Inferring Conservation Laws in Particle Physics: A Case Study in the Problem of Induction

Oliver Schulte
University of Alberta

October 3, 2001

Abstract

This paper develops a means-ends analysis of an inductive problem that arises in particle physics: how to infer from observed reactions conservation principles that govern all reactions among elementary particles. I show that there is a reliable inference procedure that is guaranteed to arrive at an empirically adequate set of conservation principles as more and more evidence is obtained. An interesting feature of reliable procedures for finding conservation principles is that in certain precisely defined circumstances they must introduce hidden particles.

Among the reliable inductive methods there is a unique procedure that minimizes convergence time as well as the number of times that the method revises its conservation principles. Thus the aims of reliable, fast and steady convergence to an empirically adequate theory single out a unique optimal inference for a given set of observed reactions—including prescriptions for when exactly to introduce hidden particles.

1 Epistemology In Practice

One of the goals of the philosophy of science is to connect general epistemology with scientific practice. In their pursuit of knowledge, scientists face epistemic problems at every step. The epistemologist can learn much from the ways in which they cope with these problems. Conversely, general epistemology should have something to say about particular problems concerning belief and knowledge, theory and evidence, that arise in scientific practice. This paper studies a particular problem of inductive inference arising in particle physics from the point of view of a general inductive epistemology. The inference problem is this: Given a set of observations of particle reactions, find a theory that makes the correct predictions about what reactions—observed and unobserved—are possible.

The central epistemological tenet of this paper is that inquiry ought to reliably lead to true beliefs. Though science may err in the short run, it ought to correct itself and eventually settle on a correct theory of the phenomena under investigation. This fallibilist vision of empirical success in the 'limit of inquiry' inspired the work of Peirce, James, Reichenbach, Putnam, and others. It is the philosophical core of a mathematical framework for analyzing inductive inference known as *formal learning theory*, which philosophers, logicians and computer scientists have elaborated in considerable detail over the past four decades or so.[1]

Since ought implies can, the precept that inquiry ought to be guaranteed to eventually arrive at a correct belief can apply only if it is in fact possible to achieve this performance. The first result in this paper is that, without further background assumptions, the problem of finding a correct theory about particle reactions is so complex that no inductive method is guaranteed to eventually arrive a correct theory. This is so even if we assume that every possible reaction can be observed in principle.[2] Thus the argument for the negative result is not that theories of particle reactions are globally underdetermined by the total body of possible evidence. Rather, the problem is *local* underdetermination: The evidence that is available at the stages of inquiry is not strong enough to sift through alternative theories of particle reactions so as to correct errors and systematically eliminate all empirically inadequate theories. (My use of the terms 'global' and 'local' underdetermination follows Kelly's thorough discussion ([1996], Ch.2).)

But that is not the end of the story. The degree to which an empirical question is underdetermined is never absolute but always relative to the assumptions that the inquirer is willing to make. Placing constraints on the theories to be considered makes an inductive problem easier because it reduces the space of alternative hypotheses to be investigated. Theories of particle reactions are constrained in just this way: physicists take recourse to *conservation principles* for formulating laws governing particle reactions. Assuming that there is an empirically adequate conservation theory of particle reactions leads to a positive result[3]: There is a reliable inference procedure that is guaranteed to find a correct conservation theory if there is one. Moreover, for a given set of n elementary particles, there is an a priori bound on the number of times that this procedure changes its conservation principles—namely at most n times. Thus in terms of the hierarchy of inductive complexity established in my article 'Means-Ends Epistemology' (Schulte [1999a]), the problem of finding a correct conservation theory ranks fairly low. Focusing on conservation principles as a special class of theories about particle reactions is a powerful antidote to local underdetermination.

[1] (Kelly's "The Logic of Reliable Inquiry" and Martin and Osherson's "Elements of Scientific Inquiry" are two recent book-length treatments of the learning-theoretic approach to the Philosophy of Science (Kelly [1996]; Martin and Osherson [1998].)

[2] An alternative interpretation of this assumption, along the lines of van Fraasen ([1980]), is that an inquirer may concern herself only with the empirical adequacy of her theories, so that she does not distinguish between impossible reactions and those that are possible but never observed.

