



Causally powerful processes

John Dupré¹

Received: 4 February 2021 / Accepted: 11 June 2021 / Published online: 29 June 2021
© The Author(s) 2021

Abstract

Processes produce changes: rivers erode their banks and thunderstorms cause floods. If I am right that organisms are a kind of process, then the causally efficacious behaviours of organisms are also examples of processes producing change. In this paper I shall try to articulate a view of how we should think of causation within a broadly processual ontology of the living world. Specifically, I shall argue that causation, at least in a central class of cases, is the interaction of processes, that such causation is the exercise of a capacity inherent in that process and, negatively, that causation should not be understood as the instantiation of universal laws. The approach I describe has substantial similarities with the process causality articulated by Wesley Salmon and Phil Dowe for physical causation, making it plausible that the basic approach can be applied equally to the non-living world. It is an approach that builds at crucial points on the criticisms of determinism and universal causality famously articulated by Elizabeth Anscombe.

Keywords Causation · Causal process · Capacity · Process ontology · Process biology · Living systems · Organism

1 Introduction

Causation, as Augustine famously remarked of time, is something I seem to understand very well until someone asks me to explain it. I experience it in my own actions. I swing an axe at a piece of wood and split it in two. Did I cause the wood

I am very grateful to Gabe Dupré for comments on an earlier draft that have led to numerous improvements, and to two anonymous referees for a number of helpful suggestions.

This article belongs to the topical collection "Causality and Determination, Powers and Agency: Anscombean Perspectives", edited by Jesse M. Mulder, Dawa Ometto, Niels vanMiltenburg, and Thomas Müller.

✉ John Dupré
j.a.dupre@exeter.ac.uk

¹ Egenis, University of Exeter, Exeter, UK

to be split? A slightly odd question perhaps, but one to which the answer is surely, yes. Similarly, I see a tree fall on my garage. Did the falling tree cause the garage to be destroyed? Again, surely yes. In these cases, I experience or perceive something (the log, the garage) being changed by something else (me, the tree). I know it when I see it.

But of course there is more. I don't split the log merely by bumping into it or passing close by. So is the cause, strictly speaking, my swinging the axe in a certain way, an event, I suppose, rather than me? And if the event caused the splitting, and I'm surely not identical to the swinging, then can it be true after all that I caused the splitting? Perhaps I caused the event, and the event caused the splitting? But surely I am part of the axe-swinging event, and if so can I also be the cause of it?

If these questions can seem puzzling, one reason, I suggest, is that whereas we are very familiar with the idea that things can, in some sense, cause stuff to happen, we have a conception of a thing that makes this surprisingly problematic. The conception I have in mind is basically passive. Things may have many dispositions to act, or behave, but dispositions need to be triggered. The default state, absent the triggering of any dispositions, is stasis. So I may have a disposition to swing axes, but what triggered the disposition, what caused me to do it at the particular time that I did? And isn't this triggering event the more genuinely explanatory cause? To do something an entity must undergo some kind of change, but it has long been recognised that it is a serious philosophical challenge to understand how a thing can change and continue to exist.¹

This very sketchy background is intended to make sense of what I propose to do in this chapter. I believe that the world is not composed of static things, but rather of more or less dynamic processes. What we think of as things, or substances, are in fact processes, actively stabilised over particular time scales. I have only attempted to argue for this in any detail with respect to living beings (Dupré, 2020; Dupré & Nicholson, 2018), but I strongly suspect that this is the right way also to think of the non-living world. If the latter claim is mistaken, then what I shall say is limited in application to the living; that is, anyhow, a very important part of the causal network in which we find ourselves.

For processes, as I think of them, change is unceasing. Stasis is death or annihilation. Consider a simple paradigm, an eddy in a river. The eddy just is a pattern of movement of parts of the water, and if the water ceases to move, the eddy ceases to exist. Moreover, I take it that the eddy is a persistent entity. As long as the water continues to move in the right way, the eddy persists. This is just the way to think of an organism, I maintain. It is a pattern in the flow of matter, maintained by countless movements and interactions of its material parts. Movement is change, so change underlies the existence of everything. As substance metaphysicians sometimes wonder why there is something rather than nothing, a process metaphysician might ask why there is change rather than stasis. But that will not be my question today. If

¹ For a survey of the history of this problem see Mortensen (2020). For an argument that it is ultimately insoluble within traditional conceptions of substance, and that we need to replace this with a process ontology, see Meincke (2019).

anything is primitive and beyond explanation, for the process metaphysician it is change.

In a world of process, the problem of causation is that of understanding the interactions of processes. This is, indeed, exactly what existing process theories of causation, notably by Salmon (1984) and Dowe (1992, 2000), have attempted to do. Unfortunately, their concerns were explicitly limited to physics, and so of limited relevance to the biological issues that are my main concern. Salmon's starting point is the concept of a causal process, one that transmits its own structure. This aligns very well with the idea of a persistent process which is central to my own position here. However, there are important respects in which the development of the idea in a physical context diverges from what is needed for a biological account. For example, much of Salmon's exposition uses as a foil the idea of a "pseudoprocess", such as a spot of light moving across the roof of a building generated by a moving light source. Unlike a genuine process, a pseudo process cannot transmit a mark. Nothing you do to the light spot on the ceiling at one point will change its state at a later point. But if there are pseudoprocesses in biology (one may indeed be afraid of shadows, but that is often for the good reason that dangerous animals cast shadows) they are not the major problem in distinguishing biological processes. The more pressing problem is that of distinguishing biological individuals from the biological processes in which they are embedded (see, e.g., several essays in Guay & Pradeu, 2016); but this is not because any part of this processual nexus is illusory. In the context of causality correct identifications of individuality may be crucial for finding the right entities over which to formulate causal generalisations, but is seldom necessary for identifying the occurrence of causal activity. If I observe a tiger sinking its teeth into the neck of an antelope, I do not doubt that there are two causal processes in my vision (more or less *sensu* Salmon) and that they have interacted. For some purposes it is important to decide whether the gut microbiome of the tiger is part of the tiger process, or another one or many engaged in a long-term interaction with it. But from the point of view of the observed causality just described this question is irrelevant. Following Anscombe (1975), I shall assume that a lot of causality is immediately observable.²

