The Oxford Handbook of Philosophy of Time

The Asymmetry of Influence

Douglas Kutach

CLARENDON PRESS • OXFORD
2010

1

THE ASYMMETRY OF INFLUENCE

1.1 Introduction

Attempts to meddle with the past are futile. While this nugget of folk wisdom serves as a respectable guide to action, its utility is standardly conceived as arising from the general inability of anything to influence the past. This explanation, though, oversimplifies the complex architecture of fact and fiction responsible for the reasonableness of not trying to affect the past. It makes little genuine progress in understanding the asymmetry of influence because it fails to distinguish two significantly different kinds of explanation. The first kind appeals to the fixity of the past as a strict, fundamental, metaphysical or scientific fact that guarantees the inefficacy of all attempts at past-directed influence. The second kind of explanation says that in virtue of the reasonableness of the folk wisdom, we interpret 'influence' in a way that ignores those senses in which the past can be influenced so that "nothing can affect the past" is rendered true largely by definition.

A fruitful analogy is the concept of solidity. One might say the reason it is useful to treat a boulder as if it were solid is because the boulder really is solid. But this platitude does not distinguish the false explanation that boulders are solid through and through from the true explanation that, although the boulder mostly consists of empty space between its atoms, its chemical bonds give rise to a cluster of complex macro-properties like rigidity and impenetrability that make it effectively solid for most ordinary practical purposes. Taking 'solid' to denote this imprecise cluster of imprecise properties makes it literally true that the boulder is solid by discounting respects in which the boulder is not solid.

One strategy for explaining the folk wisdom that affecting the past is futile follows the first kind of explanation by arguing that it is in the nature of time or somehow built into the universe's fundamental structure that the past is fixed or that any lawful dependence of the past on the present is not genuine influence. The strategy advocated in this chapter elaborates on one version of the second kind of explanation. Rather than the fixity of the past being a foundational metaphysical fact, there are numerous respectable senses in which the past can be influenced. Yet, because of a complex assortment of conceptual and physical relations, past-directed influence, also known as 'backward' influence, turns out to be critically different from future-directed influence in its practical impact. Our ordinary conceptualization of reality incorporates this pragmatic asymmetry by counting only the future-directed influence as genuine influence. Using the second kind of explanation is arguably preferable because it posits less metaphysical baggage and—more important—clarifies intricate connections between influence, causation, time, and chance, better than in accounts of the first kind.

The resulting account contributes to our understanding of time by providing a crucial component of a static theory of time: an explanation of why it is natural to think of the past and future as essentially different even though they are metaphysically on par. Dynamic theories of time virtually always enforce a fundamental difference between past and future, either by postulating a primitive ontological difference or by positing a metaphysically robust passage of time which in turn somehow vindicates the belief that the future is fundamentally different from the past. The growing block model illustrates both theses by holding that physical reality at any moment consists of a spacetime block of present and past facts but no future facts and that the block grows. The account presented here provides an alternative explanation for our stubborn conviction that the past and future are essentially different. The hard part of the explanation involves demonstrating how known physical structures make it reasonable for us to believe in the folk wisdom even if we routinely influence the past. The easy part then attributes our intuition that the past is fixed to an understandable misinterpretation. People know of the occasional practicality of trying to affect future events and the universal impracticality of trying to affect the past, but they mistake the asymmetry of *useful* influence for an asymmetry of influence simpliciter.

If successful, this account would explain the influence asymmetry using relatively innocuous resources: the fundamental laws and facts about the world's fundamental material layout. Because it dedicates no structures to the implementation of temporal passage, the only sense in which time passes is that there is something about the combination of laws and material facts that makes it reasonable to think of time as passing. This still permits the possibility that passage is fundamental. If the laws turn out to have the right properties, then the aspect of reality most responsible for the utility of conceiving time as passing might turn out to be a feature of the laws alone and not rely on features of material facts. If so, time will pass fundamentally not because there is some flow of time in the sense of dynamic passage or some fundamental arrow specifically dedicated to playing the passage role, but merely because out of all the available structures (resulting from the laws plus facts), what ends up best vindicating talk of passage is an asymmetry in the fundamental laws.

If there is nothing solely in the fundamental laws that plays the passage role, time will either not pass at all or will pass non-fundamentally, i.e. derivatively. Passing not at all means that nothing plays near enough the role of passage, and passing derivatively means there is something in the content of the laws and historical layout of facts that plays the role of passage near enough, but like temperature or solidity, is non-fundamental. Choosing between these two options depends largely on how one wants to construe passage. Much could be said about what in general constitutes bona fide temporal passage, but once one departs from theories with fundamental passage, the elements of the constitutive role of passage become rather unclear. There do seem to be some platitudes that tether the meaning of passage to other concepts: The unidirectionality of passage is tied to the idea that a number of important psychological attitudes are appropriately temporally asymmetric, for example, that the fear of death is more reasonable than the fear of birth, and that feelings of anticipation are appropriate to future but not past events. Passage is also arguably that which people perceive or seem to