[3] For brevity of expression, here and elsewhere I use the term "conservation principles" to mean additive or numeric conservation principles; physicists use other kinds of conservation principles as well. See Section 5.

Means-ends analysis requires careful attention to the details of the inference problem at issue. One of the strengths of the learning-theoretic approach is that close scrutiny often leads to fruitful questions about the logic of the scientific domain under investigation. In the particle reaction domain, learning-theoretic analysis requires us to consider the scope of conservation theories: what phenomena can they capture, and how? We will see that if the empirical phenomena are of a certain precisely specified kind, then conservation theories can make correct predictions only if they posit the presence of *undetected* particles. Thus if an inquirer wants to find an empirically adequate conservation theory, she must be prepared to posit hidden entities.

However, long-run reliability does not require a theorist to introduce a hidden particle at any particular time—she can always put off this move until later. That long-run reliability countenances this sort of arbitrariness in the short run was one of the chief criticisms that Salmon mounted against Reichenbach's 'pragmatic', means-ends vindication of his straight rule for estimating probabilities (Salmon [1991]; Reichenbach [1949]). As I showed in my paper 'Means-Ends Epistemology', the long run can lead to strong constraints in the short run if we take into account *other* epistemic aims in addition to reliable convergence to an empirically adequate theory. 'Means-Ends Epistemology' focused on two such aims, namely fast and steady convergence to a correct theory. It turns out that these criteria require a theorist to posit undetected particles at certain definite stages of inquiry. Specifically, fast and steady convergence requires a theorist to choose the 'closest fit' to the observed data; that is, the theorist should choose a conservation theory that rules out as many unobserved reactions as possible. We will see that the logic of conservation principles is such that for some possible evidence sequences the closest fit to the observed particle reactions requires introducing hidden particles.

Is does not imply ought, and normative epistemology need not take conformity with practice as its yardstick. Nonetheless, it is interesting to compare the recommendations from means-ends analysis with the theories that physicists have arrived at. The general principle of looking for the 'closest fit' to given reaction data squares well with practice; see for example (Omnes [1971], Ch.2). With regard specifically to the currently available evidence, it would be interesting to compare the conservation theory and hidden particles that my inference rule would produce on the current evidence with what physicists have developed. Since the current evidence fills a small phone book, this requires implementing my procedure in a computer program; a project to do so is underway.[4] (Valdés-Pérez has implemented another program for inferring conservation principles (Valdés-Pérez and Erdmann [1994]).) There are strong reasons to expect that the conservation theory produced by the learning-theoretic procedure will be close to what physicists have found, at least as far as its empirical content goes. This is because of the following theoretical result: Given the principles of Special Relativity—particularly conservation of energy and momentum—the number of

[4]Supported by a grant from the National Sciences and Engineering Research Council of Canada.

conservation principles is bounded by the number of stable particles (Section 10). This fact guarantees that the optimal inference procedure described below will select conservation principles that are empirically equivalent to the ones currently accepted by physicists, plus possibly one extra conservation principle. Whether or not the procedure will introduce the same undetected particles as physicists did is a complex question that requires further investigation.

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. ([1992], p. 458)

In reliabilist epistemology, one way to measure the difficulty of a set of empirical questions is by determining whether or not there is a reliable method for arriving at the correct answer to the given questions. We will see in Section 4 that the reliabilist analysis agrees with the physicists' sense that particle reactions raise considerable empirical difficulties: I show that without further background assumptions, there is no reliable inductive procedure for arriving at an empirically adequate theory of particle reactions.

As Cooper explains, physicists have dealt with the high degree of local underdetermination 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' ([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. 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). In his graduate text on particle physics Omnes explains 'once and for all' a method for inferring selection rules ([1971], p. 36).

The learning-theoretic analysis gives a mixed verdict on the difficulty of finding empirically adequate conservation principles. If all particles are detectable, there is an inference method of low complexity that is guaranteed to eventually find an empirically adequate set of conservation principles if there is one to be found (Section 6). But if we allow that undetected particles may have been present in some reactions (Section 7), the inference problem is more complicated, and further assumptions are required. In practice too we find that the combination of hard-to-detect neutrinos with conservation principles poses some complications (see for example the argument for the existence of two distinct neutrino flavours corresponding to different conserved quantities (Williams [1997], Ch. 12.3)).