Much of the metaphysical background of my account will be taken for granted. I assume without much argument that biological individuals are persistent causal processes, and that their interactions with one another and their environments are causal. By sketching some further features of this terrain, I hope to show that this is a productive way of thinking about causality, at least in biology. One thing that I think should be clear from this sketch is that much of what Anscombe argued for in her famous inaugural address makes (even more) sense in a world of causally powerful processes.

² Following Michotte (1963), there is a body of empirical work confirming the perceivability of causation. Needless to say, the philosophical interpretation of such work will depend on one's understanding of causation, and is not entirely straightforward.

2 Locating process theories in the metaphysics of causation

We have moved a long way from the orthodoxy Anscombe addresses in her inaugural lecture. First, probabilistic causality is widely acknowledged, obviously problematizing the necessity and universality that Anscombe saw as defended by her contemporaries and the Humean tradition. Second, and following from the denial of universality, it has become much more respectable to attribute causality to a single, perhaps unique occasion. These changes no doubt owe a lot to the influence of Anscombe's work. The ideas presented in this paper will be broadly congenial to Anscombe's important insights on causality. Like Anscombe, I am deeply sceptical of universality and indeed, as is now common among philosophers of biology, I am even sceptical of the laws in which universality has traditionally been inscribed. I agree that causation is something all around us, the visibility of which is only concealed by the grip of dogma. This becomes all the more obvious in the light of a metaphysical thesis I have advocated for a number of years, that our world is a world of processes rather than of substances. However it is thought that substances act, change, or do anything, activity is not at the heart of the concept of a substance. The default state of a substance is stasis. A process, on the other hand, ceases to exist if it does nothing—in fact, I take this to be the defining characteristic of a process. And to do something is to change, either oneself or some other part of the world; in short it is to be causal. There is no storm without the movement of something, even if the environment is so well battened down that only the wind and rain themselves are moving. When even these are quiet, there is no longer a storm.

What is this orthodoxy Anscombe rejects as Humean? Hume's famous ideas about causation begin with the rejection of metaphysical causal necessity, and are based on the idea that we never see the necessary connection between cause and effect; we just see the succession, for example of the impact of one billiard ball on another and the movement of the second. If the connection cannot be found in the interaction between the two, then it must be sought elsewhere. And the place Hume proposed, a solution that continues to attract many adherents, is in the regularity of the connection between events similar to these two. It identifies causation with universality—for A to cause B is for there to be universal regularity that whenever A, then B—and necessity. Although Hume is famous for his scepticism about metaphysical necessity—it is always possible for A to be followed by not-B—he was quite clear that necessity was involved in the meaning of causality, and he felt obliged to provide a psychological explanation of this, specifically via the tendency of regularities of experience to lead to ever stronger expectations of regularity. In the Humean tradition philosophers such as Mill (1843) and, roughly contemporaneously with Anscombe, Mackie (1974), laboured to describe the universal generalisations that made our everyday causal claims true. Also around the same time Lewis (1974) took up another aspect of the Humean corpus, and offered an account of causation in terms of counterfactual generalisations. More recently, Woodward's (2005) influential difference-making account illustrates one of the increasingly sophisticated moves that have been made in a broadly Humean direction.

The work of Mill and Mackie just noted, and of which Mackie's account of an individual cause as an *inus* condition—an insufficient but necessary part of an unnecessary but sufficient condition—is perhaps the most refined exemplification, elaborates the implications of a cause as a necessary and sufficient condition of its effect under the description that accommodates it to universal law. But perhaps the greatest insight of Anscombe's paper is that while the necessity of a cause for its effect is an intuitive idea³ (the cause as that without which the effect would not have happened), sufficiency is a far more dubious philosophical add-on. When I say that he would not have crashed the car had he not been drunk (his intoxication was a necessary condition for the crash), I certainly do not assume that the crash was a done deal the moment he started driving. In the search for sufficiency a whole set of additional factors are hypothesised as together bringing about the inexorable disaster. But the assumption that there is such a complete set of factors seems more a philosophical prejudice than anything implied by the original causal assertion.

