

# Inferring Conservation Principles in Particle Physics: A Case Study in Reliable and Efficient Inquiry

Oliver Schulte

Department of Philosophy, University of Alberta

August 29, 2000

## Abstract

This paper applies learning-theoretic analysis to an inductive problem that arises in particle physics: how to infer from observed reactions conservation principles that govern all reactions among elementary particles. I describe a reliable inference procedure that is guaranteed to arrive at an empirically adequate set of conservation principles as more and more evidence is gathered. In certain circumstances, finding an empirically adequate conservation theory requires positing *hidden* particles. The paper describes learning-theoretic conceptions of empirical success in addition to reliable convergence to a correct theory, which determine an essentially unique optimal inductive procedure for the particle dynamics problem.

## 1 Means-Ends Epistemology and Scientific Practice

Kelly's and Martin and Osherson's contributions to this symposium introduce many of the general philosophical and mathematical ideas behind formal learning theory. The aim of this paper is more particular: to apply learning-theoretic analysis in a specific inference problem that arises in scientific practice. The problem is how to find general conservation principles that correctly predict what reactions are possible among elementary particles. The analysis will illustrate many of the general ideas laid out in the other contributions. Moreover, elementary particle dynamics is an important domain of current physical inquiry, and illuminating its methodology is just the kind of application that our approach to the philosophy of science aims for.

I will employ this case study to elaborate several more learning-theoretic ideas: How to take account of epistemic aims *in addition to* reliable convergence to the truth; that empirical complexity, or inductive underdetermination, comes in degrees; and how to construct a hierarchy of ever more stringent conceptions of empirical success. Indeed, on my approach these ideas are intimately related.

For setting our sights on a certain epistemic aim gives rise to a standard of empirical success. And we may define the empirical complexity of an inductive problem in terms of the standards of empirical success that are feasible in the problem—the more complex the problem, the less we can expect inquiry to achieve. Specifically, I shall consider *fast and steady convergence* to a correct theory as epistemic aims in addition to reliable convergence. We might think of these additional aims as *efficiency* criteria for inquiry.

The first question that learning theory asks about an inductive problem is whether there is an inductive method that is guaranteed to find the right answer, as more and more evidence is gathered. In the particle domain it turns out that without further background assumptions, the answer is negative: particle dynamics is potentially so complicated that no method is guaranteed to settle on an empirically adequate account of it, even in the limit of inquiry.

However, the inductive complexity of an empirical problem is never absolute, but always relative to the background assumptions that the inquirer is willing to entertain. Since background assumptions depend largely on the specific context of inquiry, this is a crucial point at which the details of the given domain of inquiry enter into learning-theoretic analysis. Is there a plausible assumption that would reduce the empirical complexity of particle dynamics to the point at which reliable inquiry is feasible? Indeed there is: If we suppose that some set of conservation principles is empirically adequate, then there is an inference procedure that is guaranteed to find a correct set of conservation principles as more and more particle reactions are observed. The focus on conservation principles is a prominent feature of particle physics that many commentators have remarked on.

If we assume further that all particles are detectable, then reliable and efficient convergence to an empirically adequate conservation theory is possible. More precisely, there is an inference procedure that minimizes convergence time and succeeds with an a priori bounded number of theory changes. However, if we allow that hidden particles may be present (though not infinitely many of them) reliable inquiry into conservation laws is still feasible, but there is no a priori bound on the number of theory changes that may be required.

These results illustrate how stronger assumptions permit inquiry to realize more epistemic aims. The idea that a framework of alternative theories, such as conservation principles, provides guiding constraints to inquirers is not new to science scholars. Kuhn identifies constraints on theoretical alternatives as one of the “paradigm rules” that govern “puzzle-solving” in normal science (1970). He also speaks rather deprecatingly of “trying to force nature into boxes” (Kuhn 1970, p. 24). A more nuanced passage that fits my case study well is the following.

