The Inaccessibility of the Past is Not Statistical

Emily Adlam *

May 2, 2022

It is an inescapable reality of the human condition that, no matter how hard we try, we cannot change the past. But in most discussions of this subject in the philosophy of physics, the inaccessibility of the past is elided with thermodynamical effects: it is, we are told, a consequence of the fact that '*small, local changes produce much bigger and more diverse changes in the future than they do in the past,*' [1] which in turn is supposed to be a consequence of the thermodynamic gradient. This entails that the inaccessibility of the past is, like the second law of thermodynamics, statistical and approximate rather than a fundamental feature of reality.

Intuitively this may seem puzzling, as the inaccessibility of the past does not feel in any way approximate: it is an unforgiving, rock-solid barrier. Of course, pre-scientific intuitions have often been shown to be mistaken in the course of scientific progress. But in this article we will argue that these particular intuitions are not in fact mistaken: our inability to influence the past is a consequence of the way we ourselves are embedded into reality as temporally directed processes, and thus there is nothing statistical about the inaccessibility of the past.

We begin by making some comments about the standard account of the inaccessibility of the past. We then argue that since the universe cannot contain contradictions, the modal structure of the universe must admit only those processes which cannot give rise to contradictions; we appeal to the process matrix formalism developed in the field of quantum foundations to characterise the complete set of processes which are compatible with local free will whilst ruling out contradictions. Thus far all processes known to occur in nature have *causal* process matrices, and a guarantee that compositions of processes will always remain causal is provided by the fact that all known processes obey 'consistent chaining,' i.e. the output of a process with a certain temporal orientation can only be used as the input to another process with the same temporal orientation. Since we ourselves are in effect processes which take memories as inputs and produce actions as outputs, we too are subject to consistent chaining requirements, with the consequence that our actions can only be used as inputs to processes with the same temporal orientation as our deliberative processes. We compare our view to the perspectival account of causation advocated by Price and Ramsey, and we discuss metaphysical pictures which are compatible with this approach, ultimately advocating a 'block-process' view. Finally, we discuss the possibility that some non-causal processes may be realised in nature.

1 Can Orpheus get Eurydice Back?

The most generally accepted explanation for the temporal asymmetry of action is that it can be understood in terms of the second law of thermodynamics, which in turn is explained via something like the past hypothesis. For example, the view developed by Albert, Kutach and Loewer attributes the inaccessibility of the past to the fact that '*in virtue of (the past hypothesis), small, local changes - the kind of things we could use as 'causal handles' ... produce much bigger and more diverse changes in the future than they do in the past.*' [1] That is to say, the inaccessibility of the past is to be understood as essentially a statistical matter: 'causal handles' do produce changes in both the past and the future, but they produce a lot more changes in the future than the past, so we simply notice the changes in the future more.

But there is something odd about this account. For if the difference between the future and the past were merely a matter of averages, one would naturally expect that this difference would be only *approximate*: we should be able to change the past at least a little, particularly if we get really good at manipulating things at a microscopic level. But

*The Rotman Institute of Philosophy, 1151 Richmond Street, London N6A5B7 eadlam90@gmail.com

that is not how the difference between past and future presents itself to us. No matter how deeply Orpheus may regret looking back at Eurydice on the way up from Hades, there is absolutely no possibility he can change that by his actions in the present. This obstacle seems not a matter of approximations and averages but a deep and absolute one: even if Orpheus were to become a brilliant experimental physicist capable of exerting very fine microscopic control akin to the powers of Laplace's demon, our intuition seems to tell us that still there would be nothing he could do to change what happened that day in Hades.

Now, part of the answer to this question involves the observation that of course Orpheus cannot 'change' the past because the idea of 'changing' the past is incoherent: if we assume some kind of block universe picture in which the course of history is fixed once and forever, then the past is simply as it is and there is no hope of it ever being anything else. But of course, in this sense Orpheus cannot 'change' the future either, so this reflection does nothing to explain the difference between past and future. One might argue that the distinction arises from the fact that Orpheus has memories of looking back at Eurydice, so he already knows what happens in the past and therefore there is no point in him taking actions to attempt to alter it; whereas he doesn't know what happens in the future, so it still makes sense for him to take actions aiming to bring about the future he wants. But this too doesn't seem adequate. Could Orpheus the brilliant experimental physicist not seek to exert his fine microscopic control in order to bring it about that in the past he does *not* look back at Eurydice, but also he forms a false memory of having done so which persists up until the present day, at which point he finally remembers the truth and simultaneously Eurydice is free to return to him? This would avoid any outright contradiction between his existing memory of looking back at Eurydice and his intention of taking actions in the present to prevent himself from having looked back at Eurydice. However, common sense would suggest that this is still impossible: whether or not Orpheus has a false memory, no action he can take now will make any difference to the facts about whether or not he looked back at Eurydice.

Moreover, note that there are parts of the future which we know about with fairly high reliability. For example, your conviction that the sun will rise tomorrow is probably almost as high as your conviction that it rose yesterday. Yet you probably also believe that it is at least in principle within your power to prevent the sun from rising tomorrow (for example, by quickly constructing a very powerful missile which could travel at just under the speed of light and which would therefore reach the sun in approximately ten minutes), whereas there is nothing whatsoever that you could do which would affect the sun's rising yesterday. Of course the probability that you would actually be able to prevent the sun from rising is so low as to be effectively zero, but the obstacles standing in your way are merely a function of limited energy and resources; whereas no amount of energy and resources, we tend to think, would allow you to exert any influence on the rising of the sun yesterday. This asymmetry does not seem wholly explicable in epistemic terms, since you will assign roughly the same probability to the sun's rising yesterday and tomorrow (i.e. extremely close to 1) so at least superficially your epistemic relation to those two events seems the same.

It may of course be the case that these intuitions are just wrong. They have been formed under conditions where we do not typically have fine control over microscopic particles, so perhaps we are simply being misled by the law of large numbers into thinking that the inaccessibility of the past is absolute and incontrovertible. However, if this were simply a matter of mistaken intuitions, then one would naturally expect that these intuitions would be seen to break down once we got down to the level of individual particles, where statistical averages no longer apply. In particular, one would expect that in experiments involving the manipulation of individual particles we should find that we are able to influence either the past or the future, albeit perhaps only at a microscopic level. But this does not seem to be the case: no experiment has yet unequivocally demonstrated an intervention having an influence on the past. (Delayed choice experiments are sometimes interpreted that way, but as explained in ref [2], that interpretation is certainly not compulsory). Indeed, as shown in ref [3], quantum mechanics seems to be very carefully fine-tuned to prevent the possibility of 'signalling backwards in time.'

One possible answer would be to say that although these experiments involve microscopic particles, they are *our* experiments and thus the process as a whole is necessarily a macroscopic one, since it is mediated via our macroscopic brains and bodies. So perhaps it is only our involvement which prevents us from seeing the breakdown of the direction of time at the microscopic level. Moreover, the fact that the laws of physics appear to forbid 'signalling backwards in time' is also not conclusive, because those laws have necessarily been developed on the basis of experiments performed by observers like ourselves: it could be the case that there are other types of processes that are possible when no observer is involved which do 'signal backwards in time' (insofar as we can make sense of that notion without the involvement of a macroscopic observer). However, it's clear that trying to formulate a theory involving processes going both backwards and forward in time will quickly tie us up in all sorts of paradoxical knots - and this, we will now argue, is the true reason for the inaccessibility of the past.

2 Consistent Chaining

Suppose for a moment that it were possible for people to take actions which have an influence on the past; thus it would be possible for someone to send a 'message' back in time by taking some such action. Then a curious philosopher could perform the following construction. Let $Emily_1$ be a version of the philosopher at some time, and let $Emily_2$ be a version of her at a later time (evidently $Emily_2$ must be in the future lightcone of $Emily_1$). Suppose that $Emily_1$ sends a message to $Emily_2$ containing the value of a single bit. $Emily_2$ opens the message, flips the value of the bit, then sends a message back in time containing the flipped bit. $Emily_1$ receives that message just in time to use the bit value it contains as the content of her original message. So the value of the bit inside that message now has both value 0 and value 1 at once, which is a contradiction.

Moreover, we can take it as read that if Emily ever discovered a way to send messages back in time, she would immediately try to construct a loop of this kind in order to see what happened. Thus if the laws of nature permitted the formation of loops of this kind, it seems inevitable that eventually someone would construct a contradiction. But the universe cannot contain contradictions, and therefore loops of this kind must be impossible. That is to say, they must be ruled out not only statistically but absolutely. For if they were only statistically unlikely, then eventually, if Emily tried enough times, she would succeed in producing a contradiction; and the universe cannot contain even a single contradiction.

2.1 Consistency Requirements on Modal Structure

Let us make this argument a little more formal. We will henceforth work within the context of a block universe picture, where we take it that the universe does not come into being in some kind of temporal process but rather exists atemporally and eternally.[1] We employ this conceptual framework because the block universe approach is favoured by a significant number of modern-day philosophers and physicists, in large part as a result of the fact that it seems more hospitable to relativistic physics than other approaches to the metaphysics of time. In the context of the block universe picture, it is reasonable to require that each variable has at most one value at a given spacetime point - that is to say, we do not allow variables whose value *at a given spacetime point* undergoes some change, and we do not allow variables whose value *at a given spacetime point* may be different relative to different observers. Note that we will exclude branching worlds from consideration, so we do not allow that values at a given spacetime point could differ across different branches.

The difference between past and future is a modal one: we *can* influence the future, but we *cannot* influence the past (or at least, so it seems to us). Thus the correct way to understand this difference is to see how it arises from the underlying modal structure. In this article we will not make any assumptions about the nature of this modal structure: it could be understood within a realist approach incorporating a metaphysically robust modality which is ontologically prior to the content of physical reality [4, 5], or it could simply be understood within a Humean approach as the best systematisation of the actual Humean mosaic [6, 7]. We insist only that modality is regarded as being an objective feature of reality in some sense. We will formalise the modal structure of reality as follows. First, we model an *experiment* as a set of laboratories in each of which agents may select inputs and/or receive outputs. We then define a *process* as a class of experiments which can be implemented anywhere in spacetime, and can be composed in any way we like with other processes, such that the probabilistic relationship between inputs and outputs will always be the same if the process has been implemented correctly. That is, an experiment is a particular sequence of physical events instantiating one particular set of inputs and one particular set of outputs, whereas a process is a *family* of such sequences of physical events, encapsulating all possible inputs and the associated outputs and the probabilistic relationships between them. Note that for certain sorts of processes, it may be the case that there are restrictions on the way in which the laboratories can be arranged in spacetime - if these requirements are not met then the process has not been implemented correctly and hence the usual probabilisitic relation between inputs and outputs would not be

[1] Here and throughout this paper we will take for granted the existence of a mind-independent reality which can be described from this kind of atemporal, external point of view even in the absence of any external observer.

expected to hold. For example, typically for processes that can be regarded as transmitting signals it is necessary that the output laboratory is in the future of the input laboratory, and if this requirement is not met then the signal will not be successfully transmitted.