One way of accepting the rejection of causal sufficiency is with a move to causality as probabilistic. In the first instance this takes the form of the proposal that a cause doesn't necessitate its effect, but merely raises it probability. Smoking makes contracting lung cancer much more probable, but many lucky smokers escape this fate. But this may seem to fit easily enough into a modified Humean frame. Suppose just that for A to cause B probabilistically is for there to be a law of nature that some proportion of As be followed by B. If, given sufficient specification of the conditions, this probability can be made precise, then a certain kind of universality is even maintained, if only at the level of the population. Better still, it is often supposed that a probabilistic generalisation is merely one in which not all the relevant factors have been specified; when they have, the probability will be 1 or 0. The probability is not a contradiction of the deterministic, law-governed world, but merely a measure of our ignorance.⁴

There has, at any rate, been a tendency, at least among naturalistically inclined philosophers, to move towards positions on the anti-Humean side of this dichotomy, a tendency that has been motivated in part by the perceived decline in importance of scientific laws. Highly influential work by Cartwright (1983) has led many to doubt whether even physics could really generate true universal laws. With regard to probabilistic causality, fully general laws of the kind postulated in the last paragraph, x% of As cause Bs, have proved difficult to characterise adequately. It has become increasingly attractive to treat such non-deterministic causality as merely the outcome of capacities (As can cause Bs), made more or less likely to be exercised by the presence of further interacting conditions. And then why should there not be capacities that are just exercised sometimes but not at other times, but for which the frequency of exercise has no particular tendency to converge on a fixed value (Dupré, 1993, ch. 9; Lewontin, 1966)? While not many philosophers have

³ Though especially following Lewis (1974), it has been widely noted that causal necessity is problematized by the possibility of overdetermination. I shall say a bit more about this later in the paper.

⁴ For an argument that a Humean account of probabilistic causality is not, in the end, possible, see Dupré and Cartwright (1988).

enthusiastically embraced this last possibility, philosophers of biology have at least been increasingly ready to agree that there may be no universal laws in the life sciences (e.g. Beatty, 1995; Brandon, 1997).⁵

We should anyhow look more closely at another way forward that Anscombe does not explicitly discuss, the rejection of Hume's strictures on causal powers, capacities and so on. This is an idea that has been promoted from time to time, and increasingly of late; heretics include Harré and Madden (1975), Cartwright (1989), and Mumford and Anjum (2011). Despite important differences, all these thinkers see causation as the exercise of capacities grounded in the properties of individual entities. They also generally see causality, as did Anscombe, as something readily observable in the world, and referred to by countless causal verbs (push, pull, throw, paint, cut, sharpen, infect, vaccinate, frighten, educate, etc.), adjectives (pushy, attractive, sharp, contagious, frightening, etc.) and nouns referring to things with these capacities (tractor, knife, virus, scarecrow, vaccine, etc.). These seem, on the face of it, to be observable causal activities, properties and entities in the world. For these dissident philosophers, Hume's refusal to see causality has profoundly misrepresented a world which is in fact seething with causality for all to see. As Anscombe remarks, the Humean is "just not going to count anything as 'observation of causality'" (1975, p. 68).

Accounts of this anti-Humean kind have evolved in recent years in two directions. One that has been very widely discussed in the last 20 years focuses on mechanisms, (Craver, 2007; Glennan, 1996, 2017; Machamer et al., 2000). Ignoring a lot of detailed disagreement, mechanisms in this picture are arrangements of powerful entities that interact in ways that produce some kind of output. Mechanisms are proposed in the first instance as providers of explanations. The contraction of a muscle, for instance, is *explained* by showing how the arrangements, connections and activities of muscle cells jointly generate this effect. But of course, this is not just a once off explanation for a particular contraction. For as long as the right arrangement and connection of these cells survives, the cell will retain the capacity to contract. Mechanisms are generally explanations of causal capacities (and hence strongly anti-Humean).

Actual machines make the point easier to illustrate. Why is all the grass on my lawn roughly the same length? Because I ran my lawnmower over it. My lawnmower has the capacity to cause grass to be cut to a uniform length, and it produced just this effect on my lawn. Why does it have this capacity? Because its parts, with their own capacities—cutting, pushing, igniting, and so on—are connected together in a way designed to bring about this effect, specifically by generating the rapid movement of blades at a fixed distance above the ground. There is nothing to prevent someone from insisting that there are universal laws instantiated by the interactions of lawnmowers and lawns, but once the existence of such capacities has been conceded, there is no strong motivation for such an insistence.

⁵ This is also related to a move, especially in the life sciences, from the syntactic to the semantic view of scientific theories, the analysis of theories in terms of models rather than proposition (see, e.g., van Fraassen, 1980; Lloyd, 1988).

In a different vein I have been quite critical of mechanisms (Dupré, 2013; Dupré & Nicholson, 2018; see also Nicholson, 2012, 2013). However, I do not deny that talk of mechanisms has an important and often valuable place in the biological sciences, and this place has been illuminated by mechanistically-minded philosophers. My reservations arise at the level of metaphysics: mechanism should not be taken as providing an adequate image of what the living world is really like. Mechanisms naturally present a picture that is (a) strictly bottom-up: the properties of parts explain the behaviour of the whole; and is (b) composed of fixed, thing-like entities, as in a machine.⁶ Both these elements are generally false of biological systems. The parts are not fixed, but are constantly being created, modified or destroyed; and the properties and behaviour of the parts are often determined to some degree by the system as a whole. However, much biological explanation certainly does involve showing how parts—stable on the time scale relevant to the explanation—interact to produce the behaviour of larger scale entities. I shall sometimes refer to these as mechanisms, or mechanistic, but it should be noted that I do not thereby endorse substantial parts of (influential versions of) the contemporary mechanist programme. I take the mechanisms appealed to in explanations to be abstractions from the flow of living process; they abstract not only from the full range of causal influences, but also from the changing and always temporary natures of the parts that they do include.

