Philosophy of Science (2023), **90**, 855–879
doi:10.1017/psa.2023.55

ARTICLE

The Temporal Asymmetry of Influence Is Not Statistical

Emily Adlam[1,2]

[1]The Rotman Institute of Philosophy, Western University, London, ON, Canada and [2]Chapman University, Orange, CA, USA
Email: eadlam90@gmail.com

(Received 11 May 2022; revised 29 July 2022; accepted 09 February 2023; first published online 03 April 2023)

Abstract

I argue that the temporal asymmetry of influence is not merely the result of thermodynamics: it is a consequence of the fact that the modal structure of the universe must admit only processes that cannot give rise to contradictions. I appeal to the process matrix formalism developed in the field of quantum foundations to characterize processes that are compatible with local free will while ruling out contradictions, and I argue that this gives rise to "consistent-chaining" requirements that explain the temporal asymmetry of influence. I compare this view to the perspectival account of causation advocated by Price and Ramsey.

1. Introduction

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 temporal asymmetry of influence is identified 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" (Price and Weslake 2008), which in turn is supposed to be a consequence of the thermodynamic gradient. This entails that the temporal asymmetry of influence is, like the second law of thermodynamics, statistical and approximate rather than a fundamental feature of reality.

Intuitively this may seem puzzling, as the temporal asymmetry of influence does not feel in any way approximate: the inaccessibility of the past is an unforgiving, rock-solid barrier. Of course, prescientific intuitions have often been shown to be mistaken in the course of scientific progress. But in this article, I 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 temporal asymmetry of influence.

I begin in section 2 by arguing that because the universe cannot contain contradictions, the modal structure of the universe must admit only those processes that cannot give rise to contradictions, and I appeal to the process matrix formalism developed in the field of quantum foundations to characterize the complete set of processes that are compatible with local free will while 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"—that is, 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. Because we ourselves are, in effect, processes that 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. In section 3, I compare my view to the perspectival account of causation advocated by Price and Ramsey; finally, in section 4, I comment on the metaphysics associated with my approach, and in section 5, I discuss the possibility that some noncausal processes may be realized in nature.

2. Consistency requirements on modal structure

Although the intuition that we are able to influence the future and not the past is very widespread, the exact content of this intuition is not always clear. Sometimes it is cashed out in terms of the asymmetry of causation or by appeal to the notion that the future is "open" in a way that the past is not, but both of these analyses themselves demand substantive ontological commitments—for example, proponents of a block-universe view will not generally agree that we should think of the future as "open." In recognition of these difficulties, in this article, I will use the terminology *temporal asymmetry of influence* (following Albert 2014) to refer to this vaguely defined commonsense intuition about the difference between past and future; later, I will offer a more precise account of the content of this intuition that will be appropriate for the analytic framework that I adopt in this article.

The most generally accepted explanation for the "temporal asymmetry of influence" 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 (AKL) attributes the temporal asymmetry of influence to the fact that, due to thermodynamic asymmetries, "there are (as it were) a far wider variety of potentially available routes to influence over the future … there are a far wider variety of what you might call causal handles on the future … than there are on [the] past" (Albert 2014). This approach is based on the observation that as a result of the thermodynamic gradient, we typically have more records of the past than the future, and therefore the semantics of counterfactuals are temporally asymmetric: backtracking counterfactuals with consequents that are not true in the actual world are usually false, whereas forward counterfactuals with consequents that are not true in the actual world are often true. Thus, insofar as counterfactuals of this kind can be regarded as encoding our ability to bring about events in the world, it turns out to be the case that we can usually bring about events in the future but not in the past.

But it has often been observed that there is something odd about this account, for it is commonly agreed that the thermodynamic gradient is merely approximate, arising from statistical averages over large numbers of systems, and therefore the

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

asymmetry of records on which this account is predicated is likewise statistical and approximate. Thus the AKL account seems to suggest that we should be able to influence the past at least a little, particularly if we get really good at manipulating things at a microscopic level, because then we will able to circumvent the asymmetry of records in order to make certain backtracking counterfactuals true. Yet that is not how the difference between past and future presents itself to us—we do not typically imagine that by sufficiently careful manipulation, we will be able to influence the past. Indeed, if the difference between past and future were really statistical and epistemic, 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, and indeed, as shown by Adlam (2021), quantum mechanics seems to be very carefully fine-tuned to prevent the possibility of "signaling backward in time."

Now, AKL do have a response to this criticism: because the falsity of backtracking counterfactuals is grounded on the existence of records of the consequents, in special cases where backtracking counterfactuals are true and hence we can be said to influence the past, there cannot be any current records or memories of the consequent. Thus, whenever there are circumstances such that the past does depend counterfactually on the future, it is built into the nature of these circumstances that we "can have no way whatsoever of knowing, and [we] can have no grounds whatsoever for suspecting, when it is that they actually obtain" (Albert 2014)—that is, it may sometimes be possible to influence the past, but we can be assured that there will never be any records of such influences. But it's unclear that such epistemic considerations are adequate to explain the nature of our relationship with the past and future, for our view that we can't influence the past does not rest simply on the fact that we have no *records* of instances where we have successfully influenced the past: we also don't seem to have any idea how we could possibly go about influencing the past through our actions right now, regardless of whether or not there will be records of these influences. After all, as just noted, quantum mechanics and, indeed, all other mainstream physical theories specifically forbid signaling backward in time, and thus current physics provides us with many mechanisms by which we can bring about events in the future but no mechanisms by which we can bring about events in the past. Of course, AKL could perhaps argue that we are simply not aware of such mechanisms because we have forgotten all the evidence for them—but although this may be a coherent view, it is quite an extreme one, requiring us to postulate a whole realm of hidden physical processes for which there cannot possibly be any evidence. Before postulating a hidden sector of reality unbeknownst to modern physics, it would seem reasonable to see if our existing physics can account for the temporal asymmetry of influence on its own terms—and in this article, I will argue that indeed it can.