Thus overall, searching for empirically adequate conservation principles is relatively easy whereas searching for an arbitrary correct theory of particle reactions is very difficult. This means that the proposition that *some* set of conservation principles is empirically adequate is a strong assumption—exactly how strong I determine in Sections 6 and 7 where I characterize the particle worlds that are governed by a set of conservation principles.

It is clear that the hypothesis that some set of conservation principles constitutes the 'laws of nuclear interaction' derives much of its plausibility to scientists from the analogy between the classical conservation principles and the new ones. Thus Feynman groups the new principle of conservation of baryon number with the 'great conservation laws' ([1965], Ch.3). Ne'eman and Kirsch view the history like this: 'From the fact that there are many processes which do not contradict any of the classical conservation laws, and yet have never been observed, physicists came to the conclusion that interactions between particles obey *additional conservation laws* which are not revealed through the behaviour of macroscopic objects' ([1983], p.146), their emphasis. Omnes introduces his method for inferring selection rules with the example of the process $p \rightarrow \pi^+ + \pi$, which is not observed, and comments that 'along lines made venerable by tradition, we explain this by saying that the proton has a certain inherent property which is conserved' ([1971], p.36).

However, as epistemic support for the new conservation principles, virtue by association with the traditional ones goes only so far, since there are no 'more fundamental facts' such as space-time symmetries to justify the new principles.[5]

[5] Another disanalogy between the new and the traditional conservation principles is that it has been difficult to relate the new conserved quantities with other physical phenomena besides reactions. Thus Feynman comments "If charge is the source of a field, and baryon does the same things in other respects it ought to be the source of a field too. Too bad that so far it does not seem to be, it is possible, but we do not know enough to be sure" ([1965],

In practice, the research program of searching for selection rules has justified itself by its success so far. This observation may raise the concern that the perceived empirical success is spurious, because physicists would come up with some sort of conservation principles no matter what evidence they received. As my analysis will make clear, there is no occasion for such worries with regard to the hypothesis that the particle world is ruled by some set of conservation principles: It is indeed possible to revise conservation principles in the light of a number of contrary observations—especially if we permit theorists to posit undetected particles—but after enough false predictions, eventually all conservation theories would be inadequate.

3 Theory and Evidence about Particle Reactions

My analysis employs some basic notions from formal learning theory. Learning-theoretic models are quite simple in outline. Their basic components 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 which total body of evidence.

In the particle setting, I take the evidence items to be reactions among elementary particles. 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. For the bulk of this paper, I focus on the problem of finding an empirically adequate reaction theory; thus for the most part, correctness will be just empirical adequacy. Section 10 discusses further constraints on the adequacy of reaction theories that stem from parsimony and systematicity.

These specifications warrant comment, but before further discussion, let us make them clear with some mathematical definitions. Let P be a set of (types of) discrete elementary objects under consideration. I suppose that P is at most countable, and that the objects in P are numbered in some canonical order as $p_1, p_2, ..., p_n, ...$ A **reaction** r is a pair $r = (objects, s)$ where $objects$ is a nonempty finite sequence of elementary objects (i.e., $range(objects) \subseteq P$), and s is either 1 or 2. The number s tells us which objects in the reaction are reagents and which are products: If s is 1, the first object is the reagent and the others are products; in this case I refer to the reaction as a **decay**. If s is 2, the first two objects are the reagents and the others are products; in this case I refer to the reaction as a **collision**. Although this is my official

p.67).

definition of a reaction, I will mostly use the more familiar and more graphic *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 $p + p \to p + p + \pi$. Viewed as a sequence, this corresponds to the reaction $(\langle p, p, p, p, \pi \rangle, 2)$. For another example, $(\langle \mu^-, e^-, v_\mu, \overline{v_e} \rangle, 1)$ is the sequence corresponding to the decay $\mu^- \to e^- + \nu_\mu + \overline{\nu_e}$.

A **reaction stream** ε is an infinite sequence of reactions. The initial sequence of the first k reactions in a reaction stream is denoted as $\varepsilon|k$. A theory T of particle reactions classifies reactions as either 'possible' or 'impossible', or as physicists even say, as either 'allowed' or 'forbidden'. Mathematically I will represent a **reaction theory** T as a set of reactions, namely those that it allows. A theory T of particle reactions is empirically adequate on a reaction stream ε iff T allows all reactions that are observed on ε and T rules out all reactions that are never observed.