Evidently the claim that some physical systems instantiate a particular process is indeed a modal one: we require not only that the actual input and output are consistent with the probabilistic relationship defining the process, but also that we can make counterfactual assertions of the form 'if some other value v were chosen for the input in laboratory one, then some other value x would have been produced as output in laboratory two,' or probabilisitic counterfactuals of the form, 'if some other value v were chosen for the input in laboratory one, then with probability p some other value x would have been produced as output in laboratory two.' Indeed, another way of thinking about a process is to think of it as an implementation of a computation, in the sense of ref [8]: '*not only ... its initial and final states are those associated with the input and output states of the computation on some particular occasion, but also ... had the initial state been that corresponding to one of the other input states of the computation, the resultant physical final state would have been that associated with the appropriate output state of the computation*.' Thus processes are modal in the same sense that computations are modal, and we may take it that the set of processes available in our universe is defined by the underlying objective modal structure of our universe, including the laws of nature - for example, the literature on operational formulations of quantum mechanics [9–14] can be understood as characterising the set of processes which can be constructed from quantum-mechanical systems on a fixed spacetime background.

Now, whether modal structure is understood to determine the contents of reality (as in the realist approach) or is simply read off the actual contents of reality (as in the Humean approach), clearly one baseline condition must be met: the modal structure of our actual universe should not give rise to contradictions. For example, if the modal structure of reality includes deterministic causation, it cannot be the case that one causal process deterministically causes the value of a certain bit to be 0 and another causal process deterministically causes the value of the same bit to be 1: different parts of the modal structures of the universe must be coordinated so as to ensure that they are compatible with one another. In the Humean approach this is automatic, because the modal structure is simply the best systematisation of the actual mosaic, and since the actual mosaic does not contain contradictions, a good systematisation of it should not lead to any contradictions either - or at least it should predict only a negligibly small number of contradictions, since any time the systematization predicts a contradiction it fails to match the actual mosaic and hence performs worse on the criterion of strength, so it becomes less likely to be the best system. Conversely in the realist approach where modal structure determines the content of reality, evidently modal structure must be formulated in such a way that it does not give rise to contradictions, since after all it cannot ultimately produce a reality which contains contradictions[2]. Thus in both the realist and Humean approaches it makes sense to ask about the conditions that the modal structure must satisfy if we are to avoid contradictions.

We note that our demand for consistent modal structure is somewhat similar to the 'Principle of Self-Consistency,' which was originally proposed by Novikov to deal with problems concerning closed timelike curves (CTCs) in general relativity [15]. Novikov's principle tells us that the only solutions to the laws of physics that can occur locally in the real Universe are those which are globally self-consistent - all events happen only once and cannot be changed. However, while Novikov's principle is a constraint on the events that actually take place, our consistency requirement applies at the level of modal structure: not only the occurrent events themselves but also the modal structure associated with those events should be free from contradictions, because the modal structure either gives rise to those events or is a systematization of those events, depending on one's view of modality - either way, the modal structure should not give rise to contradictions since the occurrent events cannot contain contradictions.

What does this consistency requirement entail? First, note that processes can in general be composed to form larger processes - the output of one process can be used as the input to another process, and therefore we can think of the modal structure of our universe as giving rise to a large collection of 'chained' processes, forming something like an enormous directed graph. And the requirement that modal structure should not give rise to contradictions necessarily places some constraints on the way in which processes can be composed. For example, suppose there exists a process P_1 which can be performed in the forward direction of time, i.e. it works as normal if its input is in the past lightcone

[2]One way to guarantee this is to formulate modal structure as suggested in ref [5], in terms of constraints which are defined as sets of Humean mosaics: if a certain constraint applies to our universe, then the actual Humean mosaic must be selected from the corresponding constraint. Since no Humean mosaic contains a contradiction, it is clear that no constraint formulated in this way can individually give rise to a contradiction. And if we are combining more than one constraint of this kind, there are no contradictions provided that the intersection of all of the constraints contains at least one mosaic.

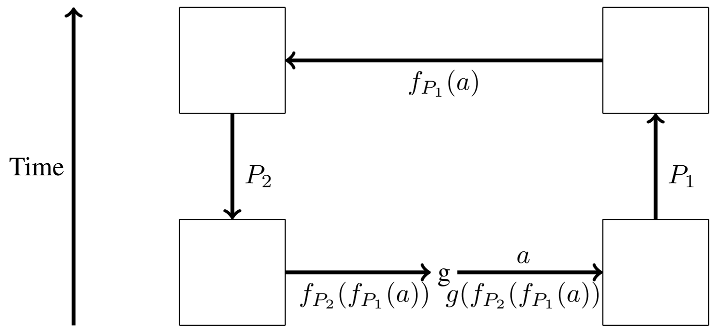

Figure 1: Schematic diagram of the composition of two processes used to demonstrate the need for consistent chaining

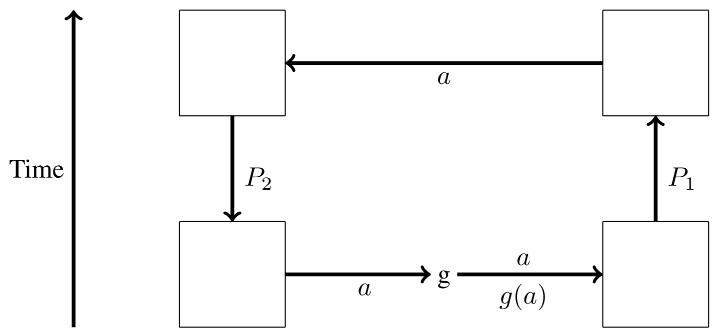

Figure 2: Simpler example of the composition of two processes used to demonstrate the need for consistent chaining

of its output, and another process P_2 which can be performed in the backwards direction of time, i.e. it works as normal if its input is in the future lightcone of its output. We can model the effect of these processes (for $i \in 1, 2$) as functions $f_{P_i} : \mathbb{I}_{P_i} \to \mathbb{O}_{P_i}$ from the set of possible inputs $\mathbb{I}_{P_i}$ to the set of possible outputs $\mathbb{O}_{P_i}$. Thus we can imagine creating the composition depicted in figure 1, where some input a is put into process P_1 and then the output $f_{P_1}(a)$ is used as the input to process P_2, producing an output $f_{P_2}(f_{P_1}(a))$. We then manually apply some function g to this output and use it as the original input to process P_1. If g is chosen such that $\forall a \in \mathbb{I}\ g(f(f(a))) \neq a$, then we have obtained a contradiction: the value of the input to process P_1 has two different values at once, in contradiction with our assumption that a given variable can have at most one value at each spacetime point. Figure 2 also shows a simple example, where the processes are just identity operations which always produce an output equal to their input. In that case we have $f_{P_1}(a) = a$, $f_{P_2}(a) = a$, and g can be any function which doesn't leave any possible input unchanged.

The thermodynamic explanation for the inaccessibility of the past would have us believe that it is merely a statistical matter that we are unable to create this kind of loop: we are prevented from acting 'backwards' in this sense only because we are trapped in the onward flow of the thermodynamic gradient. But as noted above, the absence of contradictions cannot be a statistical matter: contradictions must be ruled out absolutely, and thus the modal structure of the universe must be defined in such a way that contradictions like this simply cannot arise.

One might perhaps try to suggest that in any case like figure 1 where the objective modal structure would be expected to give rise to contradictions, the modal structure simply ceases to work in its usual way - e.g. we simply allow that in this case the output of process P_1 fails to be $f_{P_1}(a)$. However, if we were to postulate a modal structure for reality in which this kind of scenario arose frequently, that structure would evidently not be a good fit to the actual world and therefore it would seem doubtful that we could continue to maintain that this is really the right modal structure. For example, if we describe a modal structure where X deterministically causes Y, but in reality X frequently fails to be followed by Y in order to avert possible contradictions, presumably we would eventually conclude that we were wrong about the modal structure and X does not in fact deterministically cause Y. So this does not seem like a promising route if we wish to maintain that there should be a reasonably close fit between the modal structure and the physical reality associated with that modal structure, and thus we will henceforth take it that the modal structure of our universe must be such that it does not allow these kinds of contradiction to occur.

Bell made a similar observation in his comments on time travel in the Gödel universe [16], noting that the universe must be subject what he referred to as 'temporal interdicts' preventing contradictions that could arise from time travel: he argued that '*there is an important difference between the limitative principles of physics and any principles (call them "temporal interdicts") invoked to block changes of the past. In the first case it is logically possible that, for example, a body's velocity could exceed that of light or that an electron's position and momentum could be simultaneously measured with pinpoint precision. But any violation of a temporal interdict would involve a logical contradiction.*' Bell described three possibilities for these crucial 'temporal interdicts': time travel is impossible, time travel is possible but no 'changing of the past' is allowed, or time travel is possible and the universe branches whenever we change the past. Our case is more general, because we are considering not only time travel but also any other process which might be used to create a closed loop, but the possible 'temporal interdicts' are quite analogous to Bell's:

1. The modal structure of the universe does not allow processes which could give rise to contradictions.
2. The modal structure of the universe could in principle give rise to contradictions, but it arranges accidents or interferes with free choice to prevent contradictions from actually occurring.
3. The modal structure of the universe allows 'contradictions' to occur, but the universe branches so that contradictory values occur in different branches.

We have excluded branching worlds, so we can discard interdict number three. Interdict number two remains possible, but as noted in the philosophical literature on time travel [17, chapter 7], this approach would require us to accept that every attempt to create contradictions as in figure 1 is foiled by apparently inexplicable coincidences - for example by ensuring that we change our mind before carrying out the construction, or that the laboratory instrument that we are using to apply the function g mysteriously breaks every time we try to use it in this particular way. This possibility seems unappealing, and more to the point it doesn't appear to be the way things actually work in our world - we don't seem to observe large numbers of strange coincidences or random changes of mind, and indeed these features would be a contraindication for the existence of agents such as ourselves, agents who, in general, seem to have a significant amount of freedom to manipulate and compose variables as we please.

That leaves interdict number one - observers locally have freedom of choice about how they compose processes, but the modal structure of the universe does not admit any processes which could be composed in such a way as to give rise to contradictions. And indeed, this option seems to be the one that is realised in our actual world: what stops us from creating a composition as in figure 1 is not that things inexplicably go wrong every time we try, but simply that we do not know any processes which can be composed in this way - and indeed, apart from possibly General Relativity, none of our fundamental theories currently tell us that loops of this kind are possible. So there are good reasons to think that the objective modal structure of our actual world indeed limits the possible processes such that no composition of possible processes can give rise to a contradiction.