My approach is motivated by careful consideration of a specific problem faced by the AKL account. To begin, suppose there exists a process P_1 that can be performed in the forward direction of time (i.e., it works as normal if its input is in the past lightcone of its output) and another process P_2 that can be performed in the backward direction of time (i.e., it works as normal if its input is in the future lightcone of its

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

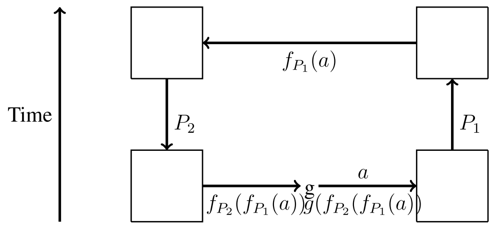

Figure 1. Schematic diagram of the composition of two processes used to demonstrate the need for consistent chaining.

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 $f_{P_2}\left(f_{P_1}(a)\right)$. 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 a single spacetime point, which most people would presumably consider to be logically and/or physically impossible.

The AKL approach 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 "backward" in this sense only because we are trapped in the onward flow of the thermodynamic gradient. But if this prohibition were really just statistical, then one would think that we should be able to circumvent it if we were to repeat this procedure enough times. Of course, the number of repetitions needed might be far too great for any real agent to actually accomplish during its lifetime, but the practicalities aren't important here—the point is that if our inability to influence the past were really just statistical, then there would be a nonzero chance of creating this kind of loop, and surely this cannot be right. If contradictions of this kind are indeed logically and/or physically impossible, their absence cannot be a statistical matter: contradictions must be ruled out absolutely, and thus the universe must function in such a way that contradictions like this simply cannot arise.

Moreover, note that it does not help here to argue, as in the AKL approach, that in the event that backward causation does occur, there will not be any records or memories of it because the requirement that the world should not contain contradictions is not epistemic—most people would presumably insist that no physical quantity can have two different values at a single spacetime point, regardless of whether or not there exist records of these values. So even if one accepts that the thermodynamic approach is consistent with all the evidence, at a higher level of analysis, it still seems to have problems because a purely statistical account of our inability to influence the past does not properly take into account the close relationship between our inability to influence the past and the absence of contradictions, as depicted in figure 1.

2.1 The process view

With this motivation in mind, I now set out to provide an alternative account of the temporal asymmetry of influence based on the simple requirement of avoiding contradictions. For concreteness, we will 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] I employ this conceptual framework because the block-universe approach is favored by a significant number of modern-day philosophers and physicists, in large part because 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.

As noted earlier, the commonsense intuition about the temporal asymmetry of influence is somewhat vaguely defined. However, it does seem clear that this intuition must be understood as having modal content: we *can* influence the future, but we *cannot* influence the past (or at least, so it seems to us). This suggests that the correct way to understand the difference between past and future is to see how it arises from some underlying modal structure. In this article, I will not make any assumptions about the nature of this modal structure: it could be understood within a metaphysically robust approach to modality that regards it as ontologically prior to the content of physical reality (Berenstain and Ladyman 2012; Adlam 2022b), or it could simply be understood within a Humean approach as the best systematization of the actual Humean mosaic (Jaag and Loew 2020; Lewis 1980). I insist only that modality is regarded as being an objective feature of reality in some sense.

The AKL approach does, in fact, proceed with an analysis of modal structure because it is grounded on an analysis of the semantics of backtracking counterfactuals. However, the relationship between this semantic approach and the way modality is modeled in physics is not always very clear, and thus in this article, the aim is to understand the modal structure of reality in a way that is more closely informed by physics. First, I model an *experiment* as a set of laboratories, in each of which agents may select inputs and/or receive outputs. I then define a *process* as a class of experiments that can be implemented anywhere in spacetime and can be composed in any desired way 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 probabilistic relation between inputs and

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

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

outputs would not be expected to hold. For example, for processes that can be regarded as transmitting signals, it is typically 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 probabilistic 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 Ladyman and Robertson (2014): "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 (Davies and Lewis 1970; Chiribella et al. 2011, 2010; Barrett 2005; Barnum et al. 2010; Masanes and Müller 2011) can be understood as characterizing the set of processes that can be constructed from quantum-mechanical systems on a fixed spacetime background.

This process-based approach allows us to finally be more precise about the nature of the intuition that we are able to influence the future and not the past. It can be reconceived as a statement about how we can use the outputs of our own decision-making processes—specifically, it seems that in general, we can use these outputs as inputs into other processes only if those processes produce outputs that are to the future of our deliberations, and not the past. Our task now is to understand how this restriction arises from fundamental requirements on modal structure.

