Synthese manuscript No.
(will be inserted by the editor)

Projection, symmetry, and natural kinds

Benjamin C. Jantzen

Received: date / Accepted: date

Acknowledgements I would like to thank the participants of the 2013 Ontology & Methodology conference at Virginia Tech, Anjan Chakravartty, and Richard Burian for helpful discussion of the ideas presented here. I am especially grateful to Kelly Trogdon and two anonymous referees for comments on earlier drafts.

B. Jantzen
Philosophy Department, 229 Major Williams Hall (0126), 220 Stanger Street, Virginia Tech, Blacksburg, VA 24061 U.S.A.
Tel.: +1-540-231-1766
Fax: +1-540-231-6367
E-mail: bjantzen@vt.edu

Synthese manuscript No.
(will be inserted by the editor)

Projection, symmetry, and natural kinds

Benjamin C. Jantzen

Received: date / Accepted: date

Abstract Scientific practice involves two kinds of induction. In one, generalizations are drawn about the states of a particular system of variables. In the other, generalizations are drawn across systems in a class. We can discern two questions of correctness about both kinds of induction: (P1) what distinguishes those systems and classes of system that are 'projectible' in Nelson Goodman's (1955) sense from those that are not, and (P2) what are the methods by which we are able to identify kinds that are likely to be projectible? In answer to the first question, numerous theories of 'natural kinds' have been advanced, but none has satisfactorily addressed both questions simultaneously. I propose a shift in perspective. Both essentialist and cluster property theories have traditionally characterized kinds directly in terms of the causally salient properties their members possess. Instead, we should focus on 'dynamical symmetries', transformations of a system to which the causal structure of that system is indifferent. I suggest that to be a member of natural kind it is necessary and sufficient to possess a particular collection of dynamical symmetries. I show that membership in such a kind is in turn necessary and sufficient for the presence of the sort of causal structure that accounts for success in both kinds of induction, thus demonstrating that (P1) has been answered satisfactorily. More dramatically, I demonstrate that this new theory of 'dynamical kinds' provides an answer to (P2) with methodological implications concerning the discovery of projectible kinds.

Keywords Discovery · Induction · Natural kinds · Projectibility · Symmetry

1 Introduction

It is a truism that some schemes for categorizing natural systems have permitted successful generalization from a sample of particulars; others have not.

Yet others, in the extension of predicates like Goodman's 'grue' are deemed defective a priori. What distinguishes the good categories or kinds from the bad, however, is far from obvious. There is a rich literature on 'natural kinds' that seeks to explain the difference. While diverse, the existing accounts of natural kinds — theories of what distinguishes kinds that admit of successful generalization from those that do not — share an approach. They all assume that kinds are to be characterized in terms of the very qualities that feature in successful generalizations involving those kinds. Essentialist theories, for example suppose that a natural system belongs to a particular kind when it possesses the properties necessary and sufficient for membership in that kind. The relevant properties are generally assumed to be theoretically salient features such as charge or proton number, the very properties that feature in scientific generalizations.

I suggest instead that characterizing the natural kinds in terms of what is causally *irrelevant* underwrites a theory of natural kinds that, while sharing important elements in common with existing views, offers better answers to questions about scientific induction. The remainder of this section spells out those questions about kinds and scientific induction, and surveys the suitability of various approaches to natural kinds as potential answers. In Section 2, I reinterpret these questions from a complementary perspective grounded in the notion of a 'dynamical symmetry', and in Section 3 I develop a theory of natural kinds in terms of this new perspective. I then argue that this new notion of natural kinds evades some problems faced by other theories while at the same time answering some questions left open by other approaches. More importantly, I show that this theory offers some novel methodological advice for discovering natural kinds. This is perhaps the strongest sort of evidence one can offer in favor of a theory of natural kinds. To borrow from Quine (1969, p129), "In induction, nothing succeeds like success." A theory of inductively useful kinds that results in successful induction is, ceteris parabus, a better theory than one that doesn't.

1.1 Science and projectibility

Empirical science involves the iteration of two sorts of induction.[1] The first concerns the behavior of particular systems, such as a specific particle collider or the chemical reagents in a particular beaker. The generalizations to be drawn involve relations amongst the variables associated with the system. These relations may be more or less quantitative, and have to do with which values of one variable are associated with which values of the others over a range of counterfactual situations. So for instance, one may generalize about the relation between oscillation period and arm length for a particular pendulum, or the initial and final concentrations of chemical species mixed in a

[1] I don't mean to suggest that this is all which empirical science involves or aims at. Nor do I intend to suggest that this is a complete taxonomy of the kinds of inductive practice. For an example of a more complete taxonomy, see (Earman, 1985).

particular apparatus. Inductions of this variety involve projecting from observed states of the system to other states that would obtain given that some of the variables take on different values. For lack of imagination, I'll call inductions of this sort *system inductions*.

The second sort of induction is over systems. Inductions of this sort involve extrapolation from the properties of a finite (typically small) set of particular systems (e.g., the orbit of the Moon about the Earth) to a potentially infinite class of actual and possible systems (e.g., all orbital systems). Given that the acceleration of *this* satellite is proportional to the inverse square of its distance to the planet, so too does it hold for *all* such gravitating systems. Given that energy is conserved in *this* mechanical system, energy is conserved for *all* such mechanical systems. Call these *kind inductions*.

Both varieties of induction pose questions of correctness. Of all the variables we can conceive, which are suitable for generalization in particular systems? Of all the classes of system we could construct, which are amenable to projection? These are, of course, just different manifestations of Goodman's 'new riddle of induction', at its heart a puzzle about distinguishing 'projectible' kinds from the non-projectible, to use Goodman's terminology (1955). In terms of scientific practice, what characterizes the systems of variables that are projectible, either in the context of system or kind induction?

There are a couple of clarifications to be made. First, when I speak of projectibility, I do not mean to refer to a predicate but to a class of things in the world, i.e., to the extension of a predicate. While reference must invariably be made via a predicate, it is the class itself that has the feature of projectibility I'm interested in. I will generally use the term 'kind' in lieu of 'predicate' to signal this intention.[2] Second, by calling a kind 'projectible', I mean that (at least some) scientifically significant generalizations over that kind are in fact true; the systems or system states not in evidence really would be found to have the projected feature if they were realized and examined. I do not mean that one is justified in drawing inferences about instances or systems not currently in evidence; whether this is true given that a kind is projectible depends upon one's available evidence and upon one's theory of inductive inference. A kind may actually be projectible and yet an epistemic agent may be unjustified in drawing inductive inferences over that kind given her epistemic context. Nor do I mean that we actually have used the kind in inductive inferences. Again, whether or not a kind is objectively projectible in the sense I intend it is independent of whether or not anyone bothers to identify that kind and draw inferences about it.

Lastly, I want to emphasize the question at issue. I am not asking about the justification for either sort of induction. Nor am I asking why there is regularity in nature, a question that I, like Quine (1969), doubt can be given a satisfying answer. Rather, I take for granted that some kinds are projectible

[2] Though Goodman (1955) frames his problem in terms of predicates—linguistic entities—his solution is in terms of the extensions of those predicates. Specifically, what he calls "entrenchment" pertains to a class that may be named by many predicates.

in that they support successful system and kind inductions. The question is what sets these kinds apart from those that do not.

1.2 Two problems of projectible kinds

For each of the two sorts of scientific induction described above, there are really two logically distinct but interrelated problems of projectibility:

P1 For system induction, what characterizes those systems of variables which will support projection across states of the system? For kind inductions, what, if anything, distinguishes the projectible kinds of system from the non-projectible?[3]

P2 What is (are) the method(s) for recognizing, discovering, or selecting systems and kinds of system of the right sort to be projectible? Or in other words, how can we efficiently identify systems and classes of system for scientific investigation?

As I said, in a given context the two questions are logically distinct. It may be the case that, aside from their projectibility, the projectible kinds of system all share some set of features in common that are not possessed by non-projectible kinds, and yet there is no method by which we can strike upon the projectible kinds at a rate better than chance. On the other hand, it is possible that there is some method for identifying projectible kinds, and yet these kinds share nothing in common (other than their projectibility) that they don't also share with at least one non-projectible kind. In other words, it may be possible to systematically generate members of the class of projectible kinds, and yet members of that class may have nothing in common that sets them apart. A good analogy is with the theorems of first-order logic for a given language and system of natural deduction. It is possible to generate theorems with the rules of the natural deductive system, but there is no property shared by all theorems other than their derivability.

One might be inclined at this point to treat the second question of projectibility as superfluous. After all, a method for finding projectible predicates seems an awful lot like a 'logic of discovery'. Popper (2007) and Laudan (1981, ch. 11), amongst others, have argued there is no such thing, and that even if there were, it would be philosophically uninteresting. It is not my aim here to take up these arguments afresh. I will, however, stress some empirical facts about the scientific enterprise itself, facts that demand an explanation. Granting that empirical science is enormously successful at discovering (and making use of) projectible kinds, it is a remarkable fact that this success comes at such comparatively low cost. I do not mean to deny that even the meanest scientific discovery is purchased at the price of enormous labor, patience, and persistence. Rather, I mean to point out that for any new empirical domain for which science has uncovered projectible kinds, the relative proportion of

[3] I mean (P1) in a thin sense: the shared feature of a kind may or may not be that in virtue of which it is projectible.

successes to attempts is astonishingly low if we are merely guessing. In fact, since there is an unlimited (and probably continuous) set of kinds one can define for any given empirical domain — an infinite number of ways to aggregate systems into classes — one would expect the probability of drawing a projectible kind from this set at random to be vanishingly small. The history of science is indeed a graveyard of discarded theories and discarded kinds, but the deceased do not outnumber the living by many orders of magnitude. Rather, scientists seem to be doing much better than random when choosing which kinds to investigate. This is a puzzle. To echo Quine (1969, p126), why should our pre-theoretical choices of variables and clustering of systems accord so well with the "...functionally relevant groupings in nature so as to make our inductions tend to come out right?"