There is a theoretical framework in quantum foundations which has been developed specifically to study the set of all possible processes which are compatible with observers locally having free choice, but which also guarantee the absence of contradictions - the process matrix framework [18–20]. In this framework, we consider a set of laboratories in which agents can freely choose to perform any operations permitted by quantum mechanics, and we then use a 'process matrix' to encode the set of all possible dynamics between the labs - including dynamics which are not permitted by any known physics, such as dynamics in 'loops' as we have described here. An equivalent formulation is available for classical physics, in which case we use process functions rather than process matrices, and agents can choose from all local operations permitted by classical physics [21]. The process matrix or function is defined in such a way that no dynamics encoded in a valid process matrix or function can ever give rise to any contradictions. So for example, it can be shown that the process used for the composition in figure 1 does not have a valid process function - the classical process function formalism forbids the identity operation on loops precisely because it could produce this kind of contradiction: '*The problem with the argument above is that it simply assumes that the identity backwards in time is a possible solution of the dynamics, based on the intuition that such evolution would be possible if a were in the future of x, without CTCs. The studies mentioned in the introduction suggest that such an assumption is typically incorrect: The system's evolution typically finds a way to 'adjust itself,' preserving the consistency of 'free interventions.*'' [21] Thus, if we assume that observers locally have free choice to perform any operations they like, and that the modal structure of the universe does not allow for exceptions or inexplicable coincidences, it follows that all possible processes must be associated with valid process matrices (if we consider quantum processes) or process functions (if we consider only classical processes). That is to say, the process formalism allows us to characterise all the processes which could possibly be allowed by the modal structure of the universe if interdict number one is adopted.

The most straightforward type of process matrix is one which is *causally ordered*, i.e. it is compatible with the existence of a strict partial order such that the probabilities for the outcome of obtained in some laboratory depend only on the settings in experiments which precede it in the partial order. We can also imagine process matrices which are not causally ordered because the order for some subset of the laboratories may depend on the settings in other laboratories, but which are still *causal* in the sense that we can assign a strict partial order dynamically in such a way that the choice of setting in one laboratory cannot affect the probabilities for outcomes in laboratories unless they come after it in the partial order, and similarly the choice of setting in one cannot affect the strict partial order for some subset of laboratories unless they all come after it in the partial order [18] - that is to say, at each individual implementation we will be able to put the experiments into a strict partial order, though we may get different orders on different implementations. Causal processes cannot violate 'causal inequalities' [22] since the correlations they produce can always be embedded into a global strict partial order.[3] And thus far, all processes known to be realised in nature (without post-selection) are causal.[4]

We will return the topic of non-causal processes in section 5, but for now we will simply follow ref [23] in postulating that, at least in the regimes with which we are familiar, the way the modal structure of the universe prevents contradictions is by limiting the allowed processes to those which are causal. Moreover, since composing processes gives rise to another process, it follows that in the regimes with which we are familiar, processes can only be composed in such a way that the resulting composite process is causal - i.e. compatible with a strict partial order. Evidently as long as all processes and compositions of processes are compatible with an underlying strict partial order, contradictions of the kind depicted in figures 1, 2 cannot occur.

As noted in ref [24], defining a general composition rule for process matrices is not straightforward - simply taking the tensor product of two process matrices can sometimes produce something which is not a valid process matrix, so it's necessary to impose some restrictions to ensure that composing processes yields another valid process. And presumably we will need some further restrictions if we also want to ensure that composing two causal processes yields another causal process. But to get a picture of what these restrictions will look like, let's focus on the case where we are composing a set of bipartite signalling causal processes - bipartite meaning the process takes an input in one lab and produces an output in another lab, and 'signalling' meaning the output depends nontrivially on the inputs, i.e. for some possible output O and input I we have $p(O|I) \neq p(O)$. Evidently if we chain processes of this kind together by using the output of one process as the input to the next, the resulting composed process will be compatible with an underlying strict partial order provided that there are no loops in the chain, because then we can obtain a strict partial order for the composed process by simply concatenating the strict partial orders for all of the individual causal processes. That is to say, the composed process is causal provided that the directed graph corresponding to the way in which the processes are chained together is acyclic.

Moreover, in the regions and regimes we have so far probed, the 'strict partial order' to which processes conform is simply temporal order: when we perform a causal bipartite signalling process whose input is in the past lightcone of its output, the output of that process can only be used as the input into another signalling process if the input of that process is also in the past lightcone of its output, as shown in the left on figure 3. Provided that we are in a spacetime which does not contain closed timelike curves (CTCs), this will ensure that all the resulting composed processes are compatible with a strict partial order, which is simply given by the temporal order of the experiments in any valid reference frame. We will use the term 'consistent chaining' to refer to the requirement that the output of a signalling process which is oriented forwards in time can only be used as input to other signalling processes

[3]More precisely, according to ref [18] a process matrix is causal iff there exists a probability distribution over possible strict partial orders on the set of local experiments and outcomes of those experiments such that for every local experiment A, every subset X of the other local experiments, and every choice of strict partial order, the probability distribution over strict partial orders such that all the experiments in X precede A in the order and the outcomes for the experiments X is independent of the settings in the experiment A.

[4]Note that we are employing the technical terms 'causally ordered' and 'causal' in the way that they are used in the process matrix literature, but we do not mean to presuppose that these processes involve causation in the broader sense in which philosophers use that term - we will return to this point in section 3, but for now these terms should be understood in a purely formal way.

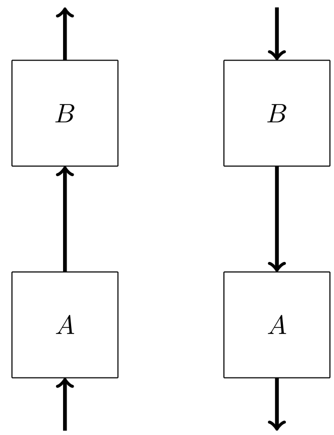

Figure 3: Switching inputs and outputs

which are also oriented forwards in time, thus giving rise to a process graph which is acyclic. If consistent chaining is satisfied, temporal order acts as a kind of regulatory device - we have complete freedom to compose bipartite signalling processes in any way that is compatible with consistent chaining, because no contradictions can ever arise from any such composition as long as there are no CTCs. It should be reinforced that we do not claim that consistent chaining is the only possible way to avoid contradictions; our argument is simply that the modal structure of the universe must implement some kind of 'temporal interdict,' and we regard it as an empirically established fact that consistent chaining is the approach instantiated by the actual modal structure of our universe, at least in our local vicinity for the kinds of processes which we are currently able to detect.

We have so far dealt strictly with bipartite signalling causal processes, but our comments generalize straightforwardly to multipartite signalling causal processes: if we chain processes of this kind together by using outputs from one process as inputs to others, the resulting composed process will be compatible with an underlying strict partial order if and only if there are no loops, i.e. there are no instances where the output O_1 of some process depends on an input I_1 of that process but also I_1 also depends on O_1 via some chain of processes, because then we can obtain a strict partial order for the composed process by simply concatenating the strict partial orders for all of the individual causal processes. And again, in the regions and regimes we have so far probed, this ordering corresponds to temporal ordering: if an output O_1 of a signalling process P_1 depends on one or more inputs which are in its past lightcone, then O_1 can only be used as an input I_2 to another process P_2 provided that all of the outputs of P_2 depending on I_2 are required to be in the future lightcone of I_2, thus giving rise to a process graph which is acyclic. So multipartite processes are likewise subject to consistent chaining requirements, and thus temporal order regulates the way in which we can compose multipartite processes in such a way as to ensure that we cannot possibly produce a contradiction.

Note that the consistent chaining requirement does not require us to postulate an *objective* distinction between the inputs and outputs of processes. We are always free to rename inputs as outputs and vice versa, and hence we could equivalently describe the case in figure 3 as a case where the input of process A is set equal to the output of process B, as shown on the right. Indeed, ref [21] demonstrates that (in the case of classical deterministic dynamics), every process function can be extended to an invertible one, which ensures that it is always possible to reverse the description so that inputs become outputs and vice versa. However, if we are going to exchange the inputs and outputs for one process, then we will also have to do so for all the other processes to which it is linked, otherwise we will be trying to set the *output* of one process equal to the output of another process. So, assuming that the process graph has no disconnected parts, this transformation will amount to changing the directions of all of the edges of the graph: clearly if it is acyclic under the original ordering it will remain acyclic under such a transformation. Thus we can equally well describe the consistent chaining requirement by saying that the output of a signalling process which is oriented 'backwards in time' can only be used as input to signalling processes which are also oriented 'backwards in time,' - from the block universe perspective, these descriptions are equivalent. So a consistent chaining mechanism does not have to rely on the existence of an objective preferred direction of time: we can make sense of this chaining requirement no matter which way we imagine the processes running. Indeed, as Ismael notes, '*nomological relationships do not have an intrinsic direction of determination ... Dynamical laws ... are constraints on the relationships between states at different times, but ... there is nothing in the law itself to say that either determines the other,*' [25] and this feature of our dynamical laws give us good reason to suppose that all of the relevant modal relations are perfectly symmetric - so the processes are not *actually* running in either direction, but nonetheless due to the consistent chaining requirement, once we choose a temporal orientation for one process that will induce an orientation for all the other processes.

2.2 Agency

Now let us try to add agents such as ourselves into this picture. The process matrix formalism treats agents as external factors, so to speak: agents are in laboratories choosing operations to performs, but they are not themselves a part of the process. However, evidently we can think of a decision made by an agent as a kind of signalling process: the agent takes an input (i.e. a set of memories and/or present perceptions) and produces an output (the final decision). The memories are themselves the output of a set of signalling processes which produce records of earlier events, and clearly the agent will regard the temporal direction recorded by their memories as the past, and the other direction as the future; hence the agent will perceive their deliberative process as being temporally oriented from the past into the future - that is, the process of making a decision will always be perceived as having an output which is later than its input relative to the direction of time perceived by the agent. A decision then typically leads the agent to manipulate some external variables (e.g. writing and posting a letter), so the output of the decision-making process is used as the input to some other process, typically a signalling process (e.g. the process of sending a message through the post). Decisions are therefore part of the process graph, so they must be 'chained' in a way that obeys the consistent chaining condition. It follows that the output of a decision can only be used as the input to a signalling process if that process has the same temporal orientation as the deliberative process - which is to say, it has an output which is later than its input and therefore later than the decision. Thus, if we take it that 'influence' necessarily requires nontrivial dependence of the event in question on the actions that supposedly influence it, so 'influence' must be mediated by signalling processes, it follows that agents will find themselves only able to influence events occurring later than their deliberation according to their own perceptual direction of time. And this is, of course, exactly how agency does work: in the words of Price and Weslake, '*we can only wiggle handles which lie in the immediate future, with respect to our own deliberations on the matter*' [1] or in Frisch's words, '*It is a striking fact about experimental interactions that we can only intervene into a system 'from the past,' as it were*' [26].

