What's So Spatial About Time Anyway?

Sam Baron and Peter W. Evans

Abstract

Skow ([2007]), and much more recently Callender ([2017]), argue that time can be distinguished from space due to the special role it plays in our laws of nature: our laws determine the behaviour of physical systems across time, but not across space. In this work we assess the claim that the laws of nature might provide the basis for distinguishing time from space. We find that there is an obvious reason to be sceptical of the argument Skow submits for distinguishing time from space: Skow fails to pay sufficient attention to the relationship between the dynamical laws and the antecedent conditions required to establish a complete solution from the laws. Callender's more sophisticated arguments in favour of distinguishing time from space by virtue of the laws of nature presents a much stronger basis to draw the distinction. By developing a radical reading of Callender's view we propose a novel approach to differentiating time and space that we call temporal perspectivalism. This is the view according to which the difference between time and space is a function of the agentive perspective.

1 Introduction

It is very natural to suppose that there is a substantial difference between time and space. But what exactly is that difference? Two of our best known models of space and time—the system derived from Newton's theory of gravitation and mechanics, and the system that arises from the special and general theories of relativity, which supersedes the former—diverge in the answer they provide to this question. In Newtonian spacetime, time is clearly and definitively distinct from the spatial dimensions, whereas in generally relativistic spacetime temporal and spatial dimensions are inexorably interwoven. The shift from Newtonian to relativistic mechanics may seem to bring with it the death of difference. But there is some distinction to be drawn between spacelike and timelike orientations in relativistic spacetime, a distinction that is ultimately encoded within both the mathematical signature of its metric and its causal structure.

There have been two attempts in recent times to clarify the exact difference between time and space. Both of these accounts, though independent, emphasise the unique role that time plays in the laws of nature: our laws determine the behaviour of physical systems across time, but not across space. The first of these accounts, due to Skow ([2007]), claims that the reason we never have evolution in a spatial direction is that such a schema would leave us with a paucity of information concerning physical behaviour in a four-dimensional spacetime region. That is, given information at, say, some location in space for all time, any physical laws resembling our current laws will not provide complete information concerning what is going on everywhere else at any time. The second of these accounts, due to Callender ([2017]), claims that time is the 'great informer' in the sense that it discriminates the direction in the manifold of events in which the laws of nature provide maximal determination of the behaviour of physical systems. That is, given certain properties of a large class of our laws, the direction in which these laws work most informatively aligns with the temporal direction.

In this work we assess the general claim that the laws of nature might provide the basis for distinguishing time from space by looking specifically at the claims of Skow and Callender. After we present Skow's account in §2, we outline in §3 an obvious reason to be sceptical of the argument Skow submits for distinguishing time from space. In short, Skow fails to pay sufficient attention to the relationship between the form of dynamical laws and the dimensionality of the antecedent conditions required to establish a complete solution from the laws. In §4 we turn to Callender's more sophisticated arguments in favour of distinguishing time from space by virtue of the laws of nature. We find these arguments to be a much stronger basis to draw the distinction. In §5 we differentiate between two readings of Callander's arguments: a conservative reading and a radical reading. In §6 we develop the radical reading of Callender's view into an appraoch that we call temporal perspectivalism. This is the view according to which the difference between time and space is a function of the agentive perspective.

2 Time as the Preferred Direction of Laws

Skow ([2007], p. 237) outlines the following position concerning the distinction between the spatial and temporal directions in spacetime:

> Timelike and spacelike directions play different roles in the laws of physics [. . .] Those laws govern the evolution of the world in timelike directions, but not in spacelike directions.

Skow takes the fundamental difference between time and space to be that the laws of physics determine behaviour by evolution in time, but they do not determine behaviour by evolution in space. This is not supposed to be a trivial claim; evolution is not to be interpreted analytically as evolution 'in time'. Rather we are to think of evolution as the determination of physical behaviour in some spatiotemporal direction, where that direction is determined by the relation between the laws of physics and the data required by the laws: the antecedent data comprises some subset of the complete information describing physical behaviour in some four-dimensional spacetime region and the laws determine the complete information given this antecedent data. Here is Skow ([2007], p. 237):

> Roughly speaking, by 'the laws govern the evolution of the world' in some direction I mean that the laws, together with complete information about what is going on in some region of spacetime, yield complete information about (or assign probabilities to complete descriptions of) what is going on in regions of spacetime that lie in that direction from the initial region.

Thus, given an amount of information concerning some subset of spacetime that corresponds to the data required to provide a well-posed initial value problem for some physical laws, we can produce enough information to describe physical behaviour in some four-dimensional spacetime region, and the direction from the initial region to the complete region is the direction of evolution. This is straightforwardly the case when the direction of evolution coincides with the direction of time. We take initial data on some spacelike surface and employ physical laws that require such data in order to determine the behaviour of some system throughout some four-dimensional region that includes the spacelike surface as an initial temporal boundary, so long as the initial data and the laws comprise a well-posed initial value problem. The laws thus govern evolution from one time to another. This is precisely the general structure that is typically characteristic of physics.

Skow's claim, then, is that the distinction between time and space consists in physical laws evolving in the direction of time, but never in the direction of space; we always take spacelike surfaces (data at a time) as initial data, our laws always govern evolution in directions normal to these surfaces, and thus evolution always occurs in the direction of timelike vectors. Skow ([2007], p. 237) says as much:

> Now, timelike vectors are not tangent to any time, on any way of partitioning any given spacetime into times [spacelike surfaces]. Rather, no matter which partitioning of spacetime into times you use, timelike vectors point from one time toward others. So timelike vectors point in the directions in which the laws govern the evolution of the world.

And then continues:

> The same is not true of points of space. If I know what is going on right here (at this location in space) for all time, the laws do not give me complete information about (or assign probabilities to complete descriptions of) what is going on anywhere else at any time.

Note that Skow's view is quite strong. Consider his definition of what it is for a direction in spacetime to be a temporal direction (Skow [2007], p. 238):

LAWS: A direction (that is either spatial or temporal according to the geometry) is timelike iff it (or its opposite) is a direction in which the laws govern the evolution of the world.

Because of the biconditional in LAWS, if the laws ever govern evolution across a given direction in spacetime then that direction must be a temporal one. So there are no possible laws that govern the evolution of the world across space.

According to Skow, the reason that there are no possible laws that evolve in a spatial direction is that such a schema would leave us with a paucity of information concerning physical behaviour in a four-dimensional spacetime region. In short, such 'laws' would never provide complete information about all of spacetime. Skow defends LAWS by considering a range of putative laws that might be thought to govern the evolution of the world across space. He argues, in each case, that these 'laws' fail to yield complete information about all of spacetime.

Consider, for instance, Skow's fourth (and most realistic, because it uses actual quantum laws) example of using laws to extrapolate information across space, rather than time (Skow [2007], p. 246):

Example 4: In this world, the laws of quantum mechanics govern the world. In an EPR-type experiment, there are two particles some distance apart, and if we measure the spin on one of them in some direction, we know with certainty the outcome of a measurement of spin of the other particle in that same direction, even if the measurement of events are spacelike separated. So these laws govern the evolution of the world in a spacelike direction.

Skow concedes that one can gain information about one part of spacetime based on information about a measured particle in another part of spacetime plus the laws, but goes on to argue that there is no way to use the same laws in such a case to extrapolate all of the information about the entire four-dimensional manifold. He extends precisely the same reasoning to each of the other three cases of a law governing the evolution of the world across space that he considers. In each case he admits that some information about another region of spacetime may be gained via the relevant laws, but denies that information about the whole manifold can be gathered.