Another anti-Humean tradition in recent decades has grounded causes in processes. This was initiated by Salmon (1984) and has been developed especially by Dowe (1992, 2000). These philosophers see causation as manifested in causal processes and interactions between causal processes. For Salmon, causal processes are defined by the ability to transmit a mark (thereby being distinguished from pseudo-processes such as moving shadows) and interacting processes may leave marks on one another. Dowe criticises this account on several grounds, and proposes instead that a causal process “is the world line of an object that manifests a conserved quantity” (1992, p. 201) such as mass-energy, linear momentum or charge; an interaction involves the exchange of such a quantity. Although my aim in this paper is to present a process-based account of causation, because of the focus of Salmon and Dowe on physics, I shall not dwell on details of this work beyond endorsing a metaphysics of causation grounded in processes and their interactions. Dowe (1992) actually admits that it would be difficult or impossible to apply his account outside physics but assumes that other sciences will at least supervene on the physical. Even if this were so, which I see no reason to believe (Dupré, 2010, pp. 44–5), it would still leave us needing to account for the apparently causal processes and interactions that abound throughout the life sciences.

⁶ It will fairly be objected that these theses present a crude picture of mechanism that ignores the subtleties that mechanists have developed to distance themselves from them. The most recent statements of the mechanist position, indeed, have much in common with the processualism I advocate. One concern is that the further mechanists move from these common associations with the view, the more it seems able to accommodate any empirical data, and the more vacuous it is in danger of becoming (Dupré, 2013). I should be clear that my position is at any rate much closer to mechanism than to the law-based accounts of biology that preceded it.

The processes that interest me particularly are persisting, or continuant, processes.⁷ A paradigm might be a storm or a horse. This categorization of a horse is of course controversial, but as I have argued for it extensively elsewhere (e.g. Dupré & Nicholson, 2018), I shall be fairly brief here. The storm and the horse are both processes, in my view, because their persistence is achieved through activity. For the storm this is the activity of circulating winds, driven by thermal energy acquired from the atmosphere. That storms persist is clear from the fact that they sometimes even have names; Hurricane Laura, for instance, visited Haiti, Cuba, Florida and Louisiana on its travels in August, 2020. For Salmon, causal processes transmit their own structure, and this is certainly a distinctive feature of many biological entities. The horse maintains its integrity through a host of metabolic activities ultimately fed by the food and oxygen the animal eats and breathes. A striking feature of processes such as organisms is that they not only include activities that maintain specific features, such as the *homeostatic* maintenance of temperature, but they also maintain characteristic trajectories, including the overall developmental course of the system. Waddington (1940) described the maintenance of such a trajectory as *homeorhetic*. A further crucial point is that it must not be assumed that the activities that stabilise such a process, whether homeostatically or homeorhetically, are all internal. In fact, organisms are deeply intertwined with many features of their biotic and their abiotic environments, and this is vital to their homeostatic and homeorhetic stability. The central point to emphasise about the persistence of a process is that it persists not by default, as is commonly supposed of persisting substances, or things, but as an achievement, a consequence of activities within or around it. I am highly doubtful whether there are any things that persist in the default way. Even the most seemingly inert objects, diamonds or lumps of gold, say, would fall apart if there were not interactions between atoms, and interactions between the subatomic particles that make up the atoms. Whether the very lowest level, if indeed there is one, is best thought of as processual is a question beyond the scope of this paper. For now I insist only that in biology nothing persists except through hard work.

So to summarise so far, I take organisms among other biological systems to be causal processes in a sense quite similar to that described by Salmon and Dowe. The persistence of such processes is made possible by many further causal processes and interactions both between parts of the system and between the system and aspects of its environment. The biological processes I am interested in have numerous causal capacities. We generally expect to be able to explain these capacities in terms of

⁷ There is a good deal of metaphysical controversy that I am passing over here. A traditional view divides persistent entities into continuants, which are said to be wholly present at all times at which they exist (they “endure”), and occurrents, which are made up of a sequence of temporal parts (they “perdure”). Processes are taken to be paradigms of occurrents. I am generally unhappy with this distinction, in part because I find it hard to understand what it is (and what it is not) to be wholly present. My own view is that all persisting entities have temporal parts, but that these are all in the past and present, not the future. Some philosophers take this view to be inconsistent with the physics of time. The categorisation of processes as occurrents is also being debated. Stout (2016) believes that processes are both occurrents and continuants. Steward (2015) holds that some processes persist, but do so in a way that is distinct both from perdurance and endurance.

interactions either among their parts or between them and other entities. An important subset of such explanations consists of those that have been described as mechanistic. I take all this to be no more than description of the place of causal language in our accounts of the biological world. But this description places the process view of causality squarely in the anti-Humean camp, and makes it clear that that universality is at best an optional part of such an account.

3 Kinds of process

This is a good point at which to make clear that a simple unitary answer to the question what is causation is unlikely to be forthcoming. Elsewhere (Dupré, 2013) I have defended this view in more detail, arguing that many or all of the various accounts of causation mentioned above have distinct uses. Here I am concerned more narrowly with the causality in biology, and more specifically still with the way that causality is inherent in, or emerges from, persistent biological processes. But even here there is an important element of pluralism, reflecting the different kinds of process that can be found in living systems. In particular, I have in mind homeorhetic processes, homeostatic processes and processes without any particular stabilisation. Homeorhetic systems are perhaps the most characteristic in biology. They are exemplified, for example, by organisms.⁸ On a much smaller scale, countless homeorhetic processes occur within the organism. Think, for instance, of the sequence of events that lead through transcription, various kinds of splicing, translation and protein folding to a mature protein molecule. This, incidentally, is neither a process that happens to something (there is no thing to which it happens) nor is it the kind of process that might have been taken for a thing or substance. It is just an “unowned” process, strongly homeorhetically constrained.