Whether the modal structure is understood to determine the contents of reality (as in the metaphysically robust 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 systematization of the actual mosaic, and because the actual mosaic does not contain contradictions, a good systematization of it should not lead to any contradictions either—or at least it should predict only a negligibly small number of contradictions because 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 metaphysically robust approach where the modal structure determines the content of reality, the modal structure must evidently be formulated in such a way that it does not give rise to contradictions because, after all, it cannot ultimately produce a reality that contains contradictions.[2] Thus, in both the metaphysically robust and Humean approaches, it makes sense to ask about the conditions that the modal structure must satisfy if we are to avoid contradictions.

The demand for a 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 (Novikov 1992). Novikov's principle tells us that the only solutions to the laws of physics that can occur locally in the real universe are those that are globally self-consistent—all events happen only once and cannot be changed. However, whereas Novikov's principle is a constraint on the events that actually take place, our consistency requirement applies at the level of the 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 because the occurrent events cannot contain contradictions.

2.2 Temporal interdicts

What does this consistency requirement entail? First, note that processes can generally 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 the modal structure should not give rise to contradictions necessarily places some constraints on the way in which processes can be composed—in particular, it must be impossible to create compositions of the kind depicted in figure 1. Bell made a similar observation in his comments on time travel in the Gödel universe (Bell n.d.), noting that the universe must be subject to 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

[2] One way to guarantee this is to formulate the modal structure, as suggested by Adlam (2022b), 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. Because 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.

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

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, but the possible temporal interdicts are quite analogous to Bell's:

1. The modal structure of the universe does not allow processes that 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 3. Interdict 2 remains viable, but as noted in the philosophical literature on time travel (Horwich 1990, chap. 7], this approach would require us to accept that every attempt to create contradictions, as in figure 1, is foiled by apparently inexplicable coincidences. 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 1—observers locally have freedom of choice about how they compose processes, but the modal structure of the universe does not admit any processes that could be composed in such a way as to give rise to contradictions. This interdict is significantly stronger than a mere prohibition on contradictions—it is, in fact, a very general constraint on the *types of processes* allowed by the modal structure of the universe. That is, assuming that we have the freedom to compose processes in any way we like, this interdict asserts that we should not have access to any processes that *could* be used to construct contradictions (regardless of whether they are *in fact* used to create contradictions). Moreover, this option seems to be the one that is realized in our actual world: what stops us from creating a composition like that in figure 1 is not that things inexplicably go wrong every time we try but simply that we do not know any processes that 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. We now turn to the question of what that entails about the set of possible processes.

2.3 The process matrix framework

There is, in fact, a theoretical framework in quantum foundations that has been developed specifically to study the set of all possible processes that are compatible with observers locally having free choice but that also guarantees the absence of contradictions—the process matrix framework (Oreshkov and Giarmatzi 2016;

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

Oreshkov et al. 2012; Araújo et al. 2015). 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 that are not permitted by any known physics, such as dynamics in "loops," as 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 (Baumeler et al. 2019). 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'" (Baumeler et al. 2019). 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 or process functions.

The most straightforward type of process matrix is one that is *causally ordered*; that is, it is compatible with the existence of a strict partial order such that the probabilities for the outcome obtained in some laboratory depend only on the settings in experiments that precede it in the partial order. We can also imagine process matrices that are not causally ordered because the order for some subset of the laboratories may depend on the settings in other laboratories but that are still *causal* in the sense that we can write down a well-defined probability distribution over possible strict partial orders such that in each possible case, the choice of setting in a local experiment does not affect the occurrence of events that are earlier in the order or the strict partial order on the set of the events and the experiment in question (Oreshkov and Giarmatzi 2016)—that is to say, at each individual implementation, we will be able to put the experiments into a strict partial order. Causal processes cannot violate "causal inequalities" (Branciard et al. 2015) because the correlations they produce can always be embedded into a global strict partial order. And thus far, all processes known to be realized in nature (without postselection) are causal.

I will return to the topic of noncausal processes in section 5, but for now, I simply follow Feix et al. (2016) 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, because 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—that is, compatible with a strict partial order. Evidently, as long as all processes and compositions of processes are compatible with

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

an underlying strict partial order, contradictions of the kind depicted in figure 1 cannot occur.

As noted by Jia and Sakharwade (2018), defining a general composition rule for process matrices is not straightforward—simply taking the tensor product of two process matrices can sometimes produce something that 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 signaling causal processes—*bipartite* meaning the process takes an input in one lab and produces an output in another lab, and *signaling* meaning the output depends nontrivially on the inputs; that is, 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 signaling 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 signaling process if the input of that process is also in the past lightcone of its output. Provided that we are in a spacetime that does not contain 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. I will use the term *consistent chaining* to refer to the requirement that the output of a signaling process that is oriented forward in time can only be used as input to other signaling processes that are also oriented forward in time, thus giving rise to a process graph that is acyclic. I emphasize that consistent chaining arises not merely from a prohibition on contradictions, but a prohibition on the existence of processes that *could* be used to create contradictions by agents with complete freedom to create arbitrary compositions of processes. Recall that the justification for taking this approach is ultimately empirical: as discussed in section 2.2, temporal interdict 2 remains a logically possible alternative to consistent chaining, but as far as we can tell from our best current physics, all known physical processes do indeed obey consistent chaining.