Previously, we had principally examined the paradigm’s role as a vehicle for scientific theory. In that role it functions by telling the scientist about the entities that nature does and does not contain and about the ways in which those entities behave. That information provides a map whose details are elucidated by mature scientific

research. And since nature is too complex and varied to be explored at random, that map is as essential as observation and experiment to science's continuing development. (Kuhn 1970, p. 109)

The result that without further background assumptions, there is no reliable method for finding an empirically adequate theory of particle reactions may be seen as a precise learning-theoretic formulation of the idea that the particle world is “too complex and varied to be explored at random”. And the current paradigm in particle physics tells scientists about “the ways in which [elementary particles] behave”, namely in ways that follow the logic of conservation principles. We thus have a space or “map” of alternative conservation theories. Reliable inquiry takes us to the part of the space, or the place on the map, where the actual world is located. The details of this paper show how exactly it guides us there.

## 2 The New Conservation Principles in Particle Physics

Whereas classical conservation principles such as conservation of energy, momentum, and electric charge rule out many reactions among elementary particles, they do not capture the full range of the phenomena. Thus physicists have sought new principles, typically also in the form of conservation principles, that give the correct predictions. Nobel Laureate Leon Cooper summarized the situation as follows.

In the analysis of events among these new particles, where the forces are unknown and the dynamical analysis, if they were known, is almost impossibly difficult, one has tried by observing what does not happen to find selection rules, quantum numbers, and thus the symmetries of the interactions that are relevant. What has gradually come out of this analysis is that in addition to the quantities, which are always thought of as being conserved (momentum, energy, heavy particle number, and so on), there are quantities such as strangeness and isotopic spin, which are conserved by some interactions and not others. No one knows why. (Cooper 1970, p. 458)

As Cooper explains, physicists have dealt with the high degree of inductive complexity in the particle domain by ‘trying to find selection rules’, that is, new conservation principles that have ‘gradually’ emerged from analyzing enough data. How difficult is it to find empirically adequate conservation principles? In contrast with many of the traditional conservation principles, there is often no a priori symmetry argument for the new selection rules. As Cooper puts it, ‘no one knows why’ some of them hold for certain classes of reactions. Similarly, Williams comments with regard to the new principle of conservation of lepton number that ‘this lepton number conservation is arbitrary and has no basis in

more fundamental ideas' (Williams 1997, p. 285). Despite the lack of theoretical guidance towards conserved quantities, physicists consider it fairly easy to find conservation principles to account for given data. In his graduate text on particle physics Omnes explains 'once and for all' a method for inferring selection rules (1971, p. 36). Feynman remarks that 'the reason why we make these tables [of conserved quantities] is that we are trying to guess at the laws of nuclear interaction, and this is one of the quick ways of guessing at nature' (1965, p. 67).

Learning theory supplies the tools for examining and sharpening such judgments about the difficulty of particle inquiry under various assumptions. To apply the framework, I cast the particle inquiry task in the mould of an inference problem. As always in building a model, different models may be appropriate depending on what aspect of the phenomenon we wish to emphasize. I will proceed with a straightforward formulation (which is essentially due to Raul Valdés-Pérez) for now and consider some variations later. More of the issues pertinent to formulating learning-theoretic models of particle inquiry are discussed in (Schulte forthcoming); see also (Valdés-Pérez and Erdmann 1994; Valdés-Pérez and Żytkow 1996).

The basic components of learning-theoretic models are (1) a set of evidence items, (2) a set of hypotheses under investigation, and (3) a correctness relation that specifies which hypotheses are correct on a given *total*, usually infinite, body of evidence. In the particle setting, I take the evidence items to be reactions among elementary particles. There are two types of reactions that are of interest in particle physics, decays and collisions. A decay is a transition from one particle into others; a collision is a transition from two particles into others. To describe reactions, I will use the familiar *arrow notation*. In the arrow notation, we write the reagents on the left side of an arrow and the products on the right. For example, consider a collision of two protons that produces two protons plus a pion. Arrow notation renders this process as  $p + p \rightarrow p + p + \pi$ . For another example,  $\mu^- \rightarrow e^- + \nu_\mu + \bar{\nu}_e$  denotes a typical decay of a muon. A **reaction stream**  $\varepsilon$  is an infinite sequence of reactions. The initial sequence of the first  $k$  reactions in a reaction stream is denoted as  $\varepsilon|k$ .