cognitively latch onto when they feel (or report that they feel) that time is passing. And, passage is linked to the thesis that the present is somehow special. The influence asymmetry too is closely aligned with the direction of passage in the sense that passage is often conceived as a conversion of the not-yet-settled possibilities of the future into settled facts of the past. None of these conceptual links are intended as necessary conditions for genuine passage, and any of them may turn out in the end to express misguided commitments. They are mentioned only because the goal of accounting for a non-fundamental asymmetry is to provide some story about why it is reasonable for people to believe in some package of beliefs that roughly resemble platitudes like these. Once one abandons the picture of passage as a metaphysically preferred presentness sweeping across the whole of history, the best account of passage will likely involve significant theoretical refinement of the gist of the passage-related platitudes. If the conceptual revision is too extensive, the question of whether the completed account possesses genuine temporal passage may well degenerate into a definitional quibble where there is no substantive dispute about whether time really passes or whether we just tend to think of time that way. The rest of the discussion will focus on the influence asymmetry and disregard other aspects of passage, but the account is definitely intended to be friendly to the conclusion that passage itself, such as it exists, is little more than the asymmetry of useful influence.

There is a tradition of attacking robust temporal passage by demonstrating the coherence and even physical possibility of backward influence, often as it appears in coherent time travel stories. Some models of general relativity with closed time-like curves (CTCs) bolster the relevance of such time travel scenarios by treating as physically possible, physical processes that locally instantiate the normal future-directed influence but which reach back around the CTC into the global past. While I believe such arguments help to undermine many dynamic (or 'moving now') conceptions of temporal passage, they fail to confront static theories of fundamental temporal passage, (e.g. Maudlin 2007), which account for passage without requiring more ontological resources than exist in standard spacetime structure plus a metaphysically privileged temporal orientation. Furthermore, arguments relying on processes or topological structures that have not been empirically confirmed might turn out to be based on models that are, on deeper investigation, not physically possible. My explanation of the influence asymmetry differs from this tradition by claiming that backward influence may well obtain amid ordinary physical processes of the kind that pervade our local environment. Regardless of whether time travel into the past successfully undermines appeal to metaphysically robust temporal passage, backward influence can exist without such time travel. In order to set aside the complications introduced by time travel, it will be assumed hereafter that the laws and contingent facts do not permit it. For instance, the topology of spacetime will be assumed to be free of CTCs, and any laws of nature governing the world's development will not allow physical states to have causal impact that skips across time. That is, the state c_1 can determine (or fix a probability for) a state c_2 only if every state temporally between them is determined by (or has some probability fixed for it by) c_1 and determines (or fixes some probability for) c_2.

The central contention of this chapter—that what we think of as the influence asymmetry is really an asymmetry of useful influence—is suggested by uncontroversial observations about ordinary future-directed influence. Our reason for believing in influence at all can be taken to result from our awareness of events at different times where the existence or character of one event, the effect, correlates with the character or existence of the other event, the cause. Some kinds of forward influence, though, are undetectable by observation of the relevant events, for example, the weak gravitational influence a person has on the motion of Jupiter or the butterfly's chaotic influence on the next decade's weather. A good reason for accepting them as cases of influence comes from believing that laws of physics are at least partially responsible for the observable influence and that these laws remain operative when the physical connection between the events becomes too weak or too chaotic for us to detect. In brief, we naturally (and correctly, I believe) accept that the fundamental laws dictate how far our paradigmatic influence extends to other kinds of influence.

Once one allows that influence should be extended to cases where the physical laws establish dependencies between events at different times, it is not outrageous to ask whether influence could be extended further to other cases of nomological dependence. Since paradigm fundamental physical theories virtually all postulate at least some kind of nomological dependence of past states on future states,[1] one should consider whether such backward nomic dependence should also count as influence. There would be no problem in doing so, I think, if it were not for our deep commitment to the folk wisdom, which is seemingly at odds with backward influence. An obvious strategy for resolving the conflict is to clarify how, despite the presence of both forward and backward influence, the folk wisdom continues to serve as a reliable practical guide.

The hypothesis under consideration is that we influence the future more or less as pre-theoretical intuition dictates but that we also influence the past in a number of different senses. Yet, for each way of affecting the past, some or other excuse always crops up for why it is unexploitable in practice.

- Some instances of past-directed influence are too weak.
- Some are too chaotic.
- Some fall under no regularities that are epistemologically accessible.
- Some are re-interpretable as ordinary, unobjectionable future-directed influence.
- Some count as influence due to ad hoc parameter settings or other ways of precisifying the influence notion that beg the question about the direction of influence.
- Some are systematically unusable due to some physical contingency that constrains events in a suitably large region to systematically fail to display the expected consequences of influence.

[1] Even theories with fundamental chanciness, like GRW-style collapse interpretations of quantum mechanics (Ghirardi, Rimini, and Weber, 1986), fix *some* probabilities for past states. There is a determinate probability at t that no collapse will have occurred at $t - \epsilon$ and that is equivalent to the probability that the state at $t - \epsilon$ implied by the deterministic backward evolution of the wave function at t will have occurred.

While these excuses will be clarified in due course, the intended conclusion is that all past-directed influence falls into at least one of the above categories and is therefore generally useless for advancing any goals one might have. Unfortunately, the argument is complicated by the fact that influence generally and past-directed influence in particular is a vague notion admitting diverse precisifications. For the argument to be successful, it must ensure that for any halfway reasonable notion of influence, some or other excuse is applicable.