The difference between homeostasis and homeorhesis may in some cases be just a matter of time scale. A reasonably long-lived adult animal is naturally treated as a homeostatic system for the purposes of its day-to-day activities and this is also the appropriate perspective for much medical science. Strictly speaking it is, of course, continuing to develop, or age, but this is of no significance for understanding its current feeding, sleeping, grooming, or whatever behaviour. More strictly homeostatic is something like the mature cardiovascular system of an animal. It will, of course, eventually fail. But this is probably not because it has features the function of which is to make it fail. Like everything else, it will eventually wear out, though the extent

⁸ One might worry that this is not the case for by far the majority of organisms, the unicellular ones. But there is, of course, the familiar cell cycle leading to and through the division of the cell. This is sometimes thought of as an oscillation that maintains a stable system. And in fact microbial cells age (Moger-Reiscer and Lennon, 2019). A condition of this is asymmetric reproduction, as if reproduction were symmetrical, all the cells would be the same age, and ageing of the cell could not be distinguished from ageing of the lineage. But, in fact, microbial cells do divide asymmetrically, and the mother cell tends to monopolise accumulated damage, so that eventually it no longer functions. This incidentally suggests a very general explanation of ageing as accumulation of damage balanced against reproductive rate.

to which biological failure is inevitable or perhaps, on the contrary, functional, is a matter of some debate.

Finally, there are countless biological processes that are neither homeostatic nor homeorhetic. Ecology and epidemiology provide good examples. Think of the development of a viral epidemic. It will shrink or grow depending on all kinds of extrinsic contingencies concerning the behaviour of host organisms and the states of their immune systems, as well as on its own random mutations. It is also possible that some processes might be classified as homeorhetic or not depending how loosely they are described. Most of vascular development appears not to follow any definite path, but the end result is typically a system that provides blood supply to the whole organism. Indeed, development generally can be seen as involving a homeorhetic core with a periphery of increasingly opportunistic or even random processes.

With these categories in mind we can look at various kinds of intersection between causal processes. One important question is whether the intersection disrupts the homeostasis or homeorhesis of either or both of the processes. The violent interactions that feature so prominently in the philosophical literature on causation tend to do this. When Jane shoots Jim, Jim will often cease to be either a homeorhetic or a homeostatic system. Similarly with the predatory interactions that form a growing proportion of David Attenborough television shows.⁹ But most interactions between or affecting organism-processes fall into one of two other categories.

One very important category comprises those interactions that exactly fall within the functional ranges of the homeostatic systems. A good example is the interaction between an organism and a microbial infection. The microbes are of course just doing what they do, replicating themselves in their natural environment. A vital homeostatic system in the host is the immune system, one function of which is to prevent or limit the damage caused by such an interaction. If all goes reasonably well a small part of the microbial lineage passes through the host and the host suffers a manageable amount of inconvenience or damage. These effects, however severe, are said to be caused by the pathogen. During the interaction, many individual microbial cells will be destroyed by the immune system, suffering a fate similar to Jim's above.

When the stabilising processes work perfectly, we tend not to think of a causal interaction at all. I don't notice the small rise in ambient temperature because I adjust for it by perspiring slightly more. It would, however, be true to say that the rise in temperature caused the increase in perspiration rate. This just points to the familiar fact that the stabilising processes work through multiple smaller scale causal interactions. Equally important in this category are the cases in which interactions between organisms just *are* parts of the stabilising system. Here I have in mind the omnipresence of symbiosis. Various chemical activities by bacteria in the cow's stomachs enable the cow to digest the cellulose in grass. Treating the bacterium and the cow as independent processes, this implies some kind of interaction between the two. But regardless of the outcome of the contemporary debate whether these really are

⁹ For non UK-based readers, Attenborough has been a presenter of outstanding nature documentaries for more than half a century, but with a possibly growing tendency to focus on predation and male sexual displays.

independent individuals or one (the so-called holobiont), this discussion will be better subsumed under a general consideration of internal processes below.

Much causal interaction between individual organisms belongs in the opportunistic or, from a biological point of view, random categories. The latter is simple enough. A coconut falls on my head and knocks me out. The impact of the coconut caused me to be unconscious. With luck, my stabilising processes will kick in and restore proper function without lasting damage. This is, no doubt, what they were designed to do, though it is unlikely that the evolutionary history of this design had much to do with coconuts.

Many opportunistic interactions are more interesting and directed. Organisms move around the world encountering more or less unexpected items. They may interact with such an item in numerous ways: eat it; mate with it; flee from it; use it to construct a shelter; play with it; and so on. These causal interactions may contribute to the stability or instability of the interacting participants, and may be more or less core parts of the stabilising processes. The capacity to interact in appropriate ways such as those mentioned above with other organisms or inanimate objects, either deliberately sought or accidentally encountered, is a characteristic feature of all living systems. The more specific capacities deployed in these interactions are central aspects of the natural history of particular organisms.

An organism, I argue, is a complex process that maintains a more or less stable complex structure by virtue of the internal metabolic processes that maintain its thermodynamic equilibrium. Parts of this structure, such as teeth, cognitive abilities to outwit the evasive strategies of prey, and suitable digestive systems, are also necessary conditions for the maintenance of the metabolic stabilisation of the overall system, but their necessity derives from their role in making possible essential interactions with other organisms. Plants do things very differently. But their ability to respond developmentally to experienced environmental features is essentially analogous. Indeed, the plant meristem is a paradigm example of a flexible and responsive process, developing into leaves, flowers, roots, and so on in adaptive response to its environment. A meristem may, for example, start to produce roots in response to its encounter with the soil (itself a massively complex process).