In the next section, we will argue that LAWS is false. Our case against LAWS will proceed in two stages. First, we shall show that the reasoning Skow uses to defend LAWS is in error. He is not considering cases of the right type. And so while he is correct that the laws he considers do not govern the evolution of the world across space, this does little to show that such laws are impossible. Following that, we will outline a direct counterexample to LAWS from classical electromagnetism. These two stages are related: seeing the error in Skow's reasoning helps to illuminate the corresponding counterexample.

3 Dynamical Laws and Antecedent Conditions

Let us begin by considering in more depth the relationship between laws and antecedent conditions—the information that is fed into a law of nature. Ordinarily our laws strike two practical balances to achieve a complete specification of physical behaviour in a four-dimensional spacetime region. The first is between the number of variables constituting the antecedent conditions and the order of the dynamical law; and the second is between the number of dimensional components of these variables and the corresponding dimension of the law.

As a quick illustration we can consider second-order dynamical equations, which comprise a large portion of our known physical laws. Second-order equations require the antecedent conditions to be constituted by two dependent variables; for particle dynamics we ordinarily take these to be the position and velocity initial conditions, and for field theories the field values and their first derivatives. Since we ordinarily take initial data to extend across three spatial dimensions in our actual physical laws, these laws must be three-dimensional, relating three components of each of the dependent variables to the independent variable (which we ordinarily take to represent time).

For a law to yield a representation of physical behaviour throughout a four-dimensional region of spacetime, there needs to be a correspondence between the dimensionality of the antecedent conditions and the dimensionality of the law. The role of pretty much all of our current physical laws is to extend full information about three-dimensions through one further dimension—time. Thus, in order for the combination of the laws and the antecedent conditions to result in a description of physical behaviour throughout a four-dimensional region of spacetime we should at least require our antecedent conditions to describe complete information across a three-dimensional sub-region of spacetime (such as a spatial configuration at some time).

For any case in which the antecedent conditions fail in combination with the dynamical laws to yield complete information about a four-dimensional sub-region, there are two possible modifications that we could make so as to achieve such complete information. One way would be to keep the antecedent conditions as they are and apply them to more wide-reaching laws. To illustrate such a law, imagine the following highly simplified scenario.

Consider a universe in which each location in Minkowski spacetime is simply a single many-valued property. We could think of this property as, say, colour: each location in spacetime is green, or red, or blue, and so on. Information concerning a single location in space for all time would consist in a linear series of colours on an everywhere-timelike curve through spacetime. We could now imagine a law that took this information as an antecedent condition and determined the value of this property throughout four-dimensional spacetime. For instance, given some foliation of the spacetime into non-overlapping surfaces of constant time (that is, Cauchy surfaces) the universe is a single colour at each time. Thus, given a linear series of colours on an everywhere-timelike curve through spacetime, and the law stipulating uniform colour on each Cauchy surface, we have complete information concerning four-dimensional spacetime.

A second way that we might achieve complete information is to keep our ordinary laws and expand the scope of the antecedent conditions instead, by feeding in more information. In the case where we take our laws to give evolution in time, we expect our antecedent conditions to contain total information about three spatial dimensions at some part of the manifold (such as a total spatial configuration at some time). Similarly, equivalent laws evolving in space

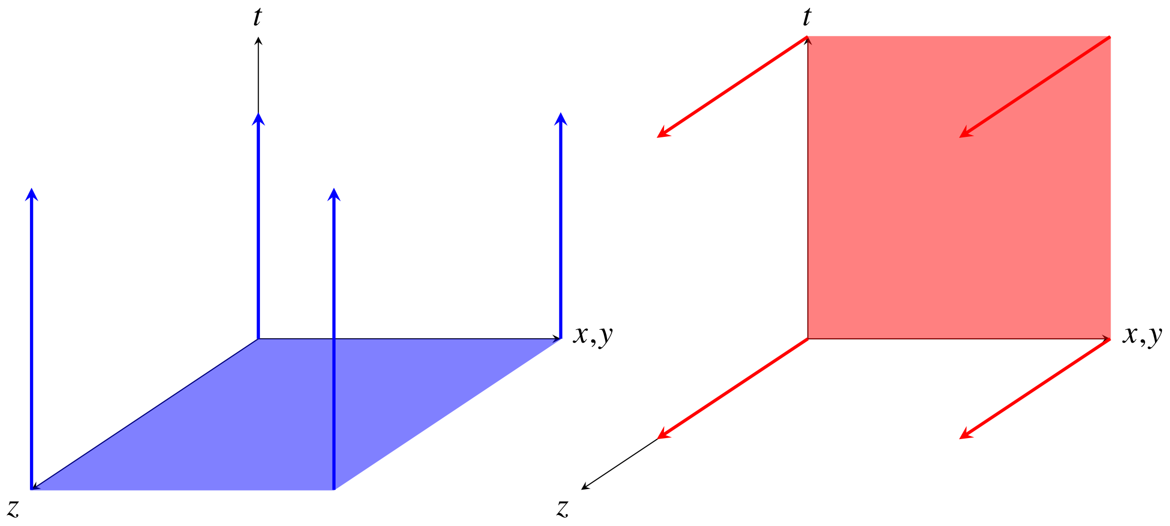

(i) (ii)

Figure 1: A representation of temporal and spatial evolution: (i) represents an extrapolation through time via the laws; (ii) represents the analogous extrapolation through space.

would also require antecedent conditions that are three-dimensional—in this case, we would need complete information concerning two dimensions of space and one dimension of time at some part of the manifold. Given such information and the requisite laws, we could compute complete information concerning four-dimensional spacetime.

When discussing the evolution over time governed by our ordinary laws, Skow considers cases in which we feed in total information about the three spatial dimensions that correspond to a particular time. Which is to say that, given some foliation of spacetime into Cauchy surfaces parameterised by t, for some such time, t_p, a complete three-dimensional specification across the associated Cauchy surface comprises the antecedent conditions. This yields total information for all $t < t_p$ and $t > t_p$. Skow argues that the analogous operation in the spatial case never yields total information. The analogous operation in the spatial case, however, would be this: given some foliation of spacetime into surfaces of constant 'location' z (with the set of surfaces parametrised by z), for some such spatial location, z_q, a complete three-dimensional specification across the associated surface of constant z, consisting of information about the remaining two spatial dimensions plus information about the temporal dimension, would constitute the antecedent conditions. The laws should then enable one to extrapolate information for all $z < z_q$ and $z > z_q$ (see Figure 1 for a representation of the two cases). In short: in the temporal case, we extrapolate from information about all three spatial dimensions at a time. In the spatial case, we extrapolate from information about two spatial dimensions and all of time at a point in space.

But this is not the operation that Skow considers when defending LAWS. Rather, Skow considers cases in which one takes partial information about some part of space, which may or may not include information about time, and then one attempts to extrapolate across all of spacetime. Each of the cases that Skow considers can thus be represented by Figure 2. Limited

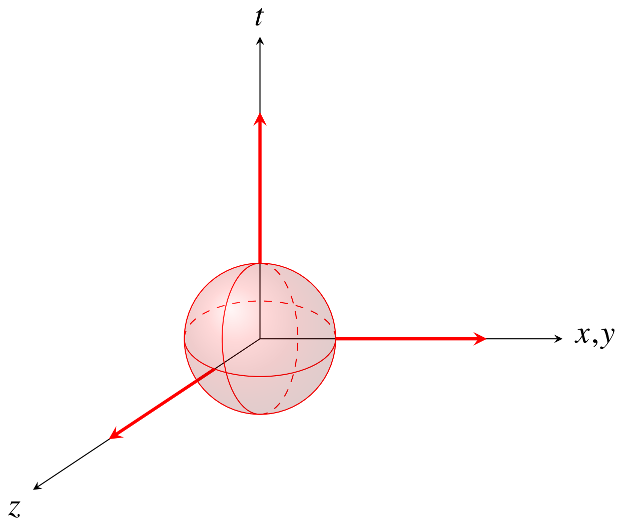

Figure 2: Skow's proposed extrapolation through space.

information about a particular spatial region is used to try to gain information about all of spacetime. This is a general feature of Skow's examples. As Skow ([2007], p. 246) puts the point:

> For laws to govern the evolution of the world in a spacelike direction, it must at least be the case that given complete information about what is going on right here, the laws give complete information about (or assign probabilities to complete description of) what is going on at other places that are spacelike separated from it. But if all we know is that (after being measured) some particle right here has spin-up in some direction, the laws don't tell us anything about what is going on elsewhere. They only give us information if we also know that there is another particle somewhere else, and that the 'system' comprising the two particles is in an entangled state. But this information is not just information about what is going on right here.