Let 'advancement asymmetry' label the claim that there exists at least one reasonable notion of influence such that future-directed influence is sometimes useful for advancing goals and that there is no reasonable notion of influence such that past-directed influence is useful for advancing goals. The purpose of this paper is to give an exhaustive account of the advancement asymmetry that (1) is connected with fundamental physics, influence, causation, counterfactual dependence and related notions in palatable ways, (2) derives the asymmetry wholly from scientifically acceptable sources, and (3) does not assume the asymmetry by fiat whether by positing an asymmetry of robust temporal passage, or by smuggling in a temporal asymmetry through the use of an asymmetric notion of chance, or by deriving the influence asymmetry from a primitive causal relation that exists in addition to the complete history of fundamental properties, relations, and laws. It is possible to have a fundamental arrow of time within a fully scientific and static account of the influence asymmetry, but I will set aside even this possibility because the asymmetry is better explained without it.

1.2 Influence and Counterfactual Dependence

The guiding idea behind one family of influence-concepts is that of some kind of counterfactual dependence. For one event C to influence an event E is for E to be different from what it would have been if C had not happened. Regimenting this idea to make it both scientifically legitimate and useful for understanding causation requires care because the use of counterfactual conditionals raises worries of arbitrariness and circularity. If we evaluate counterfactuals in some way that introduces facts that go beyond the physical condition of the world, we may beg the question about influence asymmetry or just be talking nonsense. It is therefore important that any arbitrary aspects of how influence-counterfactuals are evaluated are not misrepresented as facts about the actual world and do not rely on counterfacts—facts about what would have happened that go beyond facts about what actually happens.

Three clarifications are needed before detailing how to evaluate influence. First, one must quarantine the intuitions about influence that are at the heart of traditional struggles over free will. Incompatibilists hold that under determinism, there is no genuine free will. This might be extended into a claim that under determinism nothing really influences anything else—free or not—and that at best one has the illusion of influence. Because the kind of influence elaborated below does not take sides on the free will debate, the incompatibilist might complain that unless the laws are of a special form, one gets only pseudo-influence out of the account. Even if that is correct, the ability of the account to achieve its goal is unaffected because it can still explain why it is reasonable for us to think of ourselves as influencing the future and not the past. One need only add

that the reasonability in part arises because insofar as the utility of the folk wisdom is concerned, it is harmless to mistake pseudo-influence for the genuine article.

Second, it is important to set aside commonly accepted counterfactual logics like the kind developed by Stalnaker (1968) and Lewis (1973). Such logical systems are designed to regiment inferences and truths of natural language. To apply such systems to a scientific investigation risks contaminating one's theory of influence with linguistic quirks and artificially limiting the range of possible theories. If by some bizarre coincidence the logic of natural language conditionals shares important similarities with the physical structures that best explain influence, no harm is done, but to presume from the outset that nature must obey the rules of natural language is to restrict oneself by a fantastically implausible constraint.

Third, there is a natural temptation to balk at equating influence with counterfactual dependence because while some kinds of counterfactual dependence are causal in character, others are apparently evidential. For example, the details of a photographic portrait counterfactually vary in accord with the photographed subject. If the woman's photograph had pictured her wearing a hat, then very probably the woman would have been wearing the hat during the photo shoot. One naturally interprets this counterfactual dependence as a consequence of the ordinary future-directed influence of hats on photographs. Yet, the theory being developed here equates influence with some kinds of counterfactual dependence and is on some precisifications of 'influence' committed to the counterintuitive claim that the photograph backwardly influenced the woman's appearance. This seeming problem is partly terminological. I could align the terminology to more closely match standard usage by claiming that counterfactual dependence goes both forward and backward in time and by using the label 'influence' to mark the dependence associated with the useful-for-advancement direction. Then, the woman would counterfactually depend on the later photograph, but the photograph would not literally influence her. The more conventional usage, though, would not alter the substantive content of the theory: that some forms of backward counterfactual dependence and forward counterfactual dependence express the same kind of physical link. They both express lawful relationships between states at different times. One need not assume one direction as having a metaphysically privileged status associated with causation. I deliberately choose the unintuitive usage where 'influence' is equivalent to (my preferred versions of) counterfactual dependence, regardless of temporal direction, in order to continually reinforce the controversial character of my conclusion that the seeming temporal directionality of influence is pragmatic and not necessarily fundamental.

The model I propose for evaluating influence-counterfactuals goes as follows. The counterfactual "If C were true, E would be true," is abbreviated $C \boxdot\!\!\Rightarrow E$ and has as its "semantic" value the objective probability of E across the relevant possible worlds where C obtains. What counts for relevance is determined by developing a theory whose goal is to extract the right amount of information from the laws of nature, the actual circumstances, and the antecedent itself without being artificially restrictive or begging the question about the direction of influence. The approach I follow is to require enough structure so that the laws of physics entail determinate probabilities for E, but allow other precisifications of the relevant worlds to be left as free parameters. The value of

the counterfactual will then exist only relative to some specific setting of the parameters.

