ON RETROCAUSALITY: LOOKING FORWARD TO LOOKING BACK

Yelena Guryanova

https://doi.org/10.25364/97839033742875

1 INTRO

The allure of time-travel is overwhelming. From taking back one's words to writing a better version of this manuscript, turning back the clock proffers a simple solution to all our regrets. But why stop at righting wrongs? Trips in the TARDIS, temporal anomalies, magical machines and paradox parties have captivated the imagination of everyone from sci-fi writers to scientists. And not without reason, for in its solutions, the theory of general relativity allows for closed time-like curves in the form of wormholes, and therefore for the prospect of going backwards in time.

Our everyday experience, however, is something to the contrary; time is not reversible, which is at odds with modern physics, where no formal derivation or calculation precludes this possibility. The 'fact' that we can only send signals to the future, and not the past, is not really a fact so much as an ubiquitous observation, put in by hand by physicists in the construction of theories, which they tacitly deem 'reasonable'. Unlike Newton's laws of motion, which describe the behaviour and forces of a moving body, or Coulomb's law, which describes how electric charges repel, there is no 'law' concerning backwards-in-time signalling. Indeed, while the existing laws describe observations, the latter would be a statement that something is altogether never encountered.

Even though general relativity allows for closed time-like curves, there is currently no evidence of their existence whatsoever, and despite the appeal of time-travel, the consequences of people, objects, even information being able to propagate to the past has enormous repercussions. Travelling backwards in time to kill your own Grandfather (the so-called 'Grandfather Paradox') creates a contradiction in which you are simultaneously born, but cannot be born. This logical inconsistency can be remedied,[1] by positing that you are actually, born with probability ½. Thus, if you are born, then you travel back in time to kill your own Grandfather and therefore you are born with probability ½. No contradiction. Computer scientists were able to develop this probabilistic, logically consistent framework further and investigate how far the idea of time travel could really go. By formulating these notions in the language of programming, i.e., by reasoning in terms of computers making calculations, The Grandfather Paradox can be rephrased as, "sending the answer back to the time before the computer even started the calculation".

Even in the paradigm where logical inconsistencies are forbidden, closed time-like curves still wreak all kinds of havoc. One implication would be that quantum and classical computers would be equally and extremely powerful,[2] allowing one to efficiently find solutions to incredibly complicated problems, such as larger and larger Sudokus or piecing together enormous picture-less jigsaw puzzles. Big problems typically need big computers, but here, the presence of closed time-like curves would have the effect of rendering space and time equivalent, allowing one to 'recycle' time, instead of adding more memory space to the machine in order to find a solution.

1. David Deutsch. *Quantum mechanics near closed time-like lines.* Phys. Rev. D, 44:3197–3217, (1991).

2. Scott Aaronson and John Watrous, *Closed timelike curves make quantum and classical computing equivalent,* Proc. R. Soc. A.465631–647 (2009).

Although these arguments from computer science may pose a convincing barrier to timetravel, the barrier is not quite high enough to prohibit a little playfulness. Astonishingly, one can still construct worlds that allow the future to affect the past, nevertheless without signalling backwards-in-time. If you like, a kind of backwards-in-time influence without backwards-in-time signalling. To understand these possibilities and what it means to influence something without signalling, one must first learn to understand the world in terms of boxes.

An interesting feature of any theory is the *correlations* that it produces. Correlations are the patterns and relationships that one can analyse between at least two parties: these could be two systems, two parts of a whole, or, quite simply, anywhere where one can cut a meaningful divide between two sides. What do we mean by 'theory' and how is this connected to the correlations produced therein? In this text we will take the definition of a *physical* theory to be a mathematical formalism that: (i) provides a model/ description of a system; (ii) predicts the behaviour of system. The models and predictions in a physical theory give rise to *correlations*, i.e. the patterns and relationships *between systems*.

2 BOX-WORLD(S)

In a particular setting, classical and quantum theory differ in their correlations – an assertion which has been confirmed in numerous experiments. In this setting, two experimenters put into separate laboratories and given a task in which they were not allowed to communicate; how often they succeed at this task is related to their experimental outcomes, and thus the patterns and correlations between them – these will be different depending on whether the experimenters have access to classical or quantum systems. Here, by 'classical' system we mean all the objects that classical theory can describe – numbers, pens, paper, computers. On the other hand, 'quantum' systems, such as entangled quantum particles, are objects described in quantum but not classical theory (in some sense, quantum theory was developed because the descriptive power of classical theory was not enough to express the state and behaviour of some systems, for which an extended mathematical framework was required). Classical and quantum are both examples of physical theories because there are models for the systems and predictions for how they behave; since the models and predictions are different, it would seem reasonable that the theories give rise to different correlations.