Furthermore, part of what it is to be an 'agent' persisting over time is that the inputs to our decisions include information about our past decisions. Thus from the process point of view, an agent is simply a set of chained decisions, i.e. an ordered set S of edges of the process graph, with the input of each edge connected either directly or indirectly to the outputs of all of the previous ones. Since these processes are signalling ones, consistent chaining then mandates that all of the edges in the set S have the same orientation in time, so the agent's perceived 'past' and 'future' will be consistent throughout the lifetime of the agent. In particular, since the agent can't make decisions with outputs earlier than the time of the decision relative to the agent's perceptual direction of time, the agent will be unable to 'go backwards' or 'travel back in time' unless they encounter any CTCs.

Notice that from the external, objective point of view there is still no preferred direction here. An agent is just a signalling process, and the process can equivalently be described in either direction: in reverse, a decision would look like a map from a final decision to the set of memories and/or outputs of past decisions that produce this decision, and the memories in turn become inputs to other past-oriented processes which ultimately produce the events that the memories are of. It is not objectively the case that the memories are the inputs and the decisions are the outputs; that is simply how we experience the process, since of course we remember the content of the memories. So the asymmetry comes solely from the fact that the agent has a perspective defined by the fact that they have memories of one temporal direction and not the other: that direction is necessarily perceived as the past, and the other direction, the direction towards which agents produce decisions as outputs, is therefore perceived as the future. There are of course further questions about how there comes to be an agent with memories and thus a particular temporal perspective in the first place, and it seems likely those questions must be answered by appeal to thermodynamics. But, conditional on the existence of such an agent with memories of the past, the fact that the agent can only output decisions which affect the 'future' relative to their own orientation is an objective fact about modal structure which arises from the consistent chaining requirement. Thus for such an agent there is a very concrete distinction between past and future which has nothing to do with thermodynamics: given that the agent is oriented in a certain direction, they can only take use their actions as inputs to processes which are oriented in the same direction, and therefore there is nothing epistemic or approximate about the inaccessibility of the past.

Indeterministic processes The reasoning that led us to the need for consistent chaining requirements was concerned specifically with deterministic signalling processes. But reality also appears to contain some signalling processes that are not deterministic; indeed, if everything is ultimately made up of quantum fields undergoing scattering processes, then one might worry that really no processes are deterministic. And clearly the above argument will not go through if we compose two indeterministic signalling processes. Returning to figure 1, suppose each of the processes has a probability 0.99 of producing an output equal to its input: then we can always have one of the channels produce an output different to its input, rendering the loop consistent. This is admittedly an unlikely turn of events, but there is no contradiction here.

However, suppose that we create a large ensemble of such loops. In order to avoid inconsistencies, in every case one of the channels must produce an output different to its input, and therefore in the context of this composition, the relative frequency with which these channels produce outputs different to their inputs will be 0.5, which is very far from the expected relative frequency 0.99. For a large ensemble, the probability of such a large mismatch between the expected and predicted relative frequencies becomes extremely small, and it can be made arbitrarily close to zero by making the ensemble large enough. Moroever, this mismatch will necessarily occur across all cases of this kind of composition, and therefore there will be no meaningful sense in which it is true that the probability for these processes to produce an output matching their input remains 0.99 under this composition - regardless of one's preferred account of probability, the only reasonable conclusion would be that the probability is somewhere close to 0.5 under this composition. But we have specified that for a process the probabilistic relationship between inputs and outputs must remain the same under composition provided that the process is being implemented correctly; so under these circumstances, we would have to conclude that we have not succeeded in reproducing the same process, since the probabilities do not match the usual ones under this sort of composition.

This observation can be generalised: in any composition of processes that can be used to give rise to contradictions if the processes are deterministic, replacing the deterministic processes with indeterministic ones will give rise to inconsistencies where the processes cannot possibly exhibit their usual probabilistic relationships. As in the deterministic case, such inconsistencies can only arise if the composed process is not causal, and since all processes known to occur in nature are causal, it seems that at least in the regimes with which we are familiar, the way the modal structure of the universe prevents these sorts of inconsistencies is by limiting the allowed processes to those which are causal. And again, this is apparently achieved by means of consistent chaining: the output of an indeterministic signalling process whose input is in the past lightcone of its output can only be used as the input into another indeterministic signalling process if the input of that process is also in the past lightcone of its output. So even when we have access to indeterministic signalling processes, agents can only output decisions which influence the 'future' relative to their own orientation, and thus neither deterministic nor indeterministic signalling processes allow agents to influence the past in a spacetime without CTCs.

Nonlocality We have so far considered only *signalling* processes, i.e. processes such that something can be inferred about the input of the process from the output. However, when we move into quantum regimes we encounter novel processes which are not signalling but where nonetheless the relationship between the input and the output is not trivial. For example, using quantum mechanics we can construct a process composed of two laboratories each containing half of an entangled pair of particles, with the measurement directions as the 'inputs' and the measurement outcomes as the 'outputs.' Evidently this is indeed a process in the sense that there exists a consistent probabilistic relationship between the inputs and outputs, and moreover Bell's theorem [27] shows that the relationship between the laboratories is nontrivial: the output in one laboratory really does depend on the input in the other laboratory in a way that can't be explained by a common cause in the past lightcone. So should consistent chaining requirements apply to such processes?

Well, note that if we use non-signalling processes of this kind to create the composition shown in in figure 1, we will not obtain either a contradiction or a probability inconsistency - in fact the resulting process will still be causal, since the correlations will be compatible with a strict partial order (though the choice of order is not unique - different reference frames will lead to different quantum descriptions and thus different orders). Thus intuitively it would seem that consistent chaining need not place any restrictions on the way in which nonlocal, non-signalling processes can be arranged in spacetime. And indeed, it transpires that this is the case: if a process is constructed from measurements on entangled particles, the probabilistic relationships between inputs and outputs will be the same regardless of where the laboratories are placed in spacetime, so it doesn't matter whether they are spacelike, lightlike or timelike separated. Conversely, in order for compositions of possible processes to remain causal, clearly it is important that any such location-independent processes should not be capable of signalling, which appears to be the case - the only known nontrivial processes of this kind are those constructed from measurements on entangled particles, and the quantum no-signalling theorem [28] tells us that all such processes are indeed non-signalling. It has been observed that the existence of quantum processes of this kind which are non-local but also non-signalling is a striking feature of reality which seems in need of explanation: for example, if we try represent such a process by a causal model, it can be shown that the parameters of the model must be very carefully fine-tuned [29]. The fact that processes of this kind must be non-signalling in order to avoid contradictions might be regarded as furnishing the desired explanation: from this point of view, the universe has very carefully fine-tuned the 'toolbox' of processes available to us to ensure that we are not provided with any processes which could give rise to contradictions.

We note that Bell's theorem [27] is sometimes interpreted as saying that a measurement on one entangled quantum particle exerts a nonlocal influence on the other. Thus if we move to a broader notion of 'influence' which does not require signalling, and which includes the relationships between entangled quantum particles as a form of influence, it is no longer true that agents can only 'influence' events occurring later than their deliberation according to their own perceptual direction of time - by using an appropriately arranged pair of entangled particles an agent could 'influence' a event in their own past lightcone. However, this is not a counterexample to our argument that the past is inaccessible due to consistent chaining requirements, because the consistent chaining approach still explains why it is impossible to 'influence' the past in the stronger signalling sense, and it is this kind of influence which most people seem to have in mind when they describe the past as unchangeable or inaccessible.

Closed timelike curves Although the consistent chaining requirement is sufficient in ordinary circumstances to ensure we cannot create contradictory loops as in figure 1, this is no longer the case in spacetimes which contain CTCs; if we put processes on a CTC in such a way that the chain of processes eventually arrives back at its own beginning, the process graph associated with those processes may fail to be acyclic even though consistent chaining is obeyed everywhere, and thus we would naturally expect to be able to perform the composition of figure 1.

One possible response would simply be to rule out CTCs by fiat. There exist approaches to quantum gravity which do exactly that - for example, approaches based on canonical quantisation presuppose a globally hyperbolic spacetime, in which there cannot be CTCs [30, 31]. A somewhat more principled approach is taken by causal set theory, which builds up spacetime from a set of points with a partial order relation which is required to be reflexive, antisymmetric, transitive and locally finite [32, 33] - this necessarily entails that the spacetime thus produced will not contain any CTCs, and indeed, clearly a 'causal set' will look very much like what we have henceforth referred to as a 'process graph.'

On the other hand, the process framework does allow for non-causal processes which can be implemented on CTCs without giving rise to contradictions. So if we allow the existence of non-causal processes in certain regimes then we don't have to rule out CTCs; we simply have to require that the only possible dynamics on a CTC are dynamics which correspond to a valid process function or matrix. We will return to this topic in section 5.

Antiparticles Antiparticles are often described as being particles which are 'travelling backwards in time,' which might seem to to pose a problem for our consistent chaining requirement. However, for the purpose of consistent chaining what matters is not which temporal direction the particle is travelling in, but rather which direction we are able to *control* it in. Regardless of whether or not we think of an electron as 'a positron travelling backwards in time,' the fact is that no one has yet figured out how to exert control over an electron in such a way as to use electrons to send messages backwards in time: when we alter the state of an electron, the consequences are observed in the future and not in the past. Indeed, as we have noted, in our construction there is no need to think of processes as actually going one way or another: the relation between inputs and outputs doesn't depend on a temporal direction, and the asymmetry can be derived entirely from the way we ourselves are embedded in the process graph. Thus we agree with Price that '*the view that particles are thus intrinsically directed in time is as problematic as the corresponding doctrine about radiation. Particles and waves are simply retarded with respect to (i.e.,confined to the past light cone of) their point of emission; and advanced with respect to (i.e.,confined to the future light cone of) their point of absorption*' [34, pp. 74].

2.3 Examples

One way to clarify the distinction between our approach to the inaccessibility of the past and the usual thermodynamical account is to consider cases where the thermodynamic asymmetry fails to hold in the usual way. For example, consider a universe like ours but without the low entropy initial state. Such a universe could contain a Boltzmann brain [35]: i.e. a brain, complete with a set of (false) memories, which is produced by a large thermodynamical fluctuation. Obviously it is difficult for a lone brain to take action, so let us slightly expand the thought experiment to suppose that the fluctuation produces an entire human being called Bob, together with an environment which is adequate to keep him alive for a short period of time. What will Bob discover about the accessibility or otherwise of the past? Let us assume that Bob is to be found at the lowest entropy point of his fluctuation, so entropy is increasing in either temporal direction on each side of this point. Thus, if the inaccessibility of the past is purely thermodynamical, one would think that Bob should not be subject to any temporal restrictions; he should be able to take actions which have an impact on either the past or the future (i.e. either of the two directions of time).