These definitions specify an inference problem, with evidence items, hypotheses and a correctness relation. Clearly the model I have described captures only part of the entire complex enterprise of particle physics. For example, physicists devote much effort to establishing the validity of experimental techniques. (Franklin ([1990]) and Galison ([1997]) study the epistemology of experimental particle physics in detail.) In my model, theorists take experimental reports as data to be accounted for rather than questioned (though the extended model from Section 7 allows theorists to posit the presence of undetected particles). Also, physicists ask many more questions about particle reactions than whether they are possible or not, for example which processes occur relatively more frequently than others, and they seek to distinguish broad classes of interactions according to various criteria of interest (e.g., electromagnetic, strong and weak interactions). Section 10 discusses the relationship between my analysis and some of the aspects of scientific practice further. It would be possible to extend my model to accommodate many of these issues. In this paper I focus on the logic of inferring conservation principles, by itself a complex inductive problem with a rich structure, as we will soon see.

4 Reliable Inquiry Into Theories of 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 δ 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. (In the most general setting, we would allow inference rules not to draw any conclusions at all (cf. Schulte ([1999a],[1999b])), but for our current purposes, this simple definition suffices.) I will also refer to inference rules as 'inductive methods', 'inference methods' or simply 'theorists'. There are of course other kinds of inductive methods, for example ones that revise 'degrees of belief', as Bayesian methods do. In principle, means-ends rationality can guide agents in revising any epistemic state.[6] At any rate, it is most natural to consider theories that give a definite answer, rather than a probability, as to whether a reaction is possible or not since those are the kinds of theories that particle physicists propose.

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

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.*

For the proof I shall employ a learning-theoretic *diagonal* argument. This kind of diagonal argument entered the literature via (Gold [1965]; Putnam [1963], [1965]). Relatively non-technical expositions of learning-theoretic diagonalization may be found in (Putnam [1975], Ch.18; Kelly [1996], Ch.3; Juhl [1995]). A diagonal argument proves that every inference rule has an 'Achilles' heel', an infinite sequence of reactions on which it fails to converge to an empirically adequate theory. A vivid way to show that any theorist has an Achilles' heel is to demonstrate how an inductive demon could present data that will lead the inquirer astray. The demon does not personify any kind of 'malice' on the part of nature, or any other property of the world under investigation, 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. (I will presently examine the effects of granting more assumptions to the inquirer, and hence fewer options to the demon.) For the proof of the proposition, it suffices to consider a specific question about particle reactions: Up to

[6] Putnam showed how means-ends analysis yields a critique of Carnapian "confirmation functions" that produce "degrees of confirmation" in light of new evidence (Putnam [1963]). For learning-theoretic treatments of Bayesian updating see for example (Earman [1992], Ch.9; Martin and Osherson [1998], Ch.5; Kelly *et al.* [1997]; Juhl [1997]).

how many pions π can be produced in a collision of two protons, that is, in a reaction of the form $p + p \to p + p + \pi + \pi + \cdots + \pi$?

Consider any theorist. Suppose that the theorist observes many experiments in which no pions are produced in a collision of two protons. Eventually she must infer that no pions can be produced in such a collision; for else she would fail to infer that the collision of two protons never produces a pion. So let t be the first time at which the theorist conjectures that the collision of two protons never produces a pion. Then the demon presents the reaction $p+p \to p+p+\pi$ in which one pion is produced. The theorist has to change her mind and (eventually) infer at some later time $t' > t$ that exactly one pion is produced; or else the demon presents an infinite data sequence on which only one pion is ever produced by two protons, but the theorist fails to infer this fact. Then the demon presents the reaction $p + p \to p + p + \pi + \pi$, and so on.

What is the result of this interaction between scientist and nature? If the theorist converges to the hypothesis that up to n pions can be produced, then in fact the production of $n + 1$ pions is observed. And if she converges to the hypothesis that any number of pions may be produced, then in fact the production of only finitely many is observed. Hence either she does not converge to a theory of particle reactions at all, or she converges to one that is not empirically adequate. In either case, she does not succeed on the infinite sequence of particle data that results from the interaction between her and the demon as described.