We now draw a distinction between a physical theory and a *box-world*. A box-world is a mathematical characterisation of the correlations between (at least two) parties. It is not a physical theory because it does not satisfy requirements (i) or (ii); instead, a box-world skips over these and only describes the relationship *between* hypothetical systems, i.e., the relationship between 'boxes' in some 'world'. Box-world is therefore a universal, theory-independent, way of talking about correlations.

One can also invert the logic and ask if there is a physical theory compatible with every boxworld? The short answer is 'no'. In some cases one is able to establish that some correlations belong to a theory, but in most cases, the problem is very hard: we do not know how to construct a theory, or we cannot because it may be impossible altogether. Despite this, box-worlds are very powerful: any statements which one can make in box-world will hold in *all* worlds and theories, regardless of whether they are physical or not.

The following thought experiment will contain actors, rooms and procedures and scripts in order to *generate* correlations. In truth, these are only there to assist in the narrative, since in box-world we do not require an explanation for how the correlations are generated. For the purposes of illustration, we shall make use of imagery and analogy proceeding with pictures in mind, until it's time to obliterate them to abstraction.

To set the scene, imagine an experimenter, Alice, who has a sealed laboratory. We define an experimental round in this laboratory to consist of two steps (see Fig. 1). In the first of these, at some point, a system enters her lab; she performs a measurement on it and obtains an outcome a, which is a label that is either 0 or 1. The framework for this discourse depends on integer (whole number) outcomes, in part due to the fact that it makes it mathematically easier to compare between correlation theories. Thus, one should not think of Alice's measurement, as a measurement of height, or weight, which are more or less continuous quantities, but as a measurement of something discrete. For example, if Alice were measuring the position of a particle, she may ask: "is the particle on the left or the right?" For which the outcomes would be 0 or 1, indicating one of two delimited regions. It is important to note that 0 and 1 are arbitrary labels for anything that can be binarised (categorised into two distinguishable sets), such as on/off, left/right, and well as more imaginative pairs.

What system does Alice receive? What does it look like and what is it made of? In truth, we don't care, and nor do we necessarily know. One can imagine spheres, objects, physical items with shapes and colours, but the system need not have any of these properties – the 'particle' alluded to was a metaphor. The only demand we make, is that whatever Alice has, has a property that can take one of two values. 'System', then, is a misleading word for an abstract way of describing 'something', which we define rather minimally.

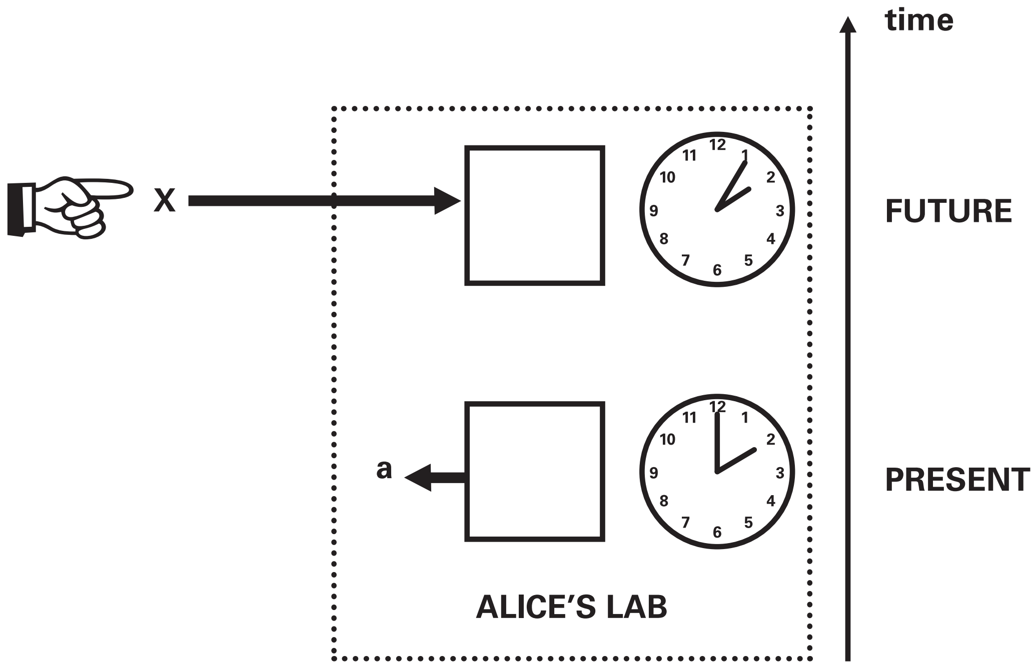

Figure 1. Schematic of Alice's lab, to be read from bottom to top; the system (box) is the same one, at two different times. At 2 PM Alice makes a measurement on her system and receives the outcome a. Five minutes later, an external referee passes her an instruction x about which transformation to perform on the system.