Of course, our success at spotting projectible kinds is only mysterious if most arbitrarily constructed kinds are not projectible. This may be mistaken — it may be that just about anything we choose to look at will turn out to support projection of one sort or another. This seem unlikely, particularly in light of Goodman's observations about 'grue' and 'bleen' and other seemingly defective predicates. But if one could provide a compelling argument that this is the case, question (P2) would be mooted. Absent such an argument, however, it seems as pressing as the first question about projectible kinds. There are certainly others who take the question seriously. Mattingly and Warwick (2009) just recently attempted to make progress on a special case of (P2). In particular, they raised the problem of ascertaining the projectible predicates for complex computer simulations that in turn are supposed to inform us about analytically intractable natural systems. They ultimately advocate an experimental approach to determining those predicates, though they provide only sketchy details. The point is that they make a strong case for the importance of answering (P2) if complex computational simulations are to yield useful conclusions about the systems they are meant to model.

I take it then, that any theory of projectibilty that can answer both (P1) and (P2) is superior, ceteris parabus, to a theory that answers only one, and a theory that answers them for both system and kind induction is superior to a theory that doesn't. My aim in this essay is to provide a new approach to answering the first question of projectibility, and to argue that this approach offers a more promising route to answering the second than any of its competitors for many if not most scientifically salient kinds.

1.3 Natural kinds and projectibility

In the search for answers to (P1) and (P2), I'll set aside Goodman's efforts at solving his own puzzle, as they have largely been superseded by theories of 'natural kinds'. Theories bearing this name have been advanced to serve a wide variety of aims, and the phrase itself has been given a diversity of meanings.[4]

[4] Quine, for instance, restricts the term 'kind' to collections defined by similarity relations. Since, in his view, mature sciences have no use for such similarity relations, there are no

Though Mill (1858), Venn (1866), Quine (1969), and Goodman (1955) all invoked a notion of natural kinds (whether or not they used that term) in order to explain aspects of inductive success, the philosophical emphasis has shifted. Since the work of Putnam (1973; 1975) and Kripke (1980), theories of natural kinds have been pursued largely in the service of semantics, in particular in attempts to understand proper names. There is also a community interested in the metaphysics of natural kinds, particularly questions of naturalness and what it is to be a 'kind' as opposed to a set, or grouping (Lewis, 1983; Lowe, 2006; Hawley and Bird, 2011). As illuminating as this work has been in its own context, my interest is strictly in answers to (P1) and (P2). It is with this in mind that I have culled from the literature on natural kinds, and it is to the relevant subset of theories I mean to refer when I mention theories of natural kinds.

So what can natural kinds tell us about projectible kinds? The idea is that a kind is projectible just when its extension coincides with an 'objective' division of the natural world, in the sense that it does not depend upon human interests or conceptual schemes; it is determined by nature, not us. The categories marked out by such divisions are called natural kinds. Of course, without a substantive account of what precisely the natural divisions are, the natural kinds solution to (P1) is not much better than noting that there exist projectible kinds. Fortunately, theories abound. In the modern literature, they can be loosely grouped into two types: essentialist and property cluster.

Essentialist theories of natural kinds are united by a common set of assumptions (Ellis, 2001, 2005):

(i) There is a set of necessary and sufficient conditions for membership in a natural kind. Generally this is taken to be a set of intrinsic properties which jointly comprise the 'essence' of the kind.
(ii) Natural kinds are categorically distinct. There is no vagueness concerning membership; either an entity possesses the requisite essence, or it does not.
(iii) Natural kinds are non-overlapping. No two natural kinds can share a member without one kind being entirely included in the other.
(iv) The essence of a kind determines or is responsible for many of the qualities associated with the kind.

According to essentialist accounts, if a putative kind shares an extension with a natural kind, then it will support successful categorical inductions, at least with respect to those properties (relations, etc.) which are held in virtue of the essential properties of the kind. As T. E. Wilkerson puts it,

> If I know that a lump of stuff is gold, or that the object in front of me is an oak, I am in a position to say what it is likely to do next, and what other things of the same kind are likely to do. I know for

'natural kinds' in science. I am only interested in notions of natural kind that stand as possible responses to (P1). Most of the views considered here take scientific categories to be paradigm examples of natural kinds.

> example that the gold cannot turn into water, and that the oak will not in due course produce tomatoes. And I know that no other piece of gold could be persuaded to turn into water, and no other oak could be persuaded to produce tomatoes. Certain outcomes are ruled in, and others are ruled out, by the real essences of gold or oaks.(1988, p30)

However, essentialism offers little hope of answering (P2). While every natural kind possesses a unique essence, essences are themselves quite inhomogenous. There are no properties that belong to the essence of every natural kind (with the exception of belonging to a natural kind), and no higher-order features of essences that distinguish them from arbitrary sets of properties, and so no particular marker of natural kindness that would let us assess whether a putative essence really belongs to a projectible kind. Furthermore, there is no obvious way to generate a list of essences. So it seems that the only way we could come to discover projectible kinds efficiently is via some sort of Aristotelian cognition of essences from particular samples. This is not a plausible route.

Currently more popular, especially amongst philosophers of biology, are property cluster theories. What unites these accounts of natural kinds is their rejection of necessary and sufficient conditions. More specifically, they assert that there is no particular property or combination of properties that each member of the kind must possess, nor any one that is sufficient. Instead, kinds are characterized by a "cluster" of properties or relations that tend to co-occur under a range of counterfactual conditions, but covary imperfectly. As N. E. Williams aptly put it, what matters for membership in the kind, "...is the *extent* to which the properties a substance instantiates overlap with the properties in the cluster, which can be satisfied to varying degrees"(2011)(emphasis in the original).

Of the property-cluster theories of natural kinds, the most widely discussed is Richard Boyd's homeostatic property cluster (HPC) theory (Boyd, 1988, 1989, 1991, 1999). On this view, the clustering properties tend to co-occur due to causal "homeostatic" mechanisms. Each token member of an HPC kind instantiates a mechanism that tends to maintain the cluster of properties in that individual, i.e., each instance of a kind instantiates a mechanism that tends to keep it an instance of the kind.[5] HPC and its various elaborations (Kornblith, 1993; Griffiths, 1999; Wilson, 1999) are designed to explain the success of inductions involving kinds, like biological species, which lack essences. The explanation is much like that of the essentialist. The causal mechanism shared

[5] There is some ambiguity as to whether these mechanisms are taken to be homeostatic with respect to the individual, as described in the text, or with respect to the kind. In the latter case, there is some causal mechanism such that the occurrence of one of the properties in the cluster tends to cause the occurrence of the others, though not necessarily in the same individual. The connection between properties is robustly maintained across many instances by the mechanism, but in any one individual the connection may be fragile over time. Boyd (1999) recognizes both possibilities, but the majority of discussion over HPC assumes that each individual belonging to a kind is a distinct instance of a type of causal mechanism. That is, each member of a kind possesses a cluster of properties maintained by an instance of the same type of homeostatic causal mechanism.

by all members of an HPC kind tends to maintain the presence of a cluster of properties, and this approximate sameness in causally relevant properties accounts for the fact that members of the same kind can be expected to behave in the same way in similar causal contexts. As Chakravartty (2007, p170) puts it, "...dispositions conferred by properties for various manifestations are present wherever such properties are found, and to the extent that the same causal properties are found in members of the same kind, their behaviors will be subject to inductive generalizations and predictions.".

Despite it's popularity, the HPC account is deeply unsatisfying. To begin with, it simply fails to capture many or most of the categories that are the focus of scientific investigation. The class of quasi-isolated gravitating systems it not an HPC kind. On the face of it, there are no constant properties held in conjunction across all such systems by a causal mechanism — the Earth-Sun system and the Jupiter-Sun system are quite different in the masses and separations of the moving bodies, orbital periods, total angular momentum, etc. And yet this is precisely the class of systems over which Newton drew what may be the most important inductive inference in the history of science.[6] To draw an example from the messier 'special sciences', the basic kinds of chemical kinetics present a similar problem for HPC. A first-order reaction, for instance, is any reacting system in which the rate at which the reaction occurs is linearly dependent on the rate-limiting reactant. Within any such system, few properties are preserved over time — it is, after all, a reaction — so there is no obvious homeostatic cluster associated with the kind. Furthermore, any given time-slice of two first-order reactions can differ dramatically in all of their chemical properties. That's because lots of different compounds can be involved in a first-order reaction. The HPC account clearly fails to capture at least some scientifically salient natural kinds.[7]

Of course, it is unfair to insist that any answer to (P1) or (P2) capture every projectible kind. An account that captures only a portion of the projectible kinds would still be significant, and the HPC account should be evaluated on how well it answers (P1) and (P2) for the subset of kinds within it's purview. However, even if we restrict our attention to the natural kinds that are supposed to be clean examples of HPC kinds, the account fails to offer a complete answer with respect to (P1), at least as far as kind induction goes. Carl Craver (2009) has pointed out a sort of conventionalism lurking in the HPC account of natural kinds. What unites two members of a kind is that they instantiate the same kind of casual mechanism maintaining similar clusters of properties. HPC thus takes for granted that mechanisms sort into kinds. But what objectively distinguishes mechanisms? What determines whether or not two particular mechanisms are of the same kind, and how could we know this? For

[6] There are complex features common to all members of the class. For instance, energy and momentum are conserved. But what is the 'causal mechanism' that maintains the association between the two properties "conserves momentum" and "conserves energy"?

[7] Slater and others have pointed out that it seems to have trouble with kinds such as chemical elements. But it's not clear that this is really the sort of induction central to scientific practice.

Craver, this is a worry about conventionality. But it should also worry anyone pursuing theories of natural kinds in order to answer questions about projectibility. Without an account of how mechanisms divide into kinds, the HPC simply fails to answer (P1). At best, it pushes the question of what constitutes a natural kind back from the level of particular collections of properties to classes of casual mechanism.

