

Dynamic and stochastic systems as a framework for metaphysics and the philosophy of science

Christian List and Marcus Pivato*

Summer 2015; final version Spring 2019; forthcoming in *Synthese*

Abstract: Scientists often think of the world (or some part of it) as a dynamical system, a stochastic process, or a generalization of such a system. Prominent examples of systems are (i) the system of planets orbiting the sun or any other classical mechanical system, (ii) a hydrogen atom or any other quantum-mechanical system, and (iii) the earth's atmosphere or any other statistical mechanical system. We introduce a general and unified framework for describing such systems and show how it can be used to examine some familiar philosophical questions, including the following: how can we define nomological possibility, necessity, determinism, and indeterminism; what are symmetries and laws; what regularities must a system display to make scientific inference possible; how might principles of parsimony such as Occam's Razor help when we make such inferences; what is the role of space and time in a system; and might they be emergent features? Our framework is intended to serve as a toolbox for the formal analysis of systems that is applicable in several areas of philosophy.

1. Introduction

For both scientific and philosophical purposes, we often find it useful to think of the world (or some part of it that we are studying) as a system evolving over time: a dynamical system, a stochastic process, or a suitable generalization of such a system. In both science and philosophy, many theories represent the world (or the part they are concerned with) in terms of such systems, with various structures and properties. Metaphysical commitments often take the form of claims about the nature of those structures and properties: which of them are real and not just artefacts of our models, which are fundamental as opposed to derivative, and which are necessary as opposed to contingent.

In this paper, we introduce a general and unified framework for describing systems, based on the theory of dynamical systems and stochastic processes, and show how this framework can be used to examine and illuminate some familiar philosophical questions. Here are some examples:

- What does it mean for a system to be deterministic or indeterministic, and which features of the system, if any, determine which others?
- Does the present determine the future? Does it determine the past? What is the smallest set of facts encoding the system's entire history? Could there be non-temporal forms of determinism?

* C. List, Department of Philosophy, Logic, and Scientific Method, London School of Economics, U.K.; M. Pivato, THEMA, Université de Cergy-Pontoise, France. Part of this paper was written when Pivato was at the Department of Mathematics at Trent University, Canada. List's work was supported by a Leverhulme Major Research Fellowship, and Pivato's work was supported by an NSERC grant #262620-2008 and by Labex MME-DII (ANR11-LBX-0023-01). We are very grateful to George Musser, Bryan Roberts, Timothy Williamson, and the anonymous reviewers of this paper for helpful comments.

- How can we define nomological possibility and necessity for a system?
- What are the laws governing a particular system, and is there a distinction between “laws” and “brute necessities”? How do laws depend on symmetries?
- What structure must a system have in order to permit generalizations from local observations to global regularities?
- How might we use principles of parsimony such as Occam’s Razor when we make such generalizations? And can we formulate a version of Occam’s Razor in terms of symmetries?
- What is the role of space and time in a system? What is the relationship between the geometry of space and time and the system’s behaviour?
- Is this spatiotemporal geometry exogenous, or is it determined by the dynamics? In other words, are space and time more fundamental than the system’s dynamics, or the other way around? Might space and time be “emergent”?
- How should we individuate systems? Should two structurally indistinguishable systems count as “the same”, or might they count as different?

For each of these questions, our framework allows us to identify in clear and precise terms what is at stake. We illustrate the generality of the framework by sketching how it can accommodate, schematically, the systems described by some standard physical theories, such as classical mechanics, electrodynamics, quantum mechanics, and special and general relativity. In principle, our framework can also be used to describe many systems studied in the special sciences, such as biological, social, and economic systems, though we do not have the space to develop these applications here. We make a few remarks about special-science systems at the end of the paper and hope that our framework will serve as a basis for future work in some of those areas.¹

The paper is structured as follows. We discuss three classes of systems, in increasing order of generality. We call the first *temporally evolving systems* (Section 2), the second *spatially extended systems* (Section 3), and the third *amorphous systems* (Section 4). We offer a conceptual toolbox for describing and analysing each class of systems, covering notions such as states and histories, determinism and indeterminism, nomological possibility and necessity, modal and probabilistic properties, symmetries and laws, ergodicity and its significance in making scientific inference possible, Occam’s Razor, and the role of time and/or space. We first explain all of these notions in the context of the simplest class of systems (in Section 2) and then generalize from there (in Sections 3 and 4). The paper also includes some more technical appendices, on factor systems (relevant to the analysis of systems at different levels of abstraction), on partial and local symmetries (relevant to “local” laws and the analysis of systems with special initial or boundary conditions), on criteria of parsimony in relation to

¹ Ideas from the theory of dynamical systems and stochastic processes have already begun to be applied to fields such as economics, biology, and cognitive science. See, e.g., van Gelder (1995), Auyang (1998), Juarrero (1999), and Silberstein and Chemero (2012).

which symmetries to postulate (relevant to Occam’s Razor), and on the definition of spatial distance in quantum-mechanical systems (which raises special challenges).

Although the paper presupposes a willingness to engage with technical material – and a basic familiarity with science will be helpful – our goal is to keep the exposition as simple and self-contained as possible. Our intended contribution is twofold: methodological and substantive. On the methodological side, we aim to offer a unified and yet accessible framework for the philosophical analysis of many of the systems studied in the sciences. While the basic ideas originate from the theory of dynamical systems and stochastic processes in mathematics and physics, and partially overlapping formalisms can be found in earlier works (e.g., by Earman 1986, van Fraassen 1989, Berkovitz, Frigg, and Kronz 2006, 2011, Werndl 2009a, 2009b, Bishop 2011, Butterfield 2012, and Yoshimi 2012), the key ideas remain underappreciated in philosophy, and to our knowledge, an equally unified (and, we think, accessible) framework is not yet available in the philosophical literature.

On the substantive side, we aim to offer a number of novel insights, for example concerning (i) the nature of nomological possibility and necessity in a system and the definition of determinism and indeterminism, (ii) the role of symmetries in distinguishing between “laws” and “brute necessities” in a system, (iii) the significance of symmetries and ergodicity as prerequisites for scientific inference, (iv) the relationship between Occam’s Razor and the symmetries of a system, and (v) the possibility that the topology and geometry of space and time may be emergent properties resulting from a system’s correlation structure. These, we hope, will be useful substantive contributions, over and above the paper’s unificatory contribution.

2. Temporally evolving systems

2.1 Basic definitions

We begin with the simplest class of systems whose states evolve over time.² To define a system in this class, we need to specify what *time* is, what the system’s *states* are, and how these states may *evolve* over time. *Time* is represented by a set of points T that is linearly ordered; we write $<$ for the “before” relation. The *state* of the system at each point in time is given by an element of some *state space* X . For the moment, we make no assumptions about the internal structure of the states in X ; they are uninterpreted primitives. A *history* of the system, capturing “state evolution”, is a path through the state space, represented by a function h from T into X . For each time t in T , $h(t)$ is the system’s state at time t . In a physical system, each state might be a completely specified microphysical state in which the system could be at a particular point in time, and histories would be possible trajectories of the system through its state space over time.

We write Ω to denote the set of all histories deemed *possible*. Histories play the role of possible worlds. Thus, the structure of Ω reflects the notion of possibility we wish to capture. If we are interested in logical possibility, then Ω is simply the set of all logically possible functions from T into X , which we call \mathcal{F} . If we are interested in some form of nomological possibility, such as physical possibility, Ω will often be a proper subset of \mathcal{F} . Our intended interpretation of possibility throughout this paper is the

² We build on the formalism in List (2014) and List and Pivato (2015).

nomological one, since we want to distinguish between histories that are permitted by the laws governing our system and histories that are not.

Subsets of Ω are called *events*. We can apply logical operations to events. The *conjunction* of two events E and E' is given by their intersection $E \cap E'$. Their *disjunction* is given by their union $E \cup E'$. The *negation* of an event E is given by its *complement* $\sim E = \Omega \setminus E$. Later we introduce *possibility* and *necessity* operators.

To complete the definition of a temporally evolving system, we must define probabilities on Ω . Formally, we introduce a *conditional probability structure*.³ This is a family of conditional probability functions $\{Pr_E\}_{E \subseteq \Omega}$, consisting of one Pr_E for each event E in Ω , where Pr_E assigns to any event in Ω the conditional probability of that event, given E .⁴ The family must satisfy certain consistency conditions, such as compatibility with Bayesian conditionalization.⁵ Now, a *temporally evolving system* is the pair consisting of the set Ω of possible histories and the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$.

For example, in a weather system, X would be the set of all possible weather states and Ω the set of all possible weather histories. For each particular weather event E , say a hot temperature on Monday, the function Pr_E assigns to every weather event D , say a thunderstorm on Tuesday, the conditional probability of its occurrence, given E .

In principle, the probability structure admits two interpretations. Under an *objectivist* interpretation, it is a feature of the system itself and thus represents *objective chance* (see, e.g., Lewis 1986, Schaffer 2007, and List and Pivato 2015). Of course, objective chance could be *degenerate*, i.e., restricted to the extremal values 0 or 1. Degenerate objective chance is a much-discussed feature of deterministic systems; we return to this point later. Under a *subjectivist* interpretation, the probability structure is not a feature of the system itself, but represents an observer's beliefs about the system, as in subjective Bayesianism (e.g., de Finetti 1972). The most natural way to read this paper is to assume the objectivist interpretation, though our formalism itself is neutral.

Familiar examples of temporally evolving systems are the system of planets orbiting the sun or any other classical mechanical system, a hydrogen atom or any other quantum-mechanical system, the earth's climate system or any other statistical mechanical system, and (arguably) the global economy or some other closed macro-economic system. Generally, any classical *dynamical system* is a special case of a temporally evolving system, as is any *stochastic process* under the standard definition.⁶

³ Conditional probability structures have previously been considered by several authors, e.g., Popper (1968), Renyi (1955), van Fraassen (1976), as reviewed in Halpern (2010).

⁴ Each Pr_E is defined on a suitable σ -algebra $\mathcal{A}(\Omega)$ on Ω ; we set the technicalities aside. For any $E \neq \emptyset$, Pr_E has the standard properties of a probability measure. But, for technical reasons, $Pr_{\emptyset}(D) = 1$ for all D .

⁵ For any subsets $C \subseteq D \subseteq E \subseteq \Omega$, we have $Pr_E(C) = Pr_E(D) \times Pr_D(C)$. Also, $Pr_D(E) = 1$ for all $D \subseteq E \subseteq \Omega$.

⁶ A classical *dynamical system* consists of a set X (the *state space*) and a function ϕ from X into itself that determines how the state changes over time (the *dynamics*). Let $T = \{0, 1, 2, 3, \dots\}$. Given any state x in X (the *initial conditions*), the *orbit* of x is the history h defined by $h(0) = x$, $h(1) = \phi(x)$, $h(2) = \phi(\phi(x))$, and so on. Let Ω be the set of all orbits determined by (X, ϕ) in this way. Let $\{Pr'_E\}_{E \subseteq X}$ be any conditional probability structure on X . For any events E and D in Ω , we define $Pr_E(D) = Pr'_E(D')$, where E' and D' are the sets of all states in X whose orbits lie in E and D , respectively. Then $\{Pr_E\}_{E \subseteq \Omega}$ is a conditional probability structure on Ω . Thus, Ω and $\{Pr_E\}_{E \subseteq \Omega}$ together form a temporally evolving system. However, not every temporally evolving system arises in this way. While classical dynamical systems are

For theoretical simplicity, we focus on *closed* systems, which are not subject to any external perturbations. However, one could also represent *open* systems in our framework, by encoding any external perturbations as additional sources of randomness in the system’s conditional probability structure (“random forcings”).⁷

2.2 Determinism and indeterminism

Conventionally, a system is called *deterministic* if, in that system, the past always determines the future. Formally, for any history h and any point in time t , let h_t be the *initial segment* of that history up to t . This is the function h restricted to the points in time up to t . History h is *deterministic* if, at any time t in T , the initial segment h_t admits only one possible continuation in Ω , where a *continuation* of h_t is a history h' such that $h'_t = h_t$. History h is *indeterministic* if, for some time t , h_t has more than one possible continuation in Ω . The system as a whole is called *deterministic* if all histories in Ω are deterministic, and *indeterministic* if some histories in Ω are indeterministic.⁸

For example, classical mechanical systems, such as the solar system on the Newtonian picture, are deterministic. By contrast, quantum-mechanical systems, such as a decaying uranium atom, are indeterministic (assuming no hidden variables). If the wave function, which encodes the state of the quantum system, collapses at time t , the initial segment h_t of the system’s history h can admit multiple continuations.

Indeterministic systems allow non-degenerate chance as we move along a given history, while deterministic systems do not.⁹ To see this, note that the chance of any event E in history h at time t is the conditional probability of E , given that the initial segment h_t has occurred. Since the event that the initial segment h_t has occurred is given by the set

deterministic, temporally evolving systems also subsume stochastic processes. Formally, a *stochastic process* is a temporally indexed collection of random variables $\{X_t : t \in T\}$ (with T as before) on some probability space $(\Omega, \mathcal{A}(\Omega), Pr)$, where Ω is some underlying (abstract) set of possible worlds, $\mathcal{A}(\Omega)$ is a σ -algebra on Ω , and Pr is a probability measure on $\mathcal{A}(\Omega)$. For each time t , we can think of the random variable X_t as expressing the state of the stochastic process at t . To see that this gives rise to a temporally evolving system, note that each world ω in Ω induces a history h in our sense, where, for each t , $h(t)$ is the realization of X_t at world ω . Most of the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ can then be derived from Pr via Bayes rule, but $\{Pr_E\}_{E \subseteq \Omega}$ may contain some additional information not encoded in Pr (namely, conditional probabilities arising from zero-probability events). In Sections 3 and 4, we extend our framework to even more general classes of systems.

⁷ The use of such random forcings does not imply that certain features of the world are genuinely random. Instead, the “randomness” of such forcings is best understood *epistemically* – as a shortcut for an explicit and detailed description of the part of the world which lies outside the model. (This is true whether we adopt an *objectivist* or *subjectivist* interpretation of the probability structure overall.)

⁸ On these definitions, see also List (2014) and List and Pivato (2015). Related definitions of determinism (broadly, in terms of a history’s unique extendibility based on an initial state or segment) can be found, for instance, in van Fraassen (1989, section 10.4) and Butterfield (2012). See also the classic overview in Earman (1986) and the discussion of varieties of determinism in Sobel (1998). Earman shares our focus on local-to-global determination, as noted below. While Sobel discusses more than 90 variants of determinism, he frames the question differently than us. First of all, his focus, unlike ours in this paper, is on the free-will debate. Secondly, for us, the central question is (roughly): does the complete specification of the world (history) for a particular space-time region (e.g., a history’s initial segment) determine a complete specification of that world (history) for other space-time regions? By contrast, for Sobel, the central question is (roughly): does an event that occurs at a particular space-time location have a cause at some antecedent space-time location? Since we do not explicitly discuss the topic of causation here, Sobel’s analysis and ours are not immediately inter-translatable.

⁹ For discussion, see, e.g., Schaffer (2007) and List and Pivato (2015).

of all continuations of h_t – call this set $[h_t]$ – the probability in question is $Pr_{[h_t]}(E)$. If history h is deterministic, the entire conditional probability function $Pr_{[h_t]}$ is *degenerate*, i.e., it assigns probability 0 or 1 to every event E . This is because the initial segment h_t has only one continuation, namely h itself, and so the specified event $[h_t]$ contains only a single history, h . Then $Pr_{[h_t]}(E)$ is 1 if h belongs to E and 0 otherwise. In contrast, if history h is indeterministic, $Pr_{[h_t]}$ may be *non-degenerate*, assigning probabilities strictly between 0 and 1 to some events E . This is because $[h_t]$ need not be singleton here, and so $Pr_{[h_t]}$ is less constrained. (For the moment, we set aside phenomena such as “higher-level” indeterminism and chance, as discussed in List and Pivato 2015. We briefly consider such phenomena at the end of this paper.)

Our framework also allows us to formulate some more general, less familiar notions of determinism. For *any* subset T' of T – not just the set of time points up to a particular time t – we can ask whether the restriction of a given history to the points in T' uniquely determines the rest of that history. Let $h_{T'}$ denote the restriction of the function h to T' . Our question, then, is whether $h_{T'}$ has a unique extension to all of T in Ω , where an *extension* of $h_{T'}$ is a history h' such that $h'_{T'} = h_{T'}$. If there is a unique extension, history h may be called *T'-deterministic*.¹⁰

We might ask, for instance, whether the entire history of a system, both past and future, is determined by its present state alone. Similarly, we might ask whether, given the states of the system at two points in time, there is a unique history connecting them. So, one can in principle consider not only the familiar idea of “past-to-future” determinism, but also other forms of “local-to-global” determinism. In Section 3.2, we develop these ideas further and consider, among other things, spatial rather than temporal forms of determinism as well as locally restricted ones.

2.3 Nomological possibility and necessity

We can explicitly define the notions of nomological necessity and possibility in our framework.¹¹ Intuitively, an event E is nomologically possible in history h at time t if the initial segment of that history up to t admits *at least one* continuation in Ω that lies in E ; and E is nomologically necessary in h at t if *every* continuation of the history’s initial segment up to t lies in E .

More formally, we say that one history, h' , is *accessible* from another, h , at time t if the initial segments of h and h' up to time t coincide, i.e., $h_t = h'_t$. We then write $hR_t h'$. The binary relation R_t on possible histories is in fact an equivalence relation (reflexive, symmetric, and transitive). Now, an event $E \subseteq \Omega$ is *nomologically possible* in history h at time t if *some* history h' in Ω that is accessible from h at t is contained in E . Similarly, an event $E \subseteq \Omega$ is *nomologically necessary* in history h at time t if *every* history h' in Ω that is accessible from h at t is contained in E .

We can thus define two modal operators, \blacklozenge_t and \blacksquare_t , to represent possibility and necessity at time t . We define each of them as a mapping from events to events. For any event $E \subseteq \Omega$,

¹⁰ For related ideas, see also van Fraassen (1989, section 10.4).

¹¹ We here employ, and subsequently generalize, a construction from List (2014).

$$\blacklozenge_t E = \{h \in \Omega : \text{for some } h' \in \Omega \text{ with } hR_t h', \text{ we have } h' \in E\},$$

$$\blacksquare_t E = \{h \in \Omega : \text{for all } h' \in \Omega \text{ with } hR_t h', \text{ we have } h' \in E\}.$$

So, $\blacklozenge_t E$ is the set of all histories in which E is possible at time t , and $\blacksquare_t E$ is the set of all histories in which E is necessary at time t . Accordingly, we say that $\blacklozenge_t E$ holds in history h if h is an element of $\blacklozenge_t E$, and $\blacksquare_t E$ holds in h if h is an element of $\blacksquare_t E$. As one would expect, the two modal operators are *duals* of each other: for any event $E \subseteq \Omega$, we have $\blacksquare_t E = \sim \blacklozenge_t \sim E$ and $\blacklozenge_t E = \sim \blacksquare_t \sim E$.

Two remarks are due. First, although we have here defined *nomological* possibility and necessity, we can analogously define *logical* possibility and necessity. To do this, we must simply replace every occurrence of the set Ω of nomologically possible histories in our definitions with the set \mathcal{H} of logically possible histories. Second, by defining the operators \blacklozenge_t and \blacksquare_t as functions from events to events, we have adopted a *semantic* definition of these modal notions. However, one could also describe them *syntactically*, by introducing an explicit modal logic. For each point in time t , the logic corresponding to the operators \blacklozenge_t and \blacksquare_t would then be an instance of a standard S5 modal logic (on S5, see, e.g., Priest 2001).

Our analysis shows how nomological possibility and necessity depend on the dynamics of the system, as evident from the time-indexed nature of the relevant modal operators. In particular, as time progresses, the notion of possibility becomes more demanding: fewer events remain possible at each time. And the notion of necessity becomes less demanding: more events become necessary at each time, for instance due to having been “settled” in the past. Formally, for any t and t' in T with $t < t'$ and any event $E \subseteq \Omega$,

$$\text{if } \blacklozenge_{t'} E \text{ then } \blacklozenge_t E,$$

$$\text{if } \blacksquare_t E \text{ then } \blacksquare_{t'} E.$$

Furthermore, in a deterministic system, for every event E and any time t , we have $\blacklozenge_t E = \blacksquare_t E$. In other words, an event is possible in any history h at time t if and only if it is necessary in h at t . In an indeterministic system, by contrast, necessity and possibility come apart.

Just as we previously discussed different notions of determinism – not just “past to future” but also “local to global” – so we can generalize the notions of possibility and necessity in a similar way. Let us say that one history, h' , is *accessible* from another, h , relative to a set T' of time points, if the restrictions of h and h' to T' coincide, i.e., $h'_{T'} = h_{T'}$. We then write $hR_{T'} h'$. Accessibility at time t is the special case where T' is the set of points in time up to time t . We can define nomological possibility and necessity relative to T' as follows. For any event $E \subseteq \Omega$,

$$\blacklozenge_{T'} E = \{h \in \Omega : \text{for some } h' \in \Omega \text{ with } hR_{T'} h', \text{ we have } h' \in E\},$$

$$\blacksquare_{T'} E = \{h \in \Omega : \text{for all } h' \in \Omega \text{ with } hR_{T'} h', \text{ we have } h' \in E\}.$$

Although these modal notions are much less familiar than the standard ones (possibility and necessity at time t), they are useful for some purposes. In particular, they allow us

to express the fact that the states of a system during a particular period of time, $T' \subseteq T$, render some events E possible or necessary.

Finally, our definitions of possibility and necessity relative to some general subset T' of T allow us to define completely “atemporal” notions of possibility and necessity. If we take T' to be the empty set, then the accessibility relation $R_{T'}$ becomes the universal relation, under which every history is related to every other. An event E is possible in this atemporal sense (i.e., $\blacklozenge_{\emptyset}E$) if and only if E is a non-empty subset of Ω , and it is necessary in this atemporal sense (i.e., $\blacksquare_{\emptyset}E$) if E coincides with all of Ω . These notions might be viewed as possibility and necessity from the perspective of some observer who has no temporal or historical location within the system and looks at it from the outside.

2.4 Modal and probabilistic properties

Ultimately, all modal properties of a temporally evolving system are encoded by the set Ω of nomologically possible histories, and all probabilistic properties are encoded by the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$. This raises the question: which, if any, of these properties qualify as “laws” of the system, and what does this mean?

One possible view is that:

- any property that is satisfied by *all* histories in Ω counts as a law of the system, specifically a “modal law”; and
- any property of the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ counts as a law of the system, specifically a “probabilistic law”.

Indeed, since the system is fully specified by Ω and $\{Pr_E\}_{E \subseteq \Omega}$, one might interpret anything that is globally true of its possible histories or its probability structure as a law of that system. A view along these lines is expressed in a classic paper by Sellars (1948, p. 309): “A natural law is a universal proposition, implicative in form, which holds of all histories of a family of possible histories; as such it is distinguished from ‘accidental’ formal implications which hold of one or more possible histories of the family, but do not hold of all.” So, the notions “being a law” and “being nomologically necessary” essentially coincide.

Against this view, however, we want to argue that even among nomologically necessary properties of a system – those that are not contingent on particular histories – one can distinguish between “laws” on the one hand and “brute necessities”, which are not law-like, on the other. Laws, we suggest, have a testable and generalizable character which brute necessities lack. To explain this, we introduce two preliminary notions, *properties of histories* and *probabilistic properties*, and then provide a criterion for identifying which of them qualify as laws.