The set of hypotheses under investigation are classifications of reactions into 'possible' and 'impossible'. I refer to such hypotheses as *theories of particle reactions*. A theory of particle reactions is empirically adequate if it (1) classifies all observed reactions as possible and (2) classifies reactions that are never observed as impossible. In addition to empirical adequacy, there are many virtues that we might desire in our theories, such as being simple, computable, or systematic (see Kelly on hypothesis virtues (1996, Ch.3.1) or Goldman on standards of epistemic evaluation (1986, Ch.1.1)). Learning-theoretic models incorporate such virtues in the correctness relation, which specifies what theories count as the right answer for a given total body of evidence. In this paper, I focus on the problem of finding an empirically adequate reaction theory. (Schulte forthcoming), Sections 9 and 10, discusses further constraints on the adequacy of reaction theories that stem from parsimony and systematicity. (Schulte forthcoming) also supplies a completely formal specification of the inference problem;

for my current purposes, the given level of detail is sufficient.

### 3 Reliable Inquiry into Particle Reactions

So far I have described the inductive problem of finding an empirically adequate theory of what transitions among elementary particles are possible. Now I formulate a notion of what it is for an investigative method to solve this problem.

An **inference rule** for the particle reaction problem is a function  $\delta$  that assigns a conjecture to each piece of evidence. Formally,  $\delta(e) = T$ , where  $e$  is some list of observed reactions and  $T$  is a reaction theory. I will also refer to inference rules as ‘inductive methods’, ‘inference methods’ or simply ‘theorists’.

An inference method **converges to an empirically adequate theory** on a given sequence of reactions **by time**  $t$  just in case the method’s theory is empirically adequate for the reaction stream  $\varepsilon$  at time  $t$  and at all times thereafter. Formally, success on a reaction stream  $\varepsilon$  by time  $t$  requires that for all times  $t' \geq t$ , it is the case that  $\delta(\varepsilon|t')$  is empirically adequate for  $\varepsilon$ . An inference method converges to an empirically adequate theory on a given reaction stream  $\varepsilon$  just in case there is a time  $t$  by which the method converges to an empirically adequate theory for  $\varepsilon$ . In that case I often simply say that the method *succeeds* on the reaction stream  $\varepsilon$ . A method is **reliable** under given background knowledge just in case the method succeeds on every reaction stream  $\varepsilon$  that is consistent with the background knowledge.

From here on, I will consider the situation of an inquirer who is investigating the dynamics of a fixed set of  $n$  observable particles (and possibly some hidden ones). We can expand the model to accommodate new particle discoveries, but this would lead to complications without much further insight into the methodology of conservation principles. Even with the restriction to a fixed set of observable particles, my first result is negative: without further background assumptions, there is no reliable method for the reaction inference problem.

**Proposition 1** *There is no inductive method that converges to an empirically adequate theory of particle reactions on every infinite sequence of particle reactions.*

There are two ways of demonstrating Proposition 1. First, let us consider how many theories of particle reactions there are. A reaction is a finite sequence of particles; since there are countably many sequences drawn from a finite set, there are potentially countably many reactions. A reaction theory is a function from the set of reactions to the set {possible, impossible}; since there are countably many reactions, it follows from basic facts of set theory (Cantor’s diagonal argument) that there are uncountably many reaction theories. But there are only countably many finite evidence sequences, so for any given inquirer, there must be (uncountably) many theories of particle reactions that the inquirer will never produce no matter what the evidence is. If one of these theories is in fact correct, then the inquirer will never produce it, much less converge to it. Thus

every inductive method will fail to converge to an empirically adequate theory on some reaction stream.