The crucial missing ingredient is temporal information: in order to extrapolate all of the information about a complete four-dimensional spacetime from two dimensions of space, full information about the temporal dimension is needed. In the cases he presents, no temporal information is used. It is unsurprising, then, that only limited information about spacetime is yielded.

What we need to look at are cases that are like those represented by Figure 1(ii). Which is to say that, in order to defend LAWS, it must be shown that there are no possible laws that govern the evolution of the world in the following manner: by feeding in full information about the temporal dimension and two spatial dimensions, full information about the remaining dimension can be obtained. Because Skow does not consider any of those cases, he has failed to provide any support for the claim that it is *impossible* for there to be laws that govern the evolution of the world across space in the same way that the laws govern the evolution of the world across time. The relevant impossibility result is crucial to Skow's account of the difference between the spatial and temporal directions of spacetime. For, as he notes, his 'view entails that it is not possible that there be laws that govern the evolution of the world [in] a spacelike direction'. Given this entailment, if it is possible that the laws govern the evolution of the world across space in exactly the manner in which they govern the evolution of the world across time, then that is sufficient to show that the essential difference between time and space cannot be to do with the fact that the laws can govern the evolution of the world across the former but cannot across the latter.

It is therefore interesting to inquire as to whether there are any plausible cases of spatial evolution in the sense of Figure 1(ii). For it is one thing to note that Skow has failed to rule such cases out as being impossible, quite another to show that there are such cases in the offing. If there are no such cases, then that would seem to strengthen Skow's conviction that the distinction between time and space ultimately resides with the manner in which the laws govern evolution across spacetime, even if his way of developing that distinction is ultimately too strong. In order to make our point more forcefully, then, let us consider an example that looks to be evolution in space of the kind that we deem to be acceptable, and the kind that we believe Skow has failed to take into consideration.

Consider classical electromagnetism, a classical relativistic field theory. Classical action principles of the sort found in electromagnetism generally require that the field boundary is fixed on a closed hypersurface, wherein the timelike parts of the boundary are just as important as the spacelike parts. Determining physical behaviour as a solution to the action integral involves integrating over the spacetime region enclosed by the boundary. It is interesting to note that neither the integral over space nor over time has any precedence over the other: the solution will comprise the same action for the same field in the same spacetime region regardless of whether we represent the integral as 'evolving' a spacelike hypersurface in time or a timelike hypersurface in space. If we imagine an electromagnetic system consisting of a perfect conductor at some location which constrains some of the electric and magnetic field components to be zero, so long as there are no other boundary conditions, an analysis of the fields near the conductor might appear as though the field is 'evolving' in a spatial direction away from the conductor, *à la* Figure 1(ii).[1]

While we take this example from electromagnetism to provide a clear counterexample to Skow's arguments regarding the distinctiveness of time, we do not wish to argue that this sort of example provides good evidence for a strong nomic symmetry between space and time (for reasons to which we will return in the context of the discussion in §4). But we do hope to have shown that this is the kind of spatially directed law that Skow needs to take into account to make his defense of LAWS convincing. Let us now turn to a second attempt to differentiate time from

[1]Callender ([2017], p. 139) offers another counterexample to Skow's view. He writes:

> [...] there are plenty of actual laws that govern in a spatial direction in contemporary phyiscs. The constraint equations of relativity seem to be an example: these operate across spatial directions and cannot be derived from the dynamical equations of relativity. For example, the ten vacuum Einstein field equations separate into six "evolution" equations $G_{ij} = 0$ and four "constraint equations," $G_{00} = 0$ and $G_{0i} = 0$, with $i = 1, 2, 3$. The latter impose nomic conditions across a spacelike slice.

space, an attempt that lies in a similar vein to Skow's but that does not face the difficulty we have outlined in this section.

4 Time as the 'Great Informer'

Similarly to Skow, Callender ([2017]) also identifies the laws of nature as the key to understanding the fundamental distinction between space and time: 'There really is a difference between space and time embodied crucially in its connection to the laws' (Callender [2017], p. 179). However, according to Callender, the feature that differentiates time from space is that time is the 'great informer': time is the direction in the manifold in which the greatest amount of information can be generated by the smallest set of antecedent conditions. This sort of informativeness—the ability to 'generate some pieces of the domain of events given other pieces' (Callender [2017], p. 141)—is a hallmark of a good balance between strength and simplicity in a best systems account of laws: the laws arise as the most accurate description of as much of the world as possible (strength) in the most succinct manner (simplicity). Thus, for Callender, 'time is that direction in spacetime in which we can tell the strongest or most informative stories' (Callender [2017], p. 142).

Callender offers two distinct but related arguments in favour of his view. The first, call it the 'informal' argument (Callender [2017], Ch.7), shows how the direction of informativeness that emerges from the process of systematisation that characterises the best systems account of laws 'binds' together a set of features ordinarily associated with time. Callender begins this argument espousing a conservative reading, in which the systematisation exposes 'an asymmetry in the distribution of events' (Callender [2017], p. 142) in the manifold, before progressing to a more radical reading, in which 'the choice of metric geometry hangs on systematizing too' such that 'the difference is not "out there" prior to systematizing' (Callender [2017], p. 151). The second argument, call it the 'formal' argument (Callender [2017], Ch.8), shows that it is a formal property of the laws that uniquely distils the direction of informativeness, and so more tightly connects the set of temporal features with this direction. We outline both arguments in this section. In §5, we consider the consequences of combining the lessons of the formal argument with the radical reading of the informal argument. We suggest that, on the radical reading, the difference between time and space in the formal analysis may be due to an underlying pragmatic choice of natural kinds.

4.1 The informal argument

The informal argument that Callender develops is based on the best systems account of laws.[2] According to this account, laws arise through a process of balancing a range of theoretical virtues amongst a set of deductive systems each correctly describing the facts of reality. The two most prominent virtues according to the view are simplicity, which is a measure of descriptive

[2] See (Lewis [1983], [1994]) for an outline of the best systems approach.

economy for fundamental properties, and strength, which is a measure of informativeness of particular matters of fact. The 'best systems' optimise the balance between simplicity and strength, and so produce simple and powerful 'algorithms' for describing reality, and the laws are those algorithms that are shared by the best systems.

Callender claims that the algorithms—that is, the laws—that best satisfy the trade-off between simplicity and strength will mark time out as special: simple and powerful laws seize upon a time parameter to generate strength from economical input such that what emerges as the temporal direction will be the direction in which the laws are as informative as possible, while being as simple as possible. Moreover, Callender argues that we can use this feature of laws as the basis for identifying the difference between time and space. Given a distribution of events in the manifold, with no presupposition concerning which of the directions are spatial or temporal, the best systems account provides a formula by which the temporal direction can be extricated from the spatial: the temporal direction will be that direction in which the laws tell the most informative stories while being as simple as possible.