To fix a determinate probability for E, the relevant worlds must obey all fundamental laws of physics. The most important kind of laws for evaluating influence are dynamical laws, rules whereby a precise physical state at one time entails an objective probability for physical states at other times. (States are events that instantiate "what happens at a single time.") It is certainly conceivable that the universe has no dynamical laws, but I will assume their existence on the hope that the fundamental laws turn out to be nice enough, like the paradigm fundamental theories of classical mechanics, relativity, and quantum mechanics. It is also possible that the dynamics permits sources of indeterminism that are unconstrained by probabilistic rules, and again I will just assume such sources either do not exist or are rare enough to be ignorable, at least for ordinary causation in mundane environments. Henceforth, only deterministic dynamics and chancy sources of indeterminism will be countenanced. If empirical evidence reveals these assumptions to be ill-founded, this account of the asymmetry will presumably fall as well.

An important quantity for evaluating influence is the counterfactual probability of the effect given the non-occurrence of the cause. To interpret this counterfactual correctly, one must distinguish between fine-grained events and coarse-grained events. A fine-grained event is by definition a perfectly precise microphysical condition inhabiting some specific spacetime region. A coarse-grained event is defined to be a (usually continuous) set of fine-grained events each of which inhabit the same spacetime region. (Some remarks on notation: Capital letters refer to coarse-grained events and lower case letters to fine-grained events. When the same letter appears in capital and lower case form, the designated fine-grained event is one element of the designated coarse-grained event. Strictly speaking, a coarse-grained event can be trivially coarse-grained, meaning that it contains a single fine-grained event as its only element. Finally, a coarse-grained event need not fall under a relatively simple natural language description, though we are typically concerned with such events.

In judging influence, the effect must not be understood in a fine-grained way because that would make its counterfactual probability zero under almost any realistic physical laws and would render it a useless quantity. By understanding the effect in a coarse-grained way, insignificantly small alterations to the physics instantiating the effect do not necessarily generate a numerically distinct event.

The characterization of the non-existence of the cause is also a coarse-grained event because when one considers what would have happened had C not occurred, one does not necessarily intend a microscopically specific instantiation of C. There exists considerable latitude in what counts as an acceptable way for the cause not to occur. If the stick had not been released, was it still being held? Broken? Burned? Because there is no systematic answer as to which states constitute the cause not happening, the non-occurrence of the cause must be left as another free parameter.

In order to get any probabilities for E, one must consider states expansive enough for the dynamical laws to be informative. This requires embedding the non-release of the stick in a large enough physical state for the laws to entail a probability for E. How one should accomplish this embedding in full generality is not perfectly systematic, but

in practice one can take an actual microstate big enough to count as "a state of the world at the time of c's occurrence" and modify it to instantiate the non-releasing of the stick. Because there are many precise ways to instantiate $\neg C$ in the larger state, and many of them fix different probabilities for E, it is useful to postulate a probability measure over all the states so that together they determine a unique counterfactual probability, the probability of E weighted over all these states. The contextualized event, $\overline{\neg C}$, is a probability distribution over a set of states each realizing $\neg C$ and some specification of the microscopic background environment. The contextualized event can treat the background in a fine-grained way by having all its elements agree on the microscopic details outside the $\neg C$ region but can also coarse-grain over the background environment. The objective probability of E, given these contextualized events is labeled $p_{\overline{\neg C}}(E)$. This quantity serves as a kind of semantic value for the counterfactual conditional $\neg C \mathrel{\Box\!\!\Rightarrow} E$ in the sense of being an objective quantity towards which belief in the counterfactual should be directed. Ordinary beliefs are directed towards truth, in that an ideally knowledgeable person will have degree of belief one in what is true and zero in what is false. Beliefs about influence counterfactuals are directed towards this probabilistic value in that an ideally knowledgeable person—one who knows what the laws entail—will have degree of belief $p_{\overline{\neg C}}(E)$ in $\neg C \mathrel{\Box\!\!\Rightarrow} E$, taking into account how $\overline{\neg C}$ fleshes out $\neg C$ in more detail. The probability is not intended to be part of a larger logical structure that one could rightly call a semantics because among other things it does not make sense for propositions generally and is not compositional.

[Three sections are omitted from this excerpt from *The Oxford Handbook of Philosophy of Time.*]

1.3 The Advancement Asymmetry

To summarize, we now have three models that have not decisively been convicted of inadmissibly presuming an asymmetry of influence.

1. A stochastic dynamics with virtually no backward transition probabilities.
2. A symmetrically deterministic dynamics.
3. A dynamics that is either stochastic or deterministic and takes the relevant counterfactual worlds to be restricted to worlds that start out in a state α that is exactly (or near enough) the universe's actual (or surrogate) initial state.

These models, it will be shown, all have the feature that advancing one's goals is asymmetric. The first case is easy to handle. Because there are virtually no probabilities for events in the past, there are virtually no ways to make past events more likely than they would have been if the present were different. Thus, there are no promoters of past events and consequently no backward advancement of goals. End of story. The other two models give probabilities that ground backward influence$_c$ and demand extensive discussion.

For comparative purposes, it is helpful to consider how to model furthering a goal towards the future, spelled out with an example. Instantiated in some agent's brain right

now is both a desire for E_f, the existence of a loaf of bread one hour from now (towards the future), and the volition to satisfy that desire. The oven contains a pan of suitable dough, and the agent's volition combined with knowledge of how ovens work leads him to turn on the oven at 220 degrees centigrade. The decision to turn on the oven, C, is in the circumstances a promoter of the oven being nearly 220 degrees not long after. Likewise, C promotes E_f because the objective probability of the bread being created 40 minutes later is very high and the probability of E_f is very low given the implicit contrast event $\neg C$, the decision to leave the oven dial in the off position.