Given the dominance of HPC views, this last point is worth pressing. The HPC account is supposed to explain projectibility by pointing to the (approximate) uniformity of the causal dispositions of members of an HPC kind. Individual things (objects or systems) possess properties, some of which confer causal dispositions to produce effects in particular contexts. For instance, all samples of reduced carbon are disposed to combust in an oxygen atmosphere if provided sufficient initial energy. If all members of a kind overlap significantly in their causally relevant properties, then we can expect them to overlap significantly in their behavior across a wide range of causal contexts. This similarity in behavior is supposed to support generalizations of the form, "This member of kind K did E in causal context C, so all K's will do E in C." However, bearing similar clusters of stably associated properties is obviously not enough. It matters which internal mechanism is responsible for the stable association of the properties. For example (and you can try this at home), put a lithium-ion and alkaline battery in a pair of flashlights and pop them in the freezer. One will stop producing light long before the other. This is because, at low temperature, the relation between current and voltage changes (though voltage is roughly maintained), and does so differently between the two batteries. For there to be uniformity in response to new causal contexts, there must be some sort of uniformity in the internal causal mechanisms shared by members of an HPC kind. But the HPC account itself, as Craver points out, doesn't give us any theory of the sameness of causal mechanism. Many different mechanisms can 'stably' associate the same properties and yet yield dramatically different behavior — different relations amongst those properties — under different external conditions.

Partly motivated by Craver's concerns,[8] Matthew Slater (2013) has recently argued for a property cluster account of natural kinds that eschews an appeal to causal mechanism. His aim is to retain those features of HPC natural kinds that are sufficient for the projectibility of the kind while purging the rest, in particular what he sees as the metaphysical baggage that comes with the notion of a causal mechanism. Membership in what he calls a stable property cluster (SPC) kind requires "cliquish stability" of the cluster. A property cluster is cliquishly stable just if the instantiation of a subcluster tends to imply the instantiation of the full cluster under a wide range of counterfactual conditions.

Whether or not it resolves worries over conventionality, Slater's proposal is a victim of its own success, at least when it comes to answering (P1) and

[8] Slater is also motivated by arguments suggesting that causal mechanisms are neither necessary nor sufficient for the sort of robust counterfactual association of properties required for successful projection.

(P2). I concede that cliquish stability is sufficient for projectibility. The problem is that projectibility seems to be sufficient for cliquish stability. Why? If a category, K, is projectible, then for some significant number of properties, P_i, generalizations of the form "x is a K and $P_i(x)$, therefore for all or most x's that are K's, $P_i(x)$" are true and would remain so over some appreciable range of counterfactual conditions (that's what distinguishes projectibility from contingent association). Put differently, there is some range of conditions other than those actually realized such that being a K is associated with possessing all or most of the P_i. But that's just to say that the properties characteristic of K are cliquishly stable.[9] Thus, to say that a kind characterized by a cluster of properties is cliquishly stable just is equivalent to saying that it is projectible. SPC therefore gives only a trivial answer to (P1), namely that the feature shared by all projectible kinds is projectibility. A trivial answer to (P1) gives us no help in answering (P2). To be fair to Slater, his aim was not to answer the epistemic questions, but rather to provide a clear accounting of the epistemic phenomena to be explained by a theory of natural kinds. That's a laudable goal, but one which gets us nowhere in actually constructing a theory of natural kinds. All we can say on the basis of the SPC account is that some kinds are projectible (i.e., cliquishly stable) and others are not.

1.4 Causal kinds

There is something right about the HPC account's appeal to causal relations and the sorts of counterfactuals they support. But the HPC is unnecessarily circumspect. If relations of causation are doing all the work in ensuring projectibility, why not simply focus on causal structure? That is, one could define natural kinds in terms of a network of causal relations: every system instantiating the same causal structure amongst instances of the same variable types is in the same kind. That way, there's no question about where to draw the line. If variable X causes variable Y such that the value of Y is a particular function of X in this system and in that system, then they belong to the same kind. For lack of an established name, I'll call this this the *causal kinds* account. According to the causal kinds account, the answer to (P1) for system induction is that systems with stable causal structures are systems with projectible states, an eminently plausible view. The folks working on causal epistemology have already answered (P2) for system induction as well. Roughly, we discover stable causal structures by detecting patterns of independence amongst variables, particularly under interventions that change some or all of them.[10]

With respect to kind induction, however, the answers are rather less impressive. If a kind is characterized by a particular causal structure specified at

[9] I am neglecting a great deal of detail in Slater's account — in particular, his use of Lange's (2009) notion of 'non-nomic stability' to spell out a precise sense of counterfactual robustness. But the rough argument I've provided is adaptable to these details.

[10] For an overview, see (Pearl, 2000; Spirtes et al, 2000).

the finest level of detail — functional relationships indicating which variables causally influence which others and according to what rule — then kind induction becomes trivial. Every member of a kind is exactly like every other. There is nothing contradictory or incoherent about defining kinds this way. However, such kinds are far more restrictive than the kinds of day to day science. In fact, they are so narrow as to be nearly sui generis. Since the functional relationships between the variables of position, momentum, and time differ between the two, our solar system and the system consisting of the star Kepler-22 and its single earth-like planet are in different causal kinds. Likewise for two first-order chemical reactions involving the same reagents but different initial concentrations — the connection between concentration and time differs between them, and so they are distinct causal kinds. Outside of human manufactured goods, it's hard to imagine finding a pair of systems in the natural world that belong to the same causal kind. This may sound pedantic; it may seem obvious that we need to take a more inclusive view of which causal systems to lump together into kinds. The question is how we can do so without running into Craver's worry.

The solution is to focus not on causal structure, but on a special property of causal structure. That property, described in detail below, is 'symmetry structure'. What I call the *dyamical kinds* theory of natural kinds is essentialist: each natural kind is defined by a symmetry structure, a property that is necessary and sufficient for membership in the kind. But it also embraces what's right about the HPC account. Instantiating a symmetry structure is necessary and sufficient for a system to exhibit a stable causal structure. This causal structure in turn supports both kind and system inductions. However, more than one causal structure can instantiate the same symmetry structure. This allows us to recognize classes like "first-order reaction" as a natural kind, despite an apparent diversity in causal structure. Importantly, symmetry structures are discoverable in straightforward ways, making it possible to answer (P2) in the affirmative — there is a method for spotting projectible kinds. I do not claim this solution is exhaustive; there may well be other sorts of projecitble kinds. But I do claim that the dynamical kinds theory is better than its competitors at capturing the projectible systems and kinds pertinent to scientific induction.

2 Looking at the ground: Classifying systems in terms of symmetry structures

My aim in the remainder of this essay is to introduce a new way of understanding natural kinds that, in some respects, offers superior answers or at least the promise of superior answers to both questions about induction. This new perspective is best introduced by way of a metaphor.

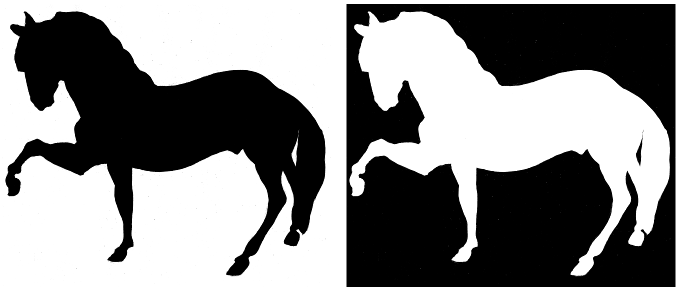

Fig. 1: In the left panel, the figure of the horse is shaded in black. On the right, the ground is shaded black.

2.1 Figure & Ground

M.C. Escher was a master of manipulating the perception of figure and ground. The term ‘figure’ refers to a region of visual space perceived as an object. Figure is always complemented in a scene by the background, or ‘ground’. In Escher’s *Metamorphosis* series, the colors perceived as figure and ground trade places as one’s gaze scans from left to right — the white gaps between black birds suddenly emerge as white fish with dark, bird-shaped gaps between them. In his *Mosaic II*, a menagerie of creatures lives behind another, each group struggling in our perception to emerge to the fore and be seen as objects. What at first appears as an odd space above a humanoid figure and at the focus of a ring of grotesque fish can be perceived a moment later as a white elephant surrounded by oddly shaped shadows. It is this duality, the possibility of describing one and the same shape as a positive figure or as the complement of the ground, that will serve as our metaphor for a new approach to natural kinds.

To have a concrete image in mind, consider the two sides of Figure 1, both of which depict the identical outline of a horse. To describe the shape of the horse, we could specify the dark shaded region in the left-hand panel. That is, we could specify what’s in the figure. Alternatively, we could characterize the very same shape — the horse — by describing the ground around it. That is, we could describe the contents of the region shaded black in the right-hand panel. As we’ll see, the content of a given natural kind whose members support system inductions can similarly be described in both positive and negative terms. It turns out that the negative characterization offers substantial benefits.

2.2 Figure: the causal connections amongst variables

Systems of causally connected variables are also amenable to dual characterization. To see how this could be, I must first say something about about what it is for two variables to be “causally connected.” What I have in mind is the interventionist account of James Woodward, a type of manipualtionist theory of causation. Since the account has been developed in great detail elsewhere (Woodward, 1997, 2000, 2001, 2003), I'll provide only enough of a sketch of the theory to suit the present purposes.

To begin with, the relation of causation is supposed to hold between variables. "...[V]ariables are properties or magnitudes that, as the name implies, are capable of taking more than one value. Values (being red, having a mass of 10kg) stand to variables (color, mass) in the relationship of determinates to determinables" (Woodward, 2003, p39). Causal relations between variables are defined, as the name of the account implies, in terms of *interventions* on variables. Briefly, "...an intervention on X (with respect to Y) is a causal process that directly changes the value of X in such a way that, if a change in the value of Y should occur, it will occur only through the change in the value of X and not in some other way "(Woodward, 2001). The general notion of cause I'll work with is what Woodward calls a *total cause*:[11] The variable X is a total cause of the variable Y if and only if there is a possible intervention on X that will change Y (Woodward, 2003, ch. 2). While precise, this notion is broad enough to capture the essence of interventionist approaches to causation, and from hereon I use"cause"as synonymous with total cause. To invoke a standard example, atmospheric pressure is a cause in this sense of both thunderstorms and the position of the dial on a barometer. Lowering the pressure results in a change in dial position. But the barometer dial position is not a cause of thunderstorms. Forcing your barometer dial to a different position while leaving atmospheric pressure unchanged does not result in a thunderstorm.