In the second step of the experiment, Alice receives an instruction x, from a referee outside the lab, labelled also by 0 or 1; to continue the allegory, one can imagine that the referee passes a post-it note with either digit written on it. Alice uses this instruction to chose between two different transformations to apply to the something-system, i.e., the *box* in box-world. This transformation could for instance be "if 0, rotate the system" or "if 1, flip the system upside down". What actually happens is again disguised in metaphor, since we do not know if the 'system' is something that can be flipped at all. The real difference between the measurement in step 1 and the transformation in step 2, is that the latter is something that results in information gain, whereas the former is something in which no information is obtained. In the first step, Alice learns whether the system belongs to the category labelled by '0' or the category labelled by '1'; in the second step she learns nothing, but may change the state of the system deterministically.

Classical theory provides us with a concrete example of the latter. Image that Alice rolls a die, and outputs a to be 0 if the number is even, or 1 if the number is odd. Then, if she receives the instruction 0 from the referee Alice changes the face of the die by rotating 180° in any direction. If she receives 1 she does nothing. Thus, if she had rolled a 4, her measurement outcome would be 0. Then, if the referee instructed 0, she would rotate the die to the number 3, or, if the instruction were 1 the same face would remain. In the first step, she learns whether the die landed on even or odd, and in the second step she learns nothing. She knows what a die looks like, and that the face with '3' is opposite the face '4' regardless of which direction one chooses to rotate by 180°.

After both steps are complete Alice sends the system out of her laboratory and the procedure is over. In every experimental round she obtains two numbers, a and x, which she duly notes down. Her logbook begins to fill up with measurement results and transformations, which she writes using the notation $a|x$:

$$0|0,\ 1|0,\ 1|1,\ 0|1,\ 0|1,\ 1|1,\ 0|1,\ 1|0,\ 1|0,\ 0|0,\ \cdots \quad (1)$$

The list is extremely long, as a rule of thumb, more than 10,000 results; Alice can look through the list and estimate the frequency of each of the four pairs. How often did, for example, $0|0$ occur? Like this, she can estimate the probability that she obtained the outcome a given that she later received the instruction x via the conditional probability distribution

$$p_A(a|x)\ , \quad (2)$$

to be read, "the probability of outcome a, given the future instruction x". The expression in (2) is notation for a list of four numbers, $p_A(0|0)$, $p_A(1|0)$, $p_A(0|1)$, $p_A(1|1)$, all between 0 and 1, with relations among them. Now, if Alice were receiving signals from the future, i.e., if the referee's instruction x at 2:05 PM were signalling to her past-self at 2 PM it would show up in her probability distribution $p_A(a|x)$. In other words, a simple mathematical test on the list in (1) would reveal if the future were signalling to the past. Specifically, if $p_A(a|x)$ depended on x, this would indicate the presence of these signals: the probability of a depends on x — the probability at present depends on the future.

From now proceed with utmost caution: we do not believe that signalling backwards-in-time is possible, or indeed 'reasonable' and choose to explicitly forbid such a phenomenon mathematically. We define the no-backwards-in-time-signalling (NBTS) principle to be:

$$p(a|x)\ =p(a)\ . \quad (3)$$

This equality states that the probability of obtaining the outcome a given the future transformation x is the *same* as the probability of simply obtaining the outcome a, i.e., that a in the present does not depend on x in the future.

Which world are we in and which theory does this correspond to? The answer is any world, in which one can embed this two-step framework (measurement then transformation). In fact, it is enough to receive a list like the one in (1), and the promise that the numbers therein were generated in two steps, with no further explanation of what went on. Perhaps a lab somewhere in London was measuring quantum spins, or perhaps somewhere in a world we've never encountered a strange printer spat out a ream of 0's and 1's on a long receipt – the models for 'what happened' are irrelevant.

Now that we have enforced, *by hand*, the NBTS principle that *a* does not depend on x (present does not depend on future), it is now very tempting to think, "Very well! Since the present cannot depend on the future, then the future should depend on the present and thus x can depend on a". This is, of course, in general completely true, but to be able to make meaningful statements about the NBTS scenarios later, we must assume that x is free and independent of a. To see why, one must consider the role of the external referee.

Let us entertain the possibility that Alice decides to completely ignore the referee and that she chooses the transformation labelled by x as desired. In particular, she could chose x to depend on a, by choosing it to be exactly the same $x = a$. Beware, that the latter does not mean that the measurement and the transformation are somehow physically identical; they are quite separate from one another. This only means that the *label* for the transformation is equivalent to the *label* for the measurement.