As we mentioned above, what ultimately renders the direction of informativeness as the temporal direction according to Callender is the set of features that we ordinarily associate with time that it binds together. The details of this element of the informal argument are not so important, but we can think of this, roughly, in terms of a functional role for time. To establish that the direction picked out by the laws is time, the relevant direction must be capable of playing the right functional role. Callender argues that there are a number of features that something must have in order to count as time; for instance, one-dimensionality and directionality. Thus Callender completes his informal argument by demonstrating that the direction in which the most informative stories are told brings together the right features, and so plays the right functional role to be time.

Why should laws that seize upon a direction of informativeness arise from the best systems account? One possibility is that the best systems expose an inherent directionality in the distribution of basic properties in the world. Thus time is picked out as special due to some asymmetry in the distribution of events on the manifold, and the laws then exploit this asymmetry to gain strength in one set of directions rather than another (Callender [2017], p. 143). More precisely, any inherent structure in the distribution of events will be captured by a metric, g, on the manifold, $\mathcal{M}$, yielding a geometry, $\langle \mathcal{M}, g \rangle$. So according to this possibility, a distribution of events with a metrical structure, g, already endowed upon it stands as an implicit input to the systematisation process for establishing the laws on $\mathcal{M}$, and is treated as part of the data that our laws aim to systematise. The direction that then emerges as time on account of the laws uncovers the signature of g. We take this to be the conservative reading of the informal argument.

Callender ([2017], pp. 149–50) goes on to differentiate this conservative reading from a more radical reading of the informal argument. The conservative and radical readings differ over exactly what it is that the best systems approach to the laws systematises; the radical reading aligns more closely with what is known as the 'better' best systems account (Cohen and Callender [2009]; Callender and Cohen [2010]; Schrenk [2014]). According to the radical reading, g is not treated as an input to the best systems account. Rather, the best systems account takes, as input, simply the worldly distribution of properties. The metric then emerges from the best systems as a package deal with the laws. Ultimately, Callender ([2017], p. 151) recommends the radical reading of his view:

> If the more radical version of the thesis is right, then the choice of metric geometry hangs on systematizing too. That would mean that the *metrical difference* between the timelike and spacelike—the centerpiece of all our physical theories—also depends on the system. To be clear, this difference will remain objective because the laws are objective on a best system theory [. . .] Nonetheless, the difference is not 'out there' prior to systematizing, and this consequence may be surprising [. . .] In many ways this radical perspective is the most natural development of the theory.

Callender views this position as 'radical' in part because the best systematisation, where 'the best system systematizes "at once" the kinds, laws, and geometry' (Callender [2017], p. 151), is beholden to a choice of natural kinds. Different choices of what the natural kinds are will therefore change the output of the best system. Given a sparse conception of the natural kinds according to which nature 'prefers' a particular group of properties, there is just one way for the natural kinds to be. And so it is unlikely that there will be much variance in the output of the best system. However, Callender goes on to express sympathy for an 'abundant' conception of natural kinds, according to which 'we' select which properties are natural kinds. Given an abundant conception of natural kinds, the output of the best system is highly contingent. For it may be that two disparate groups of scientists select disjoint groups of natural kinds to form the basis of their respective best systems analyses. The upshot being that the two groups of scientists may come up with very different laws for describing the same universe. Callender ([2017], p. 154) accepts this outcome of his view:

> the proposal is *knee deep* in relativism, but this is a virtue not a vice. If Martian scientists cared about an alternative carving of nature into kinds, then presumably they might devise a different system than we do. Let's suppose the direction of strength in that system is orthogonal to ours, so the two temporal directions don't line up. The Martians think that a direction we consider spatial is temporal. Is one of us wrong? [. . .] If an abundant conception of properties is right, and it's only *we* who pick out the natural kinds, then the Martians can pick out a distinct temporal direction without error [. . .] for what it's worth I prefer the abundant properties option.

Note that the Martian scientists cannot be correct about their 'temporal' direction if we adopt the conservative reading of the informal argument. If g is an input to the systematisation process, then the Martian scientists' direction of strength would have to produce laws that evolve in a spacelike direction (according to the metric) across $\mathcal{M}$. But the provision of any such laws will require the identification of efficient algorithms that generate the facts of reality orthogonal to the inherent asymmetry of the distribution of events. It is doubtful that such algorithms can arise from the best systems systematisation process. If, however, the metric emerges from the process of systematisation, then the Martians could plausibly establish efficient algorithms, based on their choice of natural kinds, in some direction that we consider spacelike.

4.2 The formal argument

The formal argument is in a sense an expansion on the informal argument. The formal argument shows that it is a formal property of the laws that uniquely distils the direction of informativeness, and so, by the informal argument, the direction of time. The extra feature brought to bear in the formal argument is a particular type of strength that is relevant for a very general sort of physical law.

As before, we can understand strength in terms of how many particular matters of fact some system of laws manages to imply. While a deterministic system (that is, one that generates the world completely) is the strongest such system, indeterministic systems can be strong too (in particular, Markovian systems of laws). But, as Callender points out, a formal characteristic of a system of laws that turns out to be even stronger 'for us' than determinism is that the system comprises a well-posed Cauchy problem.

Cauchy problems are characterised by seeking out a solution to a dynamical equation (usually a partial differential equation (PDE)) given antecedent data (that is, Cauchy data) specified on a hypersurface in the domain of the solution (in our case, the manifold). The problem is 'well-posed' if a solution (i) exists for any possible antecedent conditions on the Cauchy surface, (ii) is unique, and (iii) changes continuously with the antecedent conditions. This latter condition is especially relevant for understanding the qualifier 'for us': small errors in specifying the antecedent conditions lead only to small errors in our determination of the subsequent dynamics (Callender [2017], p. 162). So, for creatures such as us with the sorts of practical concerns that we have regarding what sort of information we are interested in, plus the limitations we are under with respect to the information we are capable of gathering, the most informative dynamical laws appear to be well-posed Cauchy problems. Thus, for us, well-posed Cauchy problems tell the most informative stories.

The next step in the formal argument is that, given that very many of our most important dynamical laws are second-order linear PDEs, only a specific subclass of second-order linear PDEs admit well-posed Cauchy problems: hyperbolic PDEs (Callender [2017], p. 167). This is important because, mathematically, hyperbolicity places certain restrictions on the kind of boundaries that can be used to solve such differential equations. Part of the reason for this is that any physical signals that emerge from the solution to a Cauchy problem must propagate in accordance with the hyperbolic dispersion relations of the relevant differential equation, and this prescribes a 'causal cone' at each point that the solution is defined (Callender [2017], pp. 152, 169), and this in turn endows the solution with a causal structure (in a sense, solutions 'evolve causally' from the antecedent boundary). Accordingly, physical solutions to well-posed Cauchy problems must respect this causal structure and thus only certain types of antecedent conditions will be able to accommodate such solutions. The hyperbolicity of the laws therefore establishes a clear formal distinction between the surfaces that can count as antecedent conditions and the direction normal to such surfaces that determines the direction in which the solution 'evolves'. This distinction identifies the appropriate antecedent conditions that respect this structure, and so admit solutions to well-posed Cauchy problems, as spacelike hypersurfaces.[3] So only hyperbolic PDEs admit well-posed Cauchy problems, and hyperbolic PDEs pick out a causal structure such that their admissible antecedent conditions are spacelike, and thus what we interpret as evolution must happen in a direction normal to these surfaces.

When we put all this together we see that, since well-posed Cauchy problems tell the most informative stories, the most informative stories are told in a direction normal to spacelike hypersurfaces. Since we have identified time here as the direction in which we tell the most informative stories, time must be the direction normal to spacelike hypersurfaces. Thus, 'well-posed Cauchy problems pick out temporal directions' (Callender [2017], p. 166).