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 indeed, Baumeler et al. (2019) demonstrate 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

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

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 signaling process that is oriented "backward in time" can only be used as input to signaling processes that are also oriented "backward in time"—from the block-universe perspective, these descriptions are equivalent. Thus, 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 (2016) 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," and this feature of our dynamical laws gives 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.

2.4 Agency

Now let us try to add agents such as ourselves into this picture. Evidently, we can think of a decision made by an agent as a kind of signaling 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 signaling processes that 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 that 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 signaling 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 signaling process if that process has the same temporal orientation as the deliberative process—which is to say, it has an output that 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 signaling 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 (2008), "we can only wiggle handles which lie in the immediate future, with respect to our own deliberations on the matter," or in Frisch's words, "It is a striking fact about

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

experimental interactions that we can only intervene into a system 'from the past,' as it were" (Frisch 2010).

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, that is, 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. Because these processes are signaling 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, because 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 backward" 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 signaling 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 that 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 because, 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 toward 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 that affect the "future" relative to their own orientation is an objective fact about the modal structure that arises from the consistent-chaining requirement. Thus, for such an agent, there is a very concrete distinction between past and future that has nothing to do with thermodynamics: given that the agent is oriented in a certain direction, they can only use their actions as inputs to processes that are oriented in the same direction, and therefore there is nothing epistemic or approximate about the temporal asymmetry of influence.

2.5 Example

One way to clarify the distinction between our approach to the temporal asymmetry of influence and the usual thermodynamical account is to consider a case where the thermodynamic asymmetry fails to hold in the usual way. For example, consider Price's example of the "Stargate Doughnut," involving a microscopic gate that may block the path of a photon or allow the photon to pass through (Price 2005). Because 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

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

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 temporal asymmetry of influence 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 bring it about that the photon is not 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 bring it about that the photon is not 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.

2.6 Indeterministic processes

The reasoning that led us to the need for consistent-chaining requirements was concerned specifically with deterministic signaling processes. But reality also appears to contain some signaling 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 earlier argument will not go through if we compose two indeterministic signaling 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 from its input, rendering the loop consistent. This is admittedly an unlikely turn of events, but there is no contradiction here.

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

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 from its input, and therefore in the context of this composition, the relative frequency with which these channels produce outputs different from their inputs will be 0.5, which is very far from the expected relative frequency of 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. Moreover, 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 this 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 because the probabilities do not match the usual ones under this sort of composition.

Note that as in the deterministic case, such inconsistencies can only arise if the composed process is not causal, and because 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 that are causal. So even when we have access to indeterministic signaling processes, agents can only output decisions that influence the "future" relative to their own orientation, and thus neither deterministic nor indeterministic signaling processes allow agents to influence the past in a spacetime without CTCs.

2.7 Nonlocality

We have so far considered only *signaling* processes, that is, 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 that are not signaling 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 (Bell 1990) 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 nonsignaling processes of this kind to create the composition shown in figure 1, we will not obtain either a contradiction or a probability inconsistency—in fact, the resulting process will still be causal because

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

the correlations will be compatible with a strict partial order (although 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, nonsignaling 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. It has been observed that the existence of quantum processes of this kind that are nonlocal but also nonsignaling is a striking feature of reality that seems in need of explanation: for example, if we try to represent such a process by a causal model, it can be shown that the parameters of the model must be very carefully fine-tuned (Wood and Spekkens 2015).The fact that processes of this kind must be nonsignaling 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 that could give rise to contradictions.

3. Causation

We have thus far discussed the temporal asymmetry of influence without couching that asymmetry 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 temporal asymmetry of influence 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: summarizing a common sentiment, Frisch (2010) 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." But what this way of thinking misses is that laws that 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 (2005), 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 because 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.

The account of causation that most closely aligns with the consistent-chaining approach is the perspectival analysis of Price (2005), who holds that the asymmetry of causation originates "in de facto asymmetries in our own temporal orientation, as physical structures embedded in time." 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 the words of Price (2005), "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." 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.

However, as it stands, the perspectival account is incomplete in certain ways. In particular, Price and Weslake (2008) suggest that a perspective is associated with a temporal direction by virtue of "the fact that we are always contemplating actions in the near future, with respect to the time of deliberation." But this fact in itself seems not enough to explain the temporal asymmetry of influence: for even if we can only take actions that lie in the future of our deliberations, why should those actions not have an effect that 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. And acknowledging this point seems to undermine the efficacy of the perspectival account as a way of explaining the asymmetry of causation—because now we still need to provide an explanation of why "these retroactive cases" do not actually seem to occur or at least occur very infrequently. In fact, what needs to be added to this account is exactly the consistent-chaining requirement: this requirement ensures that the output of a deliberative process can only be used as an input to another process that has the same temporal orientation, meaning that its output will also lie to the future of the deliberative process, and thus we will find ourselves unable to take actions that influence 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 temporal asymmetry of influence.