Now consider how the agent could act to further the goal of having a fresh loaf of bread an hour ago. Somehow instantiated in the brain is a desire for E_p, the past existence of a loaf of bread in the oven, and the volition to satisfy that desire. We can consider numerous ways in which C promotes E_p. The first two are far and away the most important and hold for any notion of past-directed influence, including those that employ α-chances.

1.3.1 *Backward promotion by way of the future*

First consider attempts to promote E_p by way of creating some intermediary event in the future, I_f. This is called backtracking influence because it first goes one direction in time and then backtracks by going in the other temporal direction. The strategy is to set up a situation where the volition, C, promotes I_f, which in turn promotes E_p. The problem with any such attempt is that the dynamical laws of our world apparently operate in a way where states causally contribute to effects by way of continuous transitions that screen off previous states. All paradigm theories of fundamental physics, at least, exhibit this behavior except for the CTC's allowed in general relativity. In Newtonian gravitation, for example, an arbitrary initial state S_0 defined by the relative positions and velocities of all the masses determines both the state S_1 one second later and the state S_2 two seconds later. But S_1 determines S_2 by itself as well; it encodes all the information in S_0 needed to get S_2. If the fundamental laws include stochasticity, the condition still applies. Any probability that S_1 fixes for S_2 is invariant under conditionalization on all states previous to S_1.

In a relativistically local theory, illustrated in figure 1.1, the contextualized event $\overline{C}$ that instantiates C is just barely big enough to fix a probability for I_f. Likewise, the contextualized event $\overline{I_f}$ is just barely big enough to fix a probability for the earlier E_p. But any contribution I_f makes to E_p goes through the surface where $\overline{C}$ is instantiated. Thus, any aspect of C that promotes E_p through I_f already exists in the direct connection between $\overline{C}$ and E_p. One might think that there is still some room for I_f to impact E_p by somehow affecting the details of $\overline{C}$ beyond the occurrence of C. But any such possibility would be question begging because in order to infer anything about the probability of I_f using the fundamental laws one must first have a fully detailed, complete $\overline{C}$. The fundamental laws do not ground inferences from C itself or an incompletely described $\overline{C}$ to I_f. They require fully filled out states at least as large as $\overline{C}$. Thus, any backtracking inferences in filling out how C is translated into a contextualized event does not occur in virtue of any laws and thus whatever backward influence exists by way of future events is rendered superfluous by influence that goes directly back in time to E_p.

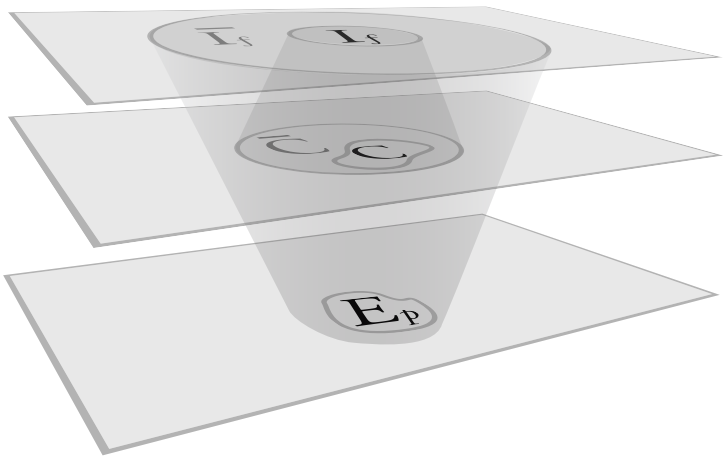

FIG. 1.1. I_f only contributes to E_p by way of $\overline{C}$.

Given this redundancy of all backtracking influence, actions to promote a past goal are useless, at least as actions are normally understood. Actions aimed at some goal can be conceived as processes roughly like the following: An action starts with one's attention being focused on some situation, evolves through a deliberative process incorporating desires and beliefs, climaxes in some decision being made, and eventuates in some (typically external) physical effects. When promoting future events, the physical development subsuming the action proceeds through various stages, each of which promotes later stages and finally external future events. But for an action to proceed toward the future through all its stages and then to go back into the past to promote past events constitutes an instance of backtracking promotion, which is guaranteed to be superfluous. Thus, the only non-trivial way an action can promote past events is to directly promote them by way of the reverse sequence of stages, that is, starting with a decision having been made, which promotes (towards the past) the agent's comparing various options, which then promotes the agent's having no idea of what to do, and finally a stage where the agent is not attending to the situation. For such actions, we have a built-in excuse for their lack of utility. Paradigmatic actions and those same actions reinterpreted as evolving towards the past are importantly different in the functional interaction of desires, beliefs, and volition. We naturally reject such reverse-actions as genuine actions directed towards a goal. Hence, we naturally reject the possibility that they could stand as a counterexample to the folk wisdom.