A *property of histories*, P , is a binary feature that a history may or may not have. Formally, it can be associated with some subset, denoted $[P]$, of the set \mathcal{H} of all logically possible histories. A history h satisfies P if h belongs to $[P]$. We call $[P]$ the *extension* of P . A property satisfied by *every* history in Ω can be called *nomologically necessary* for the system. Newton’s three laws of motion are examples of such properties in the case of a classical mechanical system.

A *probabilistic property*, \mathcal{P} , is a binary feature that a conditional probability structure may or may not have. Formally, it is associated with a subset, denoted $[\mathcal{P}]$, of the set Π of all logically possible conditional probability structures on Ω . A conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ satisfies \mathcal{P} if it belongs to $[\mathcal{P}]$. We call $[\mathcal{P}]$ the *extension* of \mathcal{P} . An example of a probabilistic property is the one that says: “The unconditional probability of event F is $1/2$.” Its extension is the set of all conditional probability structures $\{Pr_E\}_{E \subseteq \Omega}$ for which $Pr_\Omega(F) = 1/2$. Another example is the second law of thermodynamics. This is a probabilistic property that is satisfied by the conditional probability structure of a statistical mechanical system.

Our goal is to distinguish between those properties that qualify as “laws” of the system and those that do not. We capture that distinction through the notion of *symmetries*. Informally, a *symmetry* is a transformation that acts on either the state space X or the set of time points T , or both, and which can capture certain admissible changes in perspective on the system. *Laws*, we suggest, are those nomologically necessary properties which are invariant under symmetries and which therefore hold across changes in perspective. We now make this formally precise.

2.5 Symmetries

We first consider symmetries acting on the state space; we then turn to symmetries acting on time; and we finally consider more general symmetries. To introduce *state symmetries*, we begin with some preliminary definitions. Let ϕ be any function from X into itself, i.e., a transformation on the state space. We use this transformation to define a function from histories to other histories. For reasons that will become clear, we do not restrict the function to nomologically possible histories, but define it as a function on \mathcal{H} , the set of all logically possible histories. Specifically, for any history h in \mathcal{H} , we define the transformed history

$$\phi(h) = h', \text{ where, for all } t \text{ in } T, h'(t) = \phi(h(t)).$$

For example, if $X = \{a, b, c, d, \dots, z\}$, the function ϕ might shift every letter in the alphabet one place to the right, i.e., a to b , b to c , and so on, and z back to a . If we represent histories as sequences of elements in X , interpreted as the system’s states at times 1, 2, 3, ..., then applying ϕ to the history $h = (b, a, c, f, z, \dots)$ yields the history $h' = (c, b, d, g, a, \dots)$. For convenience, we use the letter ϕ to denote both the original function on the state space and the induced function on the set \mathcal{H} of histories. Note that since the set of nomologically possible histories may be a proper subset of the set of all logically possible histories, the image of a history in Ω need not be in Ω .

To define what it means for ϕ to be a symmetry, we need one further preliminary definition. For any collection of histories E in \mathcal{H} , the *inverse image* of E under ϕ is the set of all histories h in \mathcal{H} such that $\phi(h)$ lies in E .¹² For example, if E is the set of all histories whose state at time 3 is c , and ϕ is the letter-shifting transformation, then the inverse image of E under ϕ is the set of all histories whose state at time 3 is b . Now, the function ϕ is a *symmetry* of our system if

¹² Formally, we write $\phi^{-1}(E) = \{h \text{ in } \mathcal{H}: \phi(h) \text{ is in } E\}$. The use of this notation for inverse images of sets does not imply that the function ϕ is invertible. Moreover, the inverse image of an event E could be empty, namely if none of the histories in E can be “reached” as transformations of other histories.

- $\phi(\Omega) = \Omega$, i.e., (i) $\phi(h)$ is in Ω , for all h in Ω , and (ii) for any h in Ω , there is some h' in Ω such that $\phi(h') = h$; and
- for any events E and D in Ω , if E' and D' are the inverse images of E and D under ϕ , then $Pr_{E'}(D') = Pr_E(D)$.¹³

Intuitively, a symmetry is a transformation that preserves the system's modal and probabilistic structure. In our example, where $X = \{a, b, c, d, \dots, z\}$ and ϕ is the letter-shifting function, the first part of this definition implies that the set of nomologically possible histories is preserved under shifting of letters. For instance, if (b, a, c, f, z, \dots) is a nomologically possible history, then so is (c, b, d, g, a, \dots) .¹⁴ To illustrate the second part, let E be the set of all histories in Ω whose state at time 3 is c , and let D be the set of all histories in Ω whose state at time 5 is a (so that E' is a suitable set of histories whose state at time 3 is b , and D' is a suitable set of histories whose state at time 5 is z).¹⁵ The conditional probability that the state of a history at time 5 is a , given that at time 3 it is c , must then equal the conditional probability that the state at time 5 is z , given that at time 3 it is b .

Obviously, not all state transformations are symmetries. Whether there are any non-trivial state symmetries depends on the system in question, i.e., it depends on Ω and $\{Pr_E\}_{E \subseteq \Omega}$. In classical mechanical systems, state symmetries include spatial translations, which shift everything in a certain direction by a certain distance, rotations and reflections, and permutations of particles with equal mass. Those transformations preserve the modal and probabilistic structure of the relevant systems.

Similarly, we can define *time symmetries*. Again, we begin with some preliminary definitions. Let ψ be any function on T , i.e., a transformation on time. For any history h , we define the transformed history

$$\psi(h) = h', \text{ where, for all } t \text{ in } T, h'(t) = h(\psi(t)).^{16}$$

For example, if $T = \{1, 2, 3, \dots\}$, the function ψ might be given by $\psi(t) = t + 5$ for all t in T . It maps the history (x_1, x_2, x_3, \dots) (a sequence of states across time) to the history (x_6, x_7, x_8, \dots) . As in the case of state symmetries, ψ induces a function from the set \mathcal{H} to itself. In analogy to the earlier definition, ψ is a *symmetry* if

- $\psi(\Omega) = \Omega$;
- for any events E and D in Ω , if E' and D' are the inverse images of E and D under ψ , then $Pr_{E'}(D') = Pr_E(D)$.

¹³ Strictly speaking, $\{Pr_E\}_{E \subseteq \Omega}$ is only defined for subsets of Ω . But we can extend it as follows: for any $D, E \subseteq \mathcal{H}$, if $E \cap \Omega \neq \emptyset$, let $Pr_E(D) = Pr_{E \cap \Omega}(D \cap \Omega)$. Cases where $E \cap \Omega = \emptyset$ are not relevant here, because of the first bullet point in the definition of *symmetry*.

¹⁴ And further, there is a nomologically possible history, namely (a, z, b, e, y, \dots) , such that applying the letter-shifting function to it yields the history (b, a, c, f, z, \dots) .

¹⁵ Formally, E' consists of all the histories h in \mathcal{H} such that $\phi(h)$ is in Ω and $h(3) = b$. Similarly, D' consists of all the histories h in \mathcal{H} such that $\phi(h)$ is in Ω and $h(5) = z$.

¹⁶ Typically, we require ψ to be *order-preserving*, i.e., for all t and t' in T , if $t < t'$, then $\psi(t) < \psi(t')$. For example, if $T = \{1, 2, 3, \dots\}$ with the standard ordering, the functions $\psi(t) = t + 5$ and $\psi(t) = 5t$ are order-preserving. But we do not build this requirement into our definition of a time symmetry. Note that some time symmetries, such as time reversals in classical physical systems, are not order-preserving.

In our example, where $T = \{1, 2, 3, \dots\}$ and $\psi(t) = t + 5$, the first part of this definition implies that if $h = (x_1, x_2, x_3, \dots)$ is a nomologically possible history of the system, then so is $h' = (x_6, x_7, x_8, \dots)$.¹⁷ To illustrate the second part, suppose that E is the set of all histories in Ω whose state at time 3 is c , while D is the set of all histories in Ω whose state at time 4 is a (so that E' is a suitable set of histories whose state at time 8 is c , while D' is a suitable set of histories whose state at time 9 is a). The conditional probability that the state at time 9 is a , given that at time 8 it was c , must then equal the conditional probability that the state at time 4 is a , given that at time 3 it was c .¹⁸

Just as not all state transformations are symmetries, so not all time transformations are symmetries either. In most classical physical systems, time symmetries include *time translations*, such as $\psi(t) = t + 5$, but exclude *non-linear transformations*, such as $\psi(t) = t^2$. In systems where the state does not encode explicitly “kinetic” properties (such as momentum), *simple time reversals*, such as $\psi(t) = -t$, can also be time symmetries. For example, the partial differential equations describing wave propagation in an ideal medium are invariant under simple time reversals. But many other systems, such as thermodynamic ones and diffusion processes, do not admit such simple time reversals.

More general symmetries include composite functions resulting from the combination of transformations of X and transformations of T . These are best understood as functions acting on the set \mathcal{H} of logically possible histories directly, with the properties introduced above. A familiar example in classical mechanical systems is a *time reversal*, which involves *both* a negation of the time index *and* a negation of all momentum vectors in the system (not to be confused with a *simple time reversal*, as mentioned earlier).¹⁹ A more complex example is a *Galilean transformation*, which adds a constant vector to the momentum vectors of all particles and also a time-varying sequence of spatial shifts to the particle positions, thereby converting the system to a different inertial reference frame. See footnote 48 below for details.

We can think of symmetries – whether they act on the state space, on time, or on both – as transformations that encode admissible changes in perspective on a system, insofar as they preserve the system’s modal and probabilistic structure. We write Ψ to denote the set of *all* symmetries of our temporally evolving system. This set has the algebraic

¹⁷ And further, there is some nomologically possible history $h'' = (u, v, w, y, z, x_1, x_2, x_3, \dots)$ such that shifting the system’s state in h'' five time periods into the future yields the history h . (Here, the exact values of u, v, w, y, z are irrelevant, as long as h'' is nomologically possible.)

¹⁸ Note that classical dynamical systems have a particularly rich set of time symmetries. Let (X, ϕ) be a dynamical system, as defined in footnote 6. Suppose the function ϕ (which maps from X into itself) is *surjective*, i.e., for all x in X , there is some y in X such that $\phi(y) = x$. Then the set Ω of orbits is invariant under all time-shifts. Let $\{Pr'_E\}_{E \subseteq X}$ be a conditional probability structure on X , and let $\{Pr_E\}_{E \subseteq \Omega}$ be the conditional probability structure it induces on Ω . Suppose that $\{Pr'_E\}_{E \subseteq X}$ is ϕ -*invariant*, i.e., for any subsets E and D of X , if $E' = \phi^{-1}(E)$ and $D' = \phi^{-1}(D)$, then $Pr'_{E'}(D') = Pr'_E(D)$. Then *every* time shift is a temporal symmetry of the resulting temporally evolving system. The study of dynamical systems equipped with invariant probability measures is the purview of *ergodic theory*.

¹⁹ As Roberts (2013) has argued, in general, a time reversal must not only map the time coordinate t to $-t$, but also apply an appropriate transformation to the system’s state at each point in time (in special cases, this could be the identity transformation, as in the case of a simple time reversal). Generally, we can think of the state of the system at each point in time as encoding not only some “static” properties (such as each particle’s position), but also some “kinetic” properties (such as each particle’s momentum). While static properties are preserved under time reversals, kinetic properties are not generally preserved. Similarly, in quantum mechanics, time reversals involve not only a reversal of the time coordinate, but also taking the conjugate of the wave function’s values.

structure of a *monoid*. Formally, a *monoid* is a set of transformations (here of \mathcal{H}) which (i) contains the *identity transformation* (mapping every history to itself) and (ii) is *closed under composition* (i.e., for any two transformations in the set, the transformation obtained by applying first one of the two transformations and then the other is also in the set). An example of a symmetry monoid is the set of all rotations of a classical mechanical system around a fixed axis: the identity transformation obviously belongs to this set, being a rotation by an angle of zero, and the composition of any two rotations is still a rotation.²⁰

2.6 Laws and their significance

As anticipated, the laws of a system are those nomologically necessary properties within it that are invariant under symmetries. This, we show, makes laws open to testing and generalization. Laws, one might say, have a “scrutable” and “projectable” character. The close relationship between symmetries and laws has been recognized before.²¹ For instance, Wigner (1967) takes symmetries to be “a prerequisite for the very possibility of discovering the laws of nature” (as Brading and Castellani 2013 put it; see also French 2014). And van Fraassen (1989), in his classic study of symmetries in science, considers defining laws as “facts which are invariant under symmetries”, though ultimately does not endorse that definition. But none of the existing accounts clarifies the relationship between laws and symmetries in a way that we consider fully satisfactory.²² We develop this relationship in detail in the case of modal laws. We subsequently consider probabilistic laws too, but, due to space constraints, discuss those more briefly.

To define the notion of a *modal law*, consider a property of histories, P . Recall that P is *nomologically necessary* for the given system if its extension, $[P]$, includes all histories in Ω . For any symmetry ψ , we say that P is *invariant* under ψ if the set $[P]$ is

²⁰ To see that the set Ψ of *all* symmetries of a temporally evolving system forms a monoid, note that (i) the identity transformation is trivially a symmetry, and (ii) if two transformations each qualify as symmetries, by preserving the modal and probabilistic structure of the system, then so does their composition. If all symmetries are invertible, then Ψ becomes a *group*, but we need not assume this.

²¹ See, e.g., Wigner (1967), van Fraassen (1989, Part III), Mainzer (1996, Section 5.3), Brading and Castellani (2003, 2013), and Baker (2010).

²² For example, Wigner (1967) recognizes the centrality of laws to scientific discovery, but does not define laws as facts that are invariant under symmetry. Instead, he seems to regard symmetry invariance principles as “second-order laws”, which relate laws to other laws (or which relate a law to itself). So, for him, it appears that laws establish relationships between events, and symmetries establish relationships between laws. Van Fraassen’s analysis (1989) is in many ways a precursor to ours (see, e.g., *ibid.*, Section 11.2). But surprisingly he seems to reject the definition of laws as “facts invariant under symmetries” (*ibid.*, Section 11.5). His proposed counterexamples (inspired by a passage from Weyl 1952) involve symmetry-invariant facts which are not nomologically necessary (such as the number of planets in the solar system). In contrast, we define laws as *nomologically necessary* facts which are invariant under symmetries. Mainzer (1996) alludes to the connection between symmetry and simplicity (esp. on p. 580), but does not offer a precise formal analysis. Brading and Castellani (2013) discuss the importance of symmetry in modern physics, but do not propose symmetry as a criterion for scientific law. Finally, Baker (2010) discusses the importance of symmetry arguments in metaphysics, especially in relation to the “identity of indiscernables”, but again does not propose symmetry as a criterion for scientific law. That said, our analysis clearly lies in the vicinity of what others have had in mind in their accounts of laws and symmetries.

equal to its inverse image under ψ . Property P is a *law* if it is nomologically necessary for the system *and* invariant under all symmetries in Ψ .

For example, suppose $T = \{1, 2, 3, \dots\}$, and suppose that, for any non-negative integer r , the system has the time symmetry ψ_r defined by $\psi_r(t) = t+r$ for all t in T ; for simplicity, the system has no other symmetries. So, $\Psi = \{\psi_r : r = 0, 1, 2, \dots\}$. Now, suppose all histories in Ω satisfy property P which says: “If the state at time 5 is x , then at time 6 it is y .” Despite being nomologically necessary for the present system, this property falls short of being a law. The inverse image of $[P]$ under any symmetry ψ_r corresponds to the property P' which says: “If the state of the system at time $5+r$ is x , then at time $6+r$ it is y .” Clearly, unless $r = 0$, $[P']$ is not the same as $[P]$, and so P is *not* invariant under the system’s symmetries. We call such a property – nomologically necessary but not invariant under symmetries – a *brute necessity*.

By contrast, suppose all histories in Ω have the property P which says: “For any t in T , if the state of the system at time t is x , then at time $t + 1$ it is y .” It is easy to see that this property *is* invariant under all symmetries of the system: the inverse image of $[P]$ under any symmetry ψ_r is the same as $[P]$. Thus, P is a law.

For another example, consider the kinds of temporally evolving systems that arise in classical mechanics. These satisfy the law of *conservation of energy*, which says that the total energy (kinetic plus potential) remains constant over time. This can be formulated as a property P of the form: “For any times t and t' in T , the total energy of the state at time t' equals the total energy of the state at time t .” Clearly, this property is invariant under the time symmetries $\{\psi_r\}$ introduced above. As already mentioned, classical mechanical systems also have certain *state* symmetries, such as spatial translations, rotations, reflections, and the permutation of (equal-mass) particles. The total energy of a state is unchanged by such symmetries too, so the property P will also be invariant under spatial translations and (equal-mass) particle permutations. Indeed, total energy is unchanged by every symmetry of the system, and for this reason, property P is a law.²³

As we will now see, laws are testable and generalizable in a way in which properties that fall short of being laws are not, even if they are nomologically necessary. Suppose we are trying to figure out the status of some property P . Is it nomologically necessary? Does it capture a general regularity of our system? Is it a law? The first thing to note is that when we investigate a system, we are seldom able to observe all its nomologically possible histories. Conducting many “runs” of the same experiment is an attempt to observe as many histories as possible, but even the best experimental design rarely allows us to observe *all* possible histories. Furthermore, this strategy works only for smaller systems that we can isolate in laboratory conditions. When the system is the economy, the global ecosystem, or the universe as a whole, we are stuck in a single history. We cannot step outside that history and look at alternative histories. The observed history is the only evidence we have. Can we still say anything useful about the status of property P ? It is at this point that symmetries come into play.

Let us return to our simply example of a system with $T = \{1, 2, 3, \dots\}$ and time symmetries of the form ψ_r , where $\psi_r(t) = t+r$. Consider again the property P that says “if

²³ Similarly, one can define the laws of *conservation of momentum*, *of angular momentum*, and so on.

the state at time 5 is x , then at time 6 it is y ", and suppose, as before, that P is nomologically necessary, i.e., every history in Ω satisfies P . If we could observe many nomologically possible histories of the system, we would be able to verify the satisfaction of P in each case. But, as noted, we may be trapped in a single history, h . All we can do is watch this history unfold. We first see $h(1)$, then $h(2)$, then $h(3)$, and so on. Importantly, we get to observe $h(5)$ and $h(6)$ only *once*, so we get only one chance to observe whether h satisfies property P . Furthermore, even if h does satisfy P , this is only a single data point, which tells us very little about the broader status of P . Property P might as well be a contingent feature of the actual history we have observed.

However, we *do* get to observe $h(7)$, $h(8)$, $h(9)$, and so on. So, we *can* consider properties such as P' : "if the state at time 7 is x , then at time 8 it is y "; and P'' : "if the state at time 9 is x , then at time 10 it is y "; and so on. Note that P' corresponds to the inverse image of $[P]$ under ψ_2 ; and P'' corresponds to the inverse image of $[P]$ under ψ_4 ; and so forth. In other words, if we are patient, we *can* observe whether history h satisfies the properties corresponding to the inverse images of the original property under a lot of elements of the system's symmetry monoid. Similarly, in a system with spatial symmetries (of the sort we introduce in Section 3.5), we can in principle observe whether h satisfies the properties corresponding to many of the relevant inverse images simply by traveling in space.

Now, if property P was not itself invariant under symmetries, as in the case of our example, we would not learn much from this exercise. We would learn that h satisfies P ("if $h(5)=x$, then $h(6)=y$ "), that it satisfies P' ("if $h(7)=x$, then $h(8)=y$ "), and that it satisfies P'' ("if $h(9)=x$, then $h(10)=y$ "), and so on. But, strictly speaking, these are distinct properties, and on the face of it, they do not have all that much in common. By contrast, if P is symmetry-invariant, as in the case of the property which says "for all t , if $h(t)=x$, then $h(t+1)=y$ ", then P , P' , P'' , ... are all the *same* property, and thus the present exercise yields a whole series of experimental tests of the *same* law.

Moreover, in this case, the single property P picks up a general pattern, of which we can observe many instances even within a single history, h , and which lends itself to extrapolation into the future. As h unfolds, we can observe that state x is followed by state y not just once but many times. Furthermore, once we have observed this regularity a sufficient number of times, we may feel confident in hypothesizing that P is indeed a law and then predicting that, in the future, state x will also be followed by state y .

Contrast this with the case of a property that is not symmetry-invariant, such as "if $h(5)=x$, then $h(6)=y$ ". Here, there is no such general pattern, and we have no basis for making any predictions. This is the sense in which laws have a *testable* and *generalizable* character that non-symmetry-invariant properties lack, even when they are nomologically necessary.

There is another way of making the same points. Let P be some property, and let P' , P'' , P''' , and so on, be all of its inverse images under the various time (and other) symmetries of the system. Let h be the history that we observe. Suppose that, by exhaustive testing, we verify that h satisfies P , P' , P'' , P''' , and so on. (Or perhaps we only verify some *subcollection* of these properties, but then infer the rest of them through some form of "empirical induction", which is ubiquitous in science.) At this point, we have actually

verified that h satisfies an entire conjunction of properties, informally $P \wedge P' \wedge P'' \wedge P''' \wedge \dots$, or more formally, the property P^* with extension

$$[P^*] = \bigcap_{\psi \in \Psi} \psi^{-1}([P]).$$

Note that, by construction, property P^* is invariant under all symmetries in Ψ . Thus, although we get to test the initial property P only once, by testing a bunch of “ P -like” properties at various points in time (and/or positions in space etc.), we have tested not only P , but something much *stronger*, namely P^* . But note that P^* is not just any arbitrary property: it is symmetry-invariant by construction and thus qualifies as a law (provided it is also nomologically necessary). Moreover, by entailing all the various instances of P -like properties, i.e., P, P', P'', P''' , and so on, the hypothesis that property P^* is a law allows us to make predictions as to what will happen at different points in time (or in space, or after making other admissible changes corresponding to symmetries of the system).

This argument suggests that *any* property that we think we have corroborated by performing a large number of empirical tests at different times (or locations in space, or different orientations of the experimental apparatus, or different collections of otherwise identical atoms, and so on) is *ipso facto* a symmetry-invariant law, and not merely a brute necessity.²⁴

One can give a similar account of probabilistic laws. Let $\{Pr'_E\}_{E \subseteq \Omega}$ be *any* conditional probability structure, and let ψ be a symmetry of the system. We define $\psi(\{Pr'_E\}_{E \subseteq \Omega})$ to be the conditional probability structure $\{Pr''_E\}_{E \subseteq \Omega}$ such that, for any events E and D , we have $Pr''_E(D) = Pr'_{E'}(D')$, where E' and D' are, respectively, the inverse images of E and D under ψ . Let \mathcal{P} be a probabilistic property. Recall that its extension, $[\mathcal{P}]$, is a subset of the set Π of all possible conditional probability structures on Ω . We say that \mathcal{P} is *invariant* under ψ if $[\mathcal{P}]$ is equal to its inverse image under ψ . A property \mathcal{P} that is satisfied by the system’s conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ is a *law* of the system if it is invariant under *all* symmetries in Ψ .