More generally, Price and Ramsay 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 temporal asymmetry of influence: Price (2005) 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." In contrast, I 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. Thus, the consistent-chaining account strengthens the perspectival account of Price and Ramsay because it entails that the perspectival asymmetry does not arise

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

approximately and statistically but is built into the modal structure of reality at a deep level.

I do, however, acknowledge, in agreement with Price, that the thermodynamic gradient may well be a necessary precondition for the *existence* of agents having a meaningful perspective (e.g., human beings), and thus I do not deny that the thermodynamical gradient remains an important part of a perspectival approach to causation. But I emphasize that the thermodynamic gradient itself depends on an underlying consistent-chaining requirement because it is a statistical asymmetry that arises out of underlying causal processes. Indeed, if the underlying processes were not causal, there presumably could not be any well-defined thermodynamic gradient at all—generic signaling backward in time would wreak havoc with the monotonicity of the second law, so without consistent-chaining requirements, we could, at best, hope to find small regions exhibiting extremely transient thermodynamic gradients, which would not all point in the same temporal direction. And this kind of universe would not be compatible with any substantive perspectival account; indeed, it's not even clear that we could meaningfully distinguish two well-defined directions of time in such a universe, let alone observe any temporal asymmetries.

Note that I don't mean to claim that the perspectival approach requires a *universal* thermodynamic gradient—as noted by Price (2005), the perspectival account can accommodate possibilities like the Gold universe where the thermodynamic gradient goes in one direction for the first half of time and switches direction for the second half. But certainly, it requires at least a locally well-defined gradient because no meaningful perspective could arise out of a universe containing only extremely transient gradients. After all, the very existence of a perspective depends on some kind of ordering requirement because part of what it is to be a "perspective" is to have the ability to perform an ordered sequence of steps of reasoning where the output of one step is used as the input to the next, and therefore the perspectival account requires consistent chaining to hold at least locally in order that this kind of ordered reasoning can be freely performed without giving rise to contradictions. Thus, the consistent-chaining approach I have offered here is an important complement to the Price-Ramsay view: it explicates the underlying physical reasons why a well-defined temporally directed perspective is possible and answers questions left open by Price and Ramsay about why our temporally oriented deliberations slot so neatly into a more widespread causally ordered network of processes.

3.1 Retrocausality

The consistent-chaining requirement would seem to imply that retrocausality is impossible because it prohibits the existence of past-oriented processes that can take as inputs the outputs of our deliberations. However, I have noted that the consistent-chaining approach is closely related to the perspectival account of causation, and Price himself uses the perspectival account as a way of arguing that retrocausality is at least conceptually possible (Price 2005), so it seems that there is some tension here. But in fact, I have only argued that *in all regimes so far probed*, the universe seems to avoid contradictions by limiting the processes available to us in accordance with the consistent-chaining requirement—that is, temporal interdict 1. It is in these regimes that consistent chaining rules out the possibility of using the outputs of our deliberations as inputs to past-directed processes. But this empirical observation does not entail that the consistent-chaining requirement holds in all physical regimes: it remains possible that in some special regimes (e.g., regimes that are difficult to access, so that agents have limited freedom to compose processes defined in those regimes), temporal interdict 2 applies instead. So we could still find that in some special cases, we are able to use the outputs of our deliberations as inputs to past-oriented processes, and a perspectival analysis in the style of Price would presumably yield the conclusion that these sorts of cases do count as genuine retrocausality. Thus, the consistent-chaining approach offers an important caveat to Price's discussion of retrocausality because it explains why we don't see retrocausality in most ordinary regimes, but it does not contradict Price's conclusion that retrocausality is at least conceptually possible within the perspectival account.

Moreover, even if the consistent-chaining requirement were genuinely universal, this would not entirely rule out retrocausality, for it is important to keep in mind that, as detailed by Adlam (2022c), there are two very different notions of retrocausality. Some approaches to retrocausality postulate two distinct directions of dynamical causality that together determine intermediate events by forward and backward evolution, respectively, from separate and independent initial and final states—for example, the forward-evolving state and the backward-evolving state in the two-state vector interpretation (Aharonov and Vaidman 2002). Other approaches postulate an "all-at-once" picture where the laws of nature apply atemporally to the whole of history, such as, for example, in Wharton's all-at-once Lagrangian models (Wharton 2018); 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 forward and backward evolution.[3]

If we had access to dynamical retrocausality, then we would presumably be able to use the outputs of our deliberations as inputs to these past-directed processes, so processes exemplifying dynamical retrocausality would seem to violate the consistent-chaining requirement. Indeed, Maudlin (2002) and Bracken et al. (2021) describe several quantum-mechanical experiments where attempting to give an account in terms of dynamical retrocausality leads to paradoxes of a very similar kind to the one depicted in figure 1, precisely exemplifying the problems that arise in physical theories that do not obey consistent chaining. But as explained by Adlam (2022c), all-at-once retrocausality does not lead to these kinds of paradoxes. And evidently, all-at-once retrocausality is, in any case, better suited for the picture I advocate here, as I 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 that 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 because a symmetric modal relation entails that the future influences the past, and simultaneously, the past

[3] I 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, I will continue to refer to all-at-once approaches as "retrocausal" in order to maintain consistency with the literature.