It is possible, given all that has been concluded so far, that an ordinary action to promote a past goal does promote it through this highly unconventional reverse-action process. The rest of this section aims to close loopholes, so that even if one counts such reverse-actions as legitimate, it will still turn out that they cannot regularly be used to promote the goals in the past. But even if such arguments fail, the master argument in this section suffices to vindicate the folk wisdom insofar as its content is restricted to normal actions.

1.3.2 *The epistemic inaccessibility of backward promotion*

Even if there were some reverse-action influence on the past, we will be unable to become aware of it in the ordinary way we learn of future-directed promotion, that is, by observation, trials, etc. The extent to which our evidence of correlation between C's and E_p's can be properly attributed to genuine causal promotion is the extent to which the probabilistic connection exists after we have controlled for all the evidence for E_p existing in C's larger environment, $\overline{C}$. By definition, C promotes E_p to the extent that the instantiation of C in the volitional component of the agent's brain raises the probability of E_p beyond what probability is fixed by $\overline{\neg C}$. In order to gather the kind of correlational evidence relevant to establishing promotion between C and E_p, we need epistemic access to C, which we can easily get just by asking the agent to tell us what she was trying to do, and access to E_p that does not come from via C's larger environment. But, our only source of information about E_p that is independent of C itself comes by way of this larger environment, unless the laws of nature permit epistemic access to past facts that skips over intermediate physical states. Presumably, our epistemic access to E_p comes only by way of how reality evolves, which apparently does not involve any state skipping. Thus, any attempt to collect statistics about what C promotes by comparing what happens after $\overline{C}$ occurs and what happens after $\overline{\neg C}$ will leave us no access (independent of C) to previous events, and so we cannot learn by observation or trial the extent to which attempts at past-directed promotion are genuinely successful.

1.3.3 *Backward promotion that is interpreted away*

One kind of past-directed promotion sure to exist is the connection between the act of will and its immediate precursor, the volitional part of one's brain state microseconds before the act of will. Though technically the act of will influences$_c$ this precursor, such influence is useless for achieving interesting goals, and what impact it has can be reinterpreted as part of the overall volitional process.

A similar kind of past-directed promotion exists, under determinism, in one's ability to promote the state of the world ten minutes ago to be just the right kind of state to lead deterministically towards bread being baked one hour later. This is just an ordinary decision to bake the bread reinterpreted as involving past events. Such promotion is in a sense exploitable, but only to promote the existence of bread in the future. The previous argument (in §1.3.1) concluded that a present state cannot influence the past by way of the future more than what one gets by influencing the past directly. Because no asymmetry is implicated in the argument, the same holds in the other temporal direction. One cannot influence the future by way of the past more than what one gets by influencing the future directly. Any influence$_c$ over the past that is defined trivially in terms of what future goals are satisfied does not count as fulfilling goals that are about the past in any interesting sense.

1.3.4 *Backward promotion by cheating*

One's choice of contextualized events can almost always be rigged to create a contrast that generates spurious influence$_c$ or non-influence$_c$. For comparison, consider cases of spurious future-directed non-influence$_c$. One could countenance a set of C-states where

one turns on the oven but the energy supply is disconnected, or where a nemesis is lurking nearby with the aim of secretly shutting off the oven. One could also countenance a set of $\neg C$-states where one does not intentionally turn on the oven, but where a defect in the dial is set to malfunction by igniting the oven anyway. Any of these would make $C \boxdot\!\!\Rightarrow E_f$ and $\neg C \boxdot\!\!\Rightarrow E_f$ roughly equal, entailing that C does not help achieve E_f. Made precise in those ways, C does not promote E_f, but such cases are clear examples of contriving just the right kind of contrast to ensure that one gets non-influence$_c$. Thus, they do not ground substantive claims of non-influence$_c$.

The same reasoning applies to some purported cases of past-directed influence$_c$. This is clearest under determinism. Let C be the trivial coarse-graining of the precise instantiation of the volition c, and let $\overline{C}$ be its extension as far out spatially as is needed. To construct the counterfactual contrast, take the actual state when some chosen past event E_p occurs, replace it with some specific $\neg E_p$ and let that evolve toward the present, labeling that present state $\overline{\neg C}$. By construction, $C \boxdot\!\!\Rightarrow E_p$ is equal to one and $\neg C \boxdot\!\!\Rightarrow E_p$ is zero, so that C has influenced$_c$ E_p, no matter what that event is. The reason to rule out such influence$_c$ as usable for advancement is that $\overline{\neg C}$ is not plausibly an event representing a natural contrast class. It was selected only because it has coded into it all the microphysics needed to entail $\neg E_p$. Whenever the coarse-graining or contrast is selected in a biased way, that itself counts as a legitimate reason for dismissing the alleged attempt at advancing some goal for the past.

[One section is omitted from this excerpt from *The Oxford Handbook of Philosophy of Time.*]

1.4 Conclusion

The core phenomena behind the asymmetry of advancement is this:

- There exist uncontroversial, non-ad hoc ways of interpreting the counterfactual dependence implicit in influence such that there exist lots of exploitable regularities of causal promotion going from the present towards the future.
- Though some ways of understanding counterfactual dependence permit causal promotion going from the present towards the past, there are nevertheless no regularities governing such promotion that rightly count as exploitable for achieving aims.