The counterexample that we raised earlier against Skow's view—the example from classical electromagnetism from the end of §3—is no challenge here. Even for classical action principles of the sort found in electromagnetism, time and space cannot be treated as identical. The reasons for this are essentially those given by Callender ([2017], §6.2.1). For a start, the metric signature of time differs from that of space, and we can take this to be an expression of the hyperbolicity of field theories like electromagnetism. And, as we noted above, hyperbolic equations pick out a causal structure wherein the admissible antecedent conditions are spacelike. When we further note that there is no general theory for existence and uniqueness of solutions to boundary value problems with timelike antecedent conditions (which, recall, are two of the conditions for well-posedness), the possibility for timelike Cauchy data looks rather slim.[4]

5 A Difference in the Laws?

The key to Callender's argument is that there cannot be laws that inform across space that are more simple than, and just as informative as, our laws that inform across time. Put this way, we can relate Callender's account of the distinction to our critique of Skow's. Given scant antecedent data of the sort that would be required for laws to inform across space (such as is depicted in Figure 1(ii), for instance), it seems likely that strong informative laws would need

[3]By referring to these surfaces as spacelike we do not intend to imply that they are so defined with respect to the metric. Rather, it is the laws themselves that determine what their appropriate initial data surfaces are, and it just happens to be the case that these surfaces are aligned as spacelike. We follow Callender ([2017], p. 153) on this:

> informative strength doesn't merely help justify a particular set of independently existing metrically defined directions as temporal. Rather, it plus the laws help define the metrical geometry that picks out a set of directions as temporal.

[4]Although, see (Weinstein [2008]) for an exploration of existence and uniqueness of solutions to boundary value problems with timelike boundaries. In contrast, the (perhaps fatal) difficulty of such a project can be distilled from (Geroch [2011]), whereby any timelike surfaces would contain characteristic surfaces of the equations—surfaces along which causal propogation occurs—and so any initial data on such surfaces could not be freely specified whilst simultaneously obeying the laws. (This point is also made by Callender ([2017], p. 170).)

to be overly complex. If there were to be spatially directed laws simple enough to compete with temporally directed laws on a best systems account, then it had better be the case that antecedent data of a similar dimensionality to a spacelike hypersurface be required to make the spatially directed laws maximally informative. But we rarely (if ever) have access to (nor interest in gathering) this sort of antecedent data, so in this sense it is difficult to see how spatially directed laws could be more simple than temporally directed laws that arise from the best systems account.

This latter point is instructive, and will constitute at least part of our focus for the remainder of this paper. So let us be a bit more precise: it is difficult to see how spatially directed laws could be more simple than temporally directed laws, given a conservative reading of the best systems account. This conservative reading is implicit in Callander's presentation of the formal argument (see, for instance, (Callender [2017], p. 172, Fact 2) and associated discussion). The formal argument, in which the direction of informativeness, and thus the the direction of time, are formally connected to the direction of production for well-posed Cauchy problems, takes as its starting point the established laws from the systematisation process—the second-order linear PDEs. On the conservative reading of the best systems account, however, this systematisation happens over $\langle \mathcal{M}, g \rangle$ with a metric signature already endowed upon the space. Thus the signature of $\langle \mathcal{M}, g \rangle$ is embodied, through the systematisation, in the signature of the fundamental equations governing physical behaviour on $\mathcal{M}$. The formal argument consequently shows that it is a formal property of these laws, as well-posed Cauchy problems, that they pick out a unique direction of time that exploits the inherent asymmetry of the distribution of events on $\mathcal{M}$ to achieve an efficient algorithmic description. The conclusion that Callender draws from this, as we have seen, is that the 'difference between space and time [is] embodied crucially in its connection to the laws'.

In this section we suggest that if the lessons of the formal argument are combined with a radical reading of the best systems account, then the difference between time and space that inheres in the laws can be explained in terms of the pragmatics of agents undertaking the relevant systematisation. On the radical reading of the informal argument, it is not simply the laws that arise from the best systems systematisation, but the spacetime geometry $\langle \mathcal{M}, g \rangle$, particularly the metric signature.[5] Indeed, Callender ([2017], p. 152) sets out an argument from Geroch ([2011]) that, given one or more physical fields on a manifold, $\mathcal{M}$, with no metric structure, and the fact that the laws governing the fields admit of a hyperbolisation, and so admit well-posed Cauchy problems (in the same sense employed in the formal argument), the laws define a metric structure on $\mathcal{M}$. So rather than the systematisation exposing 'a particular set of independently existing metrically defined directions as temporal' (Callender [2017], p.

[5]As Callender ([2017], p. 151) puts it:

> the metrical distinction between time and space and everything that hangs on it [...] all fall on the "system" rather than "systematized" side of things. In particular, the lightcone structure that plays a crucial role in our understanding of time *emerges* from the laws that best systematize the events on $\mathcal{M}$.

153) (as in the conservative reading), systematisation on the radical reading defines the metrical geometry, and so the metrical difference between time and space. It is ‘here in [the structure of the fields] we find the principal difference between time and space’ (Callender [2017], p. 153). Recall, however, that this systematisation is carried out with respect to a particular choice of natural kinds: the previous quote continues, ‘And that difference arises thanks to the choice of kinds and form of the laws of nature’. But also recall that it is the agents conducting the systematisation who make the choice of natural kinds in the interest of optimisation. As per the quote at the end of §4.1, Callender explicitly recognises that the radical reading allows for there to be a group of Martian scientists who identify a direction in the manifold as the temporal direction that is orthogonal to the direction that we pick out.

But what is this decision about natural kinds based upon? When we carry out the systematisation process, implicit in the process is that we are not simply trying to find the laws that best meet some absolute trade-off between simplicity and strength, we are trying to find those laws that are strongest in a manner that aligns with our particular predictive practices. Given this, it is plausible that our predictive practices, and so the particular trade-off that we make, are a function of our epistemic vantage point on the world. The information about the manifold to which we have access is exclusively in our pasts, and we are interested in using such data to model our unknown futures. Since we take our past to be a predictor of our future, we are naturally predisposed (especially in a Markovian setting (Callender [2017], p. 143)) to take the boundary between our past and future as the antecedent boundary of our predictive practices, and to model our future as generated from our past in the style of a Markovian processor—and we have developed scientific practices that reflect this natural predisposition (this is exactly what our laws do). Thus when systematising over the distribution of events to which we have access, we are pragmatically constrained to identify natural kinds living on the antecedent boundary separating our past from our future—that is, spacelike hypersurfaces[6]—and that are best placed to allow efficient algorithms to take such antecedent boundaries as input and such kinds as dependent variables.

This point is not lost on Callender ([2017], p. 149):

> the natural kinds with their time/space split are essentially connected with what the laws are according to a systems perspective. The laws [. . .] are written in a special vocabulary (whether chosen by us or by the world) and the predicates of that vocabulary are the natural kinds of that world. This metaphysical connection in systems theory enshrines as metaphysical the epistemic point that we appear to choose the kinds we do in part because they are the ones with which creatures like us can make good explanations and predictions.

It is plain to see then that differently oriented agents will have precisely the same concerns according to their own epistemic vantage point in the manifold. Thus, the choice of natural kinds for Callender’s Martians will be based on their particular pragmatic concerns for undertaking their own predictive practices. To make this point more vivid, suppose that the Martians live

[6]This is despite the fact that we have as much epistemic access to such a boundary as we have to our own future.

their lives sideways, in what we consider to be the 'eastwards' direction of the manifold. These beings are not particularly interested in predicting anything across (what we would call) time, for they are smeared out in (what we would call) the temporal dimension, in just the same way that we are smeared out in (what we would call) the spatial dimension. What they want to know is what the regions 'ahead' of them are like (they 'remember' along what we would call a spatial direction, and 'predict' in the opposite direction). Just in the same way that we are pragmatically constrained to identify natural kinds on the antecedent boundary separating our past from our future, so too are the Martians pragmatically constrained: when systematising the distribution of events on $\mathcal{M}$, they will identify natural kinds on (what we would call) a timelike boundary separating their 'past' from their 'future', and devise laws that predict into their future. The kinds identified by the Martians would no doubt look grotesquely nonlocal to us, as ours would to them; indeed, one could imagine complete incommensurability between the two ways of carving up the world.