It is not difficult to show that if every inductive method fails on some reaction stream, then every inductive method fails on infinitely many reaction streams. Thus the diagonal argument just given proves that every theorist fails to arrive at the right answer about how many pions can be produced by two protons in infinitely many courses of inquiry (unless further background assumptions restrict the ways in which the evidence might come in).

The question about how many pions can be produced in a collision of two protons is just a small aspect of the terrain that a comprehensive empirically adequate theory of particle reactions has to cover. Even if physical theory could settle this particular question, or if scientists chose to simply neglect it, there are many other questions about particle reactions for which the same kind of diagonal argument shows that they cannot be settled reliably even in the limit of inquiry.

What does this negative result tell us about inquiry into particle reactions? My conclusion is not the skeptical one that the project is hopeless; rather, the argument suggests that the scientist needs some background assumptions to guide her research into what particle reactions are possible. A prima facie plausible strategy is to look for background assumptions that are just strong enough to render the research problem solvable—to defeat the inductive demon—but no stronger than what is necessary. However, Kevin Kelly has established the remarkable fact that in general, there is no such optimal background knowledge: for most typical inference problems—including the particle reaction problem—we find that background assumptions that guarantee reliable convergence to an empirically adequate theory can always be weakened further without rendering the problem unsolvable. (For more details see Kelly's discussion in ([1996], Secs. 3.7, 4.6).)

Instead, a theorist might turn to background assumptions that are strong enough to permit reliable investigation, and that are plausible to her (and her colleagues). The sources cited in Section 2 show that in the 20th century, the tradition of physics has made conservation principles plausible candidates for the 'laws of nuclear interaction'. What assumptions are plausible at a given stage of inquiry often depends on historically contingent factors such as the development of the research field in general. What is not contingent on the historical context, and subject to systematic analysis, is the methodological import of the background assumptions. I now turn to the analysis of the methodological import of conservation principles for reactions among elementary particles.

5 Conservation Principles and Selection Rules

Several kinds of conservation principles apply in particle physics: General conservation principles such as conservation of energy, momentum and electric charge, discrete space-time symmetries such as parity and CPT, and so-called numeric or additive conservation principles (Williams [1997]); Wolfenstein and Trippe [1998]). My analysis considers the third kind of conservation principle, also known as a **selection rule** (Omnes [1971], p.36). A selection rule introduces a quantity and assigns a value for that quantity to each known elementary particle. Table 1 lists 5 such quantities, namely electric charge, baryon number, lepton number, electron number, and muon number; it specifies what value each of the 20 listed particles has for a given quantity. These quantities are the same for each particle in all reactions (unlike, say, momentum). From now on, I will use the term 'quantity' to refer to such process-invariant properties of particles. I will consider only conservation principles that involve process-invariant quantities.

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 $2 \times Baryon\#(p) = 2 \times 1$, and the Baryon total of the products is $2 \times Baryon\#(p) + Baryon\#(\pi) = 2 \times 1 + 0$. A given reaction is possible according to a conservation theory only if it conserves all quantities that define selection rules.[7]

Now let us see how the use of conservation principles resolves local underdetermination. 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

[7] 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.

	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	$\overline{K^0}$	0	0	0	0	0
5	n	0	1	0	0	0
6	$\overline{n}$	0	−1	0	0	0
7	$\overline{p}$	−1	−1	0	0	0
8	p	1	1	0	0	0
9	π^-	−1	0	0	0	0
10	π^+	1	0	0	0	0
11	π^0	0	0	0	0	0
12	γ	0	0	0	0	0
13	μ^+	1	0	−1	0	−1
14	μ^-	−1	0	1	0	1
15	e^-	−1	0	1	1	0
16	e^+	1	0	−1	−1	0
17	ν_e	0	0	1	1	0
18	$\overline{\nu_e}$	0	0	−1	−1	0
19	ν_μ	0	0	1	0	1
20	$\overline{\nu_\mu}$	0	0	−1	0	−1

Table 1: Quantum Number Assignments applied in Selection Rules

whatever conservation principles govern collisions of two protons, the pion π 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 π 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 what current particle theory tells us.