Unless a viable candidate has been overlooked, there is no reasonable notion of influence such that one can advance one's goals in the past. This means it is pointless to try to influence the past and that therefore in practice it is mostly harmless to think of the past as effectively fixed and independent of one's current action. The asymmetry of influence, charitably interpreted, exists because it is nothing other than the asymmetry of advancement. Alternatively, if the influence asymmetry is construed to require that events predominantly fail to influence the past, it is a bogus but understandable simplification.

The asymmetry of advancement might strike people as too remote from the objective structure of nature to count as the correct explanation of the influence asymmetry. An explanation in terms of how agents are able to promote their goals seems implausible given that processes like stars burning and cliffs eroding possess an apparent causal or influence asymmetry that has nothing to do with goals or agency. Furthermore, some of these temporal asymmetries—for example, the radiation asymmetry and T-violations in weak nuclear decay processes—may turn out to be essential elements of fundamental causation-like relations. If so, a direction of causation would exist at the fundamental level, not—as the analogy to solidity would have it—an asymmetry that arises only macroscopically.

This charge can be rebutted on several points. First, nothing forbids a fundamental asymmetry from grounding the advancement asymmetry. If the laws of nature entail no probabilities for past events and α-chances are rejected for any reason, then there is a fundamental asymmetry that directly explains why we cannot make desired past effects more likely. Yet, however comforting such an explanation is to people who want to retain our naive conception of the influence asymmetry, there are two features of this explanation that are somewhat conceptually revisionary. For one thing, such a result would not substantiate the idea that the past is fixed because past facts would not hold regardless of counterfactual alterations to the present. Rather, the counterfactuals would just be silent about the past, which in a sense makes the past even more open than it would be with determinism or past-directed chanciness, for there would not be any laws constraining what would have happened in the past. For another thing, even if there is a fundamental explanation of the influence asymmetry, it remains true that if the laws were to have bidirectional chanciness or determinism, we would still have an advancement asymmetry. Thus, my arguments would provide a redundancy in the explanation of why it is useful to believe in asymmetric influence. Because advancement is asymmetric regardless of whether the laws themselves are asymmetric, it is not obvious that that fundamental asymmetry alone explains why we believe in the asymmetry of influence. It explains why influence is asymmetric, but if we believe in the asymmetry of influence because of the asymmetry of advancement, both the fundamental and non-fundamental explanations play a role in vindicating our beliefs. The situation is reminiscent of John Locke's (1685, II, 21, §10) character who wakes up inside a room and volunteers to stay because a friend is there, but who could not leave anyway because they are locked inside the room. We can explain why the man does not leave either by citing his volition or the locked door, but either one is redundant given the other. Similarly, we can explain why advancement is asymmetric, and thus why we think of influence the way we do, either by pointing directly to the fundamental asymmetry in the laws or to the more general fact that advancement is asymmetric regardless of any fundamental asymmetry in the laws.

Regarding the seeming oddity of accounting for the advancement asymmetry specifically in terms of achieving goals, it is worth noting that the explanation can be subsumed within a broader account of more general physical processes. The physical basis underlying the advancement asymmetry is the fact that our local environment is composed of spacetime patches with future-typical and past-atypical dynamical develop-

ment. This typicality asymmetry establishes a local condition where virtually any counterfactual difference leads dynamically to future states falling under familiar rules of thumb about how things evolve. For example, changes to the position of water molecules in the ocean probably makes little difference in at least the near term future, a change to the position of balls in the lottery bin leads to a big difference in who wins (but not to the coarse-grained probability of who wins), changes to the position of a light switch leads to a significant difference in the amount of light in the room, etc. These rules of thumbs exist because future-typicality ensures the standard connection between the chances and the distribution of chance outcomes. By contrast, past-atypicality ensures that the otherwise-to-be-expected rules of thumb do not apply. This typicality asymmetry holds whether there are agents or not and subsumes the asymmetry of advancement within an explanation of the broader physical phenomena. The explanation, briefly stated, is that a few general features of the fundamental laws together with some fact about (what we call) the early universe makes the typicality asymmetry fantastically likely throughout the known universe. The typicality asymmetry in turn explains the asymmetry of thermodynamic processes and why many macroscopic processes are asymmetric in character and are temporally aligned with each other, e.g., why the life cycle of ferns is asymmetric with all ferns sprouting, growing, reproducing and dying in the same temporal direction. The process of agency is just one of many such processes and its asymmetry is explained in the very same way.

The reason agency arises as a crucial feature of interest is due to its being central to what we naively think of as causation. If we had assembled a bunch of facts about macro-asymmetries—the temporal asymmetries of ferns or volcanos or toasters—then we could point to typicality (and features of fundamental physics) to explain the asymmetry of their functionality. That would be one kind of explanation, but there would still be a further question about what makes those macro-asymmetries *causal*. Our naive notion of cause incorporates a directionality, crudely expressed as the fact that the cause makes the effect happen and not the other way around. Merely having asymmetries in ferns and such does not vindicate this conception of causation. One could think of the 20cm high fern making itself shorter towards what we call the past, and one could think of a charred log as making a fire occur previously. Nothing disastrous would come of such talk, even if for some predictive and explanatory purposes it is inferior. What vindicates talk of 'making happen' is that among the many various macro-asymmetries, there is one called 'agency,' whose existence makes useful a concept that is useful not only for predicting on the basis of the macro-asymmetries or explaining on the basis of the macro-asymmetries, but for manipulation. My observations here closely track what Michael Dummett (1964) observed: an intelligent plant would have reason for an asymmetric notion of cause, but our concept of causation is important to us because of our agency. The intelligent plant would be able to predict and explain macro-asymmetries using the resources of the typicality asymmetry and perhaps other physical asymmetries, but it would have no need for its causal notion to incorporate the idea that present states do not bring about the past and do not influence the past. The implications of the typicality asymmetry for agency is what vindicates our conceiving of causation as essentially involving the inability of the effect to bring about the cause.