The choice to ignore the referee, is manifestly not a case of backwards-in-time signalling since Alice's actions are perfectly time ordered. By choosing x (future) based on a (present), Alice is signalling forwards-in-time, but since the labels are identical, then one can make symmetric and equivalent statements: a depends on x and x depends on a (present depends on future) and (future depends on present). The former is problematic and demonstrates why it is vital to assume that x is generated externally and independently of a and why Alice must follow the referee's instructions explicitly. In order to be able to draw consistent conclusions and enforce the principle properly, x must be free of a.

But how free can x really be? What if the referee were cheating by eavesdropping or looking at Alice's outcome and matching the instruction to it? Perhaps a subtler form of dependence is being embezzled, that is harder to detect? In general, these concerns cannot be ruled out. Here, however, it is not our objective to discuss how to guarantee the absence of such conspiracy theories, rather, in the rest of the text we will assume that x is truly independent and discuss the consequences for the setup in such a case.

3 THE TIME BETWEEN

One may wonder what interesting statements there are to make at all in such a single experimenter scenario: we have said that x is independent of a (by assumption) and also that a is independent of x (by force). Indeed, in order to see or say anything remotely interesting, one must introduce another experimenter, whom we shall refer to as Bob. Alice and Bob are now both subject to the rules laid out above. We endow each of them with a sealed laboratory and apply the same framework: a system enters Alice's laboratory and she performs a measurement on it, obtaining an outcome a. She then receives an instruction x from an external referee outside of the laboratory, which she uses to select and perform a transformation on her system; she sends the system out of her laboratory. Analogously, at some point, a system enters Bob's laboratory. He also performs a measurement on it and obtains the outcome b; a different external and independent referee passes him an instruction y; he uses y to select and perform a transformation on his system. He sends the system out of the lab.

We must now confess to another lie – the idea of the lab. The laboratory (imagined perhaps with walls and doors) is also something of a fiction, established and sealed only in order to create well-defined time-zones for Alice and Bob. In this way, one need not assume that there is any notion of time between them: in her lab, Alice has her own time, in Bob's lab he has his; outside of their respective environments anything goes. Maybe they are connected in the usual way in time just as two labs in Vienna and Paris might be, or, alternatively there may not be any time that connects them whatsoever. The 'lab' is a way of neatly making an assumption about each experimenter's local time.

In their respective 'labs' Alice and Bob perform their procedures, obtaining outcomes and performing transformations, all labelled by 0's or 1's. Fig. 2 shows the results noted down by Alice and Bob in each experimental round and the last column shows a combined way of writing the individual records. By looking down this column, one can also estimate the relative frequencies of the different combinations $a\ b\,|x\ y$ coming out of the experiments, for instance, how often did the combination $01|01$ occur? These are captured in the conditional bipartite probability distribution:

$$p(a, b|x, y), \tag{4}$$

Experiment number	Alice a | x	Bob b | y	Combined notation a b | x y
1	0 | 1	1 | 1	0 1 | 1 1
2	1 | 0	1 | 1	1 1 | 0 1
3	1 | 0	1 | 0	1 1 | 0 0
4	0 | 1	0 | 0	0 1 | 0 0
5	1 | 1	0 | 1	1 0 | 1 1
6	0 | 0	1 | 0	0 1 | 0 0
...	...	...	...

Figure 2. A record of the experimental results obtained by Alice and Bob. In many experimental rounds Alice notes down her measurement outcome and instruction $a|x$ and Bob notes down his measurement outcome and instruction $b|y$. The last column shows a combined notation for the global view of the two experiments, which can be obtained only after Alice and Bob come together.

"the probability of outcomes labelled by a and b, given that the instruction x was received in the future of a and the instruction y was received in the future of b". This joint description presents a 'bird's eye view' of both labs, and contains within it *more* information that if one were to look at Alice and Bob separately. Indeed, the probabilities that Alice and Bob see locally is not the 'global view', but rather, the probabilities they get by estimating the frequencies in their individual columns, just as Alice was doing to get $p_A(a|x)$ in the single-experimenter case. The global description $p(a, b,|x, y)$ contains all the *correlations* between Alice and Bob, but it is only possible to have this description after both experimenters have completed their procedures, met, and collated their local results into the last column of Fig. 2.