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

influences the future; that is, 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.

4. Metaphysics

It seems possible that the popularity of the statistical account of the temporal asymmetry of influence may have its origins in the popularity of a specific *metaphysical* stance. That is, it seems to be strongly associated with the metaphysical picture that sees reality as being something like a Humean mosaic—a collection of categorical properties instantiated at points with various kinds of correlations between them (Lewis 1994), for if this is the kind of picture one has in mind, then really there is no such thing as "influence" in any strong sense, and our only option is to reduce the difference between past and future to an epistemic one in terms of the "particular sorts of correlations [that] can obtain between different physical systems at different places and times" (Albert 2014), as the AKL approach aims to do. In this kind of picture, there cannot really be any very deep 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 because the value at the point *p* will typically be correlated both with categorical properties that are in the past relative to *p* and categorical properties that 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). 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 the modal structure to the best systematization of the Humean mosaic (Jaag and Loew 2020; Lewis 1980)—and indeed, as noted earlier, the Humean picture is compatible with our process-based account of the temporal asymmetry of influence because 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 that 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 (Maudlin 2007; Lange 2013; Armstrong 1983, 40), and if one is sympathetic to these objections, one might feel that the Humean approach does not give due weight to the explanatory force of the consistent-chaining requirements. After all, if a Humean accepts the arguments made here, they can still only reach the fairly weak conclusion that because the modal structure is necessarily defined in a noncontradictory way, we must *describe ourselves* as being able to influence the future and not the past—that is,

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

the Humean is compelled to think of consistent chaining as just a 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 the modal structure should be noncontradictory 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 toward a process metaphysics.

Process metaphysics has a history stretching right back to Heraclitus (Kirk 1994), and these ideas have been developed by many philosophers since, including Hegel (1970), Heidegger (1927), and Whitehead (1919, 1920). 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 1970) or the "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; sep-process-philosophy). However, this literal take on dynamical production is somewhat at odds with my 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 of causation. 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 seem 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 because 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, equivalently, produce the input by "evolution" from the output. Thus, there seems little justification for insisting that there is an underlying directedness 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 (Pooley 2013). In particular, many of the "processes" we deal with in operational approaches to quantum mechanics involve sets of laboratories that accept an input at one spacetime point and produce an output at a spacetime-separated point; although these processes are not signaling, they nonetheless 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 that looks like dynamical evolution: examples include the Einstein equations in their usual form, Lagrangian and path integral methods, the canonical quantization of gravity (which famously leads to a "timeless" model with a time-independent Schrödinger equation (Isham 1992)), and a number of interpretations of quantum mechanics that deny that the theory involves anything like a literal process of dynamical evolution (see Adlam

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

2022a,b; Chen and Goldstein 2021). 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, that is, 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 (Berenstain and Ladyman 2012) rather than a traditional process view because the realism is directed toward 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.

5. Noncausal processes

The consistent-chaining requirement and the absence of CTCs, together with the assumption that all processes that are not composed of other processes are causal, comprise a sufficient condition to ensure that the modal structure of the universe does not give rise to contradictions. However, the condition is not a necessary one—ongoing research on process matrix formalism has demonstrated the existence of possible processes that are not causal but that allow local agents to have free will while preventing contradictions, and if such noncausal processes are actually realized in nature, we could in principle have CTCs or violations of consistent chaining while still avoiding contradictions—for example, Baumeler et al. (2019) show that dynamics defined on a closed causal curve can be written in the form of a noncausal process.

However, it is not presently known if these noncausal processes are, in fact, realized in nature. At present, there is no known way of creating noncausal 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, that is, the causal ones? If the explanation for the temporal asymmetry of influence is the fact that the modal structure of our universe only allows processes that 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 noncausal processes, provided that they have process matrices and thus cannot give rise to contradictions.

As a first response, it could, in fact, be the case that nature does 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 (Oreshkov and Giarmatzi 2016; Oreshkov et al. 2012; Araújo et al. 2015). This suggests that rather than looking for a theory of quantum gravity that enforces "no closed loops" at the fundamental level, like causal set theory, we might instead choose to seek a theory that 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 that are not compatible with events having a strict partial order—although we will still not be able to "change" the past because process matrices are defined specifically so as to prevent contradictions from arising.

That said, there have also been suggestions for physical principles that might limit the set of physically possible processes to the causal ones—for example, Araújo et al. (2017) suggest 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 noncausal processes, although not all of them. In addition, it is important to note that the claim that noncausal 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 could 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, process matrix formalism tells us that this is not a possible dynamic 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 noncausal processes should not actually be physically possible. A possible compromise would be to suggest that the regimes in which noncausal processes can be implemented are precisely those regimes where the dynamics are less controllable—but note that I have argued in this article that the principles enforcing the absence of contradictions cannot be merely statistical, and therefore, if noncausal processes are possible in certain regimes, there must be some sufficiently strong and presumably nonstatistical reason why these processes cannot be performed in regimes where the dynamics become more controllable.

Although the consistent-chaining requirement is sufficient in ordinary circumstances to ensure we cannot create contradictory loops like that shown in figure 1, this is no longer the case in spacetimes that 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

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

consistent chaining is obeyed everywhere, and thus we would naturally expect to be able to perform the composition shown in figure 1 if we have access to a CTC.