For example, suppose $T = \{1, 2, 3, \dots\}$, and let the time symmetries ψ_r be as defined before. Let Y and Z be two subsets of the state space X , and suppose the system’s conditional probability structure satisfies the probabilistic property \mathcal{P} which says: “Conditional on the state being in Y at time 5, there is a 50% probability that the state will be in Z at time 6.” The inverse image of $[\mathcal{P}]$ under ψ_2 corresponds to the property \mathcal{P}' which says: “Conditional on the state being in Y at time 7, there is a 50% probability that the state will be in Z at time 8.” Clearly, $[\mathcal{P}']$ is not the same as $[\mathcal{P}]$. Thus, $[\mathcal{P}]$ is *not* invariant under ψ_2 , and so \mathcal{P} is not a probabilistic law of the system.

²⁴ One might raise the following concern. If P is nomologically necessary, and h is a possible history, then h satisfies not only P , but also all inverse images of P under all symmetries. To put it more simply, if P is nomologically necessary, then all its inverse images under all symmetries are *also* nomologically necessary. So, whenever P is nomologically necessary, we should be able to corroborate this fact via repeated testing, even if P is merely a brute necessity, rather than a symmetry-invariant law. At first sight, this may seem to challenge our claim that laws stand out in their testability. However, what have we really corroborated via the present exercise of repeated testing is not property P itself, but the much stronger (symmetry-invariant) property P^* defined above. And that property P^* is not merely nomologically necessary, but a law.

However, suppose the conditional probability structure satisfies the property \mathcal{P} which says: “For any time t in T , conditional on the state being in Y at time t , there is a 50% probability that the state will be in Z at time $t+1$.” Then it is easy to see that $[\mathcal{P}]$ is invariant under ψ_r for all positive integers r . If Ψ consists only of the time symmetries $\{\psi_r : r = 0, 1, 2, 3, \dots\}$, then \mathcal{P} is invariant under *all* elements of Ψ , and so \mathcal{P} is a probabilistic law.

As in the case of modal laws, probabilistic laws capture general and repeatable patterns. Consider again the probabilistic property \mathcal{P} which says: “Conditional on the state being in Y at time 5, there is a 50% probability that the state will be in Z at time 6.” Recall that this property is not invariant under our system’s time symmetries. Even if the system’s conditional probability structure satisfies this property, the property does not capture a general pattern. It concerns only the probabilistic transition from time 5 to time 6. If, however, the system has all the time symmetries in Ψ , then we can expect the system to satisfy the properties corresponding to the inverse images of $[\mathcal{P}]$ under the various time symmetries, for instance: \mathcal{P}' , which says: “conditional on the state being in Y at time 7, there is a 50% probability that the state will be in Z at time 8”; and \mathcal{P}'' , which says: “conditional on the state being in Y at time 9, there is a 50% probability that the state will be in Z at time 10”; and so forth. By conjoining those properties, we can deduce the more general property \mathcal{P}^* , which says: “For any t in T , if the state of the system is in the set Y at time t , there is a 50% probability that it will be in Z at time $t + 1$.” This property *is* invariant under all the time symmetries, and it does indeed qualify as a law.

The foregoing considerations show that symmetries are central to the testable and generalizable character of laws. Without suitable symmetries, generalizing from local observations to global laws or testing hypothesized laws would not be possible, especially if we can observe only a single history of a given system. Nor would it be possible to make predictions about the future based on regularities observed in the past. In a slogan, for scientific inference and prediction to work, the system must have sufficient symmetries. In effect, when we engage in scientific reasoning about some system, or even about the world at large, we rely on the auxiliary hypothesis that this system, or the world, is sufficiently symmetrical. If our system, or the world, were what Cartwright (1999) calls “dappled”, then presumably we would not be able to presuppose such symmetries, and our ability to make scientific generalizations would be compromised.²⁵

In Appendix A, we extend the present analysis to *factor systems*, which are obtained by abstracting away from certain details of the original system. In Appendix B, we extend

²⁵ In a system without symmetries (aside from the identity transformation, which is trivially a symmetry), the distinction between laws and brute necessities could not be drawn. Every nomologically necessary property would then vacuously qualify as a “law”: it could not fail to be invariant under any symmetries. But this is clearly a degenerate case. As just argued, in such a system our ability to perform science would be seriously limited. Terminologically, one might distinguish between *trivial* and *non-trivial* laws. *Trivial laws* are ones that are vacuously symmetry-invariant (because there are no non-trivial symmetries). Our interest is of course in *non-trivial laws*.

it to *partial* and *local* symmetries, which are often found in systems with special initial conditions and/or boundary conditions.²⁶

2.7 Ergodicity and its significance

We have noted that, when we scientifically investigate a system, we rely heavily on symmetries. As we may be able to observe just a single history, it is only thanks to symmetries that we can learn general features of the system from local observations. We have seen, for instance, that if we can observe that “if $h(5)=x$, then $h(6)=y$ ”, and the system has the time symmetries of the form $\psi_r(t)=t+r$, then we can infer the general law that says: “for all t , if $h(t)=x$, then $h(t+1)=y$ ”. Similarly, if we can observe that “conditional on the state being in Y at time 5, there is a 50% probability that it will be in Z at time 6”, then we can infer the general law that says: “For any t in T , if the state is in Y at time t , there is a 50% probability that it will be in Z at time $t + 1$.”

However, while the first, non-probabilistic example (where we observe that one state at time 5 is followed by another at time 6) seems unproblematic, the second, probabilistic example is trickier. If we are trapped in a single history, it is unclear how we could ever make an observation such as: “Conditional on the state being in Y at time 5, there is a 50% probability that it will be in Z at time 6.” Making this observation would seem to require looking at many repetitions of states 5 and 6. So, *even if* probabilistic properties *could* be generalized via symmetries once we have observed them, it is unclear how we could observe such properties in the first place.

The solution to this problem lies in the property of *ergodicity*. This is a property that a system may or may not have and that, if present, serves as a prerequisite for inferring probabilistic information from single histories. Indeed, it may be considered a prerequisite for scientific inference more generally. To explain this notion, let us begin with a simple example of how we learn probabilistic information from observing just a single history. Consider a system whose state at any time is the outcome of an independent coin toss, where $T = \{1, 2, 3, \dots\}$. So, the state space is $X = \{Heads, Tails\}$, and each possible history in Ω is one possible *Heads/Tails* sequence.

Suppose the true conditional probability structure on Ω is induced by the single parameter p , the probability of *Heads*. In this example, the Law of Large Numbers guarantees that, with probability 1, the limiting frequency of *Heads* in a given history (as time goes to infinity) will match p . This means that the subset of Ω consisting of “well-behaved” histories has probability 1, where a history is well-behaved if (i) there exists a limiting frequency of *Heads* for it (i.e., the proportion of *Heads* converges to a well-defined limit as time goes to infinity) and (ii) that limiting frequency is p . For this

²⁶ Finally, it is worth emphasizing that, in our analysis, laws are always defined relative to a given system. When some property qualifies as a law of the system given by Ω and $\{Pr_E\}_{E \subseteq \Omega}$, this does not imply that it will also qualify as a law of a different system, say Ω' and $\{Pr'_E\}_{E' \subseteq \Omega'}$. Different systems may be governed by different laws. Indeed, it is widely held that while our universe has certain laws of nature, other universes with distinct laws are logically possible. One might wonder whether our analysis only captures a notion of “system laws” rather than “laws of nature”. However, if we take the system given by Ω and $\{Pr_E\}_{E \subseteq \Omega}$ to represent the universe as a whole, then our account can be interpreted as an account of the laws of nature. Further, as implied by Appendix A, the laws of a larger system, such as the universe, will constrain the laws of any subsystem that can be derived from it as a “factor system” (via constraining its symmetries). Thus, the laws of the universe will constrain the laws of any smaller subsystems in it. In Section 4.4, we show that structurally equivalent systems may share the same laws.

reason, we will almost certainly (with probability 1) get arbitrarily close to the true conditional probability structure on Ω just by observing a single history and counting the number of *Heads* and *Tails* in it.

Now why does this inference work in the present example? As we will see, the system is an example of an *ergodic system*. Its ergodicity manifests itself in the fact that “almost all” histories of the system are “well-behaved”, in the sense that we can read off the desired probability parameter p from the limiting frequency of *Heads*.

To define *ergodicity* more precisely, consider again a system with $T = \{1, 2, 3, \dots\}$ which has all the time symmetries in the set $\Psi = \{\psi_r : r = 0, 1, 2, 3, \dots\}$ (and perhaps other symmetries as well, though we set these aside for now). Heuristically, the symmetries in Ψ can be interpreted as describing the evolution of the system over time.²⁷ Suppose each time-step corresponds to a day. Then the history $h = (a, b, c, d, e, \dots)$ describes a situation where today’s state is a , tomorrow’s is b , the next day’s is c , and so on. Suppose today is Monday. The transformed history $\psi_1(h) = (b, c, d, e, f, \dots)$ describes a situation where today’s state is b , tomorrow’s is c , the following day’s is d , and so on. Thus, $\psi_1(h)$ describes the same “world” as h , but *as seen from the perspective of Tuesday*. Likewise, $\psi_2(h) = (c, d, e, f, g, \dots)$ describes the same “world” as h , but *as seen from the perspective of Wednesday*, and so on.²⁸

Given the set Ψ of symmetries, an event E (a subset of Ω) is Ψ -invariant if the inverse image of E inside of Ω under ψ is E itself, for all ψ in Ψ . Formally, $\psi^{-1}(E) \cap \Omega = E$ for all such ψ . Thus, for any history h in Ω , h is an element of E if and only if $\psi(h)$ is an element of E . For example, suppose again that the elements of T represent days, and E is the event that some property P holds today. If $\psi_1, \psi_2, \psi_3, \dots$ are the symmetries that shift time by one day, by two days, by three days, and so on, then the Ψ -invariance of E implies that property P holds today if and only if it holds tomorrow, the day after tomorrow, and so on. Thus, E is a “persistent” event: an event one cannot escape from by moving forward in time. In a coin-tossing system, where Ψ is still the set of time translations, examples of Ψ -invariant events are “all *Heads*”, where E contains only the history (*Heads, Heads, Heads, ...*), and “all *Tails*”, where E contains only the history (*Tails, Tails, Tails, ...*).

Recall that symmetries preserve the unconditional probabilities of any event E . The system is *ergodic* (with respect to Ψ) if, for any Ψ -invariant event E , the unconditional probability of E , i.e., $Pr_\Omega(E)$, is either 0 or 1.²⁹ In other words, the only persistent events

²⁷ Mathematically, the pair (Ω, Ψ) can be interpreted as a classical *dynamical system*, as defined in footnote 6, with Ω playing the role of a state space (from an outside observer’s perspective) and the transformations in Ψ playing the role of state transformation rules.

²⁸ Note that, under this heuristic interpretation, the world “forgets” its past history: from the perspective of Tuesday, it is as if Monday never happened. This is just an artefact of the formal mathematical model we are using in this example and has no deeper significance. If we used the set \mathbf{Z} of all integers instead of the natural numbers to model time, it would obviate this issue.

²⁹ See, e.g., Peterson (1989, Section 2.4), Walters (2000, Section 1.5, Definition 1.4), or Krengel (1985, Section 1.1.3, Definition 1.7) for precise definitions and further discussion. Ergodicity is usually defined for a symmetry monoid generated by a single transformation ψ , in which case the pair (X, ψ) is called an *ergodic dynamical system*. But the definition generalizes immediately to arbitrary symmetry monoids. See, e.g., Krengel (1985, p. 203). At first sight, this seems unrelated to the use of the term in probability theory, where a stochastic process (such as a Markov process or random field) is called *ergodic* if it has a unique stationary measure. But the two definitions are related because any stochastic process can be

are those which occur in *almost no* history (i.e., $Pr_{\Omega}(E) = 0$) and those which occur in *almost every* history (i.e., $Pr_{\Omega}(E) = 1$).³⁰ The ergodicity of our coin-tossing system is exemplified by the fact that the Ψ -invariant events “all *Heads*” and “all *Tails*” occur with probability 0.

In an ergodic system, it is possible to estimate the probability of any event “empirically”, by counting the frequency with which that event occurs, much like the probability of *Heads* in the coin-tossing example.³¹ Frequencies are thus evidence for probabilities. The formal statement of this is the following important result from the theory of dynamical systems and stochastic processes.

Ergodic Theorem: Suppose the system is ergodic. Let E be any event and let h be any history. For all times t in T , let N_t be the number of elements r in the set $\{1, 2, \dots, t\}$ such that $\psi_r(h)$ is in E . Then, with probability 1, the ratio N_t/t will converge to $Pr_{\Omega}(E)$ as t increases towards infinity.³²

Intuitively, N_t is the number of times the event E has “occurred” in history h from time 1 up to time t . The ratio N_t/t is therefore the *frequency* of occurrence of event E (up to time t) in history h . This frequency might be measured, for example, by performing a sequence of experiments or observations at times 1, 2, ..., t . The Ergodic Theorem says that, almost certainly (i.e., with probability 1), the empirical frequency will converge to the *true* probability of E , $Pr_{\Omega}(E)$, as the number of observations becomes large. The estimation of the probability of *Heads* via the Law of Large Numbers in our coin-tossing example is a special case of this.

To understand the significance of the Ergodic Theorem, let Y and Z be two subsets of X , and suppose E is the event “ $h(1)$ is in Y ” and D is the event “ $h(2)$ is in Z ”. Then the intersection $E \cap D$ is the event “ $h(1)$ is in Y , and $h(2)$ is in Z ”. The theorem says that, by performing a sequence of observations over time, we can estimate $Pr_{\Omega}(E)$ and $Pr_{\Omega}(E \cap D)$ with arbitrarily high precision. Thus, we can compute the ratio $Pr_{\Omega}(E \cap D)/Pr_{\Omega}(E)$ (provided $Pr_{\Omega}(E) \neq 0$). But this ratio is the *conditional probability* $Pr_E(D)$. And so we are able to estimate the conditional probability that the state at time 2 will be in Z , given that at time 1 it is in Y . This illustrates that, by allowing us to estimate unconditional probabilities, the Ergodic Theorem also allows us to estimate conditional probabilities, and thereby to infer the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$. Clearly, the system’s symmetries were indispensable for this exercise. Without symmetries, the frequentist reasoning to which the Ergodic Theorem appeals would not make sense.

represented as a dynamical system (Peterson, 1983, Sections 1.2B to 1.2D). Finally, in statistical physics, Boltzmann’s Ergodic Hypothesis conjectures that the fraction of the time that a physical system spends in a particular region of its state space is proportional to the size of that region. This was the inspiration for the Birkhoff Ergodic Theorem (see below).

³⁰ If Ω is infinite, there is a subtle distinction between *almost no history* ($Pr_{\Omega}(E) = 0$) and *no history* ($E = \emptyset$). Likewise, *almost every history* ($Pr_{\Omega}(E) = 1$) is subtly distinct from *every history* ($E = \Omega$).

³¹ This insight is the basis for Reichenbach’s (1949) “straight rule”, which is to take observed frequencies as the best estimates of “true” probabilities. See, e.g., Eberhardt and Glymour (2009).

³² This result is often called the *Birkhoff (or Pointwise) Ergodic Theorem*. For simplicity, we have stated it relatively informally. For formal statements, see Krengel (1985, Section 1.2), Petersen (1989, Section 2.2), Walters (2000, Theorem 1.14), or Berkovitz, Frigg, and Kronz (2006). Further, we have defined ergodicity for a discrete set of time symmetries. More generally, the statements in this section hold for any set of symmetries that forms an *amenable monoid*. See Krengel (1985, Section 6.4).

2.8 Occam's Razor

We have seen that a system must possess a sufficiently rich set of symmetries to permit general inferences from local observations. Up to now, we have taken for granted that we know, or are justified in hypothesizing, that the system has these symmetries. But what justifies this hypothesis?

This question is crucial for the success of science. Why are we justified in assuming that the system's laws are the same at different times or in different places? Why should replicability of other scientists' experimental results be considered the norm, rather than a miraculous exception? Why is it normally safe to assume that the outcomes of experiments will be insensitive to irrelevant details such as the height of the laboratory bench, or the orientation of the apparatus relative to the planet Jupiter?

In effect, we are assuming that the phenomena under investigation are invariant under certain symmetries – both temporal, as discussed earlier, and spatial, as discussed later, including translations, rotations, and so on. But where do we get this assumption from? The answer lies in Occam's Razor.

Occam's Razor is generally a principle of parsimony. One of its best-known versions says that, when we try to explain some phenomenon, we should not postulate more entities than strictly explanatorily necessary.³³ While this version of Occam's Razor deals with the question of which *entities* to postulate, we are here focusing on another version, which concerns the question of which *regularities* to postulate. Roughly, it says that, if two hypotheses about the regularities in the world are equally consistent with our total evidence, we should prefer the simpler hypothesis.

Now the key point is that the hypothesis of a symmetry-rich system is simpler than the hypothesis of a symmetry-poor system, other things being equal.³⁴ To see why this is the case, contrast two cases. If you hypothesize that the universe has a very large set of symmetries, you are thereby postulating a very simple universe. By contrast, if you hypothesize that the universe has very few symmetries, you are postulating a very complex universe. The first universe admits a parsimonious description in light of its symmetry-induced regularity, the second does not. This suggests the following provisional formulation of Occam's Razor principle:

Occam's Razor: Always assume that a system has the largest possible set of symmetries consistent with all facts about the system that we believe to be nomologically necessary.

We must now make this more precise. We begin by explaining what we mean by “facts about the system that we believe to be nomologically necessary”. We represent this by a collection of those histories among the logically possible ones that we have not ruled out as nomologically impossible. We call this collection of histories our *total*

³³ The literature contains many proposals on how to formalize Occam's Razor. See, e.g., Baker (2013) and Fitzpatrick (2015). For an efficiency argument for Occam's Razor, see Kelly (2007).

³⁴ Mainzer (1996, Section 5.3, p. 580) also relates symmetry to simplicity and notes that scientists generally prefer simpler theories. Likewise, van Fraassen (1989, Section 10.2) notes that when constructing a scientific model, we generally assume that the model satisfies a given symmetry unless we have good reason to believe the contrary. His slogan is: “an asymmetry must always come from an asymmetry”, with some caveats. But neither of these authors connects these ideas to Occam's Razor.

nomological evidence about the system. Formally, it is a subset E of \mathcal{Z} . It could capture the “hard” constraints that we take the system to satisfy, such that, to the best of our knowledge, any history outside E is not permitted by the laws of the system. Of course, we do not *strictly* know that Ω is a subset of E . When we empirically study a system, we do not normally know what Ω is. We can at most be certain that E overlaps with Ω . We will suppose, however, that we are ready to make the auxiliary assumption that E includes, but may be logically weaker than, Ω .

Given this assumption, we are in a position to test the hypothesis that any given transformation of \mathcal{Z} is a symmetry of our system. Let ψ be such a transformation, and for any n , let ψ^n be the transformation obtained by applying ψ repeatedly, n times in a row. For example, if ψ is a rotation about some axis by angle θ , then ψ^n is the rotation by the angle $n\theta$.³⁵ For any such transformation ψ^n , we write $\psi^{-n}(E)$ to denote the inverse image in \mathcal{Z} of E under ψ^n . We say that the transformation ψ is *consistent* with the nomological evidence E if the intersection

$$E \cap \psi^{-1}(E) \cap \psi^{-2}(E) \cap \psi^{-3}(E) \cap \dots$$

is non-empty. This means that E does not falsify the hypothesis that ψ is a symmetry of the system.

For example, suppose we are interested in whether electrostatic forces work the same way at all times. We can test this hypothesis by means of Coulomb’s famous “torsion balance” experiment, which measures the electrostatic attraction or repulsion between two charged objects. Suppose we perform the experiment at time t_1 and obtain evidence E_1 , and we perform the same experiment again at time t_2 and obtain evidence E_2 . Thus, our evidence is summarized by the event $E = E_1 \cap E_2$. Let ψ be a time symmetry that shifts t_1 to t_2 . Then, focusing for simplicity just on the first two terms of the infinite intersection above, we have

$$E \cap \psi^{-1}(E) = E_1 \cap E_2 \cap \psi^{-1}(E_1) \cap \psi^{-1}(E_2).$$

If the experimental results are the same at times t_1 and t_2 , then $E_1 = \psi^{-1}(E_2)$, and the expression for $E \cap \psi^{-1}(E)$ simplifies to $E_1 \cap E_2 \cap \psi^{-1}(E_1)$. Under reasonable assumptions, this is non-empty, meaning that the evidence has not falsified time invariance of electrostatic forces. But if the experimental results at times t_1 and t_2 were different, then E_1 and $\psi^{-1}(E_2)$ would be disjoint, and so the intersection $E \cap \psi^{-1}(E)$ would be empty, which would mean that the evidence is inconsistent with time invariance. As it happens, many thousands of repetitions of Coulomb’s experiment strongly suggest that the intersection is non-empty, and so ψ is a symmetry.

Now our version of Occam’s Razor says that we should postulate as symmetries of our system a maximal monoid of transformations consistent with our evidence. Formally, a monoid Ψ of transformations (where each ψ in Ψ is a function from \mathcal{Z} into itself) is *consistent* with our total nomological evidence E if the intersection

$$\bigcap_{\psi \in \Psi} \psi^{-1}(E)$$

³⁵ In the present terms, rotations must be represented as transformations of the state space X . In Section 3.5, we represent rotations more explicitly, relying on a formal representation of space.

is non-empty. This is the generalization of the infinite intersection that appeared in our definition of an individual transformation's consistency with the evidence. Further, a monoid Ψ that is consistent with E is *maximal* if no proper superset of Ψ forms a monoid that is also consistent with E .

Occam's Razor (formal): Given our total nomological evidence E about a temporally evolving system, always assume that the set of symmetries of the system is a maximal monoid Ψ consistent with E .

What is the significance of this principle? Recall that we earlier defined Ψ to be the set of *all* symmetries of our temporally evolving system. In practice, we do not know Ψ . A monoid Ψ that passes the test of Occam's Razor, however, can be viewed as our best guess as to what the true symmetry monoid is. To disambiguate, let Ψ_{true} denote the true symmetry monoid, and let Ψ_{hyp} denote the hypothesized one.

If Ψ_{hyp} is the hypothesized symmetry monoid, and E is our total nomological evidence, the intersection

$$\bigcap_{\psi \in \Psi_{\text{hyp}}} \psi^{-1}(E)$$

can be viewed as our best guess as to what the set of nomologically possible histories is. It consists of all those histories among the logically possible ones that are not ruled out by the hypothesized symmetry monoid Ψ_{hyp} and the nomological evidence E . We call this intersection our *nomological hypothesis* and label it $\Omega(\Psi_{\text{hyp}}, E)$.

To see that this construction makes sense, note that, under certain conditions, our nomological hypothesis $\Omega(\Psi_{\text{hyp}}, E)$ will reflect the truth about nomological possibility.

Remark: If (i) the hypothesized symmetry monoid Ψ_{hyp} is a subset of the true symmetry monoid Ψ_{true} , and (ii) E is a superset of Ω , then the true set Ω of nomologically possible histories is a subset of $\Omega(\Psi_{\text{hyp}}, E)$.

Condition (i) says that we have not postulated any incorrect symmetries, which is compatible with having overlooked some correct symmetries. Condition (ii) says that we have not mistakenly ruled out any nomologically possible histories, which was our auxiliary assumption about our total nomological evidence. If these conditions hold, our nomological hypothesis will indeed be consistent with the truth and will, at most, be logically weaker than the truth.