From the global description, one can access more structure than just by looking at the individual observations of Alice and Bob. A mathematically consistent method to get from the combined column on the right in Fig. 2 and access this extra information, is known as *marginalisation*. This step is not physics, but rather mathematics, and comes from probability theory: marginalisation is true for any global view probabilities that have the form in (4). This move will reveal all the dependencies that Alice's outcome a may have, and is achieved by summing (marginalising) over all the outcomes b of Bob in the global distribution. Regardless of what Bob's measurement result is, Alice's measurement result a still has the following dependencies:

$$\sum_b p(a, b|x, y) =: p_A(a|x, y) \; , \tag{5}$$

where the colon indicates that these two sides are always true by definition. The right hand side of the equality tells us that irrespective of Bob's b, Alice's a can still depend on x and y. Likewise for Bob: all the dependencies on his measurement outcome b are learned by summing over all of Alice's outcomes a from the global probabilities:

$$\sum_a p(a, b|x, y) =: p_B(b|x, y) \; . \tag{6}$$

From here we proceed to enforce the no backwards-in-time signalling principle, just as we did in the single-experimenter scenario. In Eq. (3) we imposed that the probability of Alice's present outcome a was independent of her future transformation x. Here, we apply the same logic and enforce this rule on the marginalised probabilities. Eqs. (5) and (6) become

$$\sum_b p(a, b|x, y) =: p_A(a|x,\ y) = p_A(a|y)\ , \tag{7}$$

$$\sum_a p(a, b|x, y) =: p_B(b|x,\ y) = p_B(b|x)\ , \tag{8}$$

in each line, the first equality is by mathematical logic and the second we enforce due to our patent disbelief in time travel: these two equations enforce the NBTS principle for Alice and Bob in their respective labs.

What does it mean to mathematically enforce a principle and what does this have to do with any reality that one can relate to? Eqs. (7) and (8) exist in print but do they exist in practise?

The equations above have characterised Alice and Bob's correlations: we are in many boxworlds, in the space of all worlds which satisfy and are consistent with the no-backwards-in-time signalling principle.

4 THE TIME BETWEEN IS RELATIVE

History has shown on many occasions that whenever there is prohibition there is always dissent, and the case for box-world is no different. Even though we have precluded signals from travelling from the future, somewhat surprisingly this doesn't rule out the possibility of influences propagating in the 'forbidden direction'. These influences are not guaranteed – given an experimental list such as the one in Fig. 2, they may or may not be present, but fortunately, it is possible to systematically characterise them, in order to know where to look. The characterisation depends on the *relative* timing of Alice's and Bob's experiments, a conversation we've been careful to avoid until now, having only said that «at some point» a system enters their labs. We now pin this 'point' down and split our study into four cases – either the relative timing between Alice and Bob is known: Alice before Bob; Bob before Alice; the experiments are parallel, or, the relative timing is unknown.

The case in which the timing is known and Alice's actions are completely before Bobs we denote $A \rightarrow B$ (Fig. 3(a)). Here, the system which entered Bob's lab, was actually the same one that left Alice's, i.e., Alice sent her system to Bob. Then, the system she sent to Bob could have carried encoded information about her measurement result and transformation; thus b may depend on everything in its past, i.e., Alice's a and x. After Bob outputs b, he then receives the instruction y and performs a transformation and the experiment is over. One can check that irrespective of Alice's outcome a, i.e., on marginalising in order to access the extra information