Interventionist causation admits of degrees in more than one sense. For one, the framework comfortably accommodates probabilistic causation, such that X causes Y just if intervening on X changes the probability distribution in values over Y. For another, the causal relation itself can be stable under a greater or lesser range of changes. There are two sorts of change to consider (Woodward, 2000): (i) stability under changes of background conditions, and (ii) interventions on the variables under consideration. With the respect to the former, the causal relation between X and Y may be limited in time and space to varying degrees. With regard to the latter, an association between X and Y may be more or less fragile with respect to interventions on either variable. If X is a cause of Y, then the association between X and Y will be stable under at least some interventions on X, though not necessarily all. However, if X and Y are commonly caused by Z, the correlation between X and Y is not stable under interventions on X. Woodward calls causal relations that are stable under at least some interventions "invariant generalizations." Invariant generalizations allow for system inductions. If X causes Y, then we can infer that, were X to assume a different value, so too will Y. Insofar as science involves successful system inductions, it involves discovery of invariant generalizations amongst particular variable instances (e.g., concentration of chemicals A and B in this

[11] I won't worry here about Woodward's definitions of direct, indirect, or contributing causes. While these are useful concepts, we can get far enough with direct causation alone.

particular beaker). Insofar as these generalizations can be extended to types of variable (e.g., concentration of chemicals A and B), science involves discovery of invariant generalizations amongst variable types.

2.3 Motivational examples

Given the above characterization of causation, one might attempt to characterize natural kinds directly in terms of invariant generalizations— in other words, causal kinds. Instead, for the reasons given above in discussing causal kinds, I am going to define kinds in terms of properties of the causal structure underlying such generalizations. More specifically, I will focus on actions that are irrelevant to the casual structure in a very particular way. But before introducing the technical machinery to make this proposal precise, it will help to have in mind a couple of genuine examples from diverse scientific domains of projectible kinds that have been characterized by what is irrelevant to the causal structure of their members. That is, it will help to have a couple of actual appeals to the causal 'ground' before introducing a theory of natural kinds predicated on this notion. To this end, I'll turn first to chemistry and then to physics.

2.3.1 Chemical solutions

Chemists are great taxonomists of natural kinds in the sense I've been using the term. There are many classes of chemical system that support both types of induction. And as the metaphor of figure and ground suggests, there are two complementary ways to characterize these systems. To focus on just one corner of a vast subject, consider the subfield of qualitative analytical chemistry, which Fresenius called "...one of the main pillars upon which the entire structure of the science rests" (1913, p2). The aim of chemical analysis is to ascertain the chemical constituents of an unknown substance. Qualitative analysis proceeds using a series of interventions whose consequences can be assessed without complex instrumentation, usually by observing the presence or absence of a precipitate in solution, color in a flame, the shade of a piece of litmus paper, or sometimes by detecting the scent of a gas with a recognizable odor, such as ammonia.

For many analytical tasks, there exist well-established procedures that take the form of decision trees: first add reagent A and check to see if a precipitate forms.[12] If there is a precipitate, follow this or that sub-procedure to identify it. If no precipitate forms, add reagent B, and so on. For example, the procedure

[12] Precipitation is the sudden formation of a solid from solution that tends to sink or float as a powder or collection of small crystals.

for identifying the cations (positively charged ions) of simple compounds[13] dissolved in water runs roughly as follows:[14]

1. Add hydrochloric acid to the solution.
   (a) If a precipitate forms, then the cation in solution belongs to 'Group I' which includes lead, silver, and mercury. There is a sub-procedure to follow to determine which of these it is.
   (b) If no precipitate forms, then go on to step (2).
2. Add hydrosulfuric acid.
   (a) If a precipitate forms, then the cation in solution belongs to 'Group II' which includes copper, lead, cadmium, bismuth, and mercury amongst others. There is a sub-procedure to follow to determine which of these it is.
   (b) If no precipitate forms, then go on to step (3).
3. Neutralize the solution with ammonium chloride and ammonium hydroxide, and add ammonium sulfide.
   (a) If a precipitate forms, then the cation in solution belongs to 'Group III' which includes aluminum, chromium, iron, manganese, zinc, cobalt, and others. There is a sub-procedure to follow to determine which of these it is.
   (b) If no precipitate forms, then the cation belongs to Groups IV or V which contain barium, strontium, forms of calcium, magnesium, sodium, and potassium.

The kinds I want to focus on are the so-called analytic groups (Group I, Group II,...). These are clearly kinds that support a variety of inductions. In particular, they behave similarly with respect to the presence of particular anions (negatively charged ions) in solution. What's true of one solution containing cations from Group II concerning many chemical properties (e.g., a disposition to form a precipitate) is true of the others. Viewed this way, each group is characterized by a small set of coarsely described causal relations. So, for instance, the cations of Group I are characterized by the fact that, when included in a system also composed of dissolved hydrochloric acid, they cause a precipitate to form.

But this is to focus on the figure. We could instead focus on the ground and understand the groups as characterized by interventions to which they are indifferent, manipulations which circumscribe the causal relations within each such system. Consider the cations of Group IV and V. Or rather, consider chemical systems — aqueous solutions in this case — that will be classified as containing cations from Groups IV or V under the above scheme. These systems can be described in terms of variables such as solute concentration and precipitate mass, or perhaps more finely in terms of cation concentration, anion concentration, and mass of undissolved solids. These causal relations

[13] A 'simple compound' is just a compound that is composed of one acid and one base or one metal and one non-metallic element.

[14] The full procedure can be found in (Fresenius, 1913). For a more modern presentation, see (Vogel, 1996).

are entirely uninfluenced by the interventions described in steps (1) through (3). That is, whatever the relation between the unknown cation concentration and precipitate mass, it is uninfluenced by the addition of hydrochloric acid, hydrosulfuric acid, ammonium chloride, ammonium hydroxide, or ammonium sulfide.

2.3.2 Wigner's particles

Non-relativistic quantum mechanics is probably the single most well-confirmed physical theory in history, and the 'elementary' particle types of quantum physics are projectible kinds if anything is. In a 1939 paper that has been referred to aptly as "epochal" (Gross, 1995), Eugene Wigner precisely delineated the possible kinds of particle in a way that was both theoretically fruitful and a paradigm example of what I have in mind by appealing to ground rather than figure.

Wigner's approach rests on the assumption that all physical systems are 'invariant' under the Poincaré group of transformations. Vaguely speaking, a transformation is a change of a physical system. For instance, rigid rotation of a system through fifteen degrees about a particular axis is a transformation. Transformations of physical systems can be composed; one transformation followed by another is itself a transformation. Often, these composition relations have the structure of an algebraic group: (i) for each composition there is an inverse that takes the system back to its starting state; (ii) there is a transformation that amounts to "do nothing;" (iii) only the order of composition matters, not how transformations are grouped; and (iv) the collection of transformations is closed under the action of composition. The Poincaré group is a collection of transformations with such a structure that includes rigid rotations, rigid translations, and "Lorentz boosts" — shifts of constant velocity. To say that physical systems are invariant under the Poincaré group is to say that each transformation in the group leaves the laws governing each physical system unperturbed. This invariance is why, for instance, a pendulum swings the same whether you move it ten feet to the left, or place it on a train moving at a constant speed.

Since every physical system is presumed to consist in part of causal relations that are invariant under the Poincaré transformations, Wigner realized that this invariance could be treated as a condition or definition of a very broad natural kind: physical system. More specifically, he equated each kind of elementary particle with the smallest kinds of quantum system that respect the Poincaré group.[15] Smallest here means that the space of states of the system does not contain a subset of states that is closed under the transformations of the Poincaré group.[16]

[15] More precisely still, Wigner equated each kind of elementary particle with one of the irreducible representations of the Poincaré group (which he calls the inhomogeneous Lorentz group).

[16] To be more accurate, Wigner equated irreducible representations of the Poincaré group with what he called an "elementary system." In a later paper with Newton, he argued that

I've glossed over a great deal of detail, but my aim is simply to motivate the approach to natural kinds expounded below. There are two things to stress. First, Wigner understood particles as systems of variables — position, four-momentum — that stand in causal relations with one another, not as substances with a set of essential properties. This accords with the way I have been speaking about the targets of scientific induction. System inductions involve generalizations about particular systems, while kind inductions — those which invoke natural kinds — generalize across such systems. The account of natural kinds I am urging similarly deals in systems. Second, Wigner sorts these systems into kinds according to what is causally *irrelevant*.

2.4 Ground: Dynamical symmetries

The vague and metaphorical talk about the 'ground' of a set of causal relations can be made concrete and precise via the notion of a *dynamical symmetry*.[17] What we're after is a precise characterization of the notion that a particular change or manipulation is invisible to the causal connections amongst a collection of variables — it is irrelevant to the determination of the joint state of the system.

2.4.1 The special case of time

Consider first the special case in which time is included amongst the variables of a system. I consider the case of temporal change first because the notion of a dynamical symmetry seems particularly straightforward in the case of causal systems for which time is considered a variable. This is partly an illusion of familiarity. The most important and best known dynamical symmetries in physics involve the evolution of a system through time. But there is also a genuine simplicity in supposing that it is time, rather than some other variable, that is supposed to advance in every experiment.

In plain language, a dynamical symmetry with respect to time is any intervention that commutes with the process of time evolution. Put differently, if we arrive at the same state by applying transformation σ and then evolving the system through time, or by first evolving the system and then applying σ, then σ is a dynamical symmetry. Note that a transformation need not preserve any first-order properties of the system in order to be a dynamical symmetry—it simply must commute with the process of evolving through time. Here is the idea stated precisely:

the category of "elementary particle" is narrower, and carries the additional condition that "...it should not be useful to consider the particle as a union of other particles"(Newton and Wigner, 1949, p 400). The latter condition turns out to be problematic, so I'm sticking with the broader category. This includes such things as hydrogen atoms in their ground state. What's not an elementary system? For one, any unstable particle (Wigner and Newton cite the π-meson) that exhibits a change in state that is not relativistically invariant.

[17] I am using "dynamical" in its broad sense of "active" or "potent", not in the narrow sense of involving forces.