But clearly if we allow this we will quickly run into paradoxes: for example, what if Bob takes an action which impacts the past so as to prevent the fluctuation containing him from forming in the first place? The consistent chaining account provides an answer to this question. Like other humans, Bob will make decisions using his set of memories as input. In this case the set of memories was produced, not by a signalling process starting with the events that the memories purport to record, but by the fluctuation: the memories are the output of whatever signalling process brings about the fluctuation. Since we can always choose to redescribe processes in either temporal direction, we could in principle regard the memories as being produced by a fluctuation from either temporal direction. However, Bob's deliberative process takes these memories as input and produces an output which necessarily must be 'later' then the input according to one temporal orientation and 'earlier' according to the other orientation, so the direction of this deliberative process selects a temporal direction such that the decision is in the future of the formation of the memories according Bob's perceptual arrow. Then consistent chaining requirements entail that the decision can only be used as input to signalling processes whose output is also to the future according to Bob's perceptual arrow. So even in the absence of any thermodynamical gradient, Bob will nonetheless find that he is only able to take actions having an influence (in the signalling sense) on his perceptual future, and not his perceptual past - and he will necessarily regard his memories as having been 'produced' by a fluctuation in his perceptual past, though in an observer-independent sense we could think of the memories as having been produced in either direction. Therefore no causal paradoxes will result: Bob will not be able to act 'backwards' to interfere with the signalling process of which his very existence is an output. This demonstrates that the existence of a monotonic thermodynamic gradient is not a crucial part of the explanation for the inaccessibility of the past, because the past is just as inaccessible for Bob as it is for us, despite the absence of a monotonic thermodynamical gradient in his vicinity.

Similarly, consider Price's example of the 'Stargate Doughnut,' involving a microscopic gate which may block the path of a photon or allow the photon to pass through [36]. Since we are dealing with only microscopic entities it seems that thermodynamics is not relevant to this scenario, and yet we tend to think that the gate's being closed will prevent the photon from arriving at its destination but will not prevent it from being emitted in the first place, so there appears to be some asymmetry in our description of the situation. Price argues that '*what we bring to the case, in imagination, is the typical perspective we have as deliberating agents ... according to the perspectival view, it is this asymmetric perspective on our part that grounds the intuitive asymmetry.*' That is to say, our intuitions about this case are essentially based on an analogy - we are importing a causal asymmetry extrapolated from our experience even though there is no real asymmetry here.

However, the consistent chaining account entails that even though this scenario is microscopic, the inaccessibility of the past still applies. For example, suppose some microscopic device is oriented with respect to the Stargate so as to accept information as input from some temporal direction t_1, which we will refer to as the 'past.' Then that device can gain information on whether or not the photon is present in the 'past' before determining if the gate should be open or closed (assuming that the gate can be opened or closed instantaneously). So if this device's action in closing the gate can cause the photon to not be present in the 'past,' then this device could create a paradox by checking to see if the photon is present on the left side and closing the gate if it is. However, if *this device's* action in closing the gate can only cause the photon to not be present in the opposite temporal direction t_2, which we will refer to as the 'future,' then no paradox can arise, because the device cannot get information from the future before deciding whether or not to open the gate. (Obviously, exactly the same argument could be made in reverse about a different device configured so as to accept information as input from the temporal direction t_2 instead). So the difference between the past and future in this case doesn't just come from making analogies to more standard thermodynamic cases: it is an objective fact that if some process can gain information about the passage of the photon from one temporal direction, then the output of that process in terms of opening or closing the gate can only influence what happens to the photon in the *other* temporal direction, or else logical paradoxes would ensue. The consistent chaining account therefore entails that in such scenarios it will always be the case for any given agent or device that one direction of time is inaccessible and the other direction is accessible, so in a sense there is indeed an objective asymmetry present in this scenario, although the direction of the asymmetry is not fixed and differently oriented agents or devices will be subject to different temporal asymmetries.

3 Causation

We have thus far discussed the inaccessibility of the past without couching that inaccessibility too strongly in causal terms (apart from the technical term 'causal' inherited from the process matrix formalism), but evidently the notions are related - one way of thinking about the inaccessibility of the past is to observe that we take ourselves to be able to cause events in the future but not in the past. The temporal asymmetry of causation has often been regarded as puzzling: summarising a common sentiment, Frisch writes, '*The fundamental equations of all mature physical theories are time-reversal invariant. There is no place for an asymmetric notion of cause within a physical theory with time-reversal invariant laws. Therefore, there is no place for an asymmetric notion of cause in mature physical theories.*' [26] But what this way of thinking misses is that laws which are perfectly symmetric from the point of view of an external observer can still be experienced asymmetrically from the point of view of an internal observer. This point has been made strikingly by Price [36], who asks us to consider the example of a football field: to the external observer the field is perfectly symmetrical, but to a player there is an obvious asymmetry, since the aim is to keep the ball moving toward the opposing team's goal and therefore the football match has an overall 'orientation.' And this is exactly the kind of account suggested by the consistent chaining picture: even though the laws themselves have no temporal direction, an observer internal to the universe whose deliberative processes are oriented in one particular temporal direction will necessarily find themselves able to use the outputs of those mental processes only as inputs to other processes oriented in the same temporal direction, so such observers will under appropriate circumstances have experiences of asymmetrical causality.

Of course, asymmetry is not all that it takes for a relation to be regarded as a form of causation, and there may be features one would associate with causality which are not implemented in both directions of time in the consistent chaining picture. For example, some writers have taken the view that causation is closely tied to irreversibility, and if this accepted then presumably the direction of causation must be necessarily aligned with the thermodynamic gradient, which would imply that the perspectival asymmetry associated with consistent chaining is not the same as the causal asymmetry. This is the reason we have thus far avoided causal terminology, because we acknowledge that different analyses of causation will reach different verdicts on whether the perspectival asymmetry arising from consistent chaining is really causal in nature.

That said, the consistent chaining approach is in fact quite compatible with several popular accounts of causality. For example, it could be combined with a counterfactual account of causation [37, 38], because we have reinforced that processes must be such that they licence counterfactual assertions about what the outcome would be for various different possible inputs, and thus for most processes (particularly signalling ones) we will be able to make various statements of the form 'if the input were to be A then the output would be B, and if the input were not to be A then the output would not be B,' which according to the basic Lewisian analysis [38] entails that the event of obtaining an output B causally depends on the event of selecting an input A. That is to say, the counterfactual analysis tells us that ‘signalling processes’ as we have defined them typically involve causation (in the philosophical sense, rather than the technical process matrix sense). However, the consistent chaining approach differs from Lewis on the origin of the asymmetry of causation. For Lewis, the asymmetry of causation comes from the fact that effects are rarely overdetermined by their causes but causes are very often overdetermined by their effects [39]. Note that at the level of the most basic processes licensed by the laws of nature (such as scattering processes, or the propagation of light) there is typically a high degree of symmetry between inputs and outputs, so the ‘overdetermination’ referenced by Lewis appears only in the thermodynamic limit, so Lewis’ approach is in effect a version of the thermodynamical account of the asymmetry of the past; thus although it tells us that the asymmetry of causation is observer-independent, it also entails that this asymmetry is statistical and approximate and fails to hold at the microscopic level. Whereas we saw in section 2.3 that consistent chaining has the consequence that there must be apparent asymmetries even when we are dealing with microscopic processes to which thermodynamics does not apply, so the consistent chaining approach tells us that although the asymmetry of causation is only perspectival, it is also ubiquitous, since it is an ineliminable consequence of the way in which processes and agents must be embedded in reality.

The consistent chaining approach would also work well with an interventionist account [40, 41], because the act of choosing an input for some process can often be regarded as an intervention, and if the process is signalling we can then say that the manipulation of an input will result in the manipulation of an output, meaning that according to the interventionist criterion these processes involve causation (in the philosophical sense, rather than the technical process matrix sense). Of course there may well be some signalling processes in nature where the ‘inputs’ are not straightforwardly manipulable by human agents, so the interventionist account need not say that *all* signalling processes involve causation, but nonetheless it is clear that many signalling processes will involve causation. Indeed, the ‘process graph’ that we have discussed has much in common with the directed acyclic graphs used in interventionist analyses of causal structures [40] - except we take it that the graph is not *intrinsically* directed, and only becomes so once one adopts the perspective of an observer embedded within it. We thus agree with the interventionists that the modal structure underlying causal relations is objective, but deny that it is built out of fundamentally asymmetric relations: the modal structure and the fact that it is acyclic is an objective, perspective-free fact, but it is only the choice of a perspective which selects an orientation and thus imbues this structure with the asymmetry that turns it into something we can properly refer to as a causal structure.

This observation suggests that the account of causation which most closely aligns with the consistent chaining approach is the perspectival analysis of Price and Ramsay, who hold that the asymmetry of causation originates ‘*in de facto asymmetries in our own temporal orientation, as physical structures embedded in time*.’ [36] The consistent chaining approach agrees with the account offered by Price and Ramsay on many points, in particular through a shared emphasis on the fact that causation can be perspectival without being subjective: in Price’s words, ‘*unmasking the perspectival character of a concept does not lead to simple-minded antirealism—we may continue to use the concept, and even to affirm, in a variety of ways, the objectivity of the subject-matter concerned, despite our new understanding of what is involved (of where we ‘stand’) in doing so.*’ [36] This accords with the way things work in the consistent chaining approach, where it is an objective fact that observers are not able to use the outputs of their mental processes as inputs to other processes oriented in the opposite temporal direction, but nonetheless the direction depends on a perspective and is not written into the fundamental laws.

There are, however, some subtle distinctions between the approaches. Most importantly, Price and Ramsay seem to take the view that the perspectival asymmetry must be grounded on an objective, observer-independent asymmetry such as the thermodynamical one, and thus they partially endorse the statistical account of the inaccessibility of the past: Price writes ‘*The main candidate for a physical asymmetry that seems likely to be associated with the causal asymmetry, whether by the reductive or perspectival routes, is the asymmetry associated with the second law of thermodynamics’* [36]. Whereas we have argued that there is no need for any objective, observer-independent asymmetry to ground the perspectival one - symmetric modal relations combined with the existence of a perspective is enough. Of course we acknowledge, as does Price, that the thermodynamic gradient may well be a necessary precondition for the *existence* of agents such as ourselves, but fundamentally the fact that we can influence the future and not the past follows from the consistent chaining requirement and not from the thermodynamic gradient. Thus the consistent chaining account strengthens the perspectival account of Price and Ramsay, as it entails that the perspectival asymmetry does not arise approximately and statistically, but is built into the modal structure of reality at a deep level.