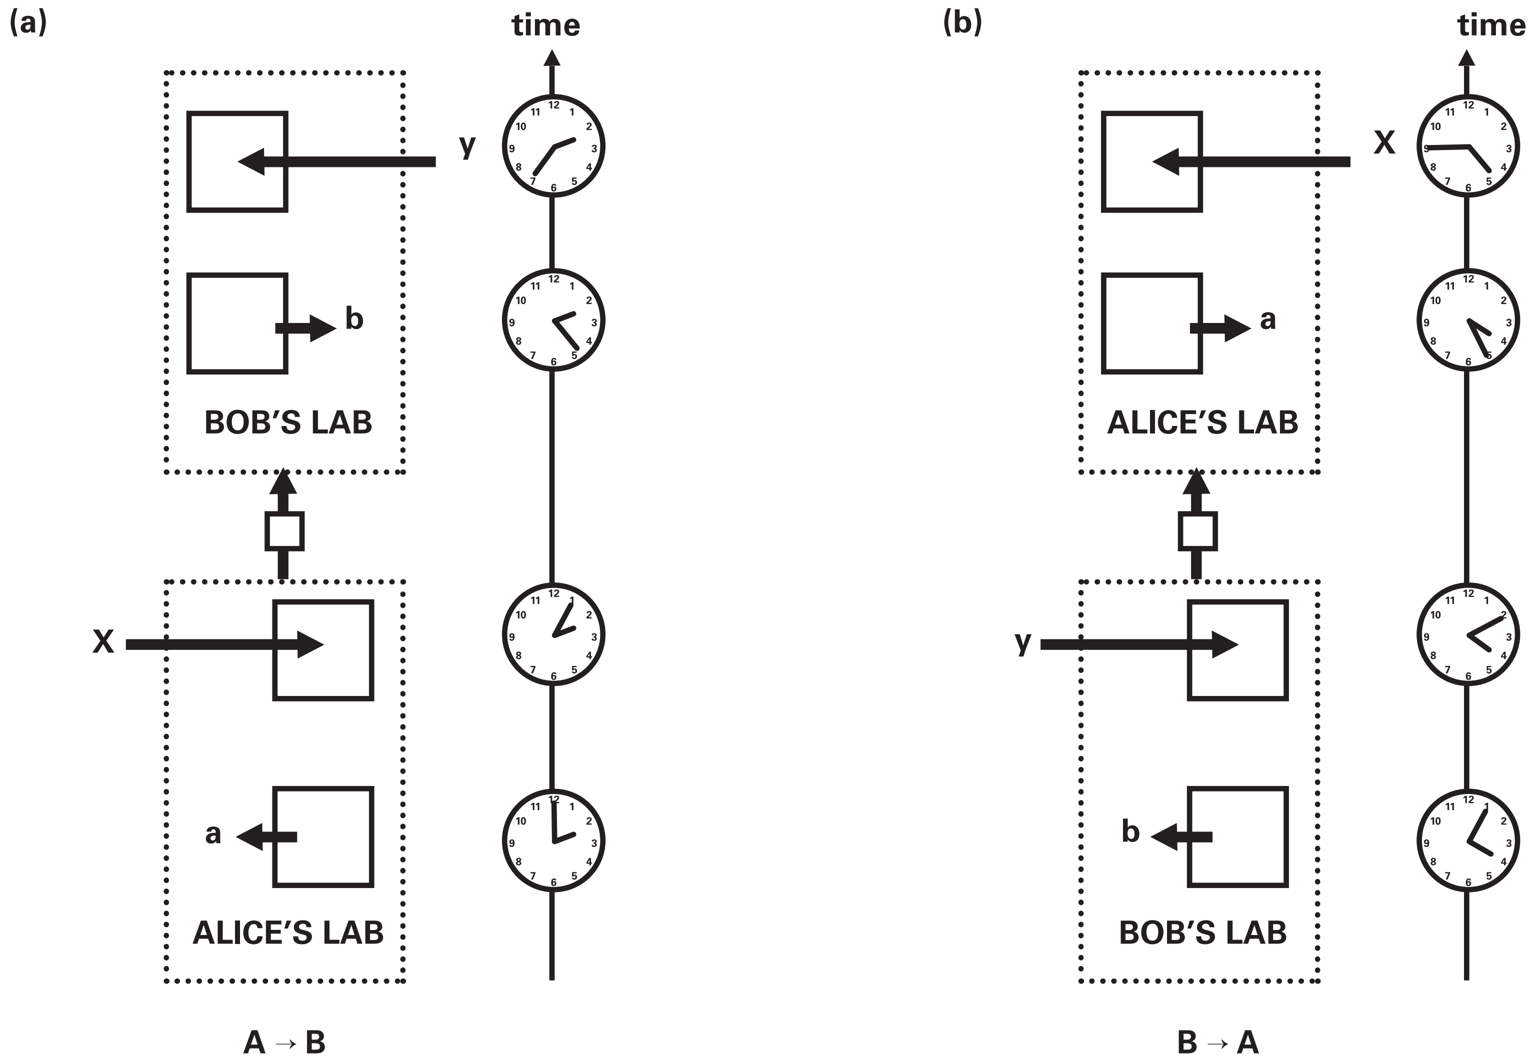

Figure 3. Known timing: the sequential scenarios. (a) $A \to B$ and (b) $B \to A$. In (a) a system enters Alice's lab at 2 PM, an external referee passes her instruction labelled x. She then send the system out of her lab to Bob, who upon receiving it performs a measurement and obtains outcome b. Later, another external referee passes him an instruction labelled y, which he uses to select and perform a transformation. In (b) the same procedure occurs, but with Bob at the start of the chain and Alice at the end.

Bob's outcome b only depends on x and not on the future y, just as in Eq. (8). Mathematically everything is consistent and all the actions are ordered in a chain. The analogous statements are true if the order were the other way around i.e., if Bob sent his system to Alice, $B \to A$ (Fig. 3(b)).

The sequential settings admit the 'intuitive' description of experimental results that follow a well-defined temporal chain, but, surprisingly, even in these ordered cases, influences may occur that allow the future to affect the past. If one considers Fig. 3(a), one can isolate instances where the following is possible: individually, Alice's outcome a at 2 PM is not affected by her transformation x at 2:05 PM. Likewise Bob's outcome b at 2:22 PM is independent of the transformation y at 2:35 PM; thus both Alice and Bob obey the NBTS constraints. However when considered *together*, it *is* possible for the outcomes a and b to depend on the future transformations x and y, i.e. a at 2 PM and b at 2:22 PM depend on the transformations x at 2:05 PM and y at 2:35 PM. This kind of strange dependence is not assured in all box-worlds, one has to check to see if it may be there.