This argument shows that the sheer *size* of the set of alternative theories is a measure, albeit crude, of the complexity of inquiry. The reach of inquiry, even if carried on indefinitely, never goes beyond countably many theories. But typically, there are uncountably many possible worlds that might be actual, as there are uncountably many conceivable particle worlds. Hence inquiry can explore only a small part of the conceivable terrain of nature, and we need further assumptions to guide its focus, to avoid “exploring at random”. This observation too suits Kuhn’s description of the constraints imposed by the paradigm: “...the range of anticipated, and thus of assimilable, results is always small compared with the range that imagination can conceive” (Kuhn 1970, p. 35). At a minimum, we need to restrict the set of alternative hypotheses among which inquiry is to decide to the same size as the set of natural numbers.

However, though necessary, restricting the set of alternative hypotheses to countable size is by no means sufficient for reliable inquiry. This is because a countable set of alternative accounts of nature can still contain great “complexity and variety”. For countable sets of alternative theories, we may establish the impossibility of reliable inquiry by means of a *diagonal* argument along the lines of Putnam (1963, 1965, 1975). A diagonal argument shows how, given a method of inquiry, an adversary—an inductive cousin of Descartes’ demon—can present evidence in such a way that the method does not succeed in settling on a correct hypothesis. The demon does not personify any kind of ‘malice’ on the part of nature, but rather the ignorance of the inquirer. The more background knowledge the inquirer has, the fewer options the inductive demon has for leading him into false beliefs. Proposition 1 does not assume any background knowledge on the theorist’s part, so the demon is free to present the scientist with any data whatsoever. In (Schulte forthcoming) I describe a typical diagonal argument showing that even a small question about particle dynamics cannot be reliably settled, even in the limit of inquiry, without further background assumptions. (The question is: how many pions can be produced in a collision of two protons?) Since particle physics addresses many questions of comparable complexity, it is clear that reliably finding the right answer to all these questions requires further background assumptions. The operative background assumption in current particle physics is that particle dynamics is governed by conservation principles (Cooper 1970, p. 458), (Ford 1963, Ch.4), (Ne’eman and Kirsch 1983). Let us see how this assumption reduces inductive complexity.

## 4 Reliable Inquiry Into Conservation Principles Without Hidden Particles

Table 1 shows a number of (types of) particles and some of their conserved properties. These quantities are the same for each type of particle in all reactions (unlike, say, momentum). From now on, I will use the term ‘quantity’ to refer to

	Particle	Charge	Baryon #	Lepton #	Electron #	Muon #
1	$K^-$	-1	0	0	0	0
2	$K^+$	1	0	0	0	0
3	$K^0$	0	0	0	0	0
4	$\bar{K}^0$	0	0	0	0	0
5	$n$	0	1	0	0	0
6	$\bar{n}$	0	-1	0	0	0
7	$\bar{p}$	-1	-1	0	0	0
8	$p$	1	1	0	0	0
9	$\pi^-$	-1	0	0	0	0
10	$\pi^+$	1	0	0	0	0
11	$\pi^0$	0	0	0	0	0
12	$\gamma$	0	0	0	0	0
13	$\mu^+$	1	0	-1	0	-1
14	$\mu^-$	-1	0	1	0	1
15	$e^-$	-1	0	1	1	0
16	$e^+$	1	0	-1	-1	0
17	$\nu_e$	0	0	1	1	0
18	$\bar{\nu}_e$	0	0	-1	-1	0
19	$\nu_\mu$	0	0	1	0	1
20	$\bar{\nu}_\mu$	0	0	-1	0	-1

Table 1: Quantum Number Assignments applied in Selection Rules

such process-invariant properties of particles. I will consider only conservation principles that involve process-invariant quantities (also called *selection rules* in particle physics).

