as most obviously in the uncontrolled cell division that constitutes cancer. These latter—disruptive interventions again—are again cases where the combination of statistical correlation with mechanistic connection described by Russo and Williamson and by the Committee on the Safety of Vaccines look about right. We should, however, be prepared for very complex answers to such questions. Events in cells such as, for instance, the splicing of an RNA transcript at a particular site, are often determined as the consensus of a considerable number of interactions predisposing for or against. Hence many aspects of the state of the cell and its wider environment will contribute to the determination of whether the event takes place. Any of these various subtle influences, in turn, will dispose the cell towards the event, but perhaps quite slightly.

On the other hand we may be interested in the question how, given that the process of spindle formation is initiated, the end result is in fact achieved. As we have seen, this is another very complex question, since a variety of different mechanisms may be employed in different circumstances. Again, also, there is a hierarchy of processes of development to consider. Countless successful processes of cell division must take place if a multicellular organism is to develop. Not only is it likely that, as in the case of microlevel processes such as cell division, there will often be redundant paths leading to the same state, but also different final states will be reached as internal and external conditions are varied. It is possible and important to identify particular circumstances that may tend to direct the process towards particular states (as in the case of parental influences on the ontogeny of violence), but these will need to be understood against the background of many other variable states of the developing process.

Finally I would like to distinguish processes of stabilisation.[8] As I have already noted, stabilisation of biological entities at many different structural levels and time scales is of great importance. Consider the genome. Stabilisation of gene sequence is accomplished at short time scales by a range of molecular processes. Organisms repair broken DNA sequences, for example, by homologous recombination, and mechanisms exist to control the activities of retrotransposons, so-called jumping genes that are disposed to reposition themselves throughout the genome. Some genomes are quite dynamic, especially among genomes that experience a lot of lateral gene transfer, but such genomes will require proportionately more work to maintain their functional integrity. An extreme case is the bacterium *Deinococcus radiodurans,* which is able to withstand extremely high doses of radiation, something it accomplishes by having very rapid editing mechanisms and multiple copies of its genome. Genomes are also stabilised at a very different timescale by natural selection. The organism that is part of the stable background from the point of view of molecular scale repair processes is a transitory constituent of the selective process operating at an evolutionary time scales.

Stabilising causes are again of great significance from the point of view of medical science, as their disruption is likely to cause malfunction of the system. Indeed one might say that a core objective of medicine is to stabilise healthy systems; those who believe that ageing should be seen as a preventable illness will be even more concerned to understand these stabilising processes, though I certainly do not personally want to endorse that premise. The stabilising causes themselves, however, are perhaps the most difficult to see as interventions. They are parts of the basic processes that maintain the living system. Nonetheless, questions about how genomes, cells, or organisms sustain their integrity, sometimes to the extent that we can usefully treat them as stable things with constant properties, is surely a causal question and a very important one.

VI *Conclusion.* I do not mean this paper to be a systematic critique of the various views about causality that I have discussed. All of them, I think, point to important features of our causal language, a fact that confirms the hypothesis that we should be pluralists about causality. Nonetheless, I do also think that none of the main current perspectives on the subject gives a quite satisfactory way of approaching causality in the context of the hierarchies of processes which, I claim, make up living beings.

There are, of course process theories of causality (Salmon 1984; Dowe 2000) but these are firmly grounded in problems in physics, and in ways that greatly limit their applicability to issues in the life sciences. The main conclusion that I hope to have supported, then, is that there is great potential value to be expected from developing a better understanding of processes in general, and thereby an account of causation properly fitted to the context of the hierarchy of processes found in living systems.

Bibliography

Bapteste, Eric and John Dupré 2012: 'Towards a Processual Microbial Ontology'. Published online, November 2012. *Biology and Philosophy*.

Bechtel, William 2011: 'Mechanism and Biological Explanation'. *Philosophy of Science* 78, pp. 533-557.

Cartwright, Nancy 1989: *Nature's Capacities and their Measurement*. Oxford: Oxford University Press.

[8] See further Bapteste and Dupré (2012).

Cartwright, Nancy 2004: 'Causation: One Word, Many Things'. *Philosophy of Science* 71, pp. 805-19

Dawkins, Richard 1982: *The Extended Phenotype*. Oxford: Oxford University Press.

Doll, Richard and Austin Bradford Hill 1950: 'Smoking and Carcinoma of the Lung'. *BMJ* 2, pp. 739–748.