It is worth explaining the significance of the auxiliary assumption that we have not mistakenly ruled out any nomologically possible histories (i.e., $E \supseteq \Omega$). Consider the simple coin-tossing system from Section 2.7, where histories are sequences of *Heads* and *Tails*, and time shifts are symmetries. Now consider the event E of getting *Heads* at time 1 and *Tails* at time 2. If we treated E as our total nomological evidence, this would exclude time shifts as symmetries: the event of getting *Heads* at time 1 and *Tails* at time 2 is not invariant under time shifts. The problem is that E , in this case, is not a superset of Ω : it excludes histories that are in fact nomologically possible. The notion of "total nomological evidence" that we require is a "cautious" one. The set E should exclude only histories that we are confident in deeming nomologically impossible. This is a subtle issue, and a full treatment is beyond the scope of this paper.

In Appendix C, we extend the present analysis by offering criteria for choosing a maximal symmetry monoid Ψ consistent with the evidence E in case more than one such monoid can be constructed. We suggest that criteria of inferential modesty and informational parsimony should guide that choice in cases of non-uniqueness.

2.9 The role of time

What is the significance of the linear order of the set T of times? Why is time ordered in one way, and not in another? Do the laws of a given system “care” about the ordering of time? To put it another way: what does it mean to say that today comes between yesterday and tomorrow? Intuitively, it means this: the events that happened yesterday cannot “directly influence” the events that will happen tomorrow; their influence must be “mediated” by the events that happen today. We now make this claim precise using a standard notion from probability theory: the *Markov property*.³⁶

To explain this property, we first introduce the notion of *conditional independence*. Let $\{Pr_E\}_{E \subseteq \Omega}$ be a conditional probability structure, and let D and E be two events (i.e., subsets of Ω). We say that D and E are *independent* if $Pr_D(E) = Pr_\Omega(E)$ and $Pr_E(D) = Pr_\Omega(D)$.³⁷ Informally, if we interpret probabilities as encoding “information”, this means that learning whether or not D has occurred provides no information about whether or not E will occur, and vice versa.

To illustrate, recall the simple coin-tossing system from Section 2.7. Let E and D be the events “the outcome at time 1 is *Heads*” and “the outcome at time 2 is *Tails*”. Then $Pr_\Omega(E) = 1/2$ and $Pr_\Omega(D) = 1/2$, assuming for simplicity that $p = 0.5$. Here, the outcome at time 1 has no effect on the outcome at time 2. So, even if we tossed *Heads* at time 1, this would not change the probability of obtaining *Tails* at time 2, and so $Pr_E(D) = 1/2$. Likewise, the outcome at time 2 tells us nothing about what happened at time 1. If we had not observed the outcome at time 1 but obtained the outcome *Tails* at time 2, we would still assign probability $1/2$ to *Heads* at time 1. So, $Pr_D(E) = 1/2$. Thus, the events E and D are independent.

Now let C , D , and E be three events. We say that C and E are *conditionally independent, given D* , if $Pr_{C \cap D}(E) = Pr_D(E)$ and $Pr_{E \cap D}(C) = Pr_D(C)$. Again, if we interpret probabilities as encoding “information”, this means the following. Suppose you already know that D has occurred. Then learning whether or not C has occurred provides no further information about whether or not E will occur, and vice versa.

To illustrate, return again to the coin-tossing example (where $T = \{1, 2, 3, \dots\}$) with $p = 0.5$, but suppose we use the tosses of the fair coin to determine the position of a token on an infinite line. We move the token after each coin toss: if we toss *Heads*, we move the token one space to the right, and if we toss *Tails*, we move it one space to the left. Let us represent the position of the token by an integer (either positive or negative); in other words, $X = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$. Let x_t denote the position of the token at time t . Then the rule becomes the following: “If you toss *Heads* at time t , then

³⁶ The importance of Markov properties in understanding causality has been emphasized by Pearl (2000) and Spirtes, Glymour, and Scheines (2000).

³⁷ If $Pr_\Omega(D) > 0$, the first equation is equivalent to $Pr_\Omega(E \cap D) = Pr_\Omega(D) Pr_\Omega(E)$. If $Pr_\Omega(E) > 0$, the second equation is equivalent to $Pr_\Omega(E \cap D) = Pr_\Omega(D) Pr_\Omega(E)$. Thus, if $Pr_\Omega(D) > 0$ and $Pr_\Omega(E) > 0$, the two equations are equivalent. But if $Pr_\Omega(D) = 0$ or $Pr_\Omega(E) = 0$, the equations must be stated separately.

$x_{t+1} = x_t + 1$; if you toss *Tails* at time t , then $x_{t+1} = x_t - 1$.” For simplicity, suppose the coin always starts at position 0 (i.e., $x_1 = 0$).³⁸

If D is an event describing the position of the token at time t , and E is an event describing its position at time $t+1$, then these two events are not independent. For example, suppose E is the event “ $x_6 = 3$ ”. Then a simple calculation shows that $Pr_{\Omega}(E) = 5/16$. If D is the event “ $x_5 = 2$ ”, then $Pr_D(E) = 1/2$, because the token now has a 50% probability of moving from position 2 to position 3 in one time step. Thus, $Pr_D(E) \neq Pr_{\Omega}(E)$. The location of the token at time 5 tells us a great deal about its probable location at time 6.

However, once we know the position at time 5, learning the position at time 4 tells us *nothing further* about the position at time 6. Continuing the previous example, let C be the event “ $x_4 = 1$ ”. Then straightforward calculations show that $Pr_{C \cap D}(E) = 1/2 = Pr_D(E)$ and $Pr_{E \cap D}(C) = 1/2 = Pr_D(C)$. In other words, if we *already knew* that the token’s position was 2 at time 5 (so that it had a 50% probability of moving to position 3 at time 6), then learning its position at time 4 tells us *nothing further* about where it might be at time 6. Likewise, if we *already knew* that the token’s position was 2 at time 5 (so that it has a 50% probability of having been at position 1 at time 4), then learning its position at time 6 tells us *nothing further* about where it might have been at time 4.

In this example, the conditional independence of the events C and E , given D , is due to the fact that D concerns the state of the system at a point in time *between* the times described by C and E and that D provides us with *complete* information about the state of the system at this intermediate time. If D provided only partial information about that state, we would not get the same result. For example, suppose D' is the event, “ $x_5 = 0, 2, \text{ or } 4$ ”, which does not fully specify the state at time 5. Then it can be shown that $Pr_{C \cap D'}(E) > Pr_{D'}(E)$. Here, learning additional information about the state at time 4 can still tell us something about where the coin is likely to be at time 6.

Now let us generalize this example. Let T be any linearly ordered set, let X be any set of states, and consider a temporally evolving system given by a collection Ω of possible histories (i.e., functions from T into X) and a conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$. For any time t in T , and any state x in X , let E_x^t denote the event “the state of the system at time t is x ”. More generally, for any subset Y of X , let E_Y^t denote the event “the state of the system at time t is an element of Y ”. We say that the system satisfies the *Markov property* if, for any times $r < s < t$ in T , any subsets Y and Z of X , and any state x in X , the events E_Y^r and E_Z^t are conditionally independent, given the event E_x^s . In other words, if you have *complete information* about the state of the system at some time s (you know that the state is x), then learning something about its state at some earlier time (e.g., that it was an element of Y at time r) tells you *nothing further* about its probable state at some later time (e.g., about how probable it is that it will fall into the set Z at time t). Roughly speaking, this means that the state of the system at time r cannot “directly influence” the state of the system at time t . It can only influence that state “indirectly”, via influencing the state at the intermediate time s . Any system with this property is called *Markovian*.

³⁸ Technically, the system just described is a *simple random walk*.

Note that the Markov property does not say that the system's future evolution is unconditionally independent of its past. It just says that the dependency of the future on the past is mediated through the present. This property is fundamental to the way we normally think about time. To see this, imagine a universe where the Markov property was *not* true. Then there would exist some times $r < s < t$ in T , some subsets Y and Z of X , and some state x in X , such that the conditional probability $Pr(E_Z^t | E_Y^r \cap E_x^s)$ is distinct from $Pr(E_Z^t | E_x^s)$.³⁹ In other words, even with a *complete* specification of the present state x , the probability of some *future* event Z would depend on whether or not some *past* event Y had occurred. This would suggest that the state specification x does *not*, in fact, contain all the information about the system's present state; somehow, information about the past is bypassing the present and "leaking" directly into the future. This, in turn, suggests that this so-called "past" is not really in the past at all; our model of the system's time structure is incorrect.

We take the Markov property to be a necessary condition for the "correct" ordering of time. To be "well-behaved", a temporally evolving system must be Markovian. What the present must do at any point in time in order to *count* as the present is "separate" the past from the future. If this property is violated, the set T does not properly play the role of time.

Three points are worth noting. First, some systems may admit multiple time orderings with respect to which they are Markovian. An extreme limiting case is given by our original coin-tossing system without the moving token, which is Markovian with respect to *every* ordering of T . Here, the precise order of time is irrelevant. By contrast, in the modified coin-tossing system with the token, the order of time matters, as we have seen. In fact, the temporal order with respect to which the system satisfies the Markov property is essentially unique; it is unique up to time reversals. This brings us to our second point. Although the Markov property says something about the linear "topology" of time, it tells us nothing about the *direction* of time. As illustrated by the modified coin-tossing system, the Markov property is completely invariant under time reversals. In other words, the Markov property only says that the present separates the past from the future. But it does not tell us on which side of the present lies the past, and on which side lies the future. And third, just as the Markov property says nothing about the direction of time, so it says nothing about its duration. There is no purely Markovian way of measuring the "length" of a time interval or saying when one time interval is longer than another.

What, then, can we say about the directionality and length of time? It turns out that symmetries are crucial for the analysis of both. In the case of length, we offer a detailed analysis in Section 3.9, showing that there is a natural way of measuring time duration, as long as the system has sufficiently rich symmetries. And in the case of directionality, we can say that a condition for time to have a direction is that time reversals are *not* symmetries of the system. Since time reversals are symmetries of classical mechanical systems (in the sense explained in footnote 19), it follows that, in those systems, there is no real direction of time: temporal orders are unique at most up to time reversal. By contrast, in thermodynamic systems, time reversals are not symmetries, and hence these systems meet the condition for time to have a direction. To the extent that the world, as

³⁹ Here, to avoid cumbersome subscripts, we are using the notation $Pr(A | B)$ to denote the conditional probability $Pr_B(A)$.

seen from our perspective, is best understood as a system in which time reversals are not symmetries, there is then a coherent basis for the directionality of time (for further discussion, see Roberts 2013).

3. Spatially extended systems

3.1 Basic definitions

We now turn to a more richly described class of systems whose states evolve over time. To define a system in this class, we still represent *time* by a linearly ordered set T , but also incorporate an explicit notion of *space*, represented by a set S of *spatial locations*. Let $S \times T$ be the set of all ordered pairs of the form (s, t) , where s is an element of S , and t is an element of T . We refer to $S \times T$ as *space-time*. Again, let X denote a set of possible states, called the *state space*. Unlike before, the elements of X are no longer “global” states, in which the system can be *at specific points in time*, but “local” states, in which the system can be *at specific points in space and time*. Again, we treat the elements of X as primitives of our model. Histories are now functions from space-time (rather than merely time) into the state space. Formally, a *spatially extended history* is a function h from $S \times T$ into X . For each point (s, t) in $S \times T$, $h(s, t)$ is the state of the system in spatial location s at time t .

In analogy to our earlier model, we write Ω to denote the set of all spatially extended histories deemed *possible*, which, as before, play the role of possible worlds. Again, this is a subset – often a proper one – of the set \mathcal{H} of all *logically possible* histories (here, all functions from $S \times T$ into X). So, membership in Ω is best interpreted as *nomological possibility*. Subsets of Ω are called *events*.

Finally, we define a *conditional probability structure* on Ω . As before, this is a family of conditional probability functions $\{Pr_E\}_{E \subseteq \Omega}$, containing one Pr_E for each event E in Ω , with standard properties. Recall that Pr_E assigns to any event in Ω the conditional probability of that event, given E . A *spatially extended system* is the pair consisting of the set Ω of possible spatially extended histories and the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$.

For example, in a classical mechanical system, T is the set \mathbf{R} of real numbers, S is the three-dimensional Euclidean space (i.e., $S = \mathbf{R}^3$), and each state $h(s, t)$ in X is given by the set of particles present at spatial location s at time t , along with their physically relevant properties (e.g., masses and momenta) and the values of any force fields (e.g., gravity) acting on these particles.⁴⁰ In a classical electrodynamical system, the state $h(s, t)$ must also specify the particles’ charges, along with the electric and magnetic field vectors at (s, t) . In that sense, electrodynamics relies on a richer ontology than classical mechanics.

In a quantum-mechanical system, it might be tempting to suppose that $S = \mathbf{R}^3$, and to suppose that $h(s, t)$ is given by the values of the wave functions of each of the particles in the system at space-time location (s, t) . But this is not correct, because the wave functions of interacting particles in a quantum system cannot generally be defined independently of each other. Instead, we must define a *joint* wave function for the entire

⁴⁰ We are *not* saying that this is the most parsimonious or computationally convenient way to represent a classical mechanical system. It is only one way of representing such a system in our framework.

multi-particle system. So, in a quantum-mechanical system with n particles, we would define space to be $S = (\mathbf{R}^3)^n$, with three coordinates representing the spatial “position” of each of the n particles in an underlying ordinary Euclidean space;⁴¹ and we would define the set X of possible states of the system to be the set of complex numbers, capturing amplitudes, whose squared absolute values behave formally like probabilities. Thus a spatially extended history h is a function from $(\mathbf{R}^3)^n \times T$ into the set of complex numbers, representing the joint wave function of the whole ensemble of particles.

For instance, if there are two particles, labelled 1 and 2, then $h(x_1, y_1, z_1, x_2, y_2, z_2, t)$ represents the joint state at time t of particles 1 and 2 at positions x_1, y_1, z_1 and x_2, y_2, z_2 in the underlying three-dimensional Euclidean space. This joint state of the two particles is a complex number whose squared absolute value can be interpreted, under some assumptions, as the probability of particles 1 and 2 being observable at positions x_1, y_1, z_1 and x_2, y_2, z_2 , respectively, at time t .

3.2 Determinism and indeterminism

As in the case of temporally evolving systems, we can define a family of notions of determinism and indeterminism for spatially extended systems. For any subset L of locations in $S \times T$, we write h_L to denote the restriction of the function h to the points in L . We can then ask for which proper subsets L of $S \times T$, if any, h_L has a unique extension to all of $S \times T$ in Ω . Again, an *extension* of h_L is a history h' such that $h'_L = h_L$. When h_L is uniquely extendible to all of $S \times T$, we say that history h is *L-deterministic*.

For example, the histories of classical mechanical systems are *L-deterministic* for any subset L of $S \times T$ that has the form $S \times T'$, where T' is any non-empty subset of T . Information about the system for even a single “time slice” of space-time, i.e., a set of the form $S \times \{t\}$ for some t in T , suffices to determine the full spatially extended history. In contrast, the histories of quantum-mechanical systems (if wave-function collapses are allowed) are not generally *L-deterministic* when L consists of time slices.

The present definitions allow us to explore some interesting possibilities not captured by standard definitions that focus exclusively on past-to-future determination.⁴² For example, some systems might encode their entire spatially extended history in each individual space-time location. Histories would then be *L-deterministic* for every singleton set $L = \{(s, t)\}$, where (s, t) is in $S \times T$. Here, we would have an extreme form of local-to-global determinism. Alternatively, some systems might encode their entire spatially extended history in some collection of “spatial slices of time”, i.e., some subset L of $S \times T$ which has the form $S' \times T$, where S' is a non-empty subset of S , possibly singleton. This would be a kind of spatial, not temporal, determinism.⁴³ Other systems might never be *L-deterministic* for any proper subset L of $S \times T$.

⁴¹ Strictly speaking, particles in quantum systems do not have “positions”, so we are using this term rather loosely. Also, there is a dual representation of the wave function (obtained via Fourier transform), where the coordinates in $(\mathbf{R}^3)^n$ represent the “momenta” (again, loosely) of the n particles. These two representations are equally valid.

⁴² For an earlier discussion of local-to-global forms of determinism, see Earman (1986, pp. 33–35).

⁴³ This sort of determinism occurs in *expansive cellular automata*, a class of spatially extended systems discussed in the theory of dynamical systems. See Pivato (2009).

There may also be some more limited, non-global forms of determination, for instance when a history restricted to some set L of locations is uniquely extendible to a history restricted to some superset L^* of L , which is still smaller than $S \times T$ in its entirety.⁴⁴ To capture this idea, we can say that a history h is L -to- L^* -deterministic if, for any history h' in Ω , if $h_L = h'_L$, then $h_{L^*} = h'_{L^*}$.

We might imagine, for instance, systems that are deterministic “across space” but not “across time”. In such a system, a history restricted to some set L of the form $S' \times \{t\}$, where S' is a non-empty subset of S and t a point in time, might determine the entire “time slice” of that history across $L^* = S \times \{t\}$, but not the rest of the history. Some crystals and other chemical or physical systems involving highly regular spatial structures might have this feature. Similarly, for suitable specifications of L and L^* , we can represent the phenomenon that, in some systems in which “information” travels with finite speed, events at particular space-time locations at time t_1 are entirely determined by the events occurring within their “backwards light cones” at some time $t_0 < t_1$. Such systems may be L -to- L^* -deterministic, but not deterministic in a more global sense.

3.3 Nomological possibility and necessity

In analogy to the case of temporally evolving systems, we can define two modal operators for each set L of space-time locations, namely nomological possibility and necessity relative to L . For each set $L \subseteq S \times T$, call one history, h' , *accessible* from another, h , relative to L , if the restrictions of h and h' to L coincide, i.e., $h'_L = h_L$. We then write $hR_L h'$. For any event $E \subseteq \Omega$, we define

$$\blacklozenge_L E = \{h \in \Omega : \text{for some } h' \in \Omega \text{ with } hR_L h', \text{ we have } h' \in E\},$$

$$\blacksquare_L E = \{h \in \Omega : \text{for all } h' \in \Omega \text{ with } hR_L h', \text{ we have } h' \in E\}.$$

Here, $\blacklozenge_L E$ and $\blacksquare_L E$ are, respectively, the sets of all histories in which E is nomologically possible and nomologically necessary *once the history in space-time region L is given*. Important special cases are (i) $L = S \times T'$, where S is all of space and T' is a particular set of time points, such as those up to time t , (ii) $L = S' \times T$, where T is all of time and S' is some spatial region, and (iii) $L = \emptyset$ for possibility and necessity in the “atemporal” sense. Since the present definitions are completely analogous to their earlier counterparts in Section 2.3, we will not say more about them here.

3.4 Modal and probabilistic properties

We now turn again to the question of how to distinguish between those properties of a system that qualify as “laws” and those that fall short of being laws. As before, our analysis is based on the notion of *symmetry*, but now with the additional ingredient that these symmetries can involve space as well as time.

In analogy to our earlier definition, a *property of histories*, P , is a binary feature that a spatially extended history may or may not have. Its *extension* is some subset $[P]$ of the set \mathcal{H} of all logically possible histories. A spatially extended history h satisfies P if h

⁴⁴ For example, this phenomenon arises frequently in the solution of *boundary value problems* in mathematical physics. See Pivato (2010).

belongs to $[P]$. Again, if $[P]$ includes all of Ω , then P can be called *nomologically necessary*. Similarly, a *probabilistic property*, \mathcal{P} , is a binary feature that a conditional probability structure may or may not have, and its extension, $[\mathcal{P}]$, is the set of all those conditional probability structures on Ω that satisfy \mathcal{P} .

3.5 Symmetries

The notion of a *state symmetry* for spatially extended systems is virtually identical to the one defined in Section 2.5 for temporally evolving systems, so we do not discuss it further.⁴⁵ Instead, we turn directly to *symmetries acting on space-time*. Let ψ be a function from $S \times T$ into itself (i.e., a transformation of space-time). Again, ψ induces a function from the set \mathcal{H} of logically possible histories into itself. For any spatially extended history h , we define the transformed history

$$\psi(h) = h', \text{ where, for all } (s,t) \text{ in } S \times T, h'(s,t) = h(\psi(s,t)).$$

As before, given any set E of histories in \mathcal{H} , the *inverse image* of E under ψ , written $\psi^{-1}(E)$, is the set of all histories h in \mathcal{H} such that $\psi(h)$ lies in E . The function ψ is a *symmetry* if

- $\psi(\Omega) = \Omega$; and
- for any events E and D in Ω , if E' and D' are the inverse images of E and D under ψ , then $Pr_{E'}(D') = Pr_E(D)$.⁴⁶

For example, if T is the set of real numbers (i.e., $T = \mathbf{R}$) and S is the three-dimensional Euclidean space (i.e., $S = \mathbf{R}^3$), we can consider a spatially extended system in classical mechanics. The following transformations of $S \times T$ are space-time symmetries of such a system, each defined for all (s,t) in $S \times T$:

- *Time translation*: $\psi(s,t) = (s, t+r)$, where r is a fixed real number;
- *Spatial translation*: $\psi(s,t) = (s+v, t)$, where v is a fixed three-dimensional vector (an element of \mathbf{R}^3); and
- *Space-time rescaling*: $\psi(s,t) = (r s, r t)$, where $r > 0$ is a fixed real number.

More general symmetries include composite functions resulting from the combination of a transformation ϕ of the state space (X) with a transformation ψ of space-time ($S \times T$).⁴⁷ Examples in classical mechanics are spatial rotations, spatial reflections, spatial rescalings, and Galilean transformations.⁴⁸ Crucially, it is possible that neither

⁴⁵ For an example, take an n -particle quantum system, where $S = (\mathbf{R}^3)^n$, X is the set of complex numbers, and a spatially extended history h is a wave function. Let ϕ be a *phase rotation map* on the complex plane; formally, there is some angle θ such that, for all x in X , $\phi(x) = e^{i\theta}x$. Then ϕ is a state symmetry.

⁴⁶ As before, for any subsets D, E of \mathcal{H} , we define $Pr_E(D) = Pr_{E \cap \Omega}(D \cap \Omega)$, provided $E \cap \Omega \neq \emptyset$.

⁴⁷ An additional property we might require of a space-time transformation ψ is *time preservation*: for any points (s_1, t_1) and (s_2, t_2) in $S \times T$, with $\psi(s_1, t_1) = (s'_1, t'_1)$ and $\psi(s_2, t_2) = (s'_2, t'_2)$, if $t_1 \leq t_2$, then $t'_1 \leq t'_2$. This implies that if $t_1 = t_2$, then $t'_1 = t'_2$. A time-preserving transformation acts on $S \times T$ such that the set of all space-time points at time t_1 gets moved *en bloc* to the set of all space-time points at time t'_1 . The transformations described above have this property, but we do not need to include it in our definition.

⁴⁸ These are defined as follows. *Spatial rotation*: Fix a line L in S and an angle θ . For any point s in S , let s' be the point obtained by rotating s by an angle of $-\theta$ around L . For all (s,t) in $S \times T$, define $\psi(s,t) = (s',t)$. Let L' be the line parallel to L , but passing through the origin. For all x in X , define $\phi(x)$ by rotating all the momentum vectors and force field vectors in x by the angle θ around L' . *Spatial reflection*: Fix a plane P in S . For any point s in S , let s' be the point obtained by reflecting s across P . For all (s,t) in $S \times T$,

the transformation ϕ of the state space nor the transformation ψ of space-time alone is a symmetry, and yet, when combined, they form a symmetry.⁴⁹

Of course, any combination of symmetries is also a symmetry. An example is a *spatiotemporal translation*, which is a combination of a time translation and a spatial translation. In a classical electrodynamical system, *only* the spatiotemporal translations and rotations are space-time symmetries. Galilean transformations are not space-time symmetries of classical electrodynamics; indeed, this was the original impetus for the development of special relativity theory.