This surprising feature is also present and even better exemplified in the other known-timing scenario: that of parallel actions. Here, Alice and Bob's experiments are completely synchronised and there are two separate systems which enter their labs, precluding the possibility of any information passing between them (Fig. 4(a)). Even in this seemingly innocent case, where nothing is exchanged or communicated, influences may still be able to escape and the following scenario is not ruled out. The results of Alice and Bob's measurements, which occur on Monday, actually depend on the external instructions x and y that they receive on Tuesday. This strange dependence can be observed on Wednesday, when Alice and Bob emerge from their labs and compare their results by creating the 'global view' column of the experimental results.

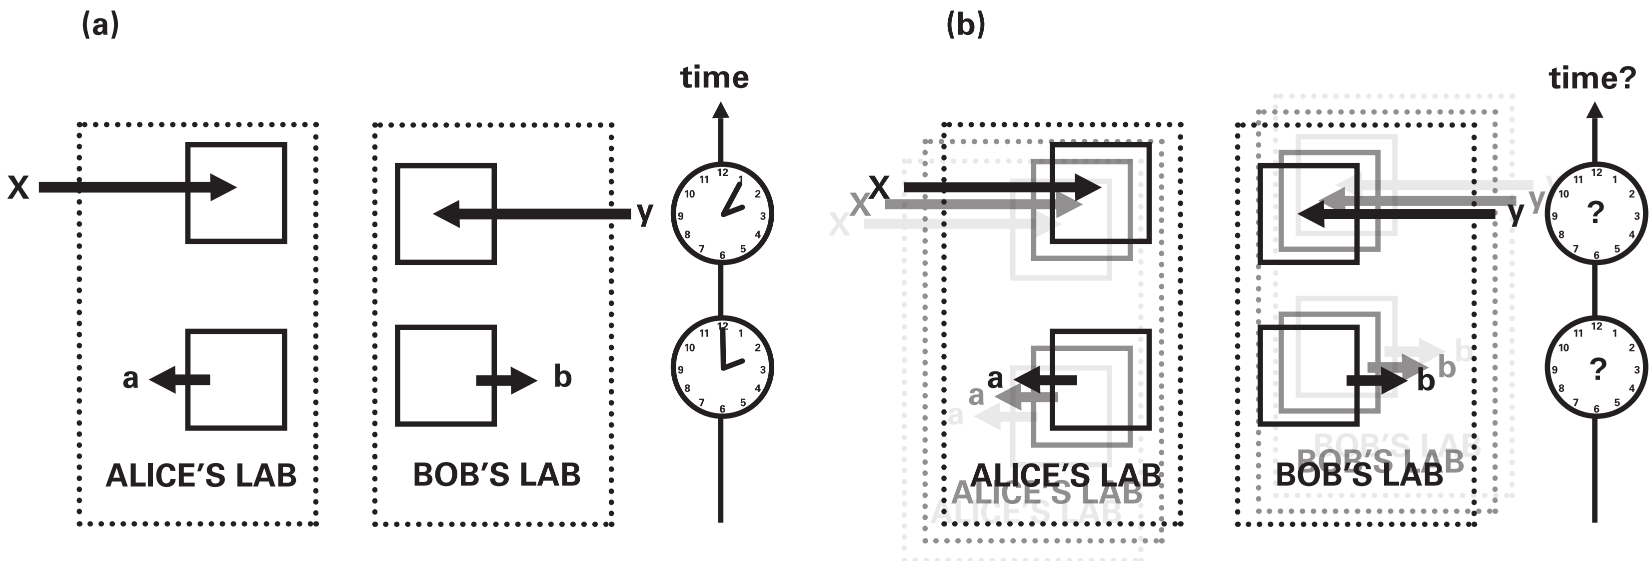

Figure 4. (a) Known timing: the parallel case. Alice and Bob's experiments are perfectly synchronised and there are two different systems in their labs, which are not exchanged. (b) Unknown relative timing: the order of Alice and Bob's experiments are completely unknown or even undefined.

Finally, the last timing scenario possible in the setup is the one in which the relative timing between the labs is unknown (Fig. 4(b)). It could be that the timing is actually well defined, but unknown to the experimenters, or that it varies from one experiment to another; it could even be the case that the timing is undefined in principle, in a strange world where no clocks can be constructed. In this case too, backwards propagating influences can be present and, intriguingly, they are 'more numerous' than in any of the other cases.