This attitude to the thermodynamic gradient is related to the fact that Price regards the perspectival asymmetry as being essentially a reflection of our epistemic situation. He analyses the structure of deliberation by dividing facts into options amongst which we deliberate and fixtures which are not held to be matters of choice in the current deliberation, and then observes that the determination of fixtures and options is not necessary objective. In particular, he holds that '*what lies in the past goes into fixtures by default*' but this is not a reflection of the objective structure of reality, merely a consequence of our epistemic situation: '*We regard the past as fixed because we regard it as knowable, at least in principle ... as information-gathering systems, we have epistemic access to things in (what we call) the past; but not, or at least not directly, to things in (what we call) the future.* This fact about our epistemic access to the past is presumably largely a function of the thermodynamic gradient, so from that point of view it makes sense for Price to argue that the causal asymmetry is grounded on the thermodynamical one. On the other hand, the consistent chaining account says that we do not simply *choose* to regard the past as fixed; the assessment of 'fixtures' and 'options' must be made relative to one's actual relationship with the process graph, and since a fact can be an 'option' only if one's actions at the present moment could exert an influence over it, it follows that 'options' must be facts which could potentially be the output of some process or chain of processes for which one could at this moment choose the input, and thus if consistent chaining is satisfied, facts which are in the past relative to one's current position cannot feature in 'options.' Our epistemic predicament is certainly relevant here, but not merely because it influences our attitude towards the past; rather the fact that we have received information in the form of inputs from the past entails that, as a matter of objective fact, we cannot produce outputs of our deliberation as actions oriented towards the past, otherwise it would be possible to form loops as in figure 1 and consistent chaining would be violated. In this sense, if consistent chaining holds it follows that facts about the past must always fall into the set of 'fixtures,' and thus the inaccessibility of the past is an objective fact about modal structure rather than simply an epistemic matter.

Price and Weslake have also suggested that the perspectival asymmetry may be related to '*the fact that we are always contemplating actions in the near future, with respect to the time of deliberation.*' [1] But this fact in itself seems not enough to explain the inaccessibility of the past: for even if we can only take actions which lie in the future of our deliberations, why should those actions not have an effect which lies in the past? Price and Weslake in fact consider this possibility, discussing a case where '*something we can choose to make the case in the future is suitably correlated with a state of affairs in the past*,' and admitting that '*disjunctive deliberation allows, at the margins, for these retroactive cases*.' So if the asymmetry of causation were really just about the fact that deliberation produces an output to the future of the deliberative process, there would seem to be no reason why the actions we take as a result of our forward-directed deliberation should not produce effects in the past. What needs to be added to this account is the consistent chaining requirement: the output of our deliberative process can only be used as an input to another process which has the same temporal orientation, meaning that its output will also lie to the future of our deliberative process, and thus we will find ourselves unable to take actions which have an impact on the past. In this sense, the consistent chaining requirement provides the final step needed in the perspectival account to rule out these retroactive cases, thus shoring up the perspectival account of the asymmetry of causation and the inaccessibility of the past.

4 Metaphysics

It seems possible that the popularity of the statistical account of the inaccessbility of the past may have its origins in the popularity of a specific *metaphysical* stance. That is, it seems to be strongly associated with the metaphysical picture which sees reality as being something like a Humean mosaic - a collection of categorical properties instantiated at points with various kinds of correlations between them [42]. For if this is the kind of picture one has in mind, then indeed it is hard to understand why there should be such a strong distinction between past and future: an agent must be understood as taking action by 'reaching into the mosaic' to set the values of the categorical properties at some point p, and since the value at the point p will typically be correlated both with categorical properties which are in the past relative to p and categorical properties which are in the future relative to p, it would seem that the agent who sets the value at that point should be understood as influencing both the past and the future (or neither, for those who couple this Humean picture with an eliminativist approach to causal talk.) In any case there does not seem to be any good reason to think of the action of setting the value of a categorical property at a point as having causal efficacy towards the future and not towards the past, so if this is the picture one has in mind, then it would seem natural to think that the difference between the past and the future can only be explained by appeal to a qualitative asymmetry such as the thermodynamical gradient.

However, one may wonder if the Humean metaphysics is really a good fit with what we know about modern physics. For the picture of reality we inherit from modern physics is not a mosaic of categorical properties: physics doesn't describe properties sitting inertly at points, it largely deals with processes, interactions and relationships. Of course, the Humean metaphysics is still capable of providing an account of processes and relationships, via the Humean reduction of modal structure to the best systematisation of the Humean mosaic [6,7] - and indeed, as we have noted, the Humean picture is compatible with our process-based account of the inaccessibility of the past, as the modal structure read off the Humean mosaic must presumably be free of contradictions. However, the Humean metaphysics may not seem like the most natural home for an approach which emphasizes process in this way. In particular, there have been persistent objections to the effect that the Humean account of modality does not really allow modal notions to actually be explanatory of anything [43, 44] [45, pp. 40], and if one is sympathetic to these objections, one might feel that the Humean approach does not give due weight to the explanatory of force of consistent chaining requirements. For after all, if a Humean accepts the arguments we have made here, they can still only reach the fairly weak conclusion that since modal structure is necessarily defined in a non-contradictory way, we must *describe ourselves* as being able to influence the future and not the past - that is, the Humean is compelled to think of consistent chaining as just fact about our modes of description which doesn't reflect any deep fact about the structure of reality. If we want to treat the requirement that modal structure should be non-contradictory as giving rise to real and meaningful constraints on our actions, we may want to take the basic ingredients of the process-based account more seriously - for example, by moving towards a process metaphysics.

Process metaphysics has a history stretching right back to Heraclitus [46] and these ideas have been developed by many philosophers since including Hegel [47], Heidegger [48] and Whitehead [49, 50]. In most cases, a process-based metaphysics is associated with a commitment to the idea that reality has a fundamentally dynamical nature - reality is the '*self-unfolding of dynamic structures or templates*' (according to Hegel) or '*growing together of the total available information of the universe at that time, according to certain principles, repeating and reinforcing certain patterns ("eternal objects") and thereby creating new ones*' (according to Whitehead) [51]. And this may at first seem like an appealing option in the context of the consistent chaining picture, for within a metaphysics which takes dynamical production literally it is straightforward to understand how consistent chaining comes about - if we think of one moment of time as being produced from the previous moment in a dynamical flow of time, it follows that outputs are really different from inputs in an objective sense, since outputs are produced from inputs by dynamical evolution and not vice versa. Thus 'outputs' present in this moment can't be used to determine facts about the previous moment from which this one was produced; they can only be used to determine facts about the moments which are yet to be generated, and thus the mechanism of time evolution ensures that consistent chaining is enforced.

However, this literal take on dynamical production is somewhat at odds with our conclusion that the consistent chaining picture supports a perspectival account of causation, because if reality has a fundamentally dynamical nature then there would in fact be a preferred direction for the process graph. Indeed, the fact that the apparent asymmetry of causation can be explained in the consistent chaining picture without any appeal to dynamic evolution significantly undercuts the usual motivations for the dynamical production metaphysics. Moreover, the known laws of nature seems to give credence to the perspectival approach rather than the dynamical production approach: when we examine the processes permitted by the laws of our current theories we typically find that there is no deep distinction between inputs and outputs at the mathematical level, as the theories from which we are constructing our 'processes' are by and large time-symmetric and hence invertible, so we can produce the output by 'evolution' from the input or we can equivalent produce the input by 'evolution' from the output. Thus there seems little justification for insisting that there is an underlying asymmetric process such that inputs are really metaphysically different from outputs. More generally, it has been noted that dynamical production approaches are somewhat in tension with relativistic physics [52]. In particular, many of the 'processes' we deal with in operational approaches to quantum mechanics involve sets of laboratories which accept an input at one spacetime point and produce an output at a spacetime separated point; although these processes are not signalling, nonetheless they exhibit nontrivial dependence of the outputs on the inputs, so if we are really committed to the view that outputs are produced by dynamical evolution from inputs and not vice versa it would seem that we are also committed to a preferred reference frame or family of reference frames in which the outputs occur in the future of the inputs, and this places us in tension with the relativistic notion that there exist no preferred reference frames. Furthermore, many theories taken seriously by modern physics do not obviously seem to incorporate anything which looks like dynamical evolution: examples include the Einstein equations in their usual form, Lagrangian and path integral methods, the canonical quantisation of gravity (which famously leads to a 'timeless' model with a time-independent Schrödinger equation [53]), and a number of interpretations of quantum mechanics which deny that the theory involves anything like a literal process of dynamical evolution. (See refs [5, 54, 55]). Thus it would seem that modern physics gives us good reason to reject the dynamical evolution picture, and indeed this may be part of the reason for the popularity of the Humean approach.

But of course, dynamical production and the Humean mosaic are not the only possible options: rejecting a literal approach to dynamical evolution does not mean we are obliged to adopt a picture of reality as a mosaic of unrelated categorical properties. Reality could instead be something like a process graph, i.e. a set of events standing in specified relations to other events as dictated by some underlying modal structure. Crucially, there is no need to think of a process graph as being generated in some particular order - for after all, a process is already, by definition, a process, and at least for our purposes there is no need to add a second processual layer in which the processes come into being in some particular direction. Because nomological relationships are not intrinsically directed, we can think of this process graph as existing eternally and atemporally, just like the block universe. In a sense this metaphysical picture is actually closer to modal structural realism [4] rather than a traditional process view, as our realism is directed towards the modal structure underlying the processes, rather than the coming-into-being of the processes. Thus rejecting the dynamical evolution picture does not compel us to revert back to the Humean approach: it is open to us to accommodate a metaphysics of processes within a block-universe picture. This approach allows us to give due weight to the inherently processual nature of reality without coming into conflict with relativistic physics or contradicting the perspectival account of causality.