The role agency plays in the asymmetry of causation can be illustrated by considering what it implies about the direction of causation in regions where the typicality asymmetry is reversed. Suppose some patch of the universe, *R*, is such that its material content is macroscopically temporally reversed from us and contains sequences that (in reverse) instantiate regularities like those on Earth. (For simplicity, set aside the kind of reasoning that motivated α-chances.) Is the arrow of causation reversed in *R*? The answer is ambiguous. If an Earthly agent were hypothetically transported to *R* in a way that preserved its temporal orientation, then actions that agent took would promote events towards what we call the future. So, the direction of causation would be the same. If we instead consider alterations to the physics in the heads of inhabitants of *R*, those changes will be of the kind that promote *their* goals in the temporal direction we call the past. So, the direction of causation in *R* would be in that sense reversed. The direction of causation is thus perspectival in the same way as the difference between up and down. On one way of thinking about it, everyone agrees that down is towards the center of the planet. On another way of thinking about it, people on the opposite side of the planet disagree about what direction counts as down.

To summarize the objectivity of the asymmetry of advancement, it is useful to return to the metaphor that initiated this chapter, the solidity of ordinary objects. There is certainly a sense in which a granite block is objectively solid; no fancy jiggering of one's subjective attitudes will negate the damage it inflicts when dropped on a stray finger. Yet, there is also a sense in which objectively the block is mostly empty space; solidity is a conceptual simplification grouping together a cluster of properties for pragmatic purposes. If we were radically different creatures—proton-sized or galaxy-sized or composed of neutrinos—we would not find the distinction between the solidity of the cinder block and the liquidity of oil to be of much significance. In this way, the distinction between solid and liquid is carved out in a way that corresponds to the interests of creatures like us. Nevertheless, one must depart rather significantly from the kinds of creatures we are to find beings for whom solidity makes no practical difference. The insensitivity to particular character of humans makes the solid/liquid distinction to that extent objective. Analogously, no adjustment of one's opinions about the direction of influence will negate the fact that people who try to bake a pie for yesterday's dessert are objectively wasting their effort. Yet the asymmetry, as it has been scrutinized and elaborated here, essentially involves relativizing facts about influence to choices about how to conceive of action, and these choices are subjective in the sense that nothing in logic or in the natural world forces one to conceive of actions, for example, as future-directed. This kind of subjectivity is innocuous because the structures of the natural world are such that the coarse-grainings and contrasts that permit useful influence over the past are ones that intuitively count as ad hoc or question-begging. Relativize to uncontrived choices and the laws ensure the unexploitability of past-directed influence. One could imagine hypothetical creatures for whom the contextualized events that ground past-directed influence count as natural ways of coarse-graining, but such creatures would arguably be so bizarre as to be nearly unrecognizable as creatures.

REFERENCES

Albert, D. *Time and Chance*. Cambridge: Harvard University Press, 2000.

Dummett, M. "Bringing About the Past," *The Philosophical Review*, **73** (3), 338–59, 1964.

Ghirardi, G. C., Rimini, A., and Weber, T. "Unified Dynamics for Microscopic and Macroscopic Systems.," *Phys. Rev. D* 34, 470–491, 1986. (doi:10.1103/ PhysRevD.34.470)

Kutach, D. "The Entropy Theory of Counterfactuals," *Philosophy of Science* **69** (1), 82–104, 2002.

Kutach, D. "The Physical Foundations of Causation," in Price, H. and Corry, R. (eds.), *Causation, Physics, and the Constitution of Reality*. Oxford: Oxford University Press, 2007.

Lewis, D. *Counterfactuals*. Oxford: Blackwell, 1973.

Locke, J. *An Essay Concerning Human Understanding*, 1685.

Loewer, B. "Counterfactuals and the Second Law," in Price, H. and Corry, R. (eds.), *Causation, Physics, and the Constitution of Reality*. Oxford: Oxford University Press, 2007.

Maudlin, T. "On the Passing of Time," in *The Metaphysics Within Physics*. Oxford: Oxford University Press, 2007.

Price, H. and Corry, R. eds. *Causation, Physics, and the Constitution of Reality*. Oxford: Oxford University Press, 2007.

Stalnaker, R. "A Theory of Conditionals," in N. Rescher (ed), *Studies in Logical Theory, American Philosophical Quarterly Monograph Series, No. 2*, Oxford: Basil Blackwell, 98–112, 1968, reprinted in E. Sosa (ed) *Causation and Conditionals*, Oxford: Oxford University Press, 165–79, and in W. L. Harper, R. Stalnaker, and G. Pearce (eds). *Ifs*, 41-55, and in Jackson, Conditionals, 28–45.