So far I have shown how, if one accepts the process view of biological systems, something quite similar to the Salmon/Dowe picture of causation as intersecting processes extends naturally to a wide range of biological examples. An organism is, more or less, the world line of a process that manifests and transmits the quality of being alive. The quality is not conserved in the way that physics describes mass/energy or momentum as being conserved, but conserving it is, nonetheless, exactly what organisms do. Organisms, in particular, bristle with capacities to interact with entities in their environments in ways that conduce to their staying alive. But they may or may not do so when they encounter suitable targets. Following Anscombe, there is no reason to suppose that these interactions instantiate laws either as instances of universal generalisations or even of statistical generalisations with determinate probabilities.

Organisms are profoundly causal processes: their persistence manifests a myriad internal causal interactions, and equally depends on constant causal interactions with entities beyond their boundaries. But there are many questions generally taken

to be causal that do not fit so easily under the rubric just outlined. In particular, we may ask either, why does this entity (process) have this capacity, and also, why does it exercise the capacity on this occasion? Let me turn now to the first of these. These are the kinds of questions most naturally emphasised by new mechanists. How can the wolf track the passage of a potential prey animal? It has sensors in its nose and an arrangement of neurons in its brain that enable it to distinguish minute quantities of distinctive odoriferous chemicals shed by the relevant prey species. This is the kind of explanation that is generally claimed as mechanistic, or at least as a mechanism sketch. Although a capacity of an organism will be sustained by internal causal processes, these are not generally the same processes that occur when the capacity is exercised.

Having said that a process ontology can readily accommodate much of the explanatory work that has been assimilated to the mechanist programme, I should make clear that there is no difficulty in reconciling a process ontology with the undoubted successes of the mechanistic explanatory strategy. Most mechanistic philosophers, following the classic paper by Machamer et al. (2000), distinguish activities from the entities that undergo them. The entities are generally understood as things, or substances. But if the entities are themselves processes, all that is required is that they be sufficiently stable over the time scale of the behaviour being explained to play the assigned role in the mechanism. A key feature of biological processes is that they are stabilised over some period of time. The appropriate period of time determines, and in a functional or evolutionary sense is no doubt equally determined by, precisely the kinds of mechanisms in which they participate. Of course, it is possible to go wrong by assuming more stability than is actually provided. One may, for instance, wrongly extrapolate the results of medical experiments to older or younger age groups. But it is certainly not necessary to make such mistakes.

A more subtle concern about mechanistic explanations is that the processes that stabilise the entities in an explanation are often higher-level processes, instances of “downward” causation. It might then appear that mechanistic explanations are incomplete. In response to this, it seems to me enough to insist that it is unimaginable that there are any complete explanations. I shall, however, have a little more to say later about downward causation.

What of the second question, why does an organism exercise a capacity on a particular occasion? A fox comes round a corner and sees a rabbit. Suppose we have established that it has the right kind of teeth and digestive system to consume rabbits, and that it can run fast enough to (sometimes) catch them. Suppose even that we have full explanations of the grounds of these capacities. Should that, in principle, tell us what happens at this crucial intersection of processes? Surely not. The rabbit almost certainly attempts to escape as soon as it notices the fox. But how soon does it notice the fox? Perhaps its attention is elsewhere. The fox may make an estimate, accurate or inaccurate, of the likelihood of catching the rabbit given their relative positions. It may be hungry, a bit hungry, or not hungry at all. Perhaps it is preoccupied with a potential mate or it has a plan for a more easily and reliably acquired meal. And so on. A complex and intelligent organism like a fox reaches a decision based on a wide range of variable factors that may be salient at the time. There is no algorithm in the fox’s brain that

somehow determines this decision by feeding in a determinate set of parameters, and absent a pre-existing commitment to determinism, no reason to suppose that there is any set of factors either necessary or sufficient for the decision.

It is an attractive supposition that the ability to respond appropriately to complex and partially unprecedented situations is a characteristic feature of living systems, and increasingly characteristic of increasingly complex and sophisticated organisms. So to the question why did the organism exercise a particular capacity on a particular occasion there may be several answers, and there may be no answer that is fully determining. The fact that it was hungry, the rabbit was looking the other way, would be cornered if it ran directly from the fox, etc., may all partly explain the fox's decision to pursue the rabbit, but even jointly may not determine it. And in a rather different vein, the bramble stem's contact with the soil, the resources available to the plant, the dampness of the soil and perhaps other chemical components of the soil, may all contribute to the plant's "decision" to put out a root. But why should we assume there must always have been a sufficient condition when such an action occurs?

It is a familiar observation that we may provide either a "proximate" or an "ultimate" explanation of a feature of a biological system (Mayr, 1993). The mechanistic explanations just discussed are proximate, and are thoroughly causal. Are ultimate explanations, explanations in terms of evolutionary history, also causal? Think of the relationship between the organism and the lineage to which it belongs. Clearly both depend on the other. The lineage generates organisms and the organisms perpetuate the lineage. Moreover, as Darwin famously emphasised, by producing a huge excess of organisms, more, that is, than the environment can sustain, the lineage can "choose" which of its varying productions will perpetuate the lineage. The lineage does not of course have a nervous system that makes such choices, or even a chemical network that serves a somewhat analogous function for a plant. Nonetheless, like an organism, it is a process that responds to a changing environment in ways that are conducive to its survival.