That said, if we can't invoke a literal process of dynamical evolution to explain how consistent chaining is enforced, we will have to come up with some other approach to explain what exactly prevents agents from manipulating outputs in the way that they manipulate inputs. The very fact that nomological relations are not intrinsically directed makes it difficult to see immediately how this could follow directly from our fundamental laws. It may be tempting to say that consistent chaining is somehow enforced by spacetime structure - for example, we might insist either that the spacetime contains no CTCS or that if it does, the geometry of the spacetime must be consistent with a topology that does not contain any CTCs (meaning that it must be time-orientable [56]); then since different possible values of the matter fields are necessarily reflected in different values of the stress-energy tensor and thus also different values of the Einstein tensor, it follows that the matter fields must also be consistent with a topology that does not contain any CTCs. However, this runs up against long-standing debates about whether or not spacetime itself can actually be explanatory of anything [57,58]. In particular, since the spacetime structure of general relativity is dynamical, we can't straightforwardly say that spacetime constrains the matter on it to behave in a certain way: the fact that the manifold as a whole has some feature like time-orientability can't fully explain the behaviour of the matter in that spacetime, because the behaviour of the matter is partly responsible for the fact that the manifold is time-orientable! It's possible that one might have to turn to a theory of quantum gravity to explain exactly how consistent chaining gets enforced - for example, we have noted that the causal set approach enforces the 'no-closed-loops' criterion directly at the level of the underlying process graph. We will not have space to address these questions in greater detail here, but we hope to elaborate further in future work.

Retrocausality The consistent chaining requirement would seem to imply that retrocausality is impossible outside of CTCs; or at least, there cannot be any processes which reliably implement retrocausality outside of CTCs. However, it is important to keep in mind that as detailed in ref [59], there are two very different notions of retrocausality. Some approaches to retrocausality postulate two distinct directions of dynamical causality which together determine intermediate events by forwards and backwards evolution respectively from separate and independent initial and final states - for example, the forwards-evolving state and the backwards-evolving state in the two-state vector interpretation [60]. Other approaches postulate an 'all-at-once,' picture where the laws of nature apply atemporally to the whole of history, as for example in Wharton's all-at-once Lagrangian models [61]; in such a picture the past and the future have a reciprocal effect on one another, so there is definitely some kind of influence from the future to the past at play, but these effects can't be separated out into separate forwards and backwards evolutions. (We will not address here the question of whether these reciprocal all-at-once influences can be properly called causal - there are certainly legitimate reasons to dispute this terminology, but nonetheless we will continue to refer to all-at-once approaches as 'retrocausal' in order to maintain consistency with the literature.)

Clearly the existence of dynamical retrocausality would violate the consistent chaining requirement, and thus it would be expected to lead to paradoxes of various kinds - indeed, refs [2, 62] describe several quantum-mechanical experiments where attempting to give an account in terms of dynamical retrocausality does lead to paradoxes. But as explained in ref [59], all-at-once retrocausality does not produce the same problems. And evidently all-at-once retrocausality is in any case better suited for the block-process metaphysics which we advocate here. In particular, we have reinforced that the relationships underlying the process picture need not be thought of as asymmetric causal relationships; there are good reasons to think that they are really symmetrical, reciprocal relationships which appear asymmetric to us only because we happen to be instantiated in some particular direction along the chain of processes. And in a way, any such symmetric processes are automatically 'retrocausal' in the all-at-once sense, as a symmetric modal relation entails that the future influences the past and simultaneously the past influences the future, i.e. past and future are on a par in terms of causal influence. Thus consistent chaining requirements do not rule out retrocausality in the more sophisticated all-at-once sense. Evidently the kind of retrocausality involved in chains of signalling processes is not particularly interesting, since we can always describe such a scenario using a straightforward directed causal order and thus no novel phenomena will arise from this kind of retrocausality, but more interesting effects may arise when this kind of 'retrocausality' is combined with quantum phenomena and/or nonlocality [61, 63].

5 Non-causal processes

The consistent chaining requirement and the absence of CTCs, together with the assumption that all processes which are not composed out of other processes are causal, comprise a sufficient condition to ensure that modal structure of the universe does not give rise to contradictions. However, the condition is not a necessary one - ongoing research on the process matrix formalism has demonstrated the existence of possible processes which are not causal but which allow local agents to have free will whilst preventing contradictions, and if such non-causal processes are actually realised in nature, we could in principle have CTCs or violations of consistent chaining whilst still avoiding contradictions - for example, ref [21] shows that dynamics defined on a closed causal curve can be written in the form of a non-causal process.

However, it is not presently known if these non-causal processes are in fact realised in nature. So far, all processes which have been implemented by real agents are not only causal but causally separable, meaning that they can be written as a probabilistic mixture of causally ordered processes. A quantum process may fail to be causally separable if it involves a superposition of two different causal orders, and some such processes have been demonstrated using quantum operations in place of laboratories [64, 65], but since quantum operations are not agents who choose freely which operation to perform, arguably these demonstrations do not count as genuine implementations of causally nonseparable processes. But if it is possible at least in principile to put human agents into superpositions then it should also be possible in principle to implement superpositions of causal orders using real agents, and thus it may well be the case that the modal structure of the universe allows genuine implementations of causally nonseparable processes[5]. However, the kinds of processes implemented in these experiments still have causal process matrices and do not violate causal inequalities. At present there is no known way of creating non-causal processes using either classical or quantum mechanics, which leaves us with a crucial question: why does nature only seem to implement a limited subset of the possible process matrices, i.e. the causal ones? If the explanation for the inaccessibility of the past is the fact that modal structure of our universe only allows processes which do not give rise to contradictions, one might naturally think that it should be possible to gain some kind of access to the past by means of non-causal processes, provided that they have process matrices and thus cannot give rise to contradictions.

As a first response, we observe that it could in fact be the case that nature does in fact implement a more general class of processes in regimes that we have not yet probed. In particular, an important motivation for the proposal of the process matrix framework was the idea that some of these indefinite causal structures might actually appear in quantum gravity [18–20]. This suggests a different kind of answer to the questions we alluded to in section 4: rather

[5] That said, as noted in ref [66] the sorts of superpositions created in experimental demonstrations utilize a quantum control system S which determines the order of operations, and thus the control system is a common cause of all the events involved; so in a sense these sorts of experiments can't really be regarded as a superposition of the two causal structures 'A causes B' and 'B causes A,' because really both cases have the same causal structure: 'S causes A and B.' And ref [67] shows that for many types of processes, 'pure' superpositions of causal order without a control are not possible within the framework of standard quantum mechanics.

than looking for a theory of quantum gravity which enforces 'no closed loops' at the fundamental level, like causal set theory, we might instead choose to seek a theory which allows all possible process matrices and then show that only the causal processes survive in some appropriate limit, so consistent chaining is enforced only in that limit. Thus it remains possible that one day, if we eventually manage to gain experimental control over quantum gravity processes, we will then actually be able to 'access' the past, in the sense that we will be able to make observations which are not compatible with events having a strict partial order - although we will still not be able to 'change' the past, as process matrices are defined specifically so as to prevent contradictions from arising.

That said, there have also been suggestions for physical principles which might limit the set of physically possible processes to the causal ones - for example, ref [68] suggests that processes are physical only if they are 'purifiable,' (i.e. they preserve the reversibility of the underlying operations) and this proposal rules out many non-causal processes, although not all of them. In addition, it is important to note that the claim that non-causal processes are conceptually possible relies on the idea that although we always have free choice over local operations, we do not necessarily have free choice about which dynamics we implement between labs. For example, if you find yourself on a CTC, prima facie there would seem to be no reason why you should not be able to construct an identity operation going all the way around the CTC - intuitively, you can just write your input down on a piece of paper and then walk all the way around the CTC until you arrive back at your starting location. But as we have seen, the process matrix formalism tells us that this is not a possible dynamics on a CTC, so in this setting it seems that there is something intrinsically uncontrollable about dynamics, whereas local operations are under our control. And yet in a sense one can think of dynamics as being just a sequence of chained local operations - for example, when you carry your piece of paper around the CTC you are essentially implementing a long chain of local identity operations. Thus one might feel that the distinction being made here is too strong: if we have free choice over local operations, we should also be able to freely construct dynamics on a CTC by means of chains of local operations, and so the kinds of restrictions placed on possible dynamics by the indefinite causal structure program are actually incompatible with free choice in a broader sense. This line of reasoning would seem to suggest that non-causal processes should not actually be physically possible.

A possible compromise would be to suggest that the regimes in which non-causal processes can be implemented are precisely those regimes where the dynamics are less controllable - so for example, presumably it is not possible to have CTCs in macroscopic regimes where we would expect to be able to 'walk around' the loop exerting direct control over the dynamics at every point, but perhaps it might be possible to have CTCs in quantum gravitational regimes where we may be able to provide 'inputs' at various points in the loop but will not necessarily be able to control the dynamics at every point in the loop. To make this approach work, we would need to show that the limit in which consistent chaining is enforced coincides with the limit in which dynamics become controllable. That said, recall that limits of this kind are often statistical - for example, this is the case for the thermodynamical limit [69] and the quantum-to-classical transition [70] - whereas we have argued in this paper that the principles enforcing the absence of contradictions cannot be merely statistical. For example, if the generalization that 'closed causal loops do not occur at macroscopic scales' were only statistically valid, then it would follow that after enough attempts we should be able to create a macroscopic closed causal loop and thus implement the identity operation on it, giving rise to a contradiction. Therefore if noncausal processes are possible in certain regimes, there must be some sufficiently strong and presumably non-statistical reason why these processes cannot be performed in regimes where the dynamics become more controllable.

6 Conclusion

We have argued that the popular account of the inaccessibility of the past as a statistical effect arising out of entropy gradients is not the full story: at the most basic level, the inaccessibility of the past follows from the fact that modal structure must be defined in such a way that it cannot give rise to contradictions, and this means that all possible processes must be associated with valid process matrices or functions. Moreover we observed that all processes observed in nature thus far have process matrices or functions which are causal, meaning that individual subprocesses must be chained together in specific ways to ensure consistency, and it turns out that this consistency requirement is deeply entwined with our usual notions of temporal order. Observers internal to the universe have the form of processes and therefore they themselves will necessarily be oriented one way or another in the chain of processes, and will find themselves capable of acting to produce effects 'downstream' but not 'upstream' - so the direction of time, for those observers, will simply be defined by their own orientation, regardless of whether or not there is any difference between the two directions from the external point of view. We have applied these ideas to the asymmetry of causation and argued that not only is causation perspectival in the sense of Price and Ramsay, in fact the asymmetry of causation does not have to be grounded in any objective asymmetry, so it does not strictly speaking depend on the thermodynamical asymmetry. Of course thermodynamics is probably a large part of the explanation for the fact that conscious agents exist and have a perspective in the first place, but conditional on the existence of such agents, it is not thermodynamics which is primarily responsible for the fact that they find themselves unable to access the past.

While this explanation of the inaccessibility of the past can feasibly be offered within any metaphysical picture which offers a sufficiently robust notion of modality, we consider that the Humean approach may not give due weight to the explanatory force of consistent chaining requirements, whilst dynamical production pictures are somewhat at odds with our view that causation should be regarded as perspectival within the consistent chaining picture. Thus we have suggested that the most appropriate metaphysics for this approach is a 'block-process' approach which reifies the process graph and thus might be regarded as a form of modal structural realism. We have also considered whether violations of consistent chaining might be possible in the form of 'non-causal' processes defined in such a way that they cannot give rise to contradictions, and have concluded that such phenomena might be possible in special regimes provided that causal processes emerge appropriately in some limit.