3.6 Laws and their significance

As in the earlier case of temporally evolving system, a *modal law* of a spatially extended system is a property of histories, P , that is nomologically necessary for the system and invariant under all of the system's symmetries. A *probabilistic law* is a probabilistic property, \mathcal{P} , that is satisfied by the system's conditional probability structure and invariant under all of its symmetries.

For example, let $S = \mathbf{R}^3$ and $T = \mathbf{R}$, and suppose the symmetry monoid Ψ contains all the spatiotemporal translations defined in the previous section. Suppose all histories of the system satisfy the property P which says: "If the state at space-time position $(3,7,2,14)$ is x , then at position $(4,8,1,17)$ it is y ." If ψ is a spatial translation by the vector $(1,2,3)$, then the inverse image of $[P]$ under ψ corresponds to the property P' which says: "If the state at $(4,9,5,14)$ is x , then at position $(5,10,4,17)$ it is y ." Clearly, $[P']$ is not the same as $[P]$, and so property P falls short of being a law.

However, suppose all histories satisfy the property P which says: "For any location (s_1, s_2, s_3) in S and any time t in T , if the state at space-time position (s_1, s_2, s_3, t) is x , then at position $(s_1+1, s_2+1, s_3-1, t+3)$ it is y ." It is easy to see that $[P]$ is invariant under all spatiotemporal translations. If Ψ consists *only* of the spatiotemporal translations, then P is invariant under *all* symmetries, and so P is a law.

An illustration is *Gauss's Law* in an electrodynamical system. This asserts, roughly, that the net "flux" of the electric field passing through the walls of any closed compartment is proportional to the net charge contained inside that compartment. This property is invariant under spatiotemporal translations, because the net flux and the net charge are unchanged by such transformations. Indeed, Gauss's Law is preserved by *every* symmetry of an electrodynamical system; that is why it is a law.

define $\psi(s,t) = (s', t)$. Let P' be the plane parallel to P , but passing through the origin. For all x in X , define $\phi(x)$ by reflecting all the momentum vectors and force field vectors in x across P' . *Spatial rescaling*: Fix some real number $r > 0$, and define $\psi(s,t) = (s/r, t)$ for all (s,t) in $S \times T$. Meanwhile, let ϕ be a transformation of X that multiplies the momentum vector of every particle by r , and also multiplies all force field vectors by r . *Galilean transformation*: For all (s,t) in $S \times T$, define $\psi(s,t) = (s-tv, t)$, where v is a fixed three-dimensional vector (an element of \mathbf{R}^3). Meanwhile, for all x in X , define $\phi(x)$ by adding the vector v to all momentum vectors in x .

⁴⁹ In all four examples in footnote 48, neither ψ nor ϕ is itself a symmetry. But when combined, they *do* form a symmetry. For another example, in classical electrodynamics, let ψ be a spatial reflection acting on $S \times T$, and let ϕ be a transformation of X which applies the corresponding reflection to all momentum vectors and field vectors, and which further *negates* the magnetic field vector. Neither one of these transformations is a symmetry by itself, but when combined, they *do* form a symmetry.

As before, the significance of laws, as opposed to properties that fall short of being laws, lies in their openness to testing and generalization. Consider again the property: “If the state at space-time position $(3,7,2,14)$ is x , then at position $(4,8,1,17)$ it is y .” This property is observable exactly once in any history, namely at space-time position $(3,7,2,14)$ alone. Taken in isolation, the observation that some history has this property tells us very little. It leaves open whether there is some broader regularity. By contrast, consider the property: “For any location (s_1, s_2, s_3) in S and any time t in T , if the state at position (s_1, s_2, s_3, t) is x , then at position $(s_1+1, s_2+1, s_3-1, t+3)$ it is y .” Recall that, if the system’s symmetry monoid consists of all spatiotemporal translations, then this property is a law. Indeed, it has many observable manifestations in each history, both at different times and in different places, and it thus picks up a pattern that we can in principle test and use as a basis for predictions, even within a single history.

3.7 Spatiotemporal ergodicity and its significance

Before turning to a more detailed analysis of the role of space in a spatially extended system, it is worth sketching how the property of ergodicity can be extended to such a system and discussing the significance of this. In the present case, too, ergodicity is the key to learning a system’s conditional probability structure, even if we are able to observe only a single history of the system.

Recall that, for some set Ψ of symmetries, an event E (a subset of Ω) is Ψ -invariant if, for every ψ in Ψ , the inverse image of E inside of Ω under ψ is E itself. For illustrative purposes, suppose Ψ consists of all spatiotemporal translations by four-dimensional vectors of integers (applying the definition from Section 3.5).⁵⁰ The system is *spatiotemporally ergodic* if the unconditional probability of any Ψ -invariant event E , $Pr_{\Omega}(E)$, is either 0 or 1.

Since Ψ consists of spatiotemporal translations, Ψ -invariant events are events from which one cannot escape by travelling through space, or by travelling forwards or backwards through time. In our example, let ψ be a spatiotemporal translation in Ψ such that, for all (s_1, s_2, s_3, t) in $S \times T$, we have $\psi(s_1, s_2, s_3, t) = (s_1+5, s_2-7, s_3+10, t+3)$. If we interpret the spatially extended history h as describing a possible world “from the perspective of position $(0,0,0,0)$ ”, then, heuristically, the transformed history $\psi(h)$ describes the same world “from the perspective of position $(5,-7,10,3)$ ”. Here a Ψ -invariant event E has the property that whenever a history h is in E , then so is $\psi(h)$. Roughly speaking, this means that the world described by h appears to be in the set E “from the perspective of position $(0,0,0,0)$ ” if and only if it appears to be in E “from the perspective of position $(5,-7,10,3)$ ”, and so on. Ergodicity requires any such event to occur either almost always (with probability 1) or almost never (with probability 0).

In a spatiotemporally ergodic system, we can estimate the probability of any event by counting the spatiotemporal frequency with which that event occurs.

Spatiotemporal Ergodic Theorem: Suppose the system is spatiotemporally ergodic. Let E be any event and let h be any history. For all $r > 0$, let Ψ_r be the set of all spatiotemporal translations by any vector (v_1, v_2, v_3, v_4) with integer coordinates between 1 and r . Let N_r be the

⁵⁰ The system might also have other spatiotemporal symmetries, but this is irrelevant here.

number of translations ψ in Ψ_r such that $\psi(h)$ is in E . Then, with probability 1, the ratio N_r / r^4 will converge to $Pr_\Omega(E)$ as r increases towards infinity.⁵¹

Intuitively, N_r is the number of times the event E has “occurred” in the spatially extended history h from time 1 to time r and inside a three-dimensional box with side-length r . The ratio N_r / r^4 is therefore the *frequency* of occurrence of event E , up to time r inside this box, in the spatially extended history h . This frequency might be measured, for example, by performing a sequence of experiments or observations inside this box. The Spatiotemporal Ergodic Theorem says that, with probability 1, the empirical frequency will converge to the true probability of E as the number of observations becomes large.⁵² As explained in Section 2.7, we can use this procedure to estimate not only unconditional probabilities but also conditional ones, and thereby to learn the properties of the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$.

A broader lesson is that whether a system is ergodic in the first place depends on the system’s symmetries. If a system is rich in symmetries, then ergodicity becomes easier to achieve than if the system has only few symmetries. To see this, note that the notion of Ψ -invariance is logically more demanding for a larger set Ψ of symmetries than for a smaller one, since an event E will need to be preserved under more symmetries in order to qualify as Ψ -invariant. As a result, there will be fewer Ψ -invariant events if Ψ is large, and hence the property of ergodicity, which constrains the probability of Ψ -invariant events, becomes less demanding. Conversely, if the set Ψ of symmetries is small, more events may qualify as Ψ -invariant. In the limit, if Ψ contains only the (trivial) identity symmetry, then *every* event E will be Ψ -invariant, and so *no* system with a non-degenerate conditional probability structure will qualify as ergodic. (Recall that, in an ergodic system, the unconditional probability of every Ψ -invariant event must be either 0 or 1. If all events are Ψ -invariant, this rules out non-degenerate probabilities.) Thus, we must conclude not only that ergodicity is a key prerequisite for inferring a system’s conditional probability structure from local observations, but also that without enough symmetries this inference would not get off the ground.⁵³

3.8 The role of space

What is the role of “space” in a spatially extended system? As we will now see, the structure of space affects the way the system evolves over time. To make this precise,

⁵¹ For simplicity, we here assume that the symmetry monoid Ψ is isomorphic to \mathbf{Z}^4 , where \mathbf{Z} is the monoid of integers. The theorem, which we have stated somewhat informally, also holds if Ψ is isomorphic to \mathbf{R}^4 , or if Ψ is any amenable monoid. For a more formal statement, see Krengel (1985, Chapter 6).

⁵² It is not necessary to average over a sequence of “boxes”; the same argument works for any sequence of sets which increase in size and thickness in an appropriate sense, technically any *Følner sequence*.

⁵³ One complication is that not all systems are ergodic. For example, in systems with conservation laws, such as conservation of energy or momentum, each value of the conserved variables determines a non-trivial invariant subset of Ω . But the “Ergodic Decomposition Theorem” shows that any non-ergodic system can be split up into “ergodic components” – informally, minimal invariant subsets of Ω , each of which (except possibly a set of measure zero) supports its own ergodic probability function (see, e.g., Glasner 2003, p. 72, Theorem 3.22). If we are part of the system, then we are already confined to one such component. Furthermore, even if the system as a whole is not ergodic, many of its factor systems may be ergodic (see Appendix A). This suggests that, by choosing the right level of description for the system (e.g., by adopting a sufficiently coarse-grained, higher-level description, as discussed in List and Pivato 2015), we may be able to reap the benefits of ergodicity. For the applicability of ergodic methods to non-ergodic Hamiltonian systems, see Berkovitz, Frigg, and Kronz (2006, Section 4).

we first introduce a formal representation of the topology of space and then discuss its role in the system's dynamics.

The topology of space can be represented by a binary relation \rightarrow between subsets of S . Heuristically, if R and R' are two subsets of S , such as two “regions” of space, then $R \rightarrow R'$ means that R and R' are “adjacent” in that information from R can flow “directly” into R' , without needing to pass through some intervening points “between” R and R' . Later, we explain exactly what we mean by “information flow”, but for our initial discussion, we leave it unexamined. We call \rightarrow the *adjacency structure* of space.⁵⁴

Adjacency structures arise naturally in many systems. For example, suppose S is ordinary three-dimensional Euclidean space, and suppose information can flow only “continuously” through this space. This would be the case, for instance, in a system consisting of particles travelling along continuous trajectories and interacting via continuous force fields, such as those found in classical mechanics, or in a system described by partial differential equations, such as those found in quantum mechanics, classical electrodynamics, or hydrodynamics. In such systems, for any subsets R and R' of S , we have $R \rightarrow R'$ if there exists a point s in R such that, for any radius $r > 0$, the ball of radius r centred at s intersects R' .⁵⁵

For another example, suppose S is the *three-dimensional integer lattice*: the set of all ordered triples $s = (s_1, s_2, s_3)$, where s_1, s_2 , and s_3 are integers. Say that two points s and s' in S are *neighbours* if they differ in only one coordinate and that difference is 1. Thus $(3, 7, 5)$ and $(3, 6, 5)$ are neighbours. Suppose information can flow only directly between neighbours in the lattice. Then, for any subsets R and R' of S , we have $R \rightarrow R'$ if some point in R is a neighbour of some point in R' .⁵⁶ Discrete spatial geometries of this kind can be found in a class of systems called *cellular automata*.⁵⁷

For a final example, consider a *directed graph*, which consists of a set of “vertices”, along with a set of “arrows” which connect pairs of vertices. Directed graphs can be used to model electric circuits, communication networks (e.g., the internet), economic and transportation networks, and biological systems (e.g., neural networks, gene regulatory networks, and epidemiological networks). Suppose S is the set of vertices. Then, for any subsets R and R' in S , we have $R \rightarrow R'$ if there is an arrow from some vertex in R to some vertex in R' .

If the sets R and R' overlap (i.e., $R \cap R' \neq \emptyset$), then clearly we have both $R \rightarrow R'$ and $R' \rightarrow R$. However, the examples above show that we can have $R \rightarrow R'$ even if R and R' do not overlap, as long as the two sets “touch” each other in some sense. Intuitively, $R \rightarrow R'$ means that it is not possible to interpose any “barrier” between R and R' ; there is no “gap” between them.

What role does the adjacency structure play in a spatially extended system? Why does space have one adjacency structure rather than another? Just as we argued earlier in the

⁵⁴ Adjacency structures are similar to proximity relations, which have been studied extensively in general topology (e.g., Willard 1970, Sections 40 and 41). But we do not assume that our adjacency structures satisfy the axioms of a proximity relation, so they are more versatile.

⁵⁵ Generally, an adjacency structure can be defined in a similar way on any metric or topological space.

⁵⁶ Generally, an adjacency structure can be defined in a similar way on any Cayley graph of any group.

⁵⁷ See Ilachinski (2001) and Moore and Mertens (2011).

case of time, we will now argue that a “correct” adjacency structure on space is one that satisfies a Markov property with respect to the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$. This Markov property is defined by considering conditional probabilities based on “partial information” about a spatially extended history.

We therefore need a precise way to talk about such “partial information”. Let R be a subset of S , and let $R \times T$ be the set of all ordered pairs (s, t) , where s is an element of R , and t is an element of T . So, $R \times T$ is the set of all time-slices restricted to the spatial region R . For any history h in Ω , recall that $h_{R \times T}$ denotes the restriction of h to the set $R \times T$. This restriction records only the part of the history h which “happens inside R ”. Let us then define the event $[h_{R \times T}]$ to be the set of all extensions of $h_{R \times T}$ to full histories in Ω , i.e., the set of all h' in Ω such that $h_{R \times T} = h'_{R \times T}$. These are precisely the histories that are accessible from h relative to the space-time region $R \times T$. The Markov property for adjacency structures will be based on conditional independence with respect to such events, in the following way.

For any event E (i.e., a subset of Ω), we say that E happens inside R if, for all histories h and h' such that $h_{R \times T} = h'_{R \times T}$, history h is an element of E if and only if history h' is an element of E . In other words, the question of whether or not a particular history is an element of E is completely determined by the restriction of that history to spatial “region” R .

A *tripartition* of S is a triple (R, R', R'') , where R, R' , and R'' are three disjoint subsets of S which together cover S (i.e., $R \cup R' \cup R'' = S$), such that it is *not* the case that $R \rightarrow R''$ or $R'' \rightarrow R$. Heuristically, this means that the set R' “separates” R from R'' . For example, suppose S is three-dimensional Euclidean space, with the adjacency structure described above. Let R be the set of all points whose distance to the origin is less than 1: the *unit ball*. Let R' be the set of all points whose distance to the origin is between 1 and 2, so R' is a sort of thick spherical “shell” around R . Finally, let R'' be the set of all points whose distance to the origin is greater than 2. Then (R, R', R'') is a tripartition of S .

We say that the adjacency structure \rightarrow satisfies the *Markov property* with respect to the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ if, for any tripartition (R, R', R'') and any history h in Ω , any event which happens inside R is conditionally independent from any event which happens inside R'' , given everything that happens in R' (i.e., given $[h_{R' \times T}]$). Heuristically, this means that there is no way for information to propagate from R into R'' , or vice versa, without first passing through R' . For example, suppose S is three-dimensional Euclidean space, and (R, R', R'') is the “concentric sphere” tripartition described above. In this case, the spherical shell R' acts as a barrier that isolates the ball-shaped compartment R from any influences coming from the “outer region” R'' . If we have complete information about the history inside R' (i.e., we know $[h_{R' \times T}]$), then we have complete control over the boundary conditions for any experiment we conduct inside R , and thus we do not need to control or even know what happens in the outer region R'' .

Scientists implicitly assume that space satisfies the Markov property every time they construct a laboratory apparatus that “isolates” some experiment from the surrounding environment. Indeed, people also implicitly assume the Markov property every time they close the doors and windows of their houses to keep out the cold. Thus, the Markov

property is central to the way we ordinarily think of space. It underpins the adjacency structure of space in the same way it underpins the order structure of time.

Just as with time, however, the Markov property does not *completely* determine the structure of space. First, there may be more than one adjacency structure on S which satisfies the Markov property with respect to $\{Pr_E\}_{E \subseteq \Omega}$, just as there may be more than one Markovian order on T . Second, the adjacency structure alone leaves many important geometric properties of S unspecified. For example, in many contexts, we would like to define a *metric* on S , which determines a notion of “distance” between points. This is obviously crucial in classical mechanics, for example. The adjacency structure does not determine a unique metric. We therefore now turn to the question of how we might arrive at such a metric.

3.9 Duration and distance

Recall that the set T of times is linearly ordered. In many contexts, we would like to define a notion of *duration* on T . That is, given four moments t_1, t_2, t_3 , and t_4 in T , with $t_1 < t_2$ and $t_3 < t_4$, we would like to determine whether the time interval between t_1 and t_2 is greater or smaller than that between t_3 and t_4 . To do this, we suppose that the monoid of temporal symmetries, Ψ , acts *freely and transitively* on T , and all symmetries in Ψ are *order-preserving*. This means that, for any times t_1 and t_2 in T , there is a *unique* symmetry ψ in Ψ such that $\psi(t_1) = t_2$, and, for any symmetry ψ in Ψ , $t_1 < t_2$ implies $\psi(t_1) < \psi(t_2)$. We can then define a formal “subtraction” operation on T as follows. Fix some reference time t_0 . Now, for any times t_1 and t_2 in T , we define

$$t_2 - t_1 = \psi(t_2), \text{ where } \psi \text{ is the unique temporal symmetry in } \Psi \text{ such that } \psi(t_1) = t_0.$$

In particular, this implies that $t - t_0 = t$, for any t in T . For any four points t_1, t_2, t_3 , and t_4 in T , we say that the time interval from t_1 to t_2 is greater than the one from t_3 to t_4 if $t_2 - t_1 > t_4 - t_3$. Similarly, we can define a formal “addition” operation on T . For any times t_1 and t_2 in T , we define

$$t_1 + t_2 = \psi(t_2), \text{ where } \psi \text{ is the unique temporal symmetry in } \Psi \text{ such that } \psi(t_0) = t_1.$$

The set T , with the ordering $<$ and the operation $+$, forms a *left-linearly ordered group*.⁵⁸

In many contexts, we would also like to define a *metric* on S , which determines a notion of “distance” between points in space. As we have noted, the adjacency structure does not determine a unique metric. But we can define a concept of distance on S by measuring how long it takes for information to travel from one point to the other. To do this, we need to use the concept of duration we have just introduced.

Given any two regions R and R' of S , and a time t in T , we define what it means for region R' to be “not reachable” from region R in time t . We begin with some preliminary definitions. For any subset R of S , and any time t in T , let $R \times \{t\}$ denote the set

⁵⁸ Formally, (i) the operation $+$ is *associative*, i.e., $(t_1 + t_2) + t_3 = t_1 + (t_2 + t_3)$ for all t_1, t_2, t_3 in T ; (ii) there is an *identity* element, namely t_0 , such that $t_0 + t = t = t + t_0$ for all t in T ; (iii) every element t in T has an *inverse* $-t$ such that $t + (-t) = t_0 = (-t) + t$; and (iv) the ordering $<$ is *left-homogeneous*, meaning that, for all t_1, t_2, t_3 in T , we have $(t_1 + t_2 < t_1 + t_3) \Leftrightarrow (t_2 < t_3)$. Left-linearly ordered groups are not generally *commutative* (“abelian”), i.e., we could have $t_1 + t_2 \neq t_2 + t_1$ for some t_1, t_2 in T . See, e.g., Fuchs (2011) for an introduction to ordered groups.

$\{(s,t): s \in R\}$. Adapting our earlier definition, we say that an event E happens inside R at time t if, for all histories h and h' such that $h_{R \times \{t\}} = h'_{R \times \{t\}}$, history h is an element of E if and only if history h' is an element of E . In other words, whether or not a particular history is an element of E is completely determined by the restriction of that history to the space-time region $R \times \{t\}$. Further, let R^C denote the *complement* of R in S , i.e., $R^C = \{s \in S: s \notin R\}$. Given any two subsets R and R' of S , and a time t in T with $t > t_0$, we now say that R' is *not reachable from R in time t* if, for any history h in Ω , any event which happens in R' at time t is conditionally independent of any event which happens in R at t_0 , given $[h_{R^C \times \{t_0\}}]$. Informally, once we have complete information about the state of the system *outside* the set R at time t_0 , learning something about the state of the system *inside* R at time t_0 gives us *no further information* about the eventual state inside R' at the later time t .⁵⁹

We now define the *distance* $d(R, R')$ between R and R' to be the *maximum* time t in T such that R' is not reachable from R in time t , if this maximum exists.⁶⁰ This can be interpreted as the *minimum time* required for information to “propagate” from R to R' . It would be natural to suppose that this notion of distance satisfies three properties:

Symmetry: For all subsets R, R' of S ,

$$d(R, R') = d(R', R).$$

Triangle Inequality: For all subsets R, R', R'' of S ,

$$d(R, R'') \leq d(R, R') + d(R', R'').$$

Non-Complementarity: For all subsets R_1, R_2, R_3 of S ,

$$d(R_1 \cup R_2, R_3) = \min\{d(R_1, R_3), d(R_2, R_3)\}.$$

However, none of these properties can be guaranteed, unless the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ has the right underlying properties. For example, if the information flow between different spatial locations is asymmetrical, such as in many communications networks, then Symmetry might not be satisfied; it might take longer for information to propagate from R to R' than vice versa. If information can be “forgotten” or “erased” at some spatial locations in the system, then the Triangle Inequality might not be satisfied; some information propagating from R to R' might be forgotten before it reaches R'' . Turning to Non-Complementarity: it is always true that $d(R_1 \cup R_2, R_3) \leq \min\{d(R_1, R_3), d(R_2, R_3)\}$. However, this inequality could be strict; i.e., we could have $d(R_1 \cup R_2, R_3) < \min\{d(R_1, R_3), d(R_2, R_3)\}$. For example, what happens in regions R_1 and R_2 at time t_1 could be like two pieces of a puzzle, which

⁵⁹ In our definition of “non-reachability”, we have referred to the reference time t_0 . However, because Ψ acts freely and transitively on T , the reference time does not matter. When region R' is not reachable from region R in time t , this implies that, for *any* times t_1 and t_2 with $t_2 - t_1 = t$, any event which happens in R' at time t_2 is conditionally independent of any event which happens in R at t_1 , given $[h_{R^C \times \{t_1\}}]$.

⁶⁰ If the maximum does not exist, we can use the supremum, provided the order of time is *supremum-complete* (i.e., any subset of T has a supremum), as it is if T is the set of real numbers. If the order of time is not supremum-complete, then the precise distance between R and R' may not be well-defined.

reveal little about what happens in region R_3 at time t_2 when considered separately, but determine it completely when put together.⁶¹

Note that our definition of distance between *regions* of space immediately entails a definition of distance between *points* in space: the distance between any two points s_1 and s_2 in S is simply the distance between the singleton regions consisting of them, i.e., $d(s_1, s_2) = d(\{s_1\}, \{s_2\})$. Clearly, $d(s, s) = 0$ for any point s in S . Thus, if our distance measure satisfies Symmetry and the Triangle Inequality, it determines a *metric* on the space S (or a *pseudo-metric* if $d(s_1, s_2) = 0$ for some $s_1 \neq s_2$). Furthermore, if it satisfies Non-Complementarity, this metric completely determines the distance between any two regions R and R' in S .⁶² However, as we have pointed out, the distance measure need not generally satisfy these properties.