The reproductive capacities of organisms, in particular, are both necessary for the perpetuation of the lineage and, a fortiori, necessary for the existence of the organism. More exactly, the organism would not exist if there were not a lineage that generated it, and the lineage would not exist if some of the organisms it generated did not have the capacity to reproduce. These necessities are causal. Just as the causal capacities of molecules and organs enable the persistence of the organism process, so the persistence and the reproductive capacities of the organisms enable the persistence of the lineage. And they do so, of course, by being exercised. Reproductive capacities, in particular, are exercised on the occasions of suitable causal interactions between organisms of opposite sexes (where the species is a sexual one). Similarly, the exercise of capacities to breathe, feed, etc., occurs in causal interaction with oxygen molecules and other organisms.

It is time to take stock. Living systems, I argue, are best seen as formed of a hierarchy of processes,¹⁰ stabilised over characteristic time periods. Such systems, and their constituent processes, interact with one another in countless ways. Some of these interactions are essential parts of the stabilisation processes of higher-level systems, whether internal interactions, of the kinds that have been much stressed by recent mechanists, or external, interactions with the wider world that the system requires to maintain itself. The cells of a multicellular organism interact with one another to maintain the stability of the organism. Organisms of the same species interact with one another to sustain lineages; in the case of organisms of different species, there is a lively debate over whether their interactions contribute to the maintenance of systems at an ecological level. All of these interactions are appropriately described as causal, and in a high proportion of cases the causal interaction involves the exercise on the part of one or both of the participants of capacities that have been exquisitely honed by natural selection precisely for the purpose their exercise serves. And, with one more nod to Elizabeth Anscombe, there is no reason to think that the exercise of such capacities is in satisfaction of any universal law or regularity. On the contrary, the better understanding of the conditions surrounding such exercise is the key to understanding the autonomy often attributed to organisms of all kinds, and most especially to humans.

I said I would return to the question of downward causation, and in fact I have already touched on the topic. The dependence of the organism on the lineage to which it belongs is an example. Now it will certainly be said that this is a dubious example. It is, after all, the organism's parents that cause its existence and they may be said to screen off the imagined influence of the lineage. But without pursuing this case here, it is easy to find less controversial examples. The cells in my liver provide various metabolic services to my body; but the persistence of my liver depends equally on its embedding within a system that contains blood vessels to carry its inputs and outputs, lungs to provide oxygen to the bloodstream, and even a brain that enables the organism to find food. The liver is not a thing that carries on functioning until something goes wrong, but a process, constantly maintained and stabilised by the complex system in which it is embedded.

Recalling the response I imagined to the case of the lineage's effect on the organism, something similar may well be said for these other cases. It is the pumping of the blood, the oxygenation of my blood in the lungs, electrical events in the brain, and so on that maintain the liver, not the whole in which all these are embedded. There are several things that might be said in reply to such a move. First, one must be cautious with such screening arguments. My throwing the brick is, in one sense, screened off from causing the window to break by the brick's impact with the window; by the time it reaches the window it is irrelevant who threw it. But this would

¹⁰ The idea of a hierarchy needs to be treated carefully. I do not assume that there is one universal hierarchy within which every biological entity has its place, just that the full understanding of a particular biological system will require attention both to systems internal to it, its parts, and further systems of which it is a part. This does not imply that there is a unique set of entities that form the parts of a system. In fact I believe that the identification of parts is to some degree dependent on the questions we want to ask of the system.

be a sorry excuse for an excuse. I did, in fact, throw the brick and initiated the causal process that led to the window breaking. I caused the breaking of the window.

But a second response is much more interesting and convincing. Screening off fails when there are alternative pathways from the initial cause to the effect. And it is a characteristic feature of biological systems that they are able to achieve important end points from multiple starting points and via multiple pathways. One kind of example is natural selection. One cannot imagine reducing a process of natural selection to all the individual matings and dyings that actually occur, as there are countless other sets of such events that would equally well have achieved the same result. The explanation is the preponderance of a feature giving a selective advantage in all of these sets. But similar redundancy is characteristic of more classically mechanistic cases.

It has been recognised for some time that mutations to the genome frequently produce little or no effect on the phenotype. It is often remarked that many gene knock outs often have no effect. Gene expression is remarkably robust. Typically robustness is achieved in machines by redundancy, that is, by duplication of parts. But it seems that this is not the main explanation of the robustness of gene expression, and that robustness is somehow distributed through the system (Wagner, 2005). The system compensates for the failure or absence of a component by changing the roles of existing parts and achieving the outcome in a different way. Hence, focusing on the particular pathway fails to identify the more robust capacity the system has to produce the observed effect. And this is just as true for the effect of the lineage on the organism as it is for the case of gene expression and countless others. In a world of universal regularity there is no question of looking for the level at which causation occurs, because as is often argued, universality can only coherently apply to one level, the lowest. But Anscombe, among others, has taught us that there is no reason to think we live in such a world; I do not believe that we do. In a world where causal regularity is local, perhaps even exceptional, we need to look for where the genuine regularity emerges. And as I have just explained, in many cases in biology this is at the higher level.¹¹