Definition 1 (Dynamical symmetry with respect to time) Let S be the set of states of a system of variables (excluding time) and let $\Lambda_{t_0,t_1} : S \to S$ be the time-evolution operator which takes the system from state s_0 at t_0 to s_1 at t_1. A *dynamical symmetry with respect to time* is any operator $\sigma : S \to S$ with the following property:

$$\forall_{s\in S}\forall_{t_1>t_0}\forall_{t_0}\left[\Lambda_{t_0,t_1}(\sigma(s)) = \sigma(\Lambda_{t_0,t_1}(s))\right].$$

What does being a dynamical symmetry in time have to do with characterizing the causal structure of a system? Joe Rosen, who has carefully articulated this temporal notion of dynamical symmetry in a number of books (Rosen, 1975, 1995, 2008), puts it this way:

> A symmetry of the laws of nature is an indifference of the laws of nature. For a transformation to be a symmetry transformation of the laws of nature, the laws of nature must ignore some aspect of physical states, and the transformation must affect that aspect only. A pair of initial states related by such a transformation are treated impartially by the laws of nature; they evolve into a pair of final states that are related by precisely the same transformation. The laws of nature are blind to the difference between the two states, which is then preserved during evolution and emerges as the difference between the final two states. (Rosen, 1995, p139)

Where Rosen speaks of "laws of nature," I would substitute "invariant generalizations" or "time-dependent causal structures," but the point is the same. However the other variables in a system are causally related to one another through time, that relationship is blind to the changes induced by a dynamical symmetry.

For an extremely simple example, consider a growing population of bacteria described by only two variables: x, the number of bacteria in the growth medium and, t, the time. Assuming the bacteria grow unchecked, these two variables are connected by a simple relation: $x = x_0e^{rt}$. This system has a family of dynamical symmetry transformations of the form kx, one for each value of a positive, real-valued parameter k. In other words, if you multiply the initial population by k (by, say, adding bacteria from another stock) and then let it grow for 10 minutes or let it grow for 10 minutes then multiply the number of bacteria present by k, you end up with the same number of bacteria in the growth medium. Thus, scaling by a constant is a dynamical symmetry of the growth system—the bacterial growth process is insensitive to the absolute number of bacteria present.

2.4.2 The general case

The case of dynamical symmetries with respect to time was special both because each system considered includes time as a variable, and because that variable is singled out as a special index for states of the system. That is, we considered the effects of a transformation on states separated by a difference in time. But in general, system inductions involve causal relations amongst all sorts of variables, oftentimes without regard to time. We thus want a general notion of dynamical symmetry that captures not just symmetries of time evolution but of general causal relations. Fortunately, the notion of a dynamical symmetry is easy to generalize in order to characterize *all* of the ways in which causal structure can be blind to certain transformations, regardless of whether the system includes time as an explicit variable. The following definition expresses this generalization:

Definition 2 (Dynamical symmetry) Let V be a set of variables. Let σ be an intervention on the variables in $Int \subset V$. The transformation σ is a dynamical symmetry with respect to some index variable $X \in V - Int$ if and only if σ has the following property: for all x_i and x_f, the final state of the system is the same whether σ is applied when $X = x_i$ and then an intervention on X makes it such that $X = x_f$, or the intervention on X is applied first, changing its value from x_i to x_f, and then σ is applied.

For a simple, atemporal example, consider a hydrostatic system consisting of two pressure gauges mounted to a vertical rail and submerged in a sealed tank of water. Suppose that the vertical distance between them is adjustable (say by sliding the upper sensor up or down) and that the pressure in the tank can itself be adjusted. We might describe such a system in terms of three variables: p_1, the pressure at the location of the upper gauge, p_2, the pressure at the lower gauge, and h, the vertical distance between the gauges. These variables are causally related to one another (in the interventionist sense), and these relations can be expressed by the simple equation:

$$p_2 - p_1 = \rho g h \tag{1}$$

where ρ is the density of water and g is the gravitational acceleration near the surface of the earth. If we consider h to be the index variable, then all transformations of the pressure p_1 of the form $\sigma(p_1) = p_1 + c$ satisfy the symmetry condition. To see this, suppose we begin with the gauges at the same elevation. That is, we start with $h = 0$, $p_2 = p_1 = P$. Suppose we then increase the temperature in the tank, thus increasing p_1 to $P + c$. The pressure at p_2 is then $p_1 + \rho g h = P + c$. Now if we increase h to a value of $h = h_f$, then the state of the system is given by $p_1 = P + c$, $p_2 = P + c + \rho g h_f$. Suppose instead we first increase h to h_f, bringing the system to the state $p_1 = P$, $p_2 = P + \rho g h_f$. Were we to then increase the pressure in the tanks such that $p_1 = P + c$, we would end up in precisely the same final state as before. However, this is not the case for transformations of the form $\tau(p_1) = kp_1$. Thus, σ is a dynamical symmetry of the system while τ is not.

2.5 From symmetry structures to dynamical kinds

Dynamical symmetries can be combined to produce other dynamical symmetries. If σ_1 and σ_2 are both symmetries with respect to index variable v, then so is the operation of first applying transformation σ_1 and then applying σ_2. We can denote this composite operation by $\sigma_2 \circ \sigma_1$. Composition may be extended indefinitely; the composite of any two symmetries is a symmetry. Of course, it is generally not the case that the composite of two dynamical symmetries is a *new* dynamical symmetry. Rigid rotation is a dynamical symmetry of, say, the Earth-Moon gravitational system. And while it is true that the composite of any two rotations is also a symmetry, it is not true that each composite is a unique operation. For example, the composite of a rotation around a particular axis by 15° and a rotation around the same axis by 30° is the same as the transformation that rotates the system by 45°. To put it more concisely, $\sigma_{15^\circ} \circ \sigma_{30^\circ} = \sigma_{45^\circ}$. There is therefore structure to the set of transformations with respect to composition. We can identify this structure with the composition function itself.[18] This leads to the following definition:

Definition 3 (Symmetry structure:) The *symmetry structure* of a collection of dynamical symmetries, $\Sigma = \{\sigma_i | i = 1, 2, \ldots\}$ is given by the composition function $\circ : \Sigma \times \Sigma \to \Sigma$.

An example of a symmetry structure is the Poincaré group invoked by Wigner (Wigner, 1939) to characterize physical systems. It includes rigid rotations and translations. Transformations of these sorts compose in familiar ways in three spatial dimensions. It also includes the Lorentz boosts which have their own rules for composition.

Where does this leave us in the development of a theory of natural kinds? So far, I've introduced the notion of a dynamical symmetry which makes precise the intuitive notion of a causally irrelevant transformation. Any one system will exhibit a characteristic collection of these symmetries, and this collection has structure, an idea captured precisely in the notion of a symmetry structure. Already it should be clear that symmetry structures are complementary to causal structures. There is an intuitive sense in which describing a system's symmetry structure is like describing the ground behind the causal figure. But we need one more technical notion to obtain the promised theory of natural kinds. That's the notion of a non-trivial symmetry structure:

Definition 4 (Non-trivial symmetry structure) A non-trivial symmetry structure is one that contains the identity (i.e., "do nothing") transformation for each variable and choice of index, and, for at least one variable X relative to some index variable T, is *not* isomorphic to the group of mappings from X to itself (i.e., the set of all transformations of X).

The condition that a non-trivial symmetry structure contain the identity transformation ensures that the variables in such a system are causally complete, i.e., there are no external causes that influence the states of the system. The second part of the condition rules out the possibility that the variables instantiating the symmetry structure are causally disconnected. Non-trivial

[18] For finite sets of transformations, the composition function is equivalent to a multiplication table. But for infinite sets, we cannot write an explicit multiplication table.

symmetry structures are the basis of the sort of kinds that have all the desired features of natural kinds. That is, they circumscribe systems of variables that are causally related in such a way as to be projectible. They constitute the essences of *dynamical kinds*:

Definition 5 (Dynamical kind) A *dynamical kind* is a class of systems of variables that share a set of dynamical symmetries that are related by a non-trivial symmetry structure.

3 Dynamical kinds as natural kinds

3.1 Dynamical kinds, causes, and induction

The dynamical kind account offers complete answers to (P1) for both system and kind induction. Specifically, it asserts that the distinguishing feature of projectible systems is possession of a non-trivial symmetry structure, and that any given non-trivial symmetry structure picks out a projectible kind. I claimed at the outset that these answers were in some ways superior to those of the causal kinds account. To see why requires us to draw out some of the features of dynamical kinds.

First, possession of a non-trivial symmetry structure is sufficient to guarantee that a system has a non-trivial causal structure, in the sense that it entails at least one relation of total causation. The proof of this is trivial. Suppose that the variables in a set V exhibit a non-trivial symmetry structure, but that no variable in V is a total cause of another. Consider any three, X, Y, Z, where Z may be the same as either X or Y. Since none of the variables is a total cause of the others, there is no intervention on X that changes Y or Z and no intervention on Z that changes X or Y (if distinct from Z). So the order in which one transforms X by σ and intervenes on the index variable Z is irrelevant. In either order, the final state is the same. As a consequence, every possible transformation is a dynamical symmetry with respect to each index variable — the symmetry structure is trivial. However, the symmetry structure is, by hypothesis, not trivial, and so there must be at least one relation of total causation in the set. The upshot is that non-trivial symmetry structure is sufficient for securing the kinds of counterfactual robustness that is needed for system induction.

Second, if the variables in a system are causally connected, then the system will exhibit a non-trivial symmetry structure. In other words, possessing a non-trivial symmetry structure is necessary for a system to contain variables that stand in a relation of total causation, and thus for supporting system induction. This claim is only slightly more taxing to prove, but I leave some details to a footnote.[19] Put roughly, it follows from the fact that at least one

[19] If the variables are casually complete in the sense that there are no latent causal influences on the variables, then the identity transformation is a symmetry of the system for every choice of variable and index. Therefore, the system possesses a symmetry structure.

change in X produces a change in Y that we can find a transformation of Y that fails to be a symmetry. The easy choice is a transformation that always takes Y back to a particular state. It then clearly makes a difference whether we apply this transformation and then change X, or change X and then apply the transformation. Causal structure entails non-trivial symmetry structure. So membership in a dynamical kind is necessary and sufficient for possessing the kind of causal structure that underwrites projectibility in systems.