A reaction conserves a quantity just in case the total sum of the quantity over the reagents is the same as the total sum over the products. For example, the reaction  $p + p \rightarrow p + p + \pi$  conserves Baryon number, since the Baryon total of the reagents is  $2 \times \text{Baryon}\#(p) = 2 \times 1$ , and the Baryon total of the products is  $2 \times \text{Baryon}\#(p) + \text{Baryon}\#(\pi) = 2 \times 1 + 0$ . A given reaction is possible according to a conservation theory only if it conserves all quantities that define selection rules.<sup>1</sup>

Now let us see how the use of conservation principles resolves the underdetermination of theory by evidence. Suppose that other resources from general physical theory, known conservation principles, etc., do not suffice to answer the question of how many pions can be produced in a collision of two protons. Suppose further that we assume that the answer must lie with conservation principles governing collisions of two protons. Now if we observe the reaction  $p + p \rightarrow p + p + \pi$ , we infer that *whatever* conservation principles govern collisions of two protons, the pion  $\pi$  must carry 0 of any conserved quantity, because clearly the two protons on the left and on the right put the same weight into the conservation balance. But if the pion  $\pi$  carries 0 of every conserved quantity, then two protons may produce any number of pions without violating a conservation law. Thus after observing one reaction such as  $p + p \rightarrow p + p + \pi$  that produces pions, we can deduce that any number of pions can be produced in a collision of two protons—which is basically what current particle theory tells us.

This example shows the power of conservation principles to resolve underdetermination: Without further assumptions, there is no reliable way for settling the pion problem, no matter how much evidence the scientist gathers (see (Schulte forthcoming, Sec.4)); under the assumption that some kind of conservation theory is empirically adequate, *one* observation entails the right answer. What about all the other questions concerning particle reactions? It can be shown that if we assume that some set of conservation principles is empirically adequate, and that no hidden particles are present in observed transitions, then there is a reliable inference method that is guaranteed to settle on an empirically adequate theory of particle reactions. Moreover, the method is of a low complexity in the sense that if we are investigating the dynamics of  $n$  particles, the method changes its conservation principles at most  $n$  times. The number of times that a method may have to change its mind is a useful measure of inductive complexity that learning theorists have employed going back to Hilary Putnam's work in the 60s (Putnam 1965). We may think of this number as measuring the extent to which methods guide us towards *stable* belief (cf. (Sklar 1975)).

---

<sup>1</sup>Some numeric conservation principles can be violated in certain kinds of interactions. For example, strangeness is not preserved in weak interactions (cf. (Feynman [1965], p. 68)). In this paper I leave aside selection rules that make distinctions among the possible reactions.

**Proposition 2** *Assume that there is an empirically adequate set of conservation principles, and that no hidden particles are present in any observed transition. Then there is a reliable inference method that is guaranteed to converge to an empirically adequate conservation theory. Furthermore, for a set of  $n$  particles under investigation, this inference method revises its conservation principles at most  $n$  times.*

The proof of Proposition 2 is a fairly straightforward argument in linear algebra. The techniques of linear algebra become available for our methodological analysis if we represent as vectors reactions and quantum properties that define conservation principles. Representing vectors as reactions is standard practice in chemistry (Aris 1969). Valdés-Pérez adapted the idea for particle reactions (Valdés-Pérez and Erdmann 1994). (Schulte forthcoming) describes the vector representation and the proof of Proposition 2.

## 5 Reliable Inquiry Into Conservation Principles With Hidden Particles

How does inquiry fare in particle worlds in which not all participants in a reaction are observed? At first glance, one might think that this will make no difference to our analysis. For so far I have discussed only how inquiry can guide us towards *empirically adequate* theories, without requiring that inquiry find *true* theories. It seems as though inquirers whose aim is just to “save the phenomena” need not concern themselves with theoretical entities. *Human* inquirers might like to make reference to theoretical entities for reasons other than saving the phenomena; they might for example feel that theoretical entities yield better explanations, or that they unify phenomena, or that they make it easier to derive predictions (see (Van Fraassen 1980), (Hempel 1965, III.8)). But nothing in my notion of an inductive method restricts methods to those that conform to our psychology and intuitions (though learning theorists have considered such restrictions in detail elsewhere). So if theoretical entities serve values other than empirical adequacy, reliable methods for finding empirically adequate theories should have no need to deal with hidden entities.