But then what would the Martians' laws look like? Given the arguments Callender marshals in favour of the generality of the laws that underlie Geroch's derivation of metrical structure from the laws and given Callender's own formal argument, we see no reason to doubt that the Martians would settle upon the same class of second-order linear PDEs that we do, just ranging over a (plausibly incommensurably) different set of dependent variables. But if this is right, then the formal argument will have exactly the same significance for the Martians as Callender points out that it has for us: it is a formal property of the Martians' laws that uniquely distils the direction of informativeness, and so the direction of time, 'for them'. Of course, whether or not the Martian's laws are second-order linear PDEs is not important for the main point we are making. The central point is that the mere possibility that laws devised from a perspective different from our own might underlie a different distinction between time and space suggests that the distinction is pragmatic in origin.

It is in this way that we can understand the lessons of the formal argument along the lines of the radical reading of the informal argument. The formal argument shows that the direction of informativeness must be normal to spacelike hypersurfaces in $\langle \mathcal{M}, g \rangle$. According to the argument the best, most efficient, and most informative algorithms developed by our scientific practices admit only antecedent data on spacelike hypersurfaces. But while a conservative reading of the formal argument renders these spacelike hypersurfaces as a property of a given metric signature inherent in the distribution of events on $\mathcal{M}$, a radical reading renders these spacelike hypersurfaces as effectively arising during the systematisation of the natural kinds and laws, which we have seen is a function of the pragmatic concerns of the systematising agents. Thus, for us, the efficient and informative algorithms that we develop that admit only antecedent data on spacelike hypersurfaces have the formal characteristics they do precisely because of our pragmatic choice of natural kinds on such spacelike boundaries, a choice that is made in the interest of modelling our future as a function of our past. The pairing of natural kinds and laws brings with it a time/space split that reflects the epistemic constraints and pragmatic interests of agents.

It should not be surprising at all, then, that the laws that are simpler are the ones that inform in what we call the temporal direction. This is because we have exclusively information about our past available to us, and because modelling our future based on our past discriminates spacelike hypersurfaces. If the signature of the Martians' fundamental equations were Lorentzian, $(-, +, +, +)$, in which their laws informed along the '$-$' direction, then the Martians might well use spacelike hypersurfaces that are distinct from our spacelike hypersurfaces without error, and in a way in which the formal argument still stands. In so far, then, as the formal argument shows that there really is a difference between time and space embodied in its connection to the laws, the radical reading opens up the possibility that this difference in the laws is due to the pragmatic concerns of the agents who developed them. If the same kinds of laws are used by two different agents from two different perspectives to carve the world in two different ways, then the difference between time and space embodied in the laws appears to reflect the differing epistemic vantage points of the two agents.

In sum, then, we are not suggesting that under some different regime of collecting antecedent data we might find ourselves in a position to trade off simplicity and strength with respect to our current laws in a manner that would pick out a different direction as temporal, and so undermine Callender's account of temporal direction. We take this to be exactly the scenario that Callender's formal argument successfully rules out. Rather, we contend that Callender's conclusion from the formal argument that the difference between time and space is 'embodied crucially in its connection to the laws' may be too quick. If we take the radical reading from the informal argument and apply it to the formal argument, then a different conclusion becomes available. On a radical reading, the very laws employed in Callender's formal argument are devised as a function of our predictive interests and, since we have no independent means to establish the virtues of any laws devised from any alternative antecedent conditions (short of us actually adopting and developing predictive practices in the alternative perspective), we have no grounds to argue that these cannot be the laws that underpin a different time/space split from another perspective. Thus, depending upon whether one reads the informal argument conservatively or radically, the difference between time and space may be fundamentally embodied in the laws, or may instead be due to a pragmatic distinction that reflects our unique epistemic vantage point on the world.[7]

Since the radical reading is the reading of the informal argument that Callender prefers, we by no means see our analysis as a problem. It is, rather, an opportunity to take Callender's analysis to its logical, albeit more radical, conclusion. And so, in the next section, we elaborate an account of the difference between time and space that is cast in explicitly pragmatic terms.

[7]Of course, as Callender notes, if nature prefers a particular way of carving the world into natural kinds then even on the radical reading, there may be just one way to select the natural kinds and so just one way to perform the best systems analysis. Then there will be just one way to isolate a temporal direction, and that isolation will be quite independent of the pragmatic interests of any scientists.

6 Temporal Perspectivalism

Our own view of the distinction between time and space is analogous to the distinction between cause and effect according to causal perspectivalism. Very roughly, causal perspectivalists argue that on an interventionist account of causation—in which we say that A is a cause of B if and only if doing A is an effective means of bringing about B (Woodward [2003])—the distinction between cause and effect can be reduced to an agent's perspective (Price [2007]; Ismael [2016]). According to causal perspectivalism, the distinction is built upon the deliberative stance of an agent towards the manipulable parts of the world, which is tied directly to an agent's epistemic vantage point. By definition, one can only deliberate to bring about unknown events, so a change in epistemic perspective can lead to a change in deliberative stance towards the world—the facts about what causes what vary accordingly. Causal perspectivalism allows for two agents to disagree about what causes what, and for that disagreement to be a no-fault disagreement. Each agent is correct, given their agentive outlook. It also allows for agents similarly placed to agree about the causal facts, so long as they occupy a broadly similar agentive perspective. Since, as it turns out, we all happen to occupy the same epistemic perspective on the world, we all generally agree about causal facts.

An analogous view is available regarding the difference between time and space. Temporal perspectivalism aims to reduce the difference between time and space to the agent's perspective.[8] A basic version of the view can be stated as follows:

Temporal perspectivalism: t is a temporal direction for an agent a iff t is the direction along which the most useful determinations can be made, from a's perspective.

We can further sharpen temporal perspectivalism by drawing on the discussion of causal perspectivalism by Ismael ([2016]). Ismael offers a broad distinction between a semantic version of causal perspectivalism and a frame-dependent version of the view. According to the semantic version of causal perspectivalism, the difference between cause and effect is to be understood in terms of some implicit or explicit contextual parameter in our semantics for causal talk. The strongest version of this view is the one where 'cause' is an indexical like 'I' and so the truth conditions for 'x causes y' make explicit reference to the context of an observer, in the same way that the truth conditions for 'I' involve a function from a context (usually a centred world) to a proposition or truth-value for a proposition. The frame-dependent view, by contrast, is not primarily a thesis about semantics. The idea is that there is some invariant fundamental structure in the world that is independent of any perspective. We then carve the invariant structure into cause and effect based on our idiosyncratic epistemic constraints and limitations concerning that structure. Thus, what is classified as a cause for one being, may be an effect for another being who has a different epistemic outlook on the world. This is roughly

[8] We know of two places where vaguely similar ideas have emerged. In the foundations of physics Rovelli ([2017]) argues that the arrow of time could be perspectival. More relevantly, Massimi ([2018]) refers explicitly to perspectivalism in the better best systems account. However, the origin of Massimi's perspectivalism is the differing of scientific virtues across different scientific communities, whereas the origin of the perspectivalism we emphasise is the epistemic limitations of agents embedded in spacetime.

analogous to the way in which, with respect to simultaneity, there is an invariant structure—a spatiotemporal manifold—that one can carve into hyperplanes of simultaneity depending on one's relative state of motion. There are many equally good ways to do the carving, and no fact of the matter about who is 'right'. Similarly, in the causal case, there are many equally good ways to carve the world into cause and effect and no fact of the matter about who is 'right'.