One notable feature of the present approach is that it measures the distance between spatial locations in units of *time*. This is, of course, consistent with the practice in modern physics of measuring distance in units such as *light seconds* or *light years*. However, the approach works only if the maximum speed of information propagation in our system is finite. In classical physics, information can propagate through space at arbitrarily high speeds. Therefore, in a classical physical system, the effective “distance” between any two spatial locations collapses to zero, according to our definition. To recover a non-trivial definition of “distance” in such a system, we must impose some restriction on the sort of “information transmission” we can use. For instance, we could consider information transmission via some messenger or signal travelling at a fixed velocity. Similarly, in Maxwell’s theory of electrodynamics, which is complementary to classical mechanics, electromagnetic waves propagate at a fixed and finite speed, namely the speed of light, even if classical-mechanical particles can exceed this speed. Thus, in the world of classical physics, we could define a non-trivial concept of “electromagnetic distance”, even if there is no non-trivial concept of “mechanical distance”. We discuss the issue of distance in quantum mechanics in Appendix D.

4. Amorphous systems: space-time as an emergent property

4.1 Basic definitions

So far, we have defined histories as functions from a set of points in time or space-time into some state space, where histories play the role of possible worlds. Time or space-time, in turn, had an exogenously given structure. In a temporally evolving system, time was some linearly ordered set (T), and in a spatially extended system, space-time was explicitly decomposed into space (S) and time (T), consistent with some fixed geometry. This picture can, and for many purposes must, be generalized. Both special and general relativity theory, for example, go against the idea that there exists a fixed temporal dimension (for a classic discussion, see Putnam 1967).

⁶¹ Technically, this means that there exist events E_1, E_2 , and E_3 in Ω which happen, respectively, in region R_1 at time t_1 , in region R_2 at time t_1 , and in region R_3 at time t_2 such that E_1, E_2 , and E_3 are *pairwise independent*, but not *jointly independent*. This situation is common in probability theory.

⁶² To be precise, $d(R, R') = \min\{d(s, s') : s \in R \text{ and } s' \in R'\}$. Strictly speaking, this only works if R and R' are *finite* sets of points. For *infinite* sets, we would need a slightly stronger version of non-complementarity, which says that $d(R, R') = \inf\{d(s, s') : s \in R \text{ and } s' \in R'\}$ (and this infimum exists).

A more general approach is to define a history as a function from some “index set”, which we call a *set of loci*, into a state space. The set of loci could be a linearly ordered set of points in time, thereby accommodating temporally evolving systems, or a set of space-time locations with an explicit decomposition into space and time, thereby accommodating spatially extended systems. But it could also be a more general four-dimensional space-time manifold without any exogenous decomposition, or even a completely abstract index set.

Formally, let I (for “index set”) be the *set of loci*, and let X denote the *state space*. A *generalized history* is a function h from I into X , where, for each locus i in I , $h(i)$ is the state of the system at locus i . As in the case of spatially extended systems, the state $h(i)$ is best interpreted, not as a “global” state in which the system is at some specific point in time (indeed, there is no exogenous notion of time), but as a “local” state in which the system is at a specific locus. We write Ω to denote the set of all generalized histories deemed *possible*, which can again be viewed as nomologically possible worlds, and subsets of Ω are called *events*.⁶³

To complete the definition of what we call an *amorphous system*, we must, once more, introduce a *conditional probability structure* on Ω . As should be clear by now, this is a family of conditional probability functions $\{Pr_E\}_{E \subseteq \Omega}$, consisting of one Pr_E for each event E in Ω . Now an *amorphous system* is the pair consisting of the set Ω of nomologically possible generalized histories and the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$.

How much of our earlier framework can be extended to amorphous systems? We might ask, for instance, whether an abstract index set, despite not being endowed with any exogenous structure, can attain some spatial and/or temporal structure as an emergent property, for instance as a byproduct of the correlations encoded in $\{Pr_E\}_{E \subseteq \Omega}$. We might also ask whether, and to what extent, the geometry of the set of loci is unique, or whether there might be multiple, equally admissible geometries.

4.2 Adjacency structure and the Markov property

Just as in Section 3.8, the topology of the set I of loci can be represented by an *adjacency structure*: a binary relation \rightarrow defined between subsets of I . For example, suppose I is a set of times, as in Section 2, i.e., $I = T$. For any subsets R and R' of I , define $R \rightarrow R'$ if there does not exist any time t such that $r < t < r'$ for all r in R and all r' in R' . For another example, let I be the four-dimensional space-time manifold of a general relativistic system. Then, for any subsets R and R' of I , we might define $R \rightarrow R'$ if there is a locus i in R such that every open neighbourhood around i intersects R' .

In Section 2.9, we related the order structure of the set T of times to the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ by means of a temporal Markov property. Likewise, in Section 3.8, we related the adjacency structure of the set S of spatial locations to the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ by means of a spatial Markov property. We now discuss a similar idea concerning a general set of loci. This will allow us to view

⁶³ Note that, in the literature on general relativity theory, the word “event” is used to refer to the objects we call “loci”. Our use of the word “event” is consistent with its use in probability theory.

the adjacency structure among loci, and thereby its topology, as an “emergent property”: something that emerges from the correlations between events encoded in $\{Pr_E\}_{E \subseteq \Omega}$.

Let R be a subset of I (i.e., a collection of loci). As before, for any generalized history h in Ω , we define h_R to be the restriction of that history to the set R . We then define the event $[h_R]$ to be the set of all histories h' in Ω such that $h_R = h'_R$. For any event E (i.e., a subset of Ω), we say that E happens inside R if, for all histories h and h' such that $h_R = h'_R$, history h is an element of E if and only if history h' is an element of E . That is, whether or not a particular history is an element of E is completely determined by the restriction of that history to R .

As in Section 3.8, we define a *tripartition* of the set I of loci as a triple (R, R', R'') , where R , R' , and R'' are disjoint subsets of I which together cover I (i.e., $R \cup R' \cup R'' = I$), such that it is *not* the case that $R \rightarrow R''$ or $R'' \rightarrow R$. Again, this means that the set R' “separates” R from R'' .

For example, let I be a set of times ($I = T$) with the adjacency structure introduced at the start of this section. Fix two times t_0 and t_1 with $t_0 \leq t_1$. Let R be the set of all times *strictly before* t_0 , let R' be the set of all times *between* t_0 and t_1 (including t_0 and t_1), and let R'' be the set of all times *strictly after* t_1 . Then (R, R', R'') is a tripartition of I .

For another example, let I be the four-dimensional Minkowski space-time of special relativity, with the “open neighbourhood” adjacency structure introduced above. Let λ be a linear time-like trajectory through I , for instance the trajectory of an “observer” traveling through space-time at a constant velocity, and let p be a point on this trajectory. In special relativity theory, there is a unique three-dimensional *simultaneity hyperplane* R' passing through p , such that all events that happen inside R' seem to occur simultaneously from the perspective of the λ -observer at p . Let R be the set of all points in I which have some part of R' in their *future* light-cone, and let R'' be the set of all points in I which have some part of R' in their *past* light-cone. Then (R, R', R'') is a tripartition of I .⁶⁴ More generally, let R and R'' be any disjoint open subsets of I ,⁶⁵ and let R' be the complement of the union $R \cup R''$. Then (R, R', R'') is a tripartition of I .

We say that the adjacency structure \rightarrow satisfies the *amorphous Markov property* with respect to the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ if, for any tripartition (R, R', R'') and any generalized history h in Ω , any event which happens inside R is conditionally independent from any event which happens inside R'' , given $[h_R]$. Again, this means, roughly, that there is no way for information to propagate from R into R'' , or vice versa, without first passing through R' . For example, suppose I is four-dimensional Minkowski space-time, and (R, R', R'') is the tripartition described above. In this case, the simultaneity hyperplane R' plays the role of the “present”, which isolates the “past” R from the “future” R'' . If we have complete information about the history inside R' (i.e., we know $[h_{R'}]$), then we have complete information about the “present state” of the world. Thus, we can predict its future evolution (in R'') without needing to know anything about its past history (in R).

⁶⁴ In a model of general relativity, a similar construction works if R' is a Cauchy surface in the four-dimensional space-time manifold.

⁶⁵ A subset R of I is *open* if, for any s in R , there is some $r > 0$ such that the ball of radius r around s is contained in R .

In Section 2.9, we argued that the temporal Markov property was the key property of time; a “correct” ordering of the set T was any ordering that satisfied this property. Likewise, in Section 3.8, we argued that the spatial Markov property was the key property of space; a “correct” adjacency structure on the set S was any adjacency structure that satisfied this property. Now we make a parallel claim for amorphous systems: a “correct” adjacency structure on I is one that satisfies the amorphous Markov property. This Markov property subsumes both the temporal Markov property of Section 2 and the spatial Markov property of Section 3.

This has an important consequence. The topology of the index set I , in the form of the adjacency structure, does not need to be imposed exogenously. Instead, this topology can emerge *endogenously* from the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$. We say that an adjacency structure \rightarrow between subsets of I is $\{Pr_E\}_{E \subseteq \Omega}$ -*admissible* if it satisfies the amorphous Markov property with respect to $\{Pr_E\}_{E \subseteq \Omega}$. If we think of I as a sort of generalized space-time, this means that the topology of space-time is an emergent property of the amorphous system.⁶⁶

4.3 Time and predictability

Both temporally evolving systems and spatially extended systems come with a set T which plays the role of “time”. What plays the role of time in an amorphous system? The adjacency structure described in the previous section tells us whether two subsets of the index set I are in “informational contact” or are “informationally separated” from one another, but it does not tell us which subset comes “before” and which comes “after”, or even whether this question makes sense. We now explain how time itself can be an emergent property of an amorphous system.

Let \rightarrow be an adjacency structure on the index set I . Let T be a linearly ordered set. A *possible time structure* on I is a function τ from I onto T (i.e., with $T = \tau(I)$) such that, for any t in T , if (i) R is the set of all points i in I such that $\tau(i) < t$, (ii) R' is the set of all points i in I such that $\tau(i) = t$, and (iii) R'' is the set of all points i in I such that $\tau(i) > t$, then (R, R', R'') is a tripartition of I . Heuristically, the function τ specifies, for each locus in I , the time at which that locus occurs, according to the given time structure.

For example, let I be four-dimensional Minkowski space-time as described in Section 4.2, and let λ be a linear time-like trajectory through I . Fix some point p_0 on the trajectory λ . Let T be the set of real numbers. Then, for every t in T , there is a unique point p_t along the trajectory λ which appears to be t seconds in the future of p_0 (or in the past, if $t < 0$), with respect to the subjective time (i.e., *proper time*) experienced by an observer traveling along the trajectory λ . Let R_t be the simultaneity hyperplane

⁶⁶ We are not the first to suggest that the geometry and/or topology of space-time could be an emergent property of more fundamental causal structures. Brown and Pooley (2001, 2006) have argued that the geometry of relativistic space-time should be seen as a consequence of the symmetries (i.e., Lorentz covariance) of the dynamical laws governing matter and electromagnetism. In their words (2006, Section 5): “space-time’s Minkowskian structure cannot be taken to explain the Lorentz covariance of the dynamical laws. From our perspective, of course, the direction of explanation goes the other way around. It is the Lorentz covariance of the laws that underwrites the fact that the geometry of space-time is Minkowskian.” See also Brown (2005). However, Brown and Pooley’s approach is very different from ours. The idea of emergent space-time geometry also appears in the literature on high-energy physics and quantum cosmology. See, e.g., Konopka et al. (2008) and Hamma et al. (2010).

passing through p_t . If we define $\tau(i) = t$ for all points i in R_t , then τ is a possible time structure on I .

As this example illustrates, an amorphous system may admit *many* possible time structures. In special relativity, there is a distinct time-structure for every inertial reference frame. All of these time structures are equally “correct”. Indeed, this is one of the key insights of special relativity theory. However, unless we impose further constraints, a system may also admit many “absurd” time structures. For example, suppose I is four-dimensional Newtonian space-time (i.e., $I = \mathbf{R}^3 \times \mathbf{R}$), with the “open neighbourhood” adjacency structure described in Section 4.2. For all points (s_1, s_2, s_3, t) in I , define $\tau(s_1, s_2, s_3, t) = s_3$. Then τ is a possible time structure on I . But if the “true” time coordinate is t , not s_3 , it seems that this time structure is not correct. So, what property of the system determines which time structures are the correct ones?

Clearly, a “correct” time structure should satisfy something like the temporal Markov property from Section 2. However, if the adjacency structure \rightarrow satisfies the amorphous Markov property with respect to the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$, then it is easy to see that *any* possible time structure will satisfy the temporal Markov property.⁶⁷ So, the Markov property alone is not enough to pick out the “correct” time structures.

Arguably, what picks out the correct time structures is *predictability*. To understand this, suppose we took a classical mechanical system with Newtonian space-time $I = \mathbf{R}^3 \times \mathbf{R}$, and applied the “absurd” time structure $\tau(s_1, s_2, s_3, t) = s_3$, as defined above. How would the system appear with respect to this time structure? It would appear very strange and unpredictable. Particles would randomly pop in and out of existence. Energy and momentum would not be conserved from one moment to the next. Events would seem to unfold over time without any rhyme or reason. This total lack of predictability would be an indication that we had picked the wrong time structure for the system.

On the other hand, if we had picked the “correct” time structure, namely $\tau(s_1, s_2, s_3, t) = t$, then the system would appear completely deterministic; its state at one “moment” in time, as defined by τ , would completely determine its “past” and “future” behaviour, as defined by τ . This total predictability is an indication that this is the correct time structure for the system.

In this example, there was a particularly stark contrast between an “incorrect” time structure, which renders the system totally unpredictable, and a “correct” one, which renders it totally predictable. This is because classical mechanical systems are deterministic. In an indeterministic system, there will not generally be such a stark contrast. Nevertheless, some time structures will render the system *more* predictable than others, and among these, we claim, the ones that render the system *most predictable* are the correct time structures for that system.

To make this idea more precise, we need a way to measure the “predictability” of a system under a given time structure. One way to do this is to use the information-

⁶⁷ To be more precise: given a possible time structure on I , we can represent the amorphous system as a temporally evolving system, and this system, in turn, will satisfy the temporal Markov property. The construction is straightforward, but to avoid too many technicalities, we set aside the details here.

theoretic notion of *entropy*.⁶⁸ For any subset R of I , let Ω_R be the set of all R -restricted histories h_R obtained from any h in Ω . For simplicity, let us assume that the underlying state space X is finite. If R' is some other finite subset of I , then $\Omega_{R'}$ is also finite.⁶⁹ Suppose we know h_R , and we want to predict $h_{R'}$. For any h_R in Ω_R , there is a quantity called the *conditional entropy of R' given h_R* , denoted by $\eta(R', h_R)$, which measures how “unpredictable” the restricted history $h_{R'}$ is, given the restricted history h_R .⁷⁰ For example, if $h_{R'}$ is entirely determined by h_R , then $\eta(R', h_R) = 0$. At the other extreme, if $h_{R'}$ is effectively as unpredictable as a collection of independent coin-tosses, even after conditioning on h_R , then $\eta(R', h_R) = 1$. Intermediate levels of entropy represent intermediate degrees of unpredictability.

Now, let τ be a time structure, mapping I into T . Let t be some time in T ; let R be the set of all points i in I such that $\tau(i) = t$; and let R^C be the set of all points i in I such that $\tau(i) \neq t$. We define $\eta(\tau, t)$, the *unpredictability* of the system under τ at t , to be the maximum value of $\eta(R', h_R)$, where h_R can be any element of Ω_R and R' is allowed to be any finite subset of R^C .⁷¹ If $\eta(\tau, t) = 0$, then this means roughly that any generalized history h in Ω is almost entirely predictable, based on its restriction h_R .⁷² If $\eta(\tau, t) > 0$, then histories in Ω are not, in general, fully predictable from their restrictions to R . The larger $\eta(\tau, t)$ is, the less predictable these histories are. We then define $\eta(\tau)$, the *unpredictability* of the system under the time structure τ , to be the maximum value of $\eta(\tau, t)$ over all times t in T .⁷³

For example, suppose I is the four-dimensional Newtonian space-time of a classical mechanical system (i.e., $I = \mathbf{R}^3 \times \mathbf{R}$), and τ is the “correct” time structure for this system, namely $\tau(s_1, s_2, s_3, t) = t$. Then $\eta(\tau) = 0$, because classical mechanics is entirely deterministic. However, if τ was an “incorrect” time structure, such as $\tau(s_1, s_2, s_3, t) = s_3$, then we would have $\eta(\tau) > 0$, because the ascription of this incorrect time structure would render the system unpredictable, as we have explained.

We now come to the key point of this section. A *correct* time structure for an amorphous system is one that *minimizes* unpredictability and thereby maximizes regularity. This definition allows that there may be many correct time structures, all of which render the system equally predictable, as in the case in special or general relativity. This has an important consequence. The correct time structure does not need to be imposed

⁶⁸ This is *not* the same as thermodynamic entropy, although it is loosely related. Thus, the discussion that follows should *not* be interpreted thermodynamically.

⁶⁹ If $|X|$ and $|R'|$ are the cardinalities of X and R' , then the cardinality of $\Omega_{R'}$ is at most $|X|^{|R'|}$.

⁷⁰ Formally, $\eta(R', h_R)$ is the sum, over all possible R' -restrictions $h_{R'}$ in $\Omega_{R'}$, of

$$-Pr([h_{R'}] | [h_R]) \log_2[Pr([h_{R'}] | [h_R])] / |R'| \log_2(|X|).$$

However, the precise formula is not important for this discussion.

⁷¹ To be more precise, it is the *supremum* of this set. The maximum is not always well-defined.

⁷² Even if $\eta(\tau, t) = 0$, there may be some “residual” unpredictability, in that Ω may contain more than one extension of h_R to all of I . However, the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ concentrates all probability on *one* of these possible extensions; the rest of the extensions get probability zero.

⁷³ Again, strictly speaking, we require the supremum of $\eta(\tau, t)$ across all t . This is not the only possible measure of the system’s unpredictability under time structure τ . We could also take the average or some other aggregate measure. For example, suppose that I is an N -dimensional integer lattice (i.e., $I = \mathbf{Z}^N$). Then we could measure the system’s unpredictability under different time structures using the theory of *entropy geometry* and *expansive subdynamics* first developed for multidimensional cellular automata by Milnor (1988) and later extended to arbitrary multidimensional symbolic dynamical systems by Boyle and Lind (1997). See the section on “Entropy” in Pivato (2009) for a summary.

exogenously. Instead, the correct time structure (or structures) could emerge *endogenously* from the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$. In other words, the structure of time itself could be an emergent property of the amorphous system. Using a more metaphysical language, it might be that space and time are grounded in the dynamics of the system, rather than the other way around.

4.4 Which features of a system are real?

A final philosophical question on which we wish to comment briefly is the following. Suppose we have described a given system using our formal framework. Should we treat all features of that system as “real”, or should we treat some features as mere artefacts of our formal description?

The debates between *relational* and *substantival* views about space and time, and between *structuralist* and *full-blown realist* views in science more generally, can be seen as attempts to answer this question.⁷⁴ Let us begin with a *relational* or *structuralist* view, which may be about space and time in particular or about the properties of a system more generally. On such a view (of which there can be several variants), only some “relational” or “structural” properties of a system count as real, while “intrinsic”, “non-structural” properties do not. It does not matter, for instance, what the nature of the system’s spatiotemporal loci in the set I is, nor what the nature of the system’s possible states in the set X is. All that matters is how these loci and/or states are related to one another and what dynamics they display. Two formally distinct systems, with formally distinct index sets I and I' and/or formally distinct state spaces X and X' , will count as the same if their nomologically possible histories and conditional probability structures are structurally indistinguishable.

By contrast, on a *substantival* or *full-blown realist* view, which may also be about space and time in particular or about the properties of a system more generally, even intrinsic, non-structural properties of a system can be real, over and above the system’s relational or structural properties. So, the system’s spatiotemporal index set I and its state space X may be significant in ways that go beyond the structures and relations in which they stand. (Again, there can be several variants of such a view.) An example of a non-structural property is the exact index of time. One can imagine two structurally identical temporally evolving systems, indexed by $T = \{0, 1, 2, 3, \dots\}$ and $T' = \{1, 2, 3, 4, \dots\}$, respectively. The only difference is that in one system history “starts at time zero”, whereas in the other it “starts at time 1”. For a relationalist or structuralist, these are “the same” system. But a substantivalist or full-blown realist might insist that there is a genuine difference between them.

The debates between these different views occur in several places in philosophy and take a variety of forms, so we cannot do justice to them here. We wish to note, however, that our formal framework can be used to express some salient positions within those debates. Specifically, different answers to the question of which features of a system are real can be expressed in terms of different criteria for individuating systems. If we begin with a very large class of systems that are formally described in our framework, there are a number of ways in which one might partition this class of systems into

⁷⁴ On a broadly “structuralist” or “relational” approach to metaphysics, see, e.g., Ladyman and Ross (2009). On “absolute” versus “relational” accounts of space and time, see, e.g., Earman (1989). On “substantivalism” and its discontents, see, e.g., Nerlich (2003).

equivalence classes that are each taken to represent the same system. Different such partitions then correspond to different answers to the question of which features of a system are real, rather than mere artefacts of our formal description. In particular, only those features that are present among *all* members of any given equivalence class count as real. Features on which there can be differences even *within* the same equivalence class count as artefacts of our formal description.

A relational or structuralist view would entail that any two systems that do not differ in any relational or structural properties count as the same and thereby fall into the same equivalence class. A substantival or full-blown realist view, by contrast, would entail that two such systems could still count as different; thus, the equivalence classes would be more fine-grained according to such a view, and might even be singleton (in which case *all* features of any given system would count as real).

Here is one way of formalizing this idea. Consider two amorphous systems, given by the pairs $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$ and $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$, where the histories in Ω are functions from the set I of loci into the state space X , and the histories in Ω' are functions from the set I' of loci into the state space X' . Let \mathcal{H} and \mathcal{H}' denote the sets of logically possible functions from I into X and from I' into X' , respectively.

Suppose there is a bijection θ from I into I' , and also a bijection ξ from X into X' (recall that a *bijection* is a one-to-one, onto function). Using θ and ξ , we can then define a bijection σ from \mathcal{H} into \mathcal{H}' which maps each history h in \mathcal{H} to the history h' in \mathcal{H}' defined as follows: for each i' in I' ,

$$h'(i') = \xi[h(i)], \text{ where } i = \theta^{-1}(i') \text{ (with } \theta^{-1} \text{ defined as the inverse of } \theta).$$

The bijection σ is an *isomorphism* between the two systems if

- $\sigma(\Omega) = \Omega'$; and
- for any events E' and D' in Ω' , if E and D are the inverse images of E' and D' under σ , then $Pr'_{E'}(D') = Pr_E(D)$.

We call two systems *isomorphic* if there exists an isomorphism between them. Isomorphic systems display the same dynamics, and they are relationally or structurally indistinguishable.⁷⁵ Moreover, any topology of space and time that is admissible for one such system can be mapped, in a structure-preserving way, onto a topology that is admissible for the other.