One possible response to this problem would simply be to rule out CTCs by fiat. There exist approaches to quantum gravity that do exactly that—for example, approaches based on canonical quantization presuppose a globally hyperbolic spacetime, in which there cannot be CTCs (Kuchař 1993; Anderson 2004). 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 that is required to be reflexive, antisymmetric, transitive, and locally finite (Sorkin 2006; Wuthrich and Callender 2017)—this necessarily entails that the spacetime thus produced will not contain any CTCs, and indeed, a "causal set" will clearly look very much like what we have henceforth referred to as a "process graph."

On the other hand, the process framework does allow for noncausal processes that can be implemented on CTCs without giving rise to contradictions. So if we allow the existence of noncausal 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 that correspond to a valid process function or matrix.

6. Conclusion

I have argued that the popular account of the temporal asymmetry of influence as a statistical effect arising out of entropy gradients is not the full story: at the most basic level, the temporal asymmetry of influence follows from the fact that the 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, all processes observed in nature thus far have process matrices or functions that 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.

I have applied these ideas to the asymmetry of causation and argued that not only is causation perspectival in the sense of Price and Ramsay, but in fact, the asymmetry of causation does not have to be grounded in any objective asymmetry, so it does not, strictly speaking, depend on 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 that is primarily responsible for the fact that they find themselves unable to access the past. I have also explored whether violations of consistent chaining might be possible in the form of "noncausal" processes defined in such a way that they cannot give rise to contradictions and concluded that such phenomena might be possible in special regimes, provided that causal processes emerge appropriately in some limit.

References

Adlam, Emily. 2021. "Contextuality, Fine-Tuning and Teleological Explanation." *Foundations of Physics* 51 (6):1–40.

Adlam, Emily. 2022a. "Determinism Beyond Time Evolution." *European Journal for Philosophy of Science* 12 (4):1–36.

Adlam, Emily. 2022b. "Laws of Nature as Constraints." *Foundations of Physics* 52:28.

Adlam, Emily. 2022c. "Two Roads to Retrocausality." *Synthese* 200 (5):422.

Aharonov, Yakir, and Lev Vaidman. 2002. "The Two-State Vector Formalism of Quantum Mechanics." In *Time in Quantum Mechanics*, edited by Gonzalo Muga, Rafael Sala Mayato, and Iñigno Egusquiza, 369–412. Berlin: Springer.

Albert, David Z. 2014. "The Sharpness of the Distinction between the Past and the Future." In *Chance and Temporal Asymmetry*, edited by Alastair Wilson, 159–74. Oxford: Oxford University Press.

Anderson, Edward. 2004. "Geometrodynamics: Spacetime or Space?" arXiv. https://arxiv.org/pdf/gr-qc/0409123.pdf.

Araújo, Mateus, Cyril Branciard, Fabio Costa, Adrien Feix, Christina Giarmatzi, and Caslav Brukner. 2015. "Witnessing Causal Nonseparability." *New Journal of Physics* 17:102001.

Araújo, Mateus, Adrien Feix, Miguel Navascués, and Časlav Brukner. 2017. "A Purification Postulate for Quantum Mechanics with Indefinite Causal Order." *Quantum* 1:10.

Armstrong, David Malet. 1983. *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Barnum, Howard, Jonathan Barrett, Lisa Orloff Clark, Matthew Leifer, Robert Spekkens, Nicholas Stepanik, Alex Wilce, and Robin Wilke. 2010. "Entropy and Information Causality in General Probabilistic Theories." *New Journal of Physics* 12 (3):033024.

Barrett, Jonathan. 2005. "Information Processing in Generalized Probabilistic Theories." arXiv. https://arxiv.org/abs/quant-ph/0508211.

Baumeler, Ämin, Fabio Costa, Timothy Ralph, Stefan Wolf, and Magdalena Zych. 2019. "Reversible Time Travel with Freedom of Choice." *Classical and Quantum Gravity* 36 (22):224002.

Bell, John. 1990. "Against 'Measurement'." *Physics World* 3 (8):33.

Bell, John. n.d. "Time and Causation in Gödel's Universe." https://publish.uwo.ca/jbell/Time.pdf.

Berenstain, Nora, and James Ladyman. 2012. "Ontic Structural Realism and Modality." In *Structural Realism: Structure, Object, and Causality*, edited by Elaine Landry and Dean Rickles, 149–68. Dordrecht, Netherlands: Springer.

Bracken, Colm, Jonte R. Hance, and Sabine Hossenfelder. 2021. "The Quantum Eraser Paradox." arXiv. https://arxiv.org/abs/2111.09347.

Branciard, Cyril, Mateus Araújo, Adrien Feix, Fabio Costa, and Časlav Brukner. 2015. "The Simplest Causal Inequalities and Their Violation." *New Journal of Physics* 18 (1):013008.

Chen, Eddy Keming, and Sheldon Goldstein. 2021. "Governing without a Fundamental Direction of Time: Minimal Primitivism about Laws of Nature." arXiv. https://arxiv.org/abs/2109.09226.

Chiribella, Giulio, Giacomo Mauro D'Ariano, and Paolo Perinotti. 2010. "Probabilistic Theories with Purification." *Physical Review A* 81 (6):062348.