The version of temporal perspectivalism we are offering here is not a semantic thesis about temporal discourse. Rather, it is analogous to Ismael's frame-dependent version of causal perspectivalism. As in the causal case, there is some invariant fundamental structure in the world. We make a distinction in that structure between time and space based on our idiosyncratic epistemic constraints and limitations and our associated predictive concerns. Temporal perspectivalism thus takes as its starting point an agent's epistemic situation. Because the epistemic limitations of different agents constrain the form of the spatiotemporal information that agents will find it natural to gather, agents with varying epistemic limitations may very well treat different directions in the manifold as temporal ones. In so far as these agents disagree, their disagreement will be no fault (as in the causal case). Temporal perspectivalism is also deeply pragmatic: the temporal direction is individuated by useful predictions. We have left 'useful' deliberately vague in order to reflect the fact that agents with different interests, projects, and goals might seek to make different kinds of predictions, and thus may carve the world into time and space differently.

Temporal perspectivalism of this kind is capable of accommodating Callender's insights regarding well-posed Cauchy problems. Because these problems give us a great many useful predictions along a certain direction, we have reason to regard this direction as time. Temporal perspectivalism treats this as a reflection of our particular epistemic vantage point on the world. Given some particular perspective, however, temporal perspectivalism renders the temporal direction uniquely determined.

In this way, temporal perspectivalism exploits another feature of causal perspectivalism: the way that causal perspectivalism informs about the objectivity of the distinction between cause and effect.[9] This provides temporal perspectivalism with the tools to provide a neat analysis concerning the objectivity of the distinction between time and space. In the causal case, it is perfectly reasonable to refer to some causal relation as 'objective', but such a claim must be understood in a deflationary, epistemic sense: there is a fact of the matter about whether the relation is causal, but only given a separation of the system containing the causal relation into system and environment and a level of grain of description that enables a series of conditions that characterise causation to be met. Both the delineation into system and environment and the level of grain of description are agent-centric specifications and are a function of a particular epistemic perspective on the world.

The same line of reasoning translates to the temporal case: it is perfectly reasonable to refer to the distinction between time and space as 'objective', so long as this is understood in a

[9]The debate concerning the objectivity of the distinction between cause and effect is ongoing. See (Menzies and Price [1993]; Price [2007], [2017]; Woodward [2007], [2009], [2016]; Ismael [2016]).

deflationary, epistemic sense. It is certainly not the case that such a distinction is subjective in the sense that an agent has some sort of control over the way that the distinction can be made. The distinction is objective in so far as there is a fact of the matter as to what counts as time and space from a particular epistemic perspective. But one should not forget the crucial role played by the epistemic perspective—differing perspectives can plausibly permit differing distinctions. Borrowing the terminology from probability theory we can call the nature of this distinction 'objective epistemic'.[10] This flavour of objectivity sufficiently extracts the agent perspective out from the objects of scientific inquiry—since all agents engaged in a common scientific practice share the same perspective—for scientific discourse to be meaningful.

To see the basic idea behind temporal perspectivalism, it is useful to consider an example. Recall from §3 the toy case we introduced involving continuously varying colour properties to demonstrate the relationship between antecedent data and laws. Recall also our elaboration of the radical reading of Callender's informal argument, in which agents with predictive practices looking 'eastwards' in the manifold could define a time direction orthogonal to our own. We can build a toy model of temporal perspectivalism by combining these two examples. Consider the two-dimensional colour space, Figure 3, depicting saturation on the x-axis and hue on the y-axis. Looking along any line parallel to the y-axis, the points along such a line depict a sequence of continuously varying tones of colour at constant saturation. Looking along any line parallel to the x-axis, the points along such a line depict a sequence of continuously varying saturations of constant tone. Both sequences have a natural linear ordering and can be described as smooth functions over a natural set of variables—frequency in the former case and intensity (at pure wavelength) in the latter.

Imagine now two sets of agents in this world, $\mathcal{H}$ and $\mathcal{S}$, with differing epistemic constraints. The first set, $\mathcal{H}$, just happen to remember, and build records of, events (saturation–hue pairs) one side of a boundary parallel to the x-axis, and evolve towards and predict the array of events to the other side of the boundary—say, parametrising changes in hue. The second set, $\mathcal{S}$, by contrast, happen to remember, and build records of, events (hue–saturation pairs) one side of a boundary parallel to the y-axis, and evolve towards and predict the array of events to the other side of the boundary—say, parametrising changes in saturation. Due to the nature of their respective interests—a function of their *de facto* situations—both $\mathcal{H}$-scientists and $\mathcal{S}$-scientists develop functional laws that take natural variables from the boundaries characterising their respective epistemic limitations (demarcating what is known to them from what is unknown) and generate predictions concerning the events they will soon encounter. Thus, $\mathcal{H}$ develop PDEs that model the 'hue field', ϕ_h, as a function of the parameter t_h, $f_{\mathcal{H}}(\phi_h(t_h), \frac{\partial \phi_h}{\partial t_h})$, which happen to inform in the direction parallel to the y-axis, and $\mathcal{S}$ develop PDEs that model the 'saturation field', ϕ_s, as a function of the parameter t_s, $f_{\mathcal{S}}(\phi_s(t_s), \frac{\partial \phi_s}{\partial t_s})$, which happen to inform in the direction parallel to the x-axis. Provided these laws have the right sort of formal properties, each of $\mathcal{H}$

[10] Callender ([2017], pp. 153–4) gets close to this position concerning the objectivity of the distinction between time and space when he says 'given certain standards of strength, a given direction on $\mathcal{M}$ objectively is or isn't temporal'. Understanding this sort of 'conditional' objectivity under the guise of temporal perspectivalism permits us to see his distinction as a case of an objective epistemic feature of the world.

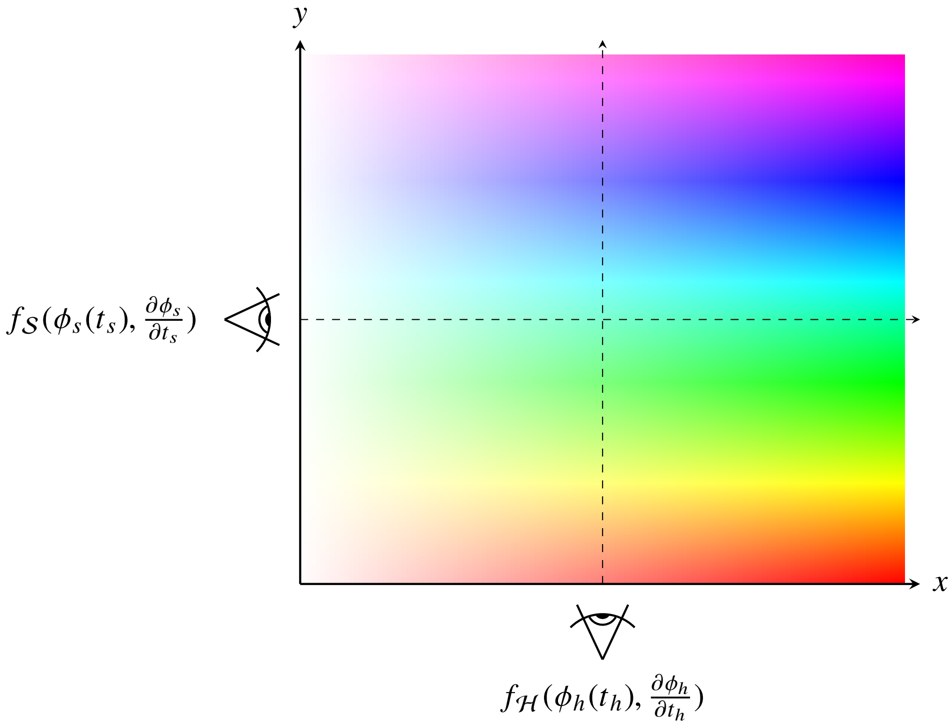

Figure 3: Two continuously varying quantities that can be used to develop differential equations from two different perspectives.

and $\mathcal{S}$ will conclude that the direction in the manifold in which their laws describe evolution is 'special'.