Thus, on a relational or structuralist view, any two isomorphic systems should be considered the same. On a substantival or full-blown realist view, they may still differ. A view of the first kind would therefore take systems to be unique only up to isomorphism, so that our initial large class of systems would be partitioned into equivalence classes of isomorphic systems. A view of the second kind would opt for a more fine-grained partition, acknowledging that even isomorphic systems may be distinct in reality.

The properties of systems on which we have focused in this paper are mainly structural and are preserved by all isomorphisms. This includes the symmetries and ergodic

⁷⁵ In fact, any bijective symmetry of a system constitutes an isomorphism from a system into itself.

properties of a system, the distinction between laws and “brute necessities”, and the topology (or topologies) and geometry (or geometries) of space and time that are compatible with the system’s correlation structure (in the sense that they satisfy the relevant Markov conditions). Thus, even a relationalist or structuralist would accept that all these properties are “real” features of the system, and not mere artefacts.

5. Concluding discussion

We have introduced a framework for describing three general classes of systems and shown how it can be used to address a number of philosophical questions. We began with the class of temporally evolving systems, of which classical dynamical systems are a special case, and then moved on to the class of spatially extended systems and the class of amorphous systems. As noted, the framework can accommodate systems as diverse as the solar system, quantum-mechanical systems, special and general relativistic systems, and the earth’s climate system.

We have discussed questions such as: how can we define nomological possibility, necessity, determinism, and indeterminism? What is special about laws, and how are laws related to symmetries? What regularities must a system display to permit global generalizations from local observations? How can we formulate principles of parsimony such as Occam’s Razor? What is the role of space and time in a system? And what is at stake in the debate between relational and substantival views about space and time, and between structuralist and full-blown realist views about systems more generally?

While our framework and what it says about these questions should already be of sufficient interest to make it worth studying, a further payoff lies arguably in the variety of applications to which the framework lends itself. Developing these is beyond the scope of this paper, but we conclude by mentioning a few.

5.1 Higher-level versus lower-level properties

Our framework can be used to explore the relationship between lower-level (“micro”) and higher-level (“macro”) properties of a system. By partitioning the system’s state space X into suitable equivalence classes, we can capture the idea that “higher-level” or “macro” states are more coarse-grained than “lower-level” or “micro” states, so that each “macro” state can be realized by different “micro” states: the phenomenon of *multiple realizability*. Consider, for example, all the different possible micro-level trajectories of a tossed coin that each correspond to the macro-property of “landing heads”. Or consider all the different possible micro-states of individual water molecules that each correspond to a macro-state such as “frozen”, “liquid”, or “gaseous”.

Suppose X is the original state space, and \mathbf{X} is the relevant set of equivalence classes, which we interpret as the higher-level state space. We can then write σ to denote the function that maps each lower-level state x in X to the corresponding higher-level state \mathbf{x} in \mathbf{X} . Note the outlined font for higher-level objects. This function can be interpreted as the *supervenience relation* connecting the two levels. We can then use σ to specify the resulting higher-level histories.⁷⁶ For each lower-level history h in the original set Ω , the corresponding higher-level history \mathbf{h} is the function from T into \mathbf{X} , where, for

⁷⁶ This construction, under the present notational conventions, was introduced in List (2014) and List and Pivato (2015). Relatedly, see also Yoshimi (2012).

each t in T , $\mathbf{h}(t) = \sigma(h(t))$. (If we are dealing with a spatially extended or amorphous system instead of a temporally evolving one, we must replace T in this definition with $S \times T$ or I .) The set of higher-level histories is therefore $\Omega = \sigma(\Omega)$. Similarly, we can use σ to arrive at a conditional probability structure defined over higher-level events, formally written $\{Pr_E\}_{E \subseteq \Omega}$. See Appendix A for details. The pair $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$ can be viewed as our system, *re-described at a higher level*. In the terminology of Appendix A, the higher-level system is a *factor system* of the original, lower-level system.

This construction allows us to study the dynamics of the higher-level system and to compare its properties with those of the lower-level system. Interestingly, the higher-level dynamics may be different from the underlying lower-level dynamics. For example, features such as determinism or indeterminism are not generally preserved under coarse-graining: the lower-level system may be deterministic, while the higher-level system is not (or vice versa). Thus indeterminism could be an emergent property (see, e.g., Butterfield 2012, List 2014, List and Pivato 2015, and relatedly Werndl 2009b).

In a similar vein, we may study the level-specificity of other properties. For instance, this approach can be used to argue that non-trivial objective chance can be an emergent phenomenon, consistently with lower-level determinism (List and Pivato 2015).⁷⁷

5.2 Laws and regularities in the special sciences

There is much debate on whether there are laws in the special sciences, as distinct from fundamental physics. The existence of laws is particularly contested in fields such as biology, ecology, geology, psychology, and the social sciences. (Chemistry, by contrast, is often viewed as a close relative of physics and thereby similar enough to it in its lawfulness.) Examples of special-science regularities that are sometimes described as laws include (i) Kleiber’s law in biology, according to which an organism’s metabolic rate is proportional to the $\frac{3}{4}$ th power of its body mass; (ii) the laws of supply and demand in economics, according to which (except for Giffen goods) the demand for a good is a decreasing function of its price, and the supply is an increasing function of price; and (iii) Duverger’s law in political science, according to which, under a first-past-the-post electoral system, the effective number of parties in the legislature will be lower than under a proportional-representation system, *ceteris paribus*. The key question is whether any of these regularities are sufficiently robust to qualify as laws.

One common view is that, as we move further away from fundamental physics, there are fewer and fewer regularities that are genuinely law-like. Kim (2010, ch. 14), for instance, argues that there are no “strict” laws in the special sciences. Among the reasons he gives for this conclusion are (i) the multiple realizability of special-science properties, which, he claims, undermines their “inductive projectibility”, and (ii) the alleged metaphysical anomalism of the mental realm, which, he suggests, undermines the existence of laws in psychology and the social sciences.

Other scholars defend the existence of laws in the special sciences. For example, focusing on the social sciences, Kincaid (1990) argues that several widely cited arguments against laws fail. He thinks that the most serious challenge to laws in the social sciences comes from the excessive *ceteris paribus* qualifications that all such

⁷⁷ For earlier work defending higher-level chance, sometimes using a strategy similar in spirit to the present one (though not fully equivalent), see, e.g., Loewer (2001), Frigg and Hoefer (2010), Glynn (2010), Strevens (2011), and Hemmo and Shenker (2012).

laws require, but argues that the procedures we routinely employ to deal with such qualifications in the natural sciences carry over to the social sciences.

Our framework might be used to make some progress in this debate. Using the framework, we can in principle describe the special-science systems in question and identify the properties these systems would have to display in order to secure the existence of laws. Those laws would then have the testable and generalizable character we have discussed. As we have seen, what laws there are in a given system depends on the system's symmetries and the properties they preserve. This is as true for a system in the special sciences as it is for a physical system. Moreover, our analysis implies that whether, given only local observations, we can gain knowledge of the probabilistic dynamics of a special-science system depends on whether the system is ergodic. The importance of ergodicity for the special sciences is much less widely recognized than its importance for physics.

Interestingly, if a special-science system arises as a higher-level description of a physical system, as discussed in Section 5.1, then it will inherit some structure from the physical system, and it will have at least as many temporal or spatiotemporal symmetries as that physical system (and possibly more), and at least as much ergodicity, for reasons explained in Appendix A. Another question is whether we are prepared to recognize weaker kinds of laws corresponding to *partial* or *local* symmetries, as defined in Appendix B. This question is particularly pertinent for the special sciences, insofar as the systems investigated in fields ranging from biology to the social sciences often have special initial or boundary conditions. While all of these issues are difficult, our framework can help us clarify what is at stake in the debate about special-science laws and thereby render the debate more tractable. For earlier applications of dynamical-systems theory to the special sciences, see Auyang (1998) and Yoshimi (2012).

5.3 *Intentional systems*

Although there has been no such thing as intentionality in our paradigmatic examples of systems, there is no barrier, in principle, to using our framework also for describing systems involving intentional agents. Indeed, van Gelder (1995) and Juarrero (1998) have argued for understanding cognitive systems as special kinds of dynamical systems (see also Spivey 2008, Hotton and Yoshimi 2010, and Silberstein and Chemero 2012); and more recently, a precursor of the present formalism has proved useful for the analysis of free will and agency (List 2014, List and Rabinowicz 2014). We can think of an agent, together with the relevant environment, as a temporally evolving system. This system can be described at different levels: at a physical level, at which we would *not* take an “intentional stance” towards the system, and at an agential level, at which we *would* take such a stance (on the notion of an “intentional stance”, see Dennett 1987). Physical-level descriptions capture the states of the agent's brain and body, while agential-level descriptions capture the agent's higher-level mental or psychological states, thereby focusing on the agent's beliefs, desires, and intentions, rather than the underlying neuronal or bodily states.

The present framework then allows us to explain, for instance, how agential-level indeterminism and an agent's *possibility of doing otherwise* can co-exist with physical-level determinism (List 2014). The framework might also shed some light on how other

psychological properties can emerge from the underlying physical dynamics of the system. In particular, as a factor system of the original physical system, the agential system may exhibit additional symmetries not present at the physical level – a point already alluded to in Section 5.2. This may, in turn, be used to explain why some higher-level regularities in an intentional system (e.g., regularities involving beliefs, desires, intentions, and norms) may qualify as “real patterns”, as Dennett (1991) has argued, and not merely as illusions due to our ignorance of the physical-level details.

Needless to say, all of these applications are challenging and raise controversial philosophical issues. We hope, however, that our framework will be a clarifying contribution to formal metaphysics and the philosophy of science and will inspire further work.

References

- Auyang, Sunny Y. (1998). *Foundations of Complex-system Theories in Economics, Evolutionary Biology, and Statistical Physics*. Cambridge: Cambridge University Press.
- Baker, Alan (2013). Simplicity. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Fall 2013 edition <<http://plato.stanford.edu/archives/fall2013/entries/simplicity/>>.
- Baker, David John (2010). Symmetry and the Metaphysics of Physics. *Philosophy Compass*, 5(12), 1157–1166.
- Berkovitz, Joseph, Roman Frigg, and Fred Kronz (2006). The ergodic hierarchy, randomness and Hamiltonian chaos. *Studies in History and Philosophy of Modern Physics*, 37(4), 661–691.
- Bishop, Robert C. (2011). Metaphysical and Epistemological Issues in Complex Systems. In Cliff Hooker (ed.), *Philosophy of Complex Systems, Volume 10 in Handbook of the Philosophy of Science* (pp. 105–136). Amsterdam: Elsevier.
- Boyle, Mike, and Douglas Lind (1997). Expansive subdynamics. *Transactions of the American Mathematical Society*, 349(1), 55–102.
- Brading, Katherine, and Elena Castellani (2013). Symmetry and Symmetry Breaking. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Spring 2013 edition <<http://plato.stanford.edu/archives/spr2013/entries/symmetry-breaking/>>.
- Brown, Harvey R. (2005). *Physical Relativity: Space-Time Structure from a Dynamical Perspective*. Oxford: Oxford University Press.
- Brown, Harvey R., and Oliver Pooley (2001). The origin of the spacetime metric: Bell’s ‘Lorentzian pedagogy’ and its significance in general relativity. In C. Callender and N. Huggett (eds.), *Physics Meets Philosophy at the Planck Length* (pp. 256–272). Cambridge: Cambridge University Press.
- Brown, Harvey R., and Oliver Pooley (2006). Minkowski space-time: a glorious non-entropy. In Dennis Dieks (ed.), *The Ontology of Spacetime* (pp. 67–89). Amsterdam: Elsevier.
- Butterfield, Jeremy (2012). Laws, Causation and Dynamics at Different Levels. *Interface Focus*, 2(1), 101–114.
- de Finetti, Bruno (1972). *Probability, induction and statistics. The art of guessing*. London: John Wiley & Sons.
- Dennett, Daniel (1987). *The Intentional Stance*. Cambridge/MA: MIT Press.
- Dennett, Daniel (1991). Real Patterns. *Journal of Philosophy*, 88(1), 27–51.
- Earman, John (1986). *A Primer on Determinism*. Dordrecht: D. Reidel.

- Earman, John (1989). *World Enough and Spacetime: Absolute and Relational Theories of Motion*. Cambridge, MA: MIT Press.
- Eberhardt, Frederick, and Clark Glymour (2009). Hans Reichenbach's probability logic. In Dov M. Gabbay, Stephan Hartmann and John Woods (eds.), *Handbook of the History of Logic. Volume 10: Inductive Logic* (pp. 357–389). Amsterdam: Elsevier.
- Fitzpatrick, Simon (2015). Simplicity in the Philosophy of Science. In *The Internet Encyclopedia of Philosophy*, ISSN 2161-0002 <<http://www.iep.utm.edu>>.
- French, Steven (2014). *The Structure of the World: Metaphysics and representation*. Oxford: Oxford University Press.
- Frigg, Roman, Joseph Berkovitz, and Fred Kronz (2011). The Ergodic Hierarchy. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Summer 2011 edition <<http://plato.stanford.edu/archives/sum2011/entries/ergodic-hierarchy/>>.
- Frigg, Roman, and Carl Hoefer (2010). Determinism and Chance from a Humean Perspective. In F. Stadler (ed.), *The Present Situation in the Philosophy of Science* (pp. 351–371). Heidelberg: Springer-Verlag.
- Fuchs, Laszlo (2011). *Partially ordered algebraic systems*. New York: Dover.
- Glasner, E. (2003). *Ergodic Theory via Joinings*. Providence, RI: American Mathematical Society.
- Glynn, Luke (2010). Deterministic Chance. *British Journal for the Philosophy of Science*, 61(1), 51–80.
- Golubitsky, Marty, Marcus Pivato, and Ian Stewart (2003). Symmetry groupoids and patterns of synchrony in coupled cell networks. *SIAM Journal of Applied Dynamical Systems*, 2(4), 609–646.
- Guay, Alexandre, and Brian Hepburn (2009). Symmetry and Its Formalisms: Mathematical Aspects. *Philosophy of Science*, 76(2), 160–178.
- Halpern, Joseph Y. (2010). Lexicographic probability, conditional probability, and nonstandard probability. *Games and Economic Behavior*, 68(1), 155–179.
- Hamma, A., F. Markopoulou, S. Lloyd, F. Caravelli, S. Severini, and K. Markström (2010). Quantum Bose-Hubbard model with an evolving graph as a toy model for emergent spacetime. *Physical Review D*, 81(10), 104032.
- Hemmo, Meir, and Orly R. Shenker (2012). *The Road to Maxwell's Demon: Conceptual Foundations of Statistical Mechanics*. Cambridge: Cambridge University Press.
- Hotton, Scott, and Jeff Yoshimi (2010). The Dynamics of Embodied Cognition. *International Journal of Bifurcation and Chaos*, 20(4), 943–972.
- Ilachinski, Andrew (2001). *Cellular Automata*. River Edge, NJ: World Scientific.
- Juarrero, Alicia (1999). *Dynamics in Action: Intentional Behavior as a Complex System*. Cambridge, MA: MIT Press.
- Kelly, Kevin T. (2007). A New Solution to the Puzzle of Simplicity. *Philosophy of Science* (Proceedings of the 2006 Biennial Meeting of the Philosophy of Science Association), 74(5), 561–573.
- Kim, Jaegwon (2010). *Essays in the Metaphysics of Mind*. Oxford: Oxford University Press.
- Kincaid, Harold (1990). Defending Laws in the Social Sciences. *Philosophy of the Social Sciences*, 20(1), 56–83.
- Konopka, T., F. Markopoulou, and S. Severini (2008). Quantum graphity: A model of emergent locality. *Physical Review D*, 77(10), 104029.
- Krengel, Ulrich (1985). *Ergodic theorems*. With a supplement by Antoine Brunel. Berlin: Walter de Gruyter & Co.
- Ladyman, James, and Don Ross (2009). *Every Thing Must Go: Metaphysics*

- Naturalized*. Oxford: Oxford University Press.
- Lewis, David (1986). A Subjectivist's Guide to Objective Chance. In *Philosophical Papers*, Vol. II (pp. 83–132). Oxford: Oxford University Press.
- List, Christian (2014). Free will, determinism, and the possibility of doing otherwise. *Noûs*, 48(1), 156–178.
- List, Christian, and Marcus Pivato (2015). Emergent Chance. *Philosophical Review*, 124(1), 119–152.
- List, Christian, and Wlodek Rabinowicz (2014). Two Intuitions about Free Will: Alternative Possibilities and Intentional Endorsement. *Philosophical Perspectives*, 28, 155–172.
- Loewer, Barry (2001). Determinism and Chance. *Studies in History and Philosophy of Science Part B: History and Philosophy of Modern Physics*, 32(4), 609–620.
- Mainzer, Klaus (1996). *Symmetries of Nature*. Berlin: Walter de Gruyter.
- Milnor, John (1988). On the entropy geometry of cellular automata. *Complex Systems*, 2(3), 357–385.
- Moore, Cristopher, and Stephan Mertens (2011). *The Nature of Computation*. Oxford: Oxford University Press.
- Nerlich, Graham (2003). Space-Time Substantivalism. In Michael J. Loux and Dean W. Zimmerman (eds.), *The Oxford Handbook of Metaphysics* (pp. 281–314). Oxford: Oxford University Press.
- Pearl, Judea (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.
- Petersen, Karl E. (1989). *Ergodic Theory*. Cambridge: Cambridge University Press.
- Popper, Karl (1968). *The Logic of Scientific Discovery*, second edition. London: Hutchinson.
- Pivato, Marcus (2009). The ergodic theory of cellular automata. In Robert A. Meyers (ed.), *The Encyclopedia of Complexity and System Science* (pp. 2980–3015). Heidelberg: Springer-Verlag.
- Pivato, Marcus (2010). *Linear Partial Differential Equations and Fourier Theory*. Cambridge: Cambridge University Press.
- Priest, Graham (2001). *An Introduction to Non-Classical Logic*. Cambridge: Cambridge University Press.
- Putnam, Hilary (1967). Time and Physical Geometry. *Journal of Philosophy*, 64(8), 240–247.
- Reichenbach, Hans (1949). *The Theory of Probability. An Inquiry into the Logical and Mathematical Foundations of the Calculus of Probability*. Transl. by E. H. Hutten and M. Reichenbach. Berkeley / Los Angeles: University of California Press.
- Rényi, Alfred (1955). On a new axiomatic theory of probability. *Acta Mathematica Acadademiae Scientiarum Hungarica*, 6(3–4), 285–335.
- Roberts, Bryan W. (2013). When We Do (and Do Not) Have a Classical Arrow of Time. *Philosophy of Science*, 80(5), 1112–1124.
- Schaffer, Jonathan (2007). Deterministic chance? *British Journal for the Philosophy of Science*, 58(2), 113–140.
- Sellars, Wilfrid (1948). Concepts as Involving Laws and Inconceivable without Them. *Philosophy of Science*, 15(4), 287–315.
- Silberstein, Michael, and Anthony Chemero (2012). Complexity and Extended Phenomenological-Cognitive Systems. *Topics in Cognitive Science*, 4(1), 35–50.
- Sobel, Jordan Howard (1998). *Puzzles for the Will: Fatalism, Newcomb and Samarra, Determinism and Omniscience*. Toronto: University of Toronto Press.
- Spirtes, Peter, Clark Glymour, and Richard Scheines (2000). *Causation, Prediction and*

- Search*, second edition. Cambridge, MA: MIT Press.
- Spivey, Michael (2008). *The Continuity of Mind*. New York: Oxford University Press.
- Strevens, Michael (2011). Probability out of Determinism. In Claus Beisbart and Stephan Hartmann (eds.), *Probabilities in Physics* (pp. 339–364). Oxford: Oxford University Press.
- van Fraassen, Bas C. (1976). Representational of conditional probabilities. *Journal of Philosophical Logic*, 5(3), 417–430.
- van Fraassen, Bas C. (1989). *Laws and Symmetry*. Oxford: Clarendon Press.
- van Gelder, Tim (1995). What Might Cognition Be, If Not Computation? *Journal of Philosophy*, 92(7), 345–381.
- Walters, Peter (2000). *An Introduction to Ergodic Theory*. Heidelberg: Springer.
- Werndl, Charlotte (2009a). What are the new implications of chaos for unpredictability? *British Journal for the Philosophy of Science*, 60(1), 195–220.
- Werndl, Charlotte (2009b). Are deterministic descriptions and indeterministic descriptions observationally equivalent? *Studies in History and Philosophy of Science Part B: History and Philosophy of Modern Physics*, 40(3), 232–242.
- Weyl, Hermann (1952). *Symmetry*. Princeton: Princeton University Press.
- Wigner, Eugene (1967). *Symmetries and reflections*. Bloomington, Indiana: Indiana University Press.
- Willard, S. (1970). *General Topology*. Reading, MA: Addison Wesley.
- Yoshimi, Jeffrey (2012). Supervenience, Dynamical Systems Theory, and Non-Reductive Physicalism, *British Journal for the Philosophy of Science*, 63(2), 373–398.

Appendix A: Factor systems

One possible objection to our framework is that it is both unrealistic and unwieldy. It is unrealistic because the actual universe is not sufficiently regular (e.g., it might lack a large enough monoid of symmetries or ergodicity). It is unwieldy because the universe as a whole is far too complex a system for us to analyze within this framework anyway.

However, there is no need to insist on applying our framework to the universe as a whole. Instead, we can apply it to a “factor” system, as we now explain. Consider a temporally evolving system, consisting of a set Ω of nomologically possible histories (each of which is a function from a set T of times into a set X of possible states), along with a conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$. Let X' be another set, and let σ be a function from X into X' . For every history h in Ω , let $\sigma(h)$ be the function h' from T into X' defined by $h'(t) = \sigma(h(t))$ for all t in T . Let $\Omega' = \{\sigma(h) : h \in \Omega\}$. For any subsets D' and E' of Ω' , let D and E be their inverse images under σ , where these are subsets of Ω , and define $Pr'_{E'}(D') = Pr_E(D)$. Then $\{Pr'_{E'}\}_{E' \subseteq \Omega'}$ is a conditional probability structure on Ω' . The system specified by Ω' and $\{Pr'_{E'}\}_{E' \subseteq \Omega'}$ is called a *factor system* of the original system specified by Ω and $\{Pr_E\}_{E \subseteq \Omega}$. The function σ from Ω into Ω' is called a *factor map*. We can define *factors* of spatially extended and amorphous systems analogously.

For a concrete example, suppose that $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$ is a classical-physics description of the entire solar system at an atomic level of detail. So X is an extremely high-dimensional space, which must specify the position and momentum of every atom in the entire solar system, along with all of their gravitational and electromagnetic

interactions.⁷⁸ Meanwhile, let $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$ be the very simple celestial-mechanical system consisting only of the Earth, Moon, and Sun, described as gravitationally interacting point masses. So $X' = \mathbf{R}^{18}$, because we must specify the three-dimensional position and momentum vectors of each of the three objects in the system, and $3 \times 6 = 18$. Let σ be the function from X into X' which translates each highly detailed atomic-level description of the solar system into the crude 18-dimensional celestial mechanical description. Then σ is a factor mapping from $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$ into $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$, and thus $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$ is a factor of $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$.