Chiribella, Giulio, Giacomo Mauro D'Ariano, and Paolo Perinotti. 2011. "Informational Derivation of Quantum Theory." *Physical Review A* 84 (1):012311.

Davies, E. Brian, and John T. Lewis. 1970. "An Operational Approach to Quantum Probability." *Communications in Mathematical Physics* 17 (3):239–60.

Feix, Adrien, Mateus Araújo, and Caslav Brukner. 2016. "Causally Nonseparable Processes Admitting a Causal Model." *New Journal of Physics* 18 (8):083040.

Frisch, Mathias. 2010. "Causal Models and the Asymmetry of State Preparation." In *EPSA Philosophical Issues in the Sciences: Launch of the European Philosophy of Science Association*, edited by Mauricio Suárez, Mauro Dorato, and Miklós Rédei, 75–85. Dordrecht, Netherlands: Springer.

Hegel, Georg Wilhelm Friedrich. 1970. *Hegel's Philosophy of Nature*. Oxford, UK: Clarendon Press.

Heidegger, Martin. 1927. *Sein Und Zeit*. Halle, Germany: M. Niemeyer.

Horwich, Paul. 1990. "Asymmetries in Time." *Noûs* 24 (5):804–6.

Isham, Chris. 1992. "Canonical Quantum Gravity and the Problem of Time." In *Integrable Systems, Quantum Groups, and Quantum Field Theories*, edited by L. A. Ibort and M. A. Rodrìguez, 157–288. Dordrecht, Netherlands: Springer.

https://doi.org/10.1017/psa.2023.55 Published online by Cambridge University Press

Ismael, Jenann. 2016. *How Physics Makes Us Free*. Oxford: Oxford University Press.

Jaag, Siegfried, and Christian Loew. 2020. "Making Best Systems Best for Us." *Synthese* 197:2525–50.

Jia, Ding, and Nitica Sakharwade. 2018. "Tensor Products of Process Matrices with Indefinite Causal Structure." *Physical Review A* 97 (3):032110.

Kirk, Geoffrey Stephen. 1994. "Natural Change in Heraclitus." In *The Pre-Socratics: A Collection of Critical Essays*, edited by Alexander P. D. Mourelatos, 189–96. Princeton, NJ: Princeton University Press.

Kuchař, Karel V. 1993. "Canonical Quantum Gravity." arXiv. https://arxiv.org/abs/gr-qc/9304012.

Ladyman, James, and Katie Robertson. 2014. "Going Round in Circles: Landauer vs. Norton on the Thermodynamics of Computation." *Entropy* 16 (4):2278–90.

Lange, Marc. 2013. "Grounding, Scientific Explanation, and Humean Laws." *Philosophical Studies* 164 (1):255–61.

Lewis, David. 1980. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, edited by Richard C. Jeffrey, 83–132. Oakland, CA: University of California Press.

Lewis, David. 1994. "Humean Supervenience Debugged." *Mind* 103 (412):473–90.

Masanes, Lluís, and Markus P. Müller. 2011. "A Derivation of Quantum Theory from Physical Requirements." *New Journal of Physics* 13 (6):063001.

Maudlin, Tim. 2002. *Quantum Non-Locality and Relativity: Metaphysical Intimations of Modern Physics*. Hoboken, NJ: Wiley-Blackwell.

Maudlin, Tim. 2007. *The Metaphysics within Physics*. Oxford: Oxford University Press.

Novikov, Igor Dmitriyevich. 1992. "Time Machine and Self-Consistent Evolution in Problems with Self-Interaction." *Physical Review D* 45 (6):1989–94.

Oreshkov, Ognyan, Fabio Costa, and Časlav Brukner. 2012. "Quantum Correlations with No Causal Order." *Nature Communications* 3:1092.

Oreshkov, Ognyan, and Christina Giarmatzi. 2016. "Causal and Causally Separable Processes." *New Journal of Physics* 18 (9):093020.

Pooley, Oliver. 2013. "Relativity, the Open Future, and the Passage of Time." *Proceedings of the Aristotelian Society* 113:321–63.

Price, Huw. 2005. "Causal Perspectivalism." In *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, edited by Huw Price and Richard Corry, 250–92. Oxford: Oxford University Press.

Price, Huw, and Brad Weslake. 2008. "The Time-Asymmetry of Causation." In *The Oxford Handbook of Causation*, edited by Helen Beebee, Peter Menzies, and Christopher Hitchcock, 414–44. Oxford: Oxford University Press.

Sorkin, Rafael. 2006. "Geometry from Order: Causal Sets." *Einstein Online* 2 (2006):02–1007.

Wharton, Ken. 2018. "A New Class of Retrocausal Models." *Entropy* 20 (6):410.

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

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

Wood, Christopher J., and Robert W. Spekkens. 2015. "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.

Wuthrich, Christian, and Craig Callender. 2017. "What Becomes of a Causal Set?" *British Journal for the Philosophy of Science* 68 (3):907–25.

Cite this article: Adlam, Emily. 2023. "The Temporal Asymmetry of Influence Is Not Statistical." *Philosophy of Science* 90 (4):855–879. https://doi.org/10.1017/psa.2023.55