According to temporal perspectivalism, $\mathcal{H}$ and $\mathcal{S}$ thus carve the world into spatial and temporal dimensions in different ways. Each will focus on the direction in which their predictive practices are oriented, $\mathcal{H}$ treating directions parallel to the y-axis as temporal, and parallel to the x-axis as spatial, $\mathcal{S}$, treating directions parallel to the x-axis as temporal, and parallel to the y-axis as spatial. The predictive asymmetry between hue and saturation depends entirely on the agents, and the knowledge that the agents happen to have about the surrounding space (that is, their 'past'). Moreover, to $\mathcal{H}$ the 'natural kind' characterised by ϕ_s is going to look grotesquely nonlocal, and to $\mathcal{S}$ the 'natural kind' characterised by ϕ_h is going to look grotesquely nonlocal. Nonetheless, it is plausible that $\mathcal{H}$ and $\mathcal{S}$ devise physical theories that share a common mathematical core, or even the very same PDEs describing evolution, given that they are both engaged in a predictive practice that is sensitive to a particular kind of linearly ordered structure in the world.

In this broad way, the toy model is compatible with the radical reading of Callender's formal argument. The relationship between $\mathcal{H}$ and $\mathcal{S}$ is of the same sort as between us and Callender's Martians. The choice of natural kinds, in the current toy example, is the choice associated with choosing the 'hue field' or the 'saturation field' as the object of scientific inquiry, and reflects the choice of natural kinds at the heart of establishing the most efficient and informative algorithms according to Callender's best systems account; and what underlies that choice is the agent perspective. But while the toy model, and the broad picture that we are offering, is compatible with the radical reading of Callender's view, it is at odds with the conservative reading. The difference between time and space is a purely pragmatic difference on our view; it is not written into an existing metric signature that is then fed into a best systems account. There are, rather, just two different properties distributed differently that lead to two different ways of developing a theoretical account of the world.

7 Final Thoughts

Temporal perspectivalism turns the issue on its head. Rather than starting from a distinction between time and space and then constructing a scientific practice to track that pretheoretic distinction, we allow the distinction to fall out of whatever our most successful scientific practices happen to be. Notice that, in so far as our current scientific practice is concerned, it is business as usual. Science can continue to produce predictions in the temporal direction, and to even imbue that direction with interesting properties. Notice also that it is relatively straightforward to determine which is the temporal and which is the spatial direction. Roughly: see what a range of agents that share an epistemic perspective are doing when they manage to carry out successful predictive practices and go from there. Such a view avoids the difficulty of unbinding the features of time that make it the great informer from those features that are merely a function of our special epistemic vantage point in spacetime. On our view, these are two sides of the same coin: time is the great informer precisely because of our special epistemic perspective, which informs who we are as agents.

Acknowledgements

For useful discussion and comments we would like to thank Sean Gryb, Alistair Wilson, Ken Wharton, Magdalena Zych, and the audience at the University of Sydney Philosophy Department's Current Projects Seminar. We would also like to thank Sara Linton, Esq. of 'Entitled Consulting' for key advice in developing the title of this paper. SB acknowledges support from the University of Western Australia and the Australian Government through the Australian Research Council (DE180100414 and DP180100105). PWE acknowledges support from the University of Queensland, the Templeton World Charity Foundation (TWCF0064/AB38), and the Australian Government through the Australian Research Council (DE170100808).

Sam Baron
Department of Philosophy
University of Western Australia
Western Australia, Australia
samuel.baron@uwa.edu.au

Peter W. Evans
School of Historical and Philosophical Inquiry

University of Queensland
Queensland, Australia
p.evans@uq.edu.au

References

Callender, C. [2017]: *What Makes Time Special?*, Oxford: Oxford University Press.

Callender, C. and Cohen, J. [2010]: 'Special Sciences, Conspiracy and the Better Best System Account of Lawhood', *Erkenntnis*, **73**(3), pp. 427–47.

Cohen, J. and Callender, C. [2009]: 'A Better Best System Account of Lawhood', *Philosophical Studies*, **145**, pp. 1–34.

Geroch, R. [2011]: 'Faster Than Light?', in M. Plaue, A. D. Rendall and M. Scherfner (*eds*), *Advances in Lorentzian Geometry: Proceedings of the Lorentzian Geometry Conference in Berlin*, vol. 49, Providence, RI: American Mathematical Society, pp. 59–70.

Ismael, J. [2016]: 'How do causes depend on us? The many faces of perspectivalism', *Synthese*, **193**, pp. 245–67.

Lewis, D. [1983]: 'New work for a theory of universals', *Australasian Journal of Philosophy*, **61**(4), pp. 343–77.

Lewis, D. [1994]: 'Humean Supervenience Debugged', *Mind*, **103**, pp. 473–90.

Massimi, M. [2018]: 'A Perspectivalist Better Best System Account of Lawhood', in W. Ott and L. Patton (*eds*), *Laws of Nature*, New York: Oxford University Press, chap. 8, pp. 139–57.

Menzies, P. and Price, H. [1993]: 'Causation as a Secondary Quality', *The British Journal for the Philosophy of Science*, **44**(2), pp. 187–203.

Price, H. [2007]: 'Causal Perspectivalism', in H. Price and R. Corry (*eds*), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, Oxford: Oxford University Press, chap. 10, pp. 250–92.

Price, H. [2017]: 'Causation, Intervention and Agency: Woodward on Menzies and Price', in H. Beebee, C. Hitchcock and H. Price (*eds*), *Making a Difference: Essays on the Philosophy of Causation*, Oxford: Oxford University Press, chap. 5, pp. 73–98.

Rovelli, C. [2017]: 'Is Time's Arrow Perspectival?', in K. Chamcham, J. Silk, J. D. Barrow and S. Saunders (*eds*), *The Philosophy of Cosmology*, Cambridge: Cambridge University Press, chap. 14, pp. 285–96.

Schrenk, M. [2014]: 'Better Best Systems and the Issue of CP-Laws', *Erkenntnis*, **79**, pp. 1787–99.

Skow, B. [2007]: ‘What Makes Time Different from Space’, *Noûs*, **41**(2), pp. 227–52.

Weinstein, S. [2008]: ‘Many Times’, Foundational Questions Institute *The Nature of Time* Essay Contest, <fqxi.org/data/essay-contest-files/Weinstein_FQXI2.pdf>.

Woodward, J. [2003]: *Making Things Happen: A Theory of Causal Explanation*, New York: Oxford University Press.

Woodward, J. [2007]: ‘Causation with a Human Face’, in H. Price and R. Corry (*eds*), *Causation, Physics, and the Constitution of Reality: Russell’s Republic Revisited*, Oxford: Oxford University Press, chap. 4, pp. 66–105.

Woodward, J. [2009]: ‘Agency and Interventionist Theories’, in H. Beebee, C. Hitchcock and P. Menzies (*eds*), *The Oxford Handbook of Causation*, Oxford: Oxford University Press, chap. 11, pp. 234–62.

Woodward, J. [2016]: ‘Causation and Manipulability’, in E. N. Zalta (*ed.*), *Stanford Encyclopedia of Philosophy*, <plato.stanford.edu/archives/win2016/entries/causation-mani/>.