The four different timing scenarios are distinguished by the size (i.e., volume) of correlations that they allow. In each case, all of these correlations do not signal backwards-in-time (they obey the NBTS principle) but among them are also influences which travel from the future to the past. The number of possible different influences varies between the four cases. From smallest to largest, the size of the spaces of influences grows from: the parallel case, the sequential cases and the unknown order case which are 6, 7 and 8 dimensional respectively. Box-worlds, it turns out, can be very accommodating places.

All of the timing scenarios allow for the possibility of the future to affect the past, without backwards-in-time signalling. How on Earth in box-world can this possibly be? The resolution to this apparent conundrum, is that Alice and Bob cannot use these influences in their real-time to do anything useful. In fact, they cannot even detect them; only afterwards, after all the experiments are complete, when they come together and create the 'global view' of their experimental results do they see, in hindsight, that these influences were present. There is absolutely no way to control them when they are happening, and nor can they be used as a resource or harnessed in any way. The fact that the effects can only be discovered later also avoids paradoxes such as killing one's own Grandfather.

5 THE BENEFIT OF HINDSIGHT

So are you, readers of this article, receiving undetectable influences from the future? Are there waves of uncontrolled impacts cursing through your body and guiding you to actions beyond your control? It unlikely. In quantum theory all of the correlations obey the NBTS principle (Eqs. (7) and (8)), but none of the correlations (in any of the timing scenarios) allow for these influences to flow – after systematically characterising these exotic effects, one can check that they do not occur in quantum theory.

Thus, if you believe that quantum theory is the correct description of your current state, then no influences can flow from the future. In fact, one can go even further. Earlier, we said that classical and quantum theory differ in their correlations, in a particular setting. It turns out that in *this* setting they do not differ at all, and are, in fact, identical. Classical theory too obeys the NBTS principle and it too does not allow for backwards-in-time influence. What this means, is that in this two-step setting (measurement and transformation) in order to reproduce the correlations of quantum theory, i.e., to obtain all the patterns and dependencies present in experimental lists like those of Fig. 2, Alice and Bob don't even need quantum particles nor any of the fabled entanglement that comes with them. Some dice and a couple of pieces of paper (objects that can be described in classical theory) will do – the patterns and correlations will be identical.[3] Correlations, it turns out, are an interesting, but not necessarily distinguishing feature of a theory.

If they are not in quantum or classical theory, where are these backwards-in-time influences to be found? Is it possible to find them in another physical theory, 'beyond' quantum? Would it be possible to find and observe these influence by looking back in the records of peculiar experiments? Perhaps. There are many more box-worlds in the 6, 7 and 8-dimensional spaces to explore that do not correspond to classical and quantum theories. One can postulate the existence of such influences in, for example, exotic contexts of quantum gravity, of which we know very little, or simply in other box-worlds, which we can define on paper, but so far have failed to encounter in experiment.

3. *Exploring the limits of no backwards in time signalling,* Yelena Guryanova, Ralph Silva, Anthony J. Short, Paul Skrzypczyk, Nicolas Brunner and Sandu Popescu, Quantum 3, 211 (2019).

We stuck with the conceptual bedrock of physics, as well as intuition, and decided defacto that time-travel and closed time-like curves do not exist by expressly forbidding them. We did so by constructing a theory-independent definition to prevent signals locally travelling from the future to the past, a principle dubbed 'no backwards-in-time signalling'. Despite this veto, we discovered that it is theoretically possible to have situations in which the future demonstrably affects the past: by performing measurements and transformations in their laboratories, or, more precisely, by gaining and not-gaining information about something-systems in individually well-defined time-zones, the experimenters were able to see, in hindsight, that the non-occurred was able to affect the occurred, i.e., that there are influences that can propagate backwards from the future without creating closed time-like curves. Unfortunately however, these influences are not present in anything described by classical or quantum theory.

6 OUTRO

It is time to destroy the last bastions upon which we constructed our story, the experimenters Alice and Bob. For box-worlds are not worlds occupied by experimenters or referees, who only serve as allegorical explanations. Box-worlds are the bare minimum: one need not have an explanation for what is going on inside the box, there is no need for specifying the details of any experimental setups, nor for creating consistent scripts with actors. The fact that we search for 'convenient explanations' in terms of physical theories is something in addendum, perhaps to appease our own unease. In any box-world there is total sensory deprivation, there are only boxes and the lists of 0's and 1's, encoding the correlations, that they produce. Despite this, they are surprisingly accommodating, occupying multidimensional spaces and allowing for novel effects to flourish.

We do not know if these influences are already out there, and neither do we know that closed time-like curves are truly forbidden. Perhaps there are particles in our universe, that behave differently in regimes that we have not encountered yet, and for which we have no theories or models. Here, it may be that backwards-in-time influences can arise, against a foreground of new physics. Thus, even if it turns out that one cannot manipulate the past, one can certainly look forward to looking back.