What of kind induction? It is the case that two systems with variables of the same type bearing the same causal relations must have identical symmetry structures. So sharing a symmetry structure is a necessary condition for sharing a causal structure. Importantly, the converse is not true; sharing a symmetry structure does *not* guarantee that two systems share the same causal structure. This is a boon, not a liability. It means that the extension of a kind defined by a symmetry structure—unlike a causal kind—can contain many specific but importantly related causal networks. How are they related? They share enough causal features to allow for limited kind inductions.

To sum up, symmetry structures are necessary and, to a limited extent, sufficient for induction, but any given dynamical kind may contain instances of a range of distinct (though related) causal structures. As an answer to (P1) then, dynamical kinds evades the excessive narrowness of the causal kinds account with respect to kind induction, while retaining its virtues vis-à-vis system induction. What remains to be shown, however, is that dynamical kinds offers something with respect to (P2) that can't be found in the HPC, SPC, essentialist, or causal kinds views. That issue will be taken up in Section 4. But first, some examples will help to give a sense of the way dynamical kinds carve up the world.

3.2 Examples of dynamical kinds

We have already encountered some dynamical kinds recognized as important by the scientific community. First, there are the analytical cation groups of qualitative chemistry. As I suggested above, these groups are in effect defined by symmetry structures on largely qualitatively valued variables (e.g., presence or absence of a precipitate). Take, for instance, Group III. This is a dynamical kind, each member of which is a system of qualitative variables with a non-trivial symmetry structure characteristic of the kind. If we treat the variable "concentration of hydrochloric acid" as the index variable (which can be in-

To show that it is not trivial, suppose that X is a total cause of Y. By definition, this means that there are some values x_0 and x_1, such that an intervention taking X from x_0 to x_1 induces the value of Y to change. In general, the dependence of Y on X can be described by a function, $f(x, y_0)$, where y_0 is the value of Y when $X = x_0$, i.e., $f(x_0, y_0) = y_0$. Since the change in X induces a change in Y, we know that $f(x_1, y_0) \neq y_0$. Consider a transformation, $\sigma(y) = y_0$. Applying σ and then changing X takes us from a state (x_0, y_0) to $f(x_1, \sigma(y_0)) = f(x_1, y_0)$. Changing X and then applying σ takes us from a state (x_0, y_0) to $\sigma(f(x_1, y_0)) = y_0 \neq f(x_1, y_0)$. Thus, σ fails to be a symmetry, and so the symmetry structure cannot be trivial.

cremented by intervention from “none” to “high”), then one of the dynamical symmetries characteristic of Group III is the transformation that changes the concentration of hydrosulfuric acid from “none” to “high”. Whether we first add hyrdosulfuric acid and then change the hydrochloric acid concentration, or add hyrdochloric acid first and then change the hydrosulfuric concentration, we obtain the same final state (a solution with no precipitate). However not all transformations are symmetries. It matters whether we first “neutralize the solution and add ammonium sulfide” and then “add hydrochloric acid” or first add the acid and only afterward neutralize and add the ammonium sulfide. The outcome (precipitate or no) will generally depend on the order.

For a paradigm example of a dynamical kind in terms of quantitative variables, one deeply entrenched in theoretical physics, consider the dynamical kind that Wigner equated with ‘physical system’. The symmetry structure that defines this kind is the Poincaré group. All rigid translations, rotations, and Lorentz boosts are dynamical symmetries of physical systems, with respect to time as the index variable. The behavior of the composition function which defines the symmetry structure in this case is complex but well studied (see, e.g., Haag, 1996). I won’t provide a full characterization, here. But, for instance, composition of rigid translations is given by simple vector addition.

Finally, consider an example from biology. Earlier I mentioned the exponential growth of some biological populations over time. Populations can be viewed as systems instantiating two variable types: population size and time. Populations exhibiting exponential growth belong to a dynamical kind that is defined by a rather simple symmetry structure. The structure contains one transformation of the form $\sigma_k(x) = kx$ for every positive, real value of k. The composition function is defined by the simple relation $\sigma_{k_2} \circ \sigma_{k_1} = \sigma_{k_1 k_2}$.[20]

3.3 Haven’t we been here before?

It’s no secret that symmetry is widely thought to provide the impetus and guide to theoretical development in modern physics (Cao, 2010; Gross, 1996, 1995; Sundermeyer, 2014). As such, the explication of symmetry has been a central concern of philosophers and philosophically-minded physicists (Brading and Castellani, 2003; Castellani, 2002; Rosen, 2008). It is also no secret that the term ‘symmetry’ covers a large number of distinct concepts (Hon and Goldstein, 2008; Mainzer, 1996). One might reasonably worry then whether the notion of a dynamical kind as I’ve introduced it—built as it is upon the idea of a dynamical symmetry—is merely a well-played idea in a shiny new case. Before exploring the putative advantages of the dynamical kinds view, I want briefly to make the case that the view is in fact novel.

Of necessity I cannot survey either the complex history of symmetry concepts, or the varied roles of symmetry in discussions of laws of nature and natural kinds, read broadly. Rather, I will restrict myself to potential answers

[20] Incidentally, the category of first-order chemical reactions mentioned earlier is a dynamical kind characterized by an analogous family of scaling transformations.

to (P1) and (P2). Relatively little has been said about what I've been calling system induction. Wigner (1967) famously stressed that laws of nature would never have been found were it not the case that they respect, at least approximately, some broad symmetries such as temporal and spatial translation. Put differently, it's unlikely we ever would have noticed sufficient stability in systems to successfully generalize over their states, let alone over groupings of such systems were it not for their indifference to absolute spatial and temporal locations. But Wigner never extended this line of reasoning beyond the space-time symmetries of physics. He certainly did not propose that we should in general recognize projectible systems and kinds on the basis of their symmetry structure.

With respect to categories appropriate for kind induction, there are only a handful of explicit taxonomies that invoke symmetry. One is the classification of crystals according to symmetries of the lattice structure, specifically the isometric transformations by which two portions of the crystal can be brought into congruence (see, e.g., Giacovazzo, 2011). But the symmetries involved in this case are not instances of what I've been calling dynamical symmetries. They concern operations that are only abstract — one never actually translates or rotates a chunk of the crystal lattice, one only transforms coordinates of a geometric representation of lattice. Another important taxonomy is Wigner's classification of particle types with irreducible representations of the Poincaré group (see Section 2.3.2 above). This ultimately sorts particles with respect to mass and spin. A finer classification based on values of the conserved quantities that correspond to so-called 'internal symmetries' of the electromagnetic, weak, and strong forces yields the taxonomy of the Standard Model. So within particle physics, it is true that something like what I've been calling symmetry structure is explicitly used to sort systems into kinds and thus answer (P1). However, as with system induction, this account of the projectible kinds—if we are to read it as such—does not extend beyond 'fundamental' physics. There has been no generalized attempt to apply this concept of grouping by symmetry outside the realm of particle physics, and thus no general answer to (P1).

It is with respect to (P2), the question of how we identify kinds that are likely to be projectible, that symmetry takes center stage in physics, albeit indirectly. A recipe often called the "gauge argument" has proven highly successful in developing quantum field theories of (most) of the fundamental interactions (electromagnetic, weak, and strong) between elementary particles (Martin, 2003; Moriyasu, 1983). Vaguely speaking, the idea is to take a non-interacting or 'free' field theory of, say, electrically charged particles, and follow a recipe for producing a theory that captures the dynamics of interaction. The first step in the recipe is to note the presence of a 'global' symmetry of the phase (a non-classical degree of freedom in quantum states). That is, scaling the phase of the field at every point in space by the same constant is a symmetry of the equations of motion. Next, we insist that this symmetry be extended to a 'local' symmetry, such that scaling the phase by a function of position is also a symmetry of the equations of motion. To satisfy this demand, we generally have to add a new field that interacts with the original free field and which transforms in a very particular way under a change of phase.The result is, after a bit of jiggering, a full theory of the dynamics of electromagnetism and charged matter. The logical structure, physical content, and philosophical importance of this gauge argument are all subject to intense debate (Belot, 2003; Earman, 2002; Healey, 2007; Martin, 2002, 2003; Redhead, 2003). Nonetheless, it is a clear example of symmetry being used to determine a set of possible causal structures. There is some similarity here with the way in which dynamical kinds relate to causal kinds in my view. Of course, there is little hope of extending the gauge argument outside of physics, and so it cannot represent a general answer to (P2).

The connection between causal laws and symmetry has not been lost on philosophers or practitioners of other sciences.[21] Charnov (1993), for instance, has used invariants—quantities that remain fixed for a system under the action of a symmetry transformation—to seek out new casual regularities in evolutionary ecology. Charnov's work, however, is the exception that proves the rule—almost no one is defending symmetry methods as general tools of discovery outside of physics. My project here thus differs in aim and scope from the bulk of the existing work on symmetry. I aim to provide a general theory of natural kinds, not a taxonomy for the objects of known or postulated laws. Nor am I trying to give methods of theory construction that apply only to particular sorts of physical theories. The scope of my account is intended to encompass inductive practices generally, not just those of fundamental physics.

Finally, I should say something about where the concept of 'dynamical symmetries' fits into the constellation of existing symmetry concepts. Wigner (1967) groups symmetries into two kinds, and his division has become standard. The first kind, which he calls the "geometrical principles of invariance," include the spacetime symmetries thought to apply to all physical laws (and thus all physical systems). These are the members of the Poincaré group. Insofar as we understand each transformation in an active sense as an actual change of location, orientation, or speed of a system, the geometric symmetries are all instances of what I've called dynamical symmetries. Symmetries of the second kind in Wigner's scheme are related to "...specific types of interaction, rather to any correlation between events" (Wigner, 1967, p17). These are the gauge principles or 'internal symmetries' sketched briefly above.[22] It is unclear whether gauge transformations should be viewed as physical transformations of a system. If not, these symmetries are *not* dynamical symmetries

[21] In particular, there is ample discussion of the use of symmetry arguments in solving physical problems, and of 'Curie's Principle'—the claim that asymmetries in an effect must be present in the cause (Ismael, 1997; Rosen, 1995; van Fraassen, 1990). These are rich topics, but well beyond the scope of this essay.