However, this argument is fallacious: an account of hidden entities and their properties can be *indispensable* for particle inquiry to arrive at an empirically adequate theory. Let’s give an example to illustrate this point for particle dynamics. The example is simple (and fictitious) but it illustrates the principle.

Consider two particles, say  $K$  and  $\mu$ . Suppose that the transitions  $K \rightarrow \mu$  and  $K+K \rightarrow K+K+\mu+\mu$  were observed. If the process  $K+K \rightarrow K+K+\mu+\mu$  conserves a quantity, then the  $\mu$  particle must carry 0 of that quantity. But then the  $K$  particle must carry 0 of the quantity too since  $K \rightarrow \mu$  conserves the quantity. Hence any conservation theory consistent with these two observations must assign the value 0 for every quantum property to both  $K$  and  $\mu$ . Then any process whatsoever involving the  $K$  and  $\mu$  particle will be consistent with a given conservation theory. Such a conservation theory would be trivial with respect

to these two particles, and empirically inadequate if there is some transition among them (e.g.,  $K + K \rightarrow \mu$ ) that is never observed.

One response to this observation is that if the two transitions  $K \rightarrow \mu$  and  $K + K \rightarrow K + K + \mu + \mu$  were observed, we would give up the supposition that some set of conservation principles governs the particle world. This possibility shows that this supposition has definite empirical content and is defeasible (for a precise characterization of its empirical content, see (Schulte forthcoming)). Kelly’s contribution analyzes in detail the role of such defeasible presuppositions in reliable inquiry.

Another response is to enlarge the set of alternative hypotheses available to the theorist and ask whether we could obtain a nontrivial conservation theory by introducing hidden particles. The answer is yes: let’s hypothesize that during the transition  $K + K \rightarrow K + K + \mu + \mu$  a neutrino  $\nu$  was present, and the reaction that actually took place was  $K + K \rightarrow K + K + \mu + \mu + \nu$ . Then if we introduce a quantity  $\mathbf{q}$  and assign  $\mathbf{q}(K) = 1, \mathbf{q}(\mu) = 1, \mathbf{q}(\nu) = -2$ , we obtain a conservation principle that is consistent with the observed transitions  $K \rightarrow \mu$  and  $K + K \rightarrow K + K + \mu + \mu$  but not with all other processes involving the  $K$  and  $\mu$  particle. For example, the selection rule is inconsistent with observing the transition  $K + K \rightarrow \mu$ : there is no way to add neutrinos to the products to make the  $\mathbf{q}$ -count balance out.

So far, so good. But perhaps the move to hidden particles trivializes the enterprise, in the sense that we can fit any data whatsoever with the right mix of hidden particles and conservation principles? This is not so: For example, no matter what hidden particles a conservation theory introduces, it cannot rule out the transition  $K + K \rightarrow \mu + \mu$  as a possible observation if it rules in the transition  $K \rightarrow \mu$ .

Using the representation of reactions as vectors mentioned in Section 4, we can characterize the range of empirical phenomena that conservation theories can capture with hidden particles but not without them. For a given set of observed transitions represented as vectors, there is the set of further reactions than can be obtained as linear combinations of the observed ones. Within the linear combinations, we may distinguish those that use only integer coefficients from those that use some fractions. Unobserved transitions that can be generated from the observed ones only by using *fractional* coefficients are exactly those that can be ruled out with hidden particles, but not without them. This result requires more linear algebra than I can go into here; for the details see (Schulte forthcoming).

So hidden particles allow a conservation theorist to express more empirical regularities, to save more phenomena. The other side of the coin is that inquiry with hidden particles is more difficult than without, because the theorist must now take account of more ways that nature might be. Indeed, there is a diagonal argument to show that there is no reliable method guaranteed to arrive at an empirically adequate conservation theory with hidden particles even assuming that there is one (Schulte forthcoming). However, if we plausibly assume that the particle world contains at most finitely many hidden particles, then reliable inquiry is again possible; but there is no a priori bound on the number of times

that a theorist may have to revise her conservation principles. Thus reliable inquiry is possible but more difficult than in the case in which the particle world contains no unobserved entities, when settling the dynamics of  $n$  particles requires at most  $n$  theory revisions.