4 Conclusion

I take this essay to be an exercise in naturalistic metaphysics. Though the account of organisms as processes may strike many as radical or revolutionary, the arguments for this view (Dupré & Nicholson, 2018) are no more than reminders of uncontroversial scientific facts about organisms. The present essay takes this conclusion for granted and attempts to show that thinking of organisms and other biological systems as partially stabilised homeostatic or homeorhetic processes and considering the various kinds of interactions they have with one another is a good way of thinking about the causal interaction between and within biological individuals. I don't

¹¹ Questions of universal lawlikeness, emergence, and downward causation, are explored in more detail in Dupré (2021) Sects. 2 and 3.

claim that it is *the* correct account of all contexts in which we speak of causation, as I doubt whether there is any such account. I do claim that it is at the very least a good way to understand a large proportion of the causal claims that are made in the life sciences. It is also, incidentally, and bearing in mind the placement of this essay in a themed collection on Elisabeth Anscombe's famous inaugural lecture, a good context in which to see that the analysis of causation has no need of the necessitation or determinism against which Anscombe so powerfully argued.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anscombe, G. E. M. (1975). Causality and determination. In E. Sosa (Ed.), *Causation and conditionals* (pp. 63–81). Oxford University Press.
- Beatty, J. (1995). The evolutionary contingency thesis. In G. Wolters & J. G. Lennox (Eds.), *Concepts, theories, and rationality in the biological sciences* (pp. 45–82). University of Pittsburgh Press.
- Brandon, R. N. (1997). Does biology have laws? The experimental evidence. *Philosophy of Science*, 64, S444–S457.
- Cartwright, N. (1983). *How the laws of physics lie*. Oxford University Press.
- Cartwright, N. (1989). *Nature's capacities and their measurement*. Oxford University Press.
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford University Press.
- Dowe, P. (1992). Wesley Salmon's process theory of causality and the conserved quantity theory. *Philosophy of Science*, 59(2), 195–216.
- Dowe, P. (2000). *Physical causation*. Cambridge University Press.
- Dupré, J. (1993). *The disorder of things: Metaphysical foundations of the disunity of science*. Harvard University Press.
- Dupré, J. (2010). It is not possible to reduce biological explanations to explanations in chemistry and/or physics? In R. Arp & F. J. Ayala (Eds.), *Contemporary debates in philosophy of biology* (pp. 32–47). Wiley-Blackwell.
- Dupré, J. (2013). Living causes. *Proceedings of the Aristotelian Society, Supplementary*, 87, 19–38.
- Dupré, J. (2020). Life as process. *Epistemology & Philosophy of Science*, 57, 96–113.
- Dupré, J. (2021). *The metaphysics of biology*. Cambridge University Press.
- Dupré, J., & Cartwright, N. (1988). Probability and causality: Why Hume and indeterminism don't mix. *Nous*, 22, 521–536.
- Dupré, J., & Nicholson, D. J. (2018). A manifesto for a processual philosophy of biology. In D. J. Nicholson & J. Dupré (Eds.), *Everything flows: Towards a processual philosophy of biology* (pp. 3–45). Oxford University Press.
- Glennan, S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, 44(1), 49–71.
- Glennan, S. (2017). *The new mechanical philosophy*. Oxford University Press.
- Guay, A., & Pradeu, T. (2016). *Individuals across the sciences*. Oxford University Press.
- Harré, R., & Madden, E. H. (1975). *Causal powers*. Oxford: Blackwell.
- Lewis, D. K. (1974). Causation. *The Journal of Philosophy*, 70(17), 556–567.
- Lewontin, R. C. (1966). Is nature probable or capricious? *BioScience*, 16, 25–27.

- Lloyd, E. A. (1988). *The structure and confirmation of evolutionary theory*. Princeton University Press.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1–25.
- Mackie, J. L. (1974). *The cement of the universe: A study of causation*. Oxford University Press.
- Mayr, E. (1993). Proximate and ultimate causation. *Biology and Philosophy*, 8, 95–98.
- Meincke, A. S. (2019). The disappearance of change: Towards a process account of persistence. *International Journal of Philosophical Studies*, 27(1), 12–30.
- Michotte, A. E. (1963). *The perception of causality*. Basic Books.
- Mill, J. S. (1843). *A System of logic, ratiocinative and inductive: Being a connected view of the principles of evidence, and the methods of scientific investigation*. J. W. Parker.
- Moger-Reischer, R. Z., & Lennon, J. T. (2019). Microbial ageing and longevity. *Nature Reviews Microbiology*, 17, 679–690.
- Mortensen, C. (2020). Change and inconsistency. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2020 Edition). <https://plato.stanford.edu/archives/spr2020/entries/change/>.
- Mumford, S., & Anjum, R. L. (2011). *Getting causes from powers*. Oxford University Press.
- Nicholson, D. J. (2012). The concept of mechanism in biology. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(1), 152–163.
- Nicholson, D. J. (2013). Organisms ≠ machines. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44(4), 669–678.
- Salmon, W. C. (1984). *Scientific explanation and the causal structure of the world*. Princeton University Press.
- Steward, H. (2015). What is a continuant?. In *Aristotelian society supplementary volume*, 89 (pp. 109–123). Oxford: Oxford University Press.
- Stout, R. (2016). The category of occurrent continuants. *Mind*, 125(497), 41–62.
- van Fraassen, B. C. (1980). *The scientific image*. Oxford University Press.
- Waddington, C. H. (1940). *Organisers and genes*. Cambridge University Press.
- Wagner, A. (2005). Distributed robustness versus redundancy as causes of mutational robustness. *BioEssays*, 27(2), 176–188.
- Woodward, J. (2005). *Making things happen: A theory of causal explanation*. Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.