[22] Unfortunately, Wigner calls these "dynamic principles of invariance," and, as with virtually every term associated with symmetry, similar phrases have been used throughout the literature to refer to a variety of distinct notions. Obviously, I am not helping the situation with my choice of terminology. However, the phrase 'dynamical symmetry' seems best to express what I have in mind and, in many of it's previous uses, is not so far off from the concept presented here. I beg the reader's indulgence for further overloading the term.

in my sense. This is not to say that my dynamical symmetries are coextensive with Wigner's geometrical transformations. to the contrary, any symmetry of the dynamics of a system that involves physical transformation—regardless of whether it is tied to a particular type of interaction—counts as a dynamical symmetry in my view. There are many more of these besides the members of Poincaré group.

It is clear that dynamical symmetries do not quit fit the standard classification. The closest concept in the literature is Rosen's notion of a "symmetry of the laws of nature" (Rosen, 1995). However, our two notions of symmetry coincide only insofar as time is taken to be the index variable. To my knowledge, no one has generalized the notion to allow for arbitrary index variables. So aside from pursuing different aims, the account of dynamical kinds given here introduces some novel concepts as well. The question remaining is what benefit this novelty brings.

4 Discovery and dynamical kinds

I suggested at the outset that what sets the dynamical kinds theory apart from other characterizations of natural kinds is its immediate utility for achieving successful kind inductions. The fact that being a member of a dynamical kind is sufficient for a system to possess non-trivial causal structure suggests a method for determining whether a particular kind is of the right sort to be projectible. In fact, it suggests a general approach to discovering new projectible kinds. The idea is simply to look for dynamical symmetries. More specifically, suppose we have at our disposal a collection of instances of a range of variable types in some new domain of investigation. For any one system of variable instances, we can determine whether or not the system possesses a non-trivial symmetry structure. This is often easier to spot than a specific causal structure; it does not take a large sample size or many experiments to notice what has no effect on the final state of the system given a manipulation of the initial state. Once we strike upon a system that does have a non-trivial symmetry structure, we can be confident that the token system will support system induction. Furthermore, the symmetry structure associated with it defines a dynamical kind that in turn will support kind induction. That is, some properties of the token system will generalize to others of the same dynamical kind. Of course, not all properties will generalize. Many specific causal structures can exhibit the same symmetry structure, and it is impossible to know a priori which properties will generalize. For example, there are lots of distinct ways in which gravitating bodies can causally influence one another. Nonetheless, there are facts that generalize across all gravitational systems.

What I'm suggesting is that one can approach the discovery of new kinds of causal system in a manner similar to the way in which causal discovery algorithms allow for the discovery of detailed causal structure. The latter include a family of algorithms that infer which variables are direct causes of which by examining the statistical relations of conditional independence amongst the variables (Spirtes et al, 2000). In this case, I am suggesting that the presence and, to a large extent structure, of a system of causally related variables can be spotted by looking for non-trivial symmetry structures. Sometimes it's easier to find the forest than the trees. In the remainder of this section, I describe two ways of finding the forest.

4.1 The more things change...

The first approach to discovering dynamical kinds is direct: look for transformations that are dynamical symmetries of a given system. To illustrate with a concrete example, suppose we are attempting to learn about heat flow. Specifically, we're interested in the relation between temperature, u, position, x, and time t in a thin rod of length, L. Suppose we keep the ends of the rod at a fixed temperature, but insulate only the sides of the rod. We also deck it out with a dense array of thermistors to measure temperature and resistive heating elements to influence the temperature profile. Once we set a particular distribution of temperature along the rod at time t_0, we wait until some later time, t_1, and then measure the new distribution of temperature over the length of the rod. The hope, of course, is that by watching the temperature change, we can learn about the desired dynamical relations. But this is usually like searching for the proverbial needle in a haystack. It's hard enough to fit a curve to a given temperature profile let alone divine the functional relation between this profile and another. The first panel on the left of Figure 2 shows two different initial temperature profiles. The data is fake—created by adding Gaussian noise to a smooth profile—but the laws used to evolve the profiles through time to generate those appearing in the middle panel are realistic. The point is that, for either the circles or the squares, it's very difficult to discern the details of the relationship between the initial temperature profile on the left and the evolved profile on the right.

But suppose we know something about the relationship between the two initial profiles. That is, suppose we know the transformation that was applied to the initial profile marked with circles to get to that marked with squares. In this case, it's a simple scaling relation. The transformed initial conditions were obtained from the original by scaling the temperature by a constant factor all along the rod. To ascertain whether this transformation is in fact a symmetry, we need only ask how the final two temperature distributions in the middle panel relate. The panel on the right shows what we get if we plot the final temperature of the experiment involving transformed initial conditions (the squares) against the final values resulting from the original initial conditions (the circles). It's not hard to spot the linear relationship; the final distributions are related to one another by a simple scaling, just like the initial conditions. We've just discovered that scaling by a constant factor is a symmetry of our unknown dynamics. This puts significant constraints on what the details of those dynamics can look like. Furthermore, this success does not depend on some unusual feature of heat dynamics. Quite generally, symmetries of a given

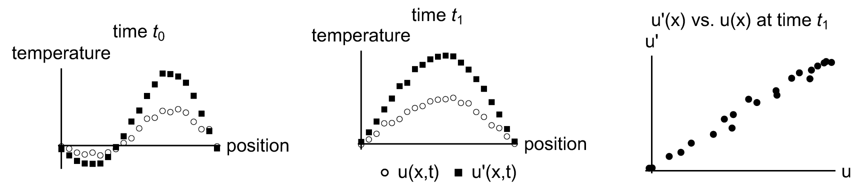

Fig. 2: *Left*: Plots of initial temperature at time t_0 as a function of position along a heated rod for two different experiments. The profile marked by squares is the result of a transformation that scales the temperature along the rod by a constant factor relative to the profile marked with circles. *Middle*: The temperature profiles that result at a later time, t_1, given the two initial profiles on the left. *Right*: A plot of the final temperatures in the experiment with transformed initial conditions (squares) versus the final temperatures resulting from the unperturbed experiment (circles). A straight line indicates that the final states are also related by a scaling transformation.

dynamics depend upon fewer free parameters than the dynamics themselves. This reduced complexity makes them easier to spot in systematic search, just like the linear relationship between the final temperature distributions in our two heat experiments.

I haven't the space for detailed case studies, but I can at least point to a couple of important discoveries made through the direct symmetry approach. The first is the role of 'isospin' in working out the dynamics of the strong force. The discovery of isospin was the discovery that whatever held nucleons together was approximately indifferent to their electrical charge—swapping a proton for a neutron is a dynamical symmetry of the then unknown strong force. This observation led to the theoretically fruitful approach of treating the neutron and proton as two states of the same particle.[23] Another example is the seminal work of Osborne Reynolds on the inception of turbulence in pipe flows (Reynolds, 1883). In a series of delicate experiments involving, amongst other things, streams of dye injected into the center of the flow of water through a glass pipe, Reynolds found that the velocity at which the flow became turbulent (v_c) was proportional to the kinematic viscosity of the water (ν) and inversely proportional to the diameter of the pipe (D). In other words, the number $\mathrm{Re} = \frac{v_c D}{\nu}$ is a constant. This relation encodes an entire family of symmetries of the unknown solutions to the equation of motion. One could, for instance, scale up the velocity while reducing the diameter, and end up with a flow that evolves in the same way. The symmetries encoded in Reynolds's scaling relation have been enormously useful in working out the details of viscous fluid flow in a wide range of circumstances.

[23] This story is recounted in Pais (1986).

4.2 ...the more they stay the same

The second approach to discovery through symmetry is somewhat indirect. This approach to discovery is built upon the concept of an invariant of the motion. An invariant of the motion is a function of the variables in a causal system the value of which does not change through time. For instance, linear momentum, angular momentum, and the Lagrangian are all invariants of the motion of systems obeying classical mechanics—they are all conserved through time. There is an intimate connection between invariants and dynamical symmetries. Emmy Noether[24] proved, coarsely speaking, that if a system exhibits a continuous (differentiable) family of dynamical symmetries with respect to time, then it possess an invariant of the motion. More broadly, the conservation of a non-trivial function of the dynamical variables implies the presence of a dynamical symmetry. So an alternate approach to discovering symmetry structures is to look for conserved quantities, for what stays the same.

This approach has already met with significant success. This is essentially the approach of Charnov (1993) in his work on invariants of organism life-histories. It's also the basis for some recent work in artificial intelligence. The automated discovery algorithm of Schmidt and Lipson (2009) takes time-series data and performs a search through a space of symbolic mathematical expressions, looking for non-trivial invariants of the motion. Not every invariant is associated with a family of dynamical symmetries. The most innovative aspects of Schmidt and Lipson's work involve techniques for sifting invariants to isolate those indicative of a dynamical symmetry, and thus informative about the causal structure of the system under investigation. The salient point is that they actually built a system for the automated discovery of projectible kinds on the basis of symmetries. Given that their method looks only for dynamical symmetries with respect to time and only for symmetries connected with invariants of the motion, this is only the beginning of what is possible.

I concede that the symmetry approaches to discovery I've described are more like sketchy promises than demonstrations. Only development and implementation of the methods will demonstrate conclusively that symmetry gives us a handle on projectibility, that dynamical kinds are the natural kinds. That is work for another time.

5 Conclusion

The questions at issue are simple: what distinguishes the projectible kinds from the non-projectible, and is there a method for identifying such kinds? The answers, of course, are notoriously obscure. I have tried in this essay to rough out a different perspective on the matter. With regard to the first question, I suggest that the projectible kinds are dynamical kinds. While related to previous accounts of so-called natural kinds, the central innovation of the dynamical kinds approach is a shift of focus from causal relations to properties of

[24] See (Neuenschwander, 2011) for a thorough discussion.

causal relations. Specifically, the emphasis is placed on dynamical symmetries — transformations to which a given causal structure is blind or indifferent. My emphasis on symmetry is new to the natural kinds literature, but hardly original in the art and philosophy of science. As Wigner (1967), Rosen (1995), and others have stressed, the presence and recognition of symmetries is essential to scientific discovery. Mechanics would be intractable if the laws of physics were not indifferent to the position, orientation, or temporal location of physical systems. Chemistry would be a chaos if every reaction depended upon every detail of the environment or the time of day.