We have seen that an inquirer who seeks an empirically adequate conservation theory, whether or not she also wants her theory to be simple, and to unify or explain particle dynamics, may obtain information about the observed particle world that forces her to consider the inhabitants of the unobserved particle world and their properties. This is not to deny that an inquirer may take different attitudes towards different parts of her theory. Along the lines of (Van Fraassen 1980), for example, she may believe only what her theory tells her about observable particle transitions, and merely “accept” what her theory tells her about the presence of undetected particles. From the point of view of the current analysis, such a distinction, while coherent, appears rather artificial, since theories about unobserved particles make definite predictions about observable particle transitions, and in some possible particle worlds conservation theories that introduce unobserved particles are the *only* ones that make correct predictions.

## 6 Reliable and Efficient Inquiry Into Particle Reactions

From a reliabilist point of view, the first question to ask about an empirical problem is whether there is a reliable method for solving it. But that is not the only question. As we have seen, considering the feasibility of other cognitive aims, such as minimizing theory changes and convergence time, reveals fine gradations in the complexity of inductive problem. Such considerations also yield more constraints on what inductive methods ought to infer in the short run. Sometimes these constraints are very strong. For example, in the particle reaction problem, methods that infer that a reaction is disallowed until it is observed settle on the right answer about the possibility of each reaction with at most one change of conjecture, no matter what the evidence is. It is possible to show that such methods are the only ones that achieve this performance while also minimizing convergence time (Schulte forthcoming). Thus steady and fast reliable inquiry requires choosing the “closest fit” to reaction data, ruling out as many unseen reactions as is possible with conservation principles. This procedure accords well with practice (Omnes 1971, p. 36).

One of the familiar criticisms of long-run, means-ends epistemology in the Philosophy of Science is that long-run success is consistent with any behaviour in the short run. This was the criticism that Salmon levelled against Reichenbach’s “pragmatic vindication” of induction (Salmon 1991). The example above points the way to a response to this criticism: if we take into account epistemic aims *in addition to* reliable convergence to a correct theory, then there may be very few or even just one way of drawing inferences that meets *all* of the desiderata.

This is the case in the particle reaction problem for the desiderata of fast, steady and reliable inquiry.

In other problems, the same goals apply, but may sanction quite different rules (Schulte 1999a). For example, if we are interested in the existence of various observable entities, reliable, fast and steady inquiry sanctions a variant of Occam’s Razor: “Do not posit the existence of observable entities until you have observed them.” And in a Goodmanian Riddle of Induction, reliable, fast and steady inquiry singles out the natural projection rule as the best one (project that all emeralds are green as long as that hypothesis is consistent with the evidence), regardless of the language in which hypotheses and evidence are framed. It is remarkable that such disparate methodological directives stem from the same cognitive values. The reason for this confluence is that the particle reaction problem, the Goodmanian Riddle, and the Occam problem share a common structure of “temporal entanglement” (Kelly forthcoming, Sec.3) among the various possible ways in which the evidence might come in. This structure is a topological structure in the space of evidence sequences. A learning-theoretic characterization theorem states that reliable, fast and steady inquiry is feasible if and only if the problem under investigation has the topological structure in question. There is no room here for further details (they are in (Schulte 1999b); see also (Kelly 1996, Ch.9.7)), except to note that the particle reaction problem, the Goodmanian Riddle and the Occam problem exhibit precisely the topology of evidence sequences that characterizes reliable inquiry with a bounded number of mind changes. Uncovering common structure among seemingly disparate problems is one of the great strengths of learning theory.

## 7 Conclusion

Like any good theory, learning theory seeks generality and abstraction. This does not mean that the theory abstracts away from important details. On the contrary, when we come to apply the learning-theoretic analysis of inquiry in specific domains, the analysis requires careful attention to the details of the empirical problem, and indeed often guides us towards the methodologically pivotal aspects of a domain of inquiry.