As this example illustrates, a factor system can be seen as a sort of “abstraction” or “simplification” of the original system, obtained by discarding some properties. Now, suppose ψ is a function from T into itself (e.g., a time shift) which is a temporal symmetry of the original system $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$. Then it is easy to verify that ψ will also be a temporal symmetry of the factor system $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$. Thus, the temporal symmetry monoid of the factor system $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$ is *at least as large* as the temporal symmetry monoid of the original system $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$. In a spatially extended system, the exact same statement applies to spatiotemporal symmetries. Furthermore, if Ψ is an amenable monoid of temporal (or spatiotemporal) symmetries, and $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$ is ergodic relative to Ψ , then $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$ will *also* be ergodic relative to Ψ . In other words, $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$ is *at least as ergodic* as $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$.

This means that, even if the original system $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$ lacks certain symmetries or ergodic properties, the factor system $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$ may well possess these properties. Furthermore, even if the original system $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$ is too complicated to analyze using the formal tools we have described, the system $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$ may well be simple enough. To illustrate this, consider our example of the solar system. The original system $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$ describes the entire solar system at an atomic level of detail. Whether or not the system possesses the desired symmetries or ergodic properties, it is certainly too complex to analyze. In contrast, the abstract Earth-Moon-Sun system $(\Omega', \{Pr'_{E'}\}_{E' \subseteq \Omega'})$ is very simple. In fact, it is an example of a *quasiperiodic* dynamical system: it can be described as two independently rotating “wheels”, one describing the orbit of the Moon around the Earth, and the other describing the orbit of the Earth around the Sun. This is a prototypical example of an ergodic dynamical system.

Appendix B: Partial symmetries and local symmetries

An important assumption of this paper has been that there is a fairly large monoid Ψ of symmetries acting on the set Ω of nomologically possible histories. We have argued that a property of the system qualifies as a “law” only if it is invariant under all of these symmetries. But this argument runs into a problem: many systems studied in the sciences lack sufficient symmetries to account for all of their “law-like” features.

For example, suppose space is represented by the set of all integers, while time is represented by the set of positive integers, i.e., $S = \{\dots, -1, 0, 1, 2, \dots\}$ and $T = \{1, 2, 3, \dots\}$, and consider the simple random-walk system described in Section 2.9. Nomologically speaking, the token could begin at any spatial location at time one. But suppose the conditional probability structure $\{Pr_E\}_{E \subseteq \Omega}$ is such that, with probability one, the token

⁷⁸ For simplicity, we eschew a quantum-mechanical description in this example.

begins at spatial location zero at time one.⁷⁹ In that case, the probability distribution of its location at time t is a $(t-1, 1/2)$ -binomial distribution.⁸⁰ Evidently, this distribution is not invariant under spatial translations, since it is centred around zero. Furthermore, it changes over time. Thus, spatiotemporal translations are not symmetries of this system. But this contradicts our intuition that the motion of the token is highly “law-like”: it can be described by a simple rule which is the same everywhere in space and time.

To solve this problem, we now introduce the notion of “partial” symmetries. We use the framework of spatially extended systems. Recall that \mathcal{H} is the set all logically possible spatially extended histories. A *partial symmetry monoid* of a spatially extended system is a collection Ψ of transformations of \mathcal{H} , along with a collection \mathcal{E} of ordered pairs of events (E, D) , such that:

- $\psi(\Omega) = \Omega$ for all ψ in Ψ ; and
- for any event pair (E', D') in \mathcal{E} and any ψ in Ψ , if E and D are the inverse images (in Ω) of E' and D' under ψ , then (E, D) is also in \mathcal{E} , and $Pr_E(D) = Pr_{E'}(D')$.⁸¹

For example, in our random-walk example (re-constructed in the framework of spatially extended systems), let \mathcal{E} be the set of all ordered pairs of events (E, D) such that event E exactly specifies the location of the particle at some time t , while event D happens at some later time t' . Thus, $Pr_E(D)$ is the conditional probability that the token satisfies such-and-such property at time t' , *given* that it was at such-and-such location at time t . Let Ψ be the monoid of all spatiotemporal translations of $S \times T$. If ψ is any element of Γ , then the set \mathcal{E} of pairs of events is invariant under ψ , and the conditional probability $Pr_E(D)$ is preserved by ψ for any (E, D) in \mathcal{E} , in the sense described above. Thus, the pair (Ψ, \mathcal{E}) is a partial symmetry monoid for the random-walk system. Crucially, the set \mathcal{E} does not include pairs of the form (Ω, D) , so we do not require unconditional probabilities of the form $Pr_\Omega(D)$ to be preserved by spatiotemporal translations.

Seen from this perspective, the transition probabilities of the random walk are “law-like”, because they are preserved by all the transformations in Ψ . In contrast, the initial probability distribution of the system is merely a brute necessity of the present system, since it is not preserved by any symmetries.

For another example suppose that the temporally evolving (or spatially extended) system $(\Omega', \{Pr_{E'}\}_{E' \subseteq \Omega'})$ is a *factor* of the system $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$, via some factor map σ , as described in Appendix A. For any event $E' \subseteq \Omega'$, let $\sigma^{-1}(E')$ denote its inverse image under σ (here defined as a subset of Ω). Then define $\mathcal{E} = \{(\sigma^{-1}(E'), \sigma^{-1}(D')) : E', D' \subseteq \Omega'\}$. Let Ψ be a monoid of spatiotemporal symmetries of the factor system $(\Omega', \{Pr_{E'}\}_{E' \subseteq \Omega'})$. The elements of Ψ might not be symmetries of the original system $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$. However, they will be *partial* symmetries, with respect to the set \mathcal{E} . So (Ψ, \mathcal{E}) is a partial symmetry monoid for $(\Omega, \{Pr_E\}_{E \subseteq \Omega})$. As already explained in Appendix A, one can greatly extend the scope of our framework by focusing attention

⁷⁹ The following argument does not depend on this assumption. Indeed, our argument would apply to *any* initial probability distribution for the token. Note that there is no such thing as a uniform probability distribution over the set of integers.

⁸⁰ To be precise: if t is odd, and $t' = t-1$, then for any even s between $-t'/2$ and $t'/2$, the probability that the token will be at spatial location s at time t is $2^{-t'} B((t'+s)/2, t')$, where B is the binomial coefficient function. The formula for even times is similar, but more complicated.

⁸¹ Formally, $E = \psi^{-1}(E') \cap \Omega$ and $D = \psi^{-1}(D') \cap \Omega$.

on a factor system rather than the original system. We now see that this is a special case of the broader concept of a partial symmetry monoid.

However, partial symmetry monoids cannot accommodate another feature of many systems. To illustrate this, consider a temporally evolving system where time is a *finite* sequence of integers, such as $T = \{1, 2, \dots, 100\}$. For such a system, time translations are not even *well-defined*.⁸² But in most such systems, we still want to say that the system obeys the same causal laws at all times, except perhaps at times 0 and 100. A similar problem arises in a spatially extended system where the space S is bounded (e.g., a partial differential equation defined on a cube, with specified boundary conditions) or a finite set of points (e.g., a cellular automaton defined on a 100×100 grid, with specified boundary conditions). In such a system, spatial translations are not well-defined. But in most such systems, we still want to say that the system obeys the same causal laws everywhere in the “interior” of the spatial domain.

To solve this problem, we now introduce “local” symmetries. We begin with some preliminary definitions. Let N be a subset of $S \times T$; call this a *neighbourhood* of space-time. Extending our earlier terminology, we say that an event $E \subseteq \Omega$ *happens inside* N if, for all histories h and h' in Ω , if $h_N = h'_N$, then h_N is in E if and only if h'_N is in E . Let $\Omega_N = \{h_N : h \text{ in } \Omega\}$; this is the set of all nomologically possible histories restricted to N (a set of functions from N into X). Let N' be another neighbourhood of $S \times T$, and suppose ψ is a function from N' into N . We use this to construct a function from histories restricted to N into histories restricted to N' . Specifically, for any h_N in Ω_N , we define $\psi(h_N)$ to be the function $h'_{N'}$ from N' into X given by $h'_{N'}(n') = h_N(\psi(n'))$, for all n' in N' . Note that while $h'_{N'}$ is a logically possible history restricted to N' , it is not necessarily an element of $\Omega_{N'}$.

We now define a *local symmetry groupoid* of a spatially extended system to be a combination of three components:

- a collection \mathcal{N} of subsets of $S \times T$ (called *neighbourhoods*);
- for each neighbourhood N in \mathcal{N} , a set \mathcal{E}_N of ordered pairs of events (E, D) which happen inside N (called *local events*); and
- for each pair of neighbourhoods N and N' in \mathcal{N} , a collection $\Psi_{N,N'}$ of bijections from N' into N (called *local symmetries*).

We refer to the collection $\{\Psi_{N,N'} : N, N' \in \mathcal{N}\}$ as a *groupoid* because it must satisfy two algebraic closure properties:

- for all neighbourhoods M and N in \mathcal{N} , and any local symmetries ψ in $\Psi_{M,N}$, its inverse ψ^{-1} is in $\Psi_{N,M}$; and
- for all neighbourhoods L, N , and M in \mathcal{N} , and all local symmetries α in $\Psi_{L,M}$ and β in $\Psi_{M,N}$, the composition $\alpha \circ \beta$ is in $\Psi_{L,N}$.

For any neighbourhoods N and N' in \mathcal{N} , and any ψ in $\Psi_{N,N'}$, we call ψ a *local symmetry* because it must preserve the modal and probabilistic structure of the system in the following sense:

⁸² For example, suppose we try to define $\psi(t) = t+1$ for all t in T ; then $\psi(100)$ is not well-defined, because 101 is not an element of T .

- $\psi(\Omega_N) = \Omega_{N'}$; and
- for all event pairs (E', D') in $\mathcal{E}_{N'}$, if E and D are the inverse images (in Ω_N) of E' and D' under ψ , then (E, D) is in \mathcal{E}_N , and $Pr_E(D) = Pr_{E'}(D')$.⁸³

For example, suppose that $S = \{1, 2, \dots, 10\}$ and $T = \{1, 2, \dots, 100\}$. For any s in $\{2, \dots, 9\}$ and t in $\{2, \dots, 99\}$, let $N_{s,t}$ be the 3×2 “space-time rectangle” of the form $N_{s,t} = \{s-1, s, s+1\} \times \{t, t+1\}$. Let \mathcal{N} be the set of all such space-time rectangles. For any s and s' in $\{2, \dots, 9\}$, and any t and t' in $\{2, \dots, 99\}$, if $N = N_{s,t}$ and $N' = N_{s',t'}$, then we define $\Psi_{N,N'} = \{\psi_{s',t' \rightarrow s,t}\}$, where $\psi_{s',t' \rightarrow s,t}$ is the function from N' into N which sends each space-time point (s_0, t_0) in N' to the point $(s_0 - s' + s, t_0 - t' + t)$ in N . Then, with a suitable specification of the local event sets \mathcal{E}_N for all N in \mathcal{N} , we could construct a local symmetry groupoid for many of the spatially extended systems (such as cellular automata) that one might define on $S \times T$. However, a fully worked out example would be rather technically involved and is beyond the scope of this paper; see Golubitsky, Pivato, and Stewart (2003) and Guay and Hepburn (2009).

Most of the ideas we have developed in this paper for the monoid of “full” symmetries can be generalized to partial symmetries and local symmetries. However, this is also beyond the scope of this paper.

Appendix C: Inferential modesty, informational parsimony, and the nomological hypothesis

Our version of Occam’s Razor requires us to assume that the symmetries of our system are given by a maximal symmetry monoid consistent with our total nomological evidence E (a superset of Ω). Under natural assumptions, at least one maximal consistent monoid will indeed exist.⁸⁴ However, there may be more than one. In this case, we need a criterion to choose one maximal symmetry monoid rather than another. We now develop such a criterion.

Let us begin with an example. Consider a very simple temporally evolving system, where the set T of times contains only a single element. So, histories can be identified with states at that single time; this expositional simplification has no substantive consequences. Suppose that the state of the system is described by a two-dimensional grid of zeros and ones, which is infinite in every direction. Let X be the set of all logically possible grids of this kind. Then the set \mathcal{H} of all logically possible histories can be identified with X . In this system, one elementary kind of nomological constraint is one that constrains the values of one or more cells, for example the constraint “in any possible history, the cell (2,3) must have the value zero”.⁸⁵ Suppose we have obtained evidence that any possible history must satisfy the constraints shown in Figure 1. This evidence would be represented by the subset E of \mathcal{H} consisting of all single-period

⁸³ Formally, $E = \psi^{-1}(E') \cap \Omega_N$ and $D = \psi^{-1}(D') \cap \Omega_N$.

⁸⁴ For example, if (i) \mathcal{H} has a topology, (ii) E is a compact subset of \mathcal{H} , and (iii) all the transformations in question are continuous, then Zorn’s lemma implies the existence of a maximal consistent monoid.

⁸⁵ Formally, this constraint corresponds to the set $E = \{h \in \mathcal{H} : h(2,3) = 0\}$. Of course, we have chosen this rather artificial example only for expositional simplicity. Typically, we would be interested not so much in nomological constraints on single coordinates, but in constraints on the relationships between two or more coordinates, such as the constraint “No two adjacent cells can both contain a zero”.

histories in which the grid coincides with Figure 1 in all non-empty cells. For the sake of argument, let us treat E as our total nomological evidence about the system.

	:	:	:	:	:	:
...		1				...
...						...
...		1				...
...		1				...
...				0	0	...
...		1				...
	:	:	:	:	:	:

Figure 1

Now, for any integer n , let ψ_n^{\rightarrow} be the transformation that shifts the entire grid to the *right* by n spaces.⁸⁶ Let $\Psi^{\rightarrow} := \{ \dots, \psi_{-1}^{\rightarrow}, \psi_0^{\rightarrow}, \psi_1^{\rightarrow}, \psi_2^{\rightarrow}, \dots \}$ denote the monoid of all such horizontal shifts. Meanwhile, let ψ_n^{\uparrow} be the transformation that shifts the entire grid *upwards* by n spaces, and let $\Psi^{\uparrow} := \{ \dots, \psi_{-1}^{\uparrow}, \psi_0^{\uparrow}, \psi_1^{\uparrow}, \psi_2^{\uparrow}, \dots \}$ denote the monoid of all such vertical shifts. Consider two hypotheses:

Hypothesis 1: All transformations in Ψ^{\rightarrow} are symmetries of the system.

Hypothesis 2: All transformations in Ψ^{\uparrow} are symmetries of the system.

The evidence represented in Figure 1 is consistent with either of these hypotheses. However, it cannot accommodate both simultaneously. If Hypothesis 1 were true, then the constraints in Figure 1 would entail the constraints shown in Figure 2. If Hypothesis 2 were true, then they would entail the constraints shown in Figure 3. In each figure, the constraints that were part of the initial nomological evidence are highlighted in boldface; extrapolated constraints (based on the postulated symmetries) appear in non-bold font. Clearly, Hypotheses 1 and 2 cannot *both* be true, since they yield mutually contradictory constraints on the values of the grey cells.

	:	:	:	:	:	:
...	1	1	1	1	1	...
...						...
...	1	1	1	1	1	...
...	1	1	1	1	1	...
...	0	0	0	0	0	0
...	1	1	1	1	1	...
	:	:	:	:	:	:

Figure 2

⁸⁶ Of course, if n is *negative*, then ψ_n^{\rightarrow} is actually a shift to the *left*.

	:	:	:	:	:	:
...		1		0	0	...
...		1		0	0	...
...		1		0	0	...
...		1		0	0	...
...		1		0	0	...
...		1		0	0	...
...		1		0	0	...
	:	:	:	:	:	:

Figure 3

Let Ψ be some maximal consistent monoid of transformations that we postulate as the symmetry monoid, in accordance with Occam's Razor. Hypothesis 1 then asserts that $\Psi^{\rightarrow} \subseteq \Psi$, while Hypothesis 2 asserts that $\Psi^{\uparrow} \subseteq \Psi$. Since both hypotheses cannot simultaneously be true, it follows that there are at least two distinct ways in which we could specify Ψ : one including Ψ^{\rightarrow} and another including Ψ^{\uparrow} . So even in this very simple example, there is no unique maximal consistent monoid.

At first sight, the choice between these two maximal symmetry monoids seems arbitrary. But it is not. To see this, note that both hypotheses could have entailed the same constraints they did, using less initial evidence. For example, Hypothesis 1 would have entailed the same constraints from the weaker evidence represented in Figure 4.

	:	:	:	:	:	:
...		1				...
...						...
...		1				...
...		1				...
...				0		...
...		1				...
	:	:	:	:	:	:

Figure 4

The original evidence in Figure 1 constrained *six* cell values (i.e., six "bits" of information). But Hypothesis 1 can make do with only *five* of them (in particular, the second zero is redundant). Meanwhile, Hypothesis 2 would have entailed the same constraints from only *three* bits of information, as represented in Figure 5.

In this example, neither $\Omega(\Psi^\rightarrow, E)$ nor $\Omega(\Psi^\uparrow, E)$ includes the other, so neither monoid is more inferentially modest than the other, according to our definition. Likewise, neither E^\rightarrow nor E^\uparrow includes the other, so neither monoid is more informationally parsimonious. So our formal definitions up to this point are not sensitive enough to capture the plausible intuition that Ψ^\uparrow is both more inferentially modest and more informationally parsimonious than Ψ^\rightarrow .

One possible way of capturing this intuition is to use concepts from information theory, such as entropy. To do this, we must introduce a prior probability distribution Pr_0 on the set \mathcal{H} of all logically possible histories. In the example with the infinite grid, this could be the *uniform Bernoulli* distribution, which treats all the cells in the grid as independent, identically distributed random variables, where zero and one each appear with probability $\frac{1}{2}$. Given two different symmetry monoids Ψ_1 and Ψ_2 that are consistent with the same nomological evidence E , we can use Pr_0 to compare them:

Inferential modesty (relative to Pr_0): If $\text{Pr}_0(\Omega(\Psi_2, E)) \leq \text{Pr}_0(\Omega(\Psi_1, E))$, then we say that Ψ_1 is (at least weakly) more *inferentially modest* than Ψ_2 , relative to Pr_0 .

Informational Parsimony (relative to Pr_0): Let E_1 be the largest superset of E such that $\Omega(\Psi_1, E_1) = \Omega(\Psi_1, E)$. Let E_2 be the largest superset of E such that $\Omega(\Psi_2, E_2) = \Omega(\Psi_2, E)$. If $\text{Pr}_0(E_2) \leq \text{Pr}_0(E_1)$, then we say that Ψ_1 is (at least weakly) more *informationally parsimonious* than Ψ_2 , relative to Pr_0 .

Do these criteria enable us to prefer Ψ^\uparrow to Ψ^\rightarrow , as intuition suggests? Let us begin with the second criterion. Comparing Figures 4 and 5, we see that $\text{Pr}_0(E^\rightarrow) = 2^{-5}$, whereas $\text{Pr}_0(E^\uparrow) = 2^{-3}$, and so Ψ^\uparrow is indeed more informationally parsimonious than Ψ^\rightarrow , relative to the uniform Bernoulli distribution. The first criterion, by contrast, does not help. Comparing Figures 2 and 3, we see that $\text{Pr}_0(\Omega(\Psi^\rightarrow, E))$ and $\text{Pr}_0(\Omega(\Psi^\uparrow, E))$ are each zero, because they constrain an infinite number of cells. So, they do not differ in inferential modesty relative to Pr_0 . They do differ in more sensitive measures of inferential modesty, computed using more advanced notions from information theory, such as “entropy density”. But the details are beyond the scope of this paper.

Note that, if Ψ_1 is more inferentially modest than Ψ_2 in the original sense, which did not refer to any prior probability, then Ψ_1 is more inferentially modest than Ψ_2 in the information-theoretic sense, relative to *any* prior Pr_0 . This is because if $\Omega(\Psi_2, E) \subseteq \Omega(\Psi_1, E)$, then $\text{Pr}_0(\Omega(\Psi_2, E)) \leq \text{Pr}_0(\Omega(\Psi_1, E))$. Likewise, if Ψ_1 is more informationally parsimonious than Ψ_2 in the original sense, then Ψ_1 is more informationally parsimonious than Ψ_2 in the new sense, relative to *any* prior Pr_0 . The reason is that if $E_2 \subseteq E_1$, then $\text{Pr}_0(E_2) \leq \text{Pr}_0(E_1)$.

Appendix D: Spatial distance in quantum mechanical systems

In Section 3.9, we proposed a definition of the *distance* between two regions R and R' in space, based on the minimum time required for a “signal” to travel from R to R' . As we have already observed, this definition is not entirely satisfactory in systems where signals can travel at arbitrarily high speeds (such as classical mechanics). This is particularly problematic in quantum mechanics, for two reasons.

First, there is the well-known phenomenon of *entanglement*, where two particles, perhaps separated by a large spatial distance, can apparently correlate their behaviour. But, in fact, this is less of a problem than it first appears. Rather than interpreting entanglement as “spooky action at a distance”, we can interpret it as a sign that we have not correctly specified the space S for this spatially extended system. A three-dimensional quantum system with n particles is not a collection of n wave functions on a three-dimensional space; rather, it should be viewed as a *single* wave function on a $3n$ -dimensional space. So, we should define $S = \mathbf{R}^{3n}$. Even if two particles appear widely “separated” from our three-dimensional perspective, their joint location is described by a single “hump” of the wave function in a six-dimensional space.⁸⁸ From this perspective, the entangled behaviour of the two particles does *not* appear as a non-local phenomenon.

However, there is a more fundamental problem, which affects even a single-particle quantum system. Solutions to the Schrödinger equation on unbounded domains generally have *full support*: they give non-zero probability to every part of the space. This means, in effect, that the particle has a non-zero probability (albeit tiny) of “jumping” arbitrarily large distances through space.⁸⁹ Thus, no two regions of space are ever unreachable from one another in any time duration, no matter how short, and so the distance between any two regions will be zero, according to the definition given in Section 3.9.

To address this problem, we must introduce a slightly more nuanced version of “unreachability”. Let $\varepsilon > 0$ be some small “error tolerance”. Given three events E , F , and G in Ω , we say that E and G are ε -*conditionally independent given F* , if $1 - \varepsilon < \Pr_F(E \cap G) / \Pr_F(E) \cdot \Pr_F(G) < 1 + \varepsilon$. In other words, the conditional probability $\Pr_F(E \cap G)$ is “almost” the same as the product $\Pr_F(E) \cdot \Pr_F(G)$, which means that E and G are “almost” conditionally independent, given F . If R and R' are two regions of S , and $t > t_0$, then we say that R' is ε -*unreachable* from R in time t if, for any h in Ω , any event which happens in R' at time t is ε -conditionally independent of any event which happens in R at time t_0 , given the event $[h_{R \times \{t_0\}}]$. If ε is small, this means that, *with very high probability*, a signal which originates in R at time t_0 cannot reach R' before time t . We then define the ε -*distance* between R and R' to be the supremum of the set of all t such that R' is ε -unreachable from R in time t (if this supremum exists).

By using a small but non-zero ε , we can thus define a non-trivial notion of ε -distance between different regions of space, even in a quantum-mechanical system. This measure of distance will obviously depend on the value of ε , but it will roughly approximate the “classical” notion of distance. However, a detailed development of this approach is beyond the scope of this paper.

⁸⁸ We are being slightly imprecise here; in quantum mechanics, particles do not even *have* precisely specified locations.

⁸⁹ Again, we are being somewhat imprecise in even ascribing a particular “location” to the particle. The Fourier transform of the wave function also has full support; this means that any velocity for the particle, no matter how large, has a tiny but non-zero probability of being realized.