7 Acknowledgements

This publication was made possible through the support of the ID 61466 grant from the John Templeton Foundation, as part of the "The Quantum Information Structure of Spacetime (QISS)" Project (qiss.fr). The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the John Templeton Foundation.

References

[1] Huw Price and Brad Weslake. The time-asymmetry of causation. In Helen Beebee, Peter Menzies, and Christopher Hitchcock, editors, *The Oxford Handbook of Causation*. Oxford University Press, 2008.

[2] Colm Bracken, Jonte R. Hance, and Sabine Hossenfelder. The quantum eraser paradox, 2021.

[3] E. Adlam. Contextuality, fine-tuning and teleological explanation. *Found Phys*, 51, 2021.

[4] Nora Berenstain and James Ladyman. Ontic structural realism and modality. In Elaine Landry and Dean Rickles, editors, *Structural Realism: Structure, Object, and Causality*. Springer, 2012.

[5] Emily Adlam. Laws of nature as constraints. *Foundations of Physics*, 52, 2022.

[6] Siegfried Jaag and Christian Loew. Making best systems best for us. *Synthese*, 197:2525–2550, 2020.

[7] David Lewis. A subjectivist's guide to objective chance. In Richard C. Jeffrey, editor, *Studies in Inductive Logic and Probability*, pages 83–132. University of California Press, 1980.

[8] James Ladyman and Katie Robertson. Going round in circles: Landauer vs. norton on the thermodynamics of computation. *Entropy*, 16(4):2278–2290, 2014.

[9] E Brian Davies and John T Lewis. An operational approach to quantum probability. *Communications in Mathematical Physics*, 17(3):239–260, 1970.

[10] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Informational derivation of quantum theory. *Physical Review A*, 84(1), Jul 2011.

[11] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Probabilistic theories with purification. *Phys. Rev. A*, 81:062348, Jun 2010.

[12] J. Barrett. Information processing in generalized probabilistic theories. *eprint arXiv:quant-ph/0508211*, August 2005.

[13] H. Barnum, J. Barrett, L. Orloff Clark, M. Leifer, R. Spekkens, N. Stepanik, A. Wilce, and R. Wilke. Entropy and information causality in general probabilistic theories. *New Journal of Physics*, 12(3):033024, March 2010.

[14] Lluís Masanes and Markus P Müller. A derivation of quantum theory from physical requirements. *New Journal of Physics*, 13(6):063001, Jun 2011.

[15] I. D. Novikov. Time machine and self-consistent evolution in problems with self-interaction. *Phys. Rev. D*, 45:1989–1994, Mar 1992.

[16] John Bell. Time and causation in gödel's universe.

[17] Paul Horwich. Asymmetries in time. *Noûs*, 24(5):804–806, 1990.

[18] Ognyan Oreshkov and Christina Giarmatzi. Causal and causally separable processes. *New Journal of Physics*, 18, 09 2016.

[19] O. Oreshkov, F. Costa, and Č. Brukner. Quantum correlations with no causal order. *Nature Communications*, 3:1092, October 2012.

[20] Mateus Araújo, Cyril Branciard, Fabio Costa, Adrien Feix, Christina Giarmatzi, and Časlav Brukner. Witnessing causal nonseparability. *New Journal of Physics*, 17, 10 2015.

[21] Ämin Baumeler, Fabio Costa, Timothy Ralph, Stefan Wolf, and Magdalena Zych. Reversible time travel with freedom of choice. *Classical and Quantum Gravity*, 36, 11 2019.

[22] Cyril Branciard, Mateus Araú jo, Adrien Feix, Fabio Costa, and Časlav Brukner. The simplest causal inequalities and their violation. *New Journal of Physics*, 18(1):013008, dec 2015.

[23] Adrien Feix, Mateus Araú jo, and Caslav Brukner. Causally nonseparable processes admitting a causal model. *New Journal of Physics*, 18(8):083040, aug 2016.

[24] Ding Jia and Nitica Sakharwade. Tensor products of process matrices with indefinite causal structure. *Phys. Rev. A*, 97:032110, Mar 2018.

[25] J.T. Ismael. *How Physics Makes Us Free*. Oxford University Press, 2016.

[26] Mathias Frisch. *Causal Models and the Asymmetry of State Preparation*, pages 75–85. Springer Netherlands, Dordrecht, 2010.

[27] J. Bell. Against 'measurement'. *Physics World*, August 1990.

[28] Abner Shimony. *Controllable and uncontrollable non-locality*, volume 2, page 130–139. Cambridge University Press, 1993.

[29] C. J. Wood and R. W. Spekkens. The lesson of causal discovery algorithms for quantum correlations: causal explanations of Bell-inequality violations require fine-tuning. *New Journal of Physics*, 17(3):033002, March 2015.

[30] Karel V. Kuchař. Canonical quantum gravity. *arXiv: General Relativity and Quantum Cosmology*, 1993.

[31] Edward Anderson. Geometrodynamics: Spacetime or space ? *arXiv: General Relativity and Quantum Cosmology*, 2004.

[32] R. Sorkin. Geometry from order: causal sets. *Einstein Online*, 2006.

[33] C. Wuthrich and C. Callender. What becomes of a causal set. *ArXiv e-prints*, January 2015.

[34] H. Price. *Time's Arrow & Archimedes' Point: New Directions for the Physics of Time*. Oxford Paperbacks: Philosophy. Oxford University Press, 1996.

[35] Sean M. Carroll. Why boltzmann brains are bad, 2017.

[36] Huw Price. Causal perspectivalism. In Huw Price and Richard Corry, editors, *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*. Oxford University Press, 2005.

[37] Ned Hall. Two concepts of causation. In John Collins, Ned Hall, and Laurie Paul, editors, *Causation and Counterfactuals*, pages 225–276. MIT Press, 2004.

[38] David Lewis. Causation. *Journal of Philosophy*, 70(17):556–567, 1973.

[39] David Lewis. Counterfactual dependence and time?s arrow. *Noûs*, 13(4):455–476, 1979.

[40] Judea Pearl. *Causality*. Cambridge university press, 2009.

[41] James Woodward. *Making things happen: A theory of causal explanation*. Oxford university press, 2005.

[42] David Lewis. Humean supervenience debugged. *Mind*, 103(412):473–490, 1994.

[43] Tim Maudlin. *The Metaphysics Within Physics*. Oxford University Press, 2007.

[44] Marc Lange. Grounding, scientific explanation, and humean laws. *Philosophical Studies*, 164(1):255–261, 2013.

[45] D. M. Armstrong. *What is a Law of Nature?* Cambridge University Press, 1983.

[46] G. S. Kirk. 7. natural change in heraclitus. In Alexander P. D. Mourelatos, editor, *The Pre-Socratics: A Collection of Critical Essays*, pages 189–196. Princeton University Press, 1994.

[47] Georg Wilhelm Friedrich Hegel. *Hegel's Philosophy of Nature*. Oxford, Clarendon Press, 1970.

[48] Martin Heidegger. *Sein Und Zeit*. M. Niemeyer, 1927.

[49] Alfred North Whitehead. *An Enquiry Concerning the Principles of Natural Knowledge*. Cambridge University Press, 1919.

[50] Alfred North Whitehead. *The Concept of Nature: Tarner Lectures*. Cambridge University Press, 1920.

[51] Johanna Seibt. Process Philosophy. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2022 edition, 2022.

[52] Oliver Pooley. Relativity, the open future, and the passage of time. *Proceedings of the Aristotelian Society*, 113:321–363, 2013.

[53] Chris Isham. Canonical quantum gravity and the problem of time. *Integrable Systems, Quantum Groups, and Quantum Field Theories*, 11 1992.

[54] Emily Adlam. Determinism beyond time evolution, 2021.

[55] Eddy Keming Chen and Sheldon Goldstein. Governing without a fundamental direction of time: Minimal primitivism about laws of nature, 2021.

[56] E. Minguzzi and Miguel Sánchez. *The causal hierarchy of spacetimes*, pages 299–358. 01 2008.

[57] H.R. Brown and Oxford University Press. *Physical Relativity: Space-time Structure from a Dynamical Perspective*. Oxford Scholarship Online. Philosophy module. Clarendon Press, 2005.

[58] Harvey R. Brown and James Read. The dynamical approach to spacetime theories, April 2018. Forthcoming in E. Knox and A. Wilson (eds.), "The Routledge Companion to Philosophy of Physics", London: Routledge, 2019.

[59] Emily Adlam. Two roads to retrocausality, 2022.

[60] Yakir Aharonov and Lev Vaidman. *The Two-State Vector Formalism of Quantum Mechanics*, pages 369–412. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.

[61] Ken Wharton. A new class of retrocausal models. *Entropy*, 20(6):410, May 2018.

[62] Tim Maudlin. *Quantum Non-Locality and Relativity: Metaphysical Intimations of Modern Physics*. Blackwell, 2002.

[63] H. Price. A Neglected Route to Realism About Quantum Mechanics. *ArXiv General Relativity and Quantum Cosmology e-prints*, June 1994.

[64] K. Goswami, C. Giarmatzi, M. Kewming, F. Costa, C. Branciard, J. Romero, and A. G. White. Indefinite causal order in a quantum switch. *Physical Review Letters*, 121(9), aug 2018.

[65] Giulia Rubino, Lee Rozema, Adrien Feix, Mateus Araújo, Jonas Zeuner, Lorenzo Procopio, Časlav Brukner, and Philip Walther. Experimental verification of an indefinite causal order. *Science Advances*, 3, 08 2016.

[66] Jean-Philippe W. MacLean, Katja Ried, Robert W. Spekkens, and Kevin J. Resch. Quantum-coherent mixtures of causal relations. *Nature Communications*, 8(1), may 2017.

[67] Fabio Costa. A no-go theorem for superpositions of causal orders. *Quantum*, 6:663, March 2022.

[68] Mateus Araú jo, Adrien Feix, Miguel Navascués, and Časlav Brukner. A purification postulate for quantum mechanics with indefinite causal order. *Quantum*, 1:10, apr 2017.

[69] Wayne C. Myrvold. Statistical mechanics and thermodynamics: A maxwellian view. *Studies in History and Philosophy of Science Part A*, 42(4):237–243, 2011.

[70] M.A. Schlosshauer and M.S.S.L.O. Maximilian Schlosshauer-Selbach. *Decoherence: And the Quantum-To-Classical Transition*. The Frontiers Collection. Springer, 2007.