I have argued, however, that symmetries can do more for us than merely set the stage for discovery. An explicit focus on dynamical symmetries illuminates otherwise mysterious aspects of inductive success. Dynamical symmetries hang together in symmetry structures. Non-trivial symmetry structures are both necessary and sufficient for the presence of the kinds of causal relations that support the counterfactuals needed for successful induction. Thus, it is no surprise that dynamical kinds — each of which is defined by a particular symmetry structure — are projectible kinds. What is a bit surprising is that once we recognize that the projectible kinds are dynamical kinds, this immediately suggests a number of methodological innovations. That is, adopting the view that induction is successful when carried out in terms of dynamical kinds suggests ways of modifying inductive practice. As I demonstrated in Section 4, they offer some promise for discovery. In many contexts it is easier to spot a symmetry or the invariant quantity associated with one than to discern a specific causal connection. Aside from the role symmetry has played in the development of modern physics, automated discovery systems have already been deployed that demonstrate the effectiveness of seeking the causal figure by looking at the ground.

To bring us back around to the questions with which we began, the dynamical kinds account provides answers to both (P1) and (P2). At least some projectible kinds are dynamical kinds, and we can identify them by means of the discovery methods described above. As a consequence, I claim that the dynamical kinds approach does a better job than its competitors at capturing many or most scientifically salient kinds. In a sense, this account of natural kinds is testable; it suggests adopting particular methods for efficient induction in scientific practice. Insofar as these methodological consequences bear fruit, the dynamical kinds approach is that much more plausible as an answer to (P2). Conversely, failure looks bad for the account. I do not have significant evidence along these lines to present here. The point I want to stress is that what evidence is already available favors the view.

References

Belot G (2003) Symmetry and gauge freedom. Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics 34(2):189–225

Boyd R (1988) How to be a moral realist. In: Sayre-McCord G (ed) Moral Realism, Cornell University Press, Ithaca, NY

Boyd R (1989) What realism implies and what it does not. Dialectica 43(1-2):5–29

Boyd R (1991) Realism, anti-foundationalism and the enthusiasm for natural kinds. Philosophical Studies 61(1/2):127–148

Boyd R (1999) Homeostasis, species, and higher taxa. In: Wilson RA (ed) Species : New Interdisciplinary Essays, MIT Press, Cambridge, Mass, pp 141–185

Brading K, Castellani E (eds) (2003) Symmetries in physics: philosophical reflections. Cambridge University Press, Cambridge, UK ; New York

Cao TY (2010) From current algebra to quantum chromodynamics: a case for structural realism. Cambridge University Press, Cambridge ; New York

Castellani E (2002) Symmetry, quantum mechanics, and beyond. Foundations of Science 7(1-2):181–196

Chakravartty A (2007) A metaphysics for scientific realism: knowing the unobservable. Cambridge University Press, Cambridge

Charnov EL (1993) Life history invariants: some explorations of symmetry in evolutionary ecology. Oxford series in ecology and evolution, Oxford University Press, Oxford [England] ; New York

Craver CF (2009) Mechanisms and natural kinds. Philosophical Psychology 22(5):575–594

Earman J (1985) Concepts of projectibility and the problems of induction. Noûs 19(4):521

Earman J (2002) Gauge matters. Philosophy of Science 69(S3):S209–S220

Ellis B (2005) Physical realism. Ratio 18(4):371–384

Ellis BD (2001) Scientific essentialism. Cambridge studies in philosophy, Cambridge University Press, New York

van Fraassen BC (1990) Laws and Symmetry. Oxford University Press, Oxford ; New York

Fresenius CR (1913) Manual of Qualitative Chemical Analysis, 16th edn. John Wiley & Sons, New York

Giacovazzo C (ed) (2011) Fundamentals of crystallography, 3rd edn. No. 15 in IUCr texts on crystallography, Oxford University Press, Oxford ; New York

Goodman N (1955) Fact, fiction and forecast. Harvard University Press, Cambridge

Griffiths PE (1999) Squaring the circle: Natural kinds with historical essences. In: Wilson RA (ed) Species: New Interdisciplinary Essays, MIT Press, Cambridge, Mass, pp 209–228

Gross D (1995) Symmetry in physics: Wigner's legacy. Physics Today 48:46–50, 12

Gross DJ (1996) The role of symmetry in fundamental physics. Proceedings of the National Academy of Sciences of the United States of America 93(25):14,256–14,259

Haag R (1996) Local quantum physics: fields, particles, algebras, 2nd edn. Texts and monographs in physics, Springer, Berlin ; New York

Hawley K, Bird A (2011) What are natural kinds? Philosophical Perspectives 25(1):205–221

Healey R (2007) Gauging what's real: the conceptual foundations of contemporary gauge theories. Oxford University Press, Oxford ; New York

Hon G, Goldstein BR (2008) From Summetria to Symmetry: The Making of a Revolutionary Scientific Concept. Springer

Ismael J (1997) Curie's principle. Synthese 110(2):167–190

Kornblith H (1993) Inductive inference and its natural ground: an essay in naturalistic epistemology. MIT Press, Cambridge, Mass

Kripke SA (1980) Naming and necessity. Blackwell, Oxford

Lange M (2009) Laws and lawmakers: science, metaphysics, and the laws of nature. Oxford University Press, Oxford ; New York

Laudan L (1981) Science and Hypothesis. D. Reidel Publishing Company, Dordrecht, Holland

Lewis D (1983) New work for a theory of universals. Australasian Journal of Philosophy 61(4):343–377

Lowe EJ (2006) The four-category ontology: a metaphysical foundation for natural science. Clarendon Press, Oxford

Mainzer K (1996) Symmetries of nature: a handbook for philosophy of nature and science. Walter de Gruyter, New York

Martin C (2002) Gauge principles, gauge arguments and the logic of nature. Philosophy of Science 69(S3):S221–S234

Martin CA (2003) On continuous symmetries and the foundations of modern physics. In: Brading K, Castellani E (eds) Symmetries in physics: philosophical reflections, Cambridge University Press, Cambridge, UK ; New York, pp 29–60

Mattingly J, Warwick W (2009) Projectible predicates in analogue and simulated systems. Synthese 169(3):465–482, ArticleType: research-article / Full publication date: Aug., 2009 / Copyright © 2009 Springer

Mill JS (1858) A system of logic, ratiocinative and inductive. Harper & brothers, New York

Moriyasu K (1983) An elementary primer for gauge theory. World Scientific, Singapore

Neuenschwander DE (2011) Emmy Noether's wonderful theorem. Johns Hopkins University Press, Baltimore, Md.

Newton TD, Wigner EP (1949) Localized states for elementary systems. Reviews of Modern Physics 21:400–406, 3

Pais A (1986) Inward bound: of matter and forces in the physical world. Clarendon Press ; Oxford University Press, Oxford [Oxfordshire] : New York

Pearl J (2000) Causality : models, reasoning, and inference. Cambridge University Press, Cambridge, U.K.; New York

Popper KR (2007) The logic of scientific discovery. International Society for Science and Religion, Cambridge

Putnam H (1973) Meaning and reference. The Journal of Philosophy 70(19):699–711

Putnam H (1975) Mind, language, and reality. No. v. 2 in His Philosophical papers, Cambridge University Press, Cambridge [Eng.] ; New York

Quine WV (1969) Ontological relativity, and other essays. No. no. 1 in The John Dewey essays in philosophy, Columbia University Press, New York

Redhead M (2003) The interpretation of gauge symmetry. In: Brading K, Castellani E (eds) Symmetries in physics: philosophical reflections, Cambridge University Press, Cambridge, UK ; New York, pp 124–139

Reynolds O (1883) An experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and of the law of resistance in parallel channels. Philosophical Transactions of the Royal Society of London 174:935–982

Rosen J (1975) Symmetry discovered: concepts and applications in natureand science. Cambridge University Press, Cambridge ; New York

Rosen J (1995) Symmetry in science : an introduction to the general theory. Springer-Verlag, New York

Rosen J (2008) Symmetry rules how science and nature are founded on symmetry. Springer, Berlin

Schmidt M, Lipson H (2009) Distilling free-form natural laws from experimental data. Science 324(5923):81–85

Slater M (2013) Natural kindness. The British Journal for the Philosophy of Science (forthcoming)

Spirtes P, Glymour CN, Scheines R (2000) Causation, prediction, and search, 2nd edn. Adaptive computation and machine learning, MIT Press, Cambridge, Mass

Sundermeyer K (2014) Symmetries in fundamental physics, Fundamental Theories of Physics, vol 176. Springer, Berlin

Venn J (1866) The logic of chance. Macmillan and co., London and Cambridge

Vogel AI (1996) Vogel's qualitative inorganic analysis, 7th edn. Longman, Harlow

Wigner E (1939) On unitary representations of the inhomogeneous lorentz group. The Annals of Mathematics 40:149–204, 1

Wigner EP (1967) Symmetries and reflections: scientific essays of Eugene P. Wigner. Indiana University Press, Bloomington

Wilkerson TE (1988) Natural kinds. Philosophy 63(243):29–42, ArticleType: research-article / Full publication date: Jan., 1988 / Copyright © 1988 Royal Institute of Philosophy

Williams NE (2011) Arthritis and nature's joints. In: Campbell JK, O'Rourke M, Slater MH (eds) Carving nature at its joints: natural kinds in metaphysics and science, Topics in contemporary philosophy, MIT Press, Cambridge, Mass

Wilson RA (1999) Realism, essentialism, and kind: Resuscitating species essentialism? In: Wilson RA (ed) Species: New Interdisciplinary Essays, MIT Press, Cambridge, Mass, pp 187–207

Woodward J (1997) Explanation, invariance, and intervention. Philosophy of Science 64:S26–S41

Woodward J (2000) Explanation and invariance in the special sciences. The British journal for the philosophy of science 51:197–254, 2
Woodward J (2001) Law and explanation in biology: Invariance is the kind of stability that matters. Philosophy of Science 68:1–20, 1
Woodward J (2003) Making things happen : a theory of causal explanation. Oxford University Press, New York