The domain of inquiry under consideration in this paper was elementary particle dynamics. The learning-theoretic analysis vindicates the central role of conservation principles in reducing the inductive complexity of particle dynamics. We saw how the search for empirically adequate conservation principles can require the introduction of hidden particles. I analyzed the implications of epistemic values in addition to long-run convergence to a correct hypothesis, specifically stable belief and fast convergence to a correct hypothesis. It turns out that these aims impose strong constraints on efficient inquiry into particle dynamics in the short run (namely ruling out as many unobserved reactions as is possible with conservation principles). These constraints are plausible and accord with scientific practice.

The tools of learning theory are in place for further applications. Appli-

cations such as the one outlined in this paper require careful attention to the details of a domain of scientific inquiry; if the learning-theoretic approach to the philosophy of science is correct, they will be worth the effort.

## References

- Aris, R. (1969), *Elementary Chemical Reactor Analysis*. Prentice Hall: Englewood Cliffs, N.J.
- Cooper, L. (1970), *Physics: Structure and Meaning*. Hanover: University Press of New England.
- Feynman, R. (1965; 19th ed. 1990) *The Character of Physical Law*. Cambridge, Mass.: MIT Press.
- Ford, K.W. (1963), *The World of Elementary Particles*. New York: Blaisdell.
- Goldman, A. (1986), *Epistemology and Cognition*. Cambridge, Mass.: Harvard University Press.
- Hempel, C.G. (1965), *Aspects of Scientific Explanation*. New York: Free Press.
- Kelly, K. (2000), "The Logic of Success", *The British Journal for the Philosophy of Science*. Forthcoming.
- Kelly, K. (1996), *The Logic of Reliable Inquiry*. Oxford: Oxford University Press.
- Kelly, K., Schulte, O. and Juhl, C. (1997), "Learning Theory and the Philosophy of Science", *Philosophy of Science*. 64: 245–267.
- Kuhn, T. (1970), *The Structure of Scientific Revolutions*, 2nd edition. Chicago: University of Chicago Press.
- Ne'eman, Y. and Kirsh, Y. (1983), *The Particle Hunters*. Cambridge: Cambridge University Press.
- Omnes, R. (1971), *Introduction to Particle Physics*. London, New York: Wiley Interscience.

- Putnam, H. (1963), “Degree of Confirmation’ and Inductive Logic”, in *The Philosophy of Rudolf Carnap*, ed. A. Schilpp. La Salle, Ill: Open Court.
- Putnam, H. (1965), “Trial and Error Predicates and a Solution to a Problem of Mostowski”, *Journal of Symbolic Logic* 30: 49–57.
- Putnam, H. (1975), “Probability and Confirmation”, in *Mathematics, Matter and Method, Philosophical Papers Vol. 1*, London: Cambridge University Press.
- Salmon, W. (1991), “Hans Reichenbach’s Vindication of Induction”, *Erkenntnis*. 35:99–122.
- Schulte, O. “Inferring Conservation Laws in Particle Physics: A Case Study in the Problem of Induction”. Forthcoming in *The British Journal for the Philosophy of Science*.
- Schulte, O. (1999), “Means-Ends Epistemology”, *The British Journal for the Philosophy of Science*, 50:1–31.
- Schulte, O. (1999), “The Logic of Reliable and Efficient Inquiry”, *The Journal of Philosophical Logic* 28:399–438.
- Sklar, L. (1975), “Methodological Conservatism”, *Philosophical Review* LXXXIV: 374–400.
- Valdés-Pérez, R. and Erdmann, M. (1994), “Systematic Induction and Parsimony of Phenomenological Conservation Laws”, *Computer Physics Communications* 83: 171–180.
- Valdés-Pérez, R. and Żytkow, J. (1996), “Systematic Generation of Constituent Models of Particle Families”, *Physical Review E* 54:2102–2110.
- Van Fraassen, B. (1980), *The Scientific Image*. Oxford: Clarendon Press.
- Williams, W.S.C. (1997), *Nuclear and Particle Physics*. New York: Oxford University Press.