Dowe, Phil 2000: *Physical Causation*. New York: Cambridge University Press.

Duncan, T. & J. G. Wakefield 2011: '50 Ways to Build a Spindle: The Complexity of Microtubule Generation During Mitosis'. *Chromosome Research* 19, pp. 321–333.

Dupré, John 1984: 'Probabilistic Causality Emancipated'. *Midwest Studies in Philosophy* 9, pp. 169-175.

Dupré, John 1993: *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.

Dupré, John 2012: *Processes of Life: Essays in the Philosophy of Biology*. Oxford: Oxford University Press.

Dupré, John and Maureen A. O'Malley 2007: 'Metagenomics and Biological Ontology'. *Studies in the History and Philosophy of the Biological and Biomedical Sciences* 38: 834-846. (Reprinted in Dupré 2012, pp. 188-205.)

Dupré, John and Maureen A. O'Malley 2009: 'Varieties of Living Things: Life at the Intersection of Lineage and Metabolism'. *Philosophy and Theory in Biology*, 2009. (Reprinted in Dupré 2012, pp. 206-229.)

Eells, Ellery 1991: *Probabilistic Causality*. Cambridge: Cambridge University Press.

Eells, Ellery and Elliott Sober 1983: 'Probabilistic Causality and the Question of Transitivity'. *Philosophy of Science* 50, pp. 35-57.

Glennan, Stuart 1996: 'Mechanisms and the Nature of Causation'. *Erkenntnis* 44, pp. 49-71.

Godfrey-Smith, Peter 2010: 'Causal Pluralism'. In Helen Beebee, Christopher Hitchcock and Peter Menzies, (eds.), *Oxford Handbook of Causation*. Oxford: Oxford University Press. Pp. 326-337.

Hitchcock, Christopher 2001: 'A Tale of Two Effects'. *Philosophical Review* 110, pp. 361–396.

Hitchcock, Christopher 2003: 'Of Humean Bondage'. *British Journal of Philosophy of Science* 54, pp. 1-25.

Huang, Sui, Gabriel Eichler, Yaneer Bar-Yam, and Donald E. Ingber 2005: 'Cell Fates as High-Dimensional Attractor States of a Complex Gene Regulatory Network'. *Physical Review Letters* 94, p. 128701

Hull, David L. 1974: *The Philosophy of Biological Science*. Englewood Cliffs, NJ: Prentice-Hall.

Illari, Phyllis M. and Jon Williamson 2012: 'What is a Mechanism: Thinking about Mechanisms Across the Sciences'. *European Journal for Philosophy of Science* 2, pp. 119-135.

Kitcher, Philip 1984: '1953 and All That. A Tale of Two Sciences". *The Philosophical Review* 43, pp. 335-371.

Lewis, David K. 1973: 'Causation'. *Journal of Philosophy* 70, pp. 556-67.

Lewis, David K. 1983: 'Survival and Identity'. In *Philosophical Papers, Vol. 1*. Oxford: Oxford University Press, pp. 55-77

Machamer, Peter K., Lindley Darden and Carl F. Craver 2000: 'Thinking about Mechanisms'. *Philosophy of Science* 67, pp. 1–25.

Rescher, Nicholas 1996: *Process Metaphysics: An Introduction to Process Philosophy*. Albany: SUNY Press.

Russo, Federica and Jon Williamson 2007: 'Interpreting Causality in the Health Sciences'. *International Studies in the Philosophy of Science* 21, pp. 157—170.

Mumford, Stephen and Rani Lill Anjum 2011: *Getting Causes from Powers*. Oxford: Oxford University Press.

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

Sider, Theodore 2001: *Four-Dimensionalism*. Oxford: Oxford University Press 2001.

US Institute of Medicine Report 2011: 'Adverse Effects of Vaccines: Evidence and Causality'. http://www.iom.edu/Reports/2011/Adverse-Effects-of-Vaccines-Evidence-and-Causality.aspx

Whitehead, Alfred North (1979/1929): *Process and Reality: An Essay in Cosmology.* Corrected edition, edited by D. R. Griffin and D. W. Sherburne. New York: The Free Press.

Williamson, Jon 2006: 'Causal Pluralism versus Epistemic Causality'. *Philosophica* 77, pp. 69-96

Woodward, James 2003: *Making Things Happen: a Theory of Causal Explanation*. New York: Oxford University Press.