This example shows the power of conservation principles to resolve local underdetermination: Without further assumptions, there is no reliable way for settling the pion problem, no matter how much evidence the scientist gathers, as we saw in Section 4; under the assumption that some kind of conservation theory is empirically adequate, *one* observation entails the right answer.

The question about the number of pions concerns but a small aspect of what particle reactions are possible. Is the focus on conservation principles similarly successful in resolving underdetermination with respect to all other aspects? To answer this question, we must investigate the logic of selection rules to determine what possible particle worlds they can and cannot describe.

6 Reliable Inquiry Into Conservation Principles Without Hidden Particles

At a given stage of inquiry into particles, there is a set of known elementary particles whose behaviour we wish to account for. Let us examine the problem of finding an empirically adequate conservation theory of reactions involving a finite set of n known elementary particles, which we may assume are numbered in some order as $p_1, p_2, ..., p_n$.

As we saw in the pion production example, if one reaction is possible (such as $p+p \rightarrow p+p+\pi$), then selection rules must permit other reactions as well (such as $p+p \rightarrow p+p+\pi+\pi$). Thus focusing on conservation principles allows us to infer that certain unobserved reactions are possible (and hence will be observed or at least would be if the experimenter went to the trouble of creating the necessary conditions). The question is: which unobserved reactions can we thus conclude to be possible? The key to the answer is that we can combine reactions in certain ways to obtain a new reaction such that if each of the combined reactions separately conserves a set of quantities, then their combination also conserves all quantities in that set. Thus to the extent that numeric conservation principles govern the reactions in question, we can infer that any combination of observed reactions is possible.

With some care, we can take the operations for forming new reactions to be the operations of vector sum and scalar multiplication familiar from linear algebra. This requires viewing reactions as vectors. Representing reactions as vectors is standard procedure in chemistry (see for example (Aris [1969])); this scheme allows us to apply the powerful and elegant apparatus of linear algebra to the methodology of particle physics. The basic idea is to associate with each reaction r an n-ary vector $\mathbf{r}$, where n is the number of known elementary particles. The i-th entry $\mathbf{r}(i)$ tells us how often particle p_i occurs in reaction r. To distinguish products from reagents, we write the entries for products as negative numbers. For example, let's consider the 6 elementary particles $p, \pi, \mu^-, e, v_\mu, \overline{v_e}$, numbered in that order. We represent the decay $\mu^- \to e^- + \nu_\mu + \overline{\nu_e}$ by the vector $(0, 0, 1, -1, -1, -1)$.

A difficulty with this scheme is that it isn't clear how to represent a reaction in which the same type of particle occurs both as a reagent and as a product (like the proton in the reaction $p+p \to p+p+\pi$). Valdés-Pérez has solved this problem elegantly by noting that, as far as selection rules go, we can cancel occurrences of a particle on different sides of the arrow without losing information about what reactions are and are not possible (Valdés-Pérez and Erdmann [1994]). More precisely, if in a reaction r a particle p occurs among both the products and the reagents, then cancelling this occurrence leads to a process r' that conserves exactly the same quantities as r does. Thus for example the transition $p+p \to p+p+\pi$ conserves exactly the same (process-invariant) quantities as the reaction $p \to p+\pi$. Following Valdés-Pérez, let the **net occurrence** of a particle p_i in a reaction r be the number of times that p_i occurs among the reagents of r minus the number of times that p_i occurs among the products of r. Then with each reaction r we associate the n-ary vector $\mathbf{r}$ whose i-th entry $\mathbf{r}(i)$ is the net occurrence of particle p_i in r. So both the collision $p + p \to p + p + \pi$ and the decay $p \to p + \pi$ are represented by the 6-dimensional vector $(0, -1, 0, 0, 0, 0)$.

We think of n-ary vectors representing reactions as members of R^n, the vector space whose elements are n-tuples of real numbers with the usual operations of vector addition and scalar multiplication. With this representation scheme, linear algebra serves as a neat calculus for deriving predictions about what reactions are possible given a set of observed reactions. To illustrate, let's see how we can apply linear operations to conclude that if $p + p \to p + p + \pi$ is possible, then so is $p + p \to p + p + \pi + \pi + \cdots$ for any number of pions. As noted above, the first reaction is represented by the vector $\mathbf{r}_1 = (0, -1, 0, 0, 0, 0)$. A collision of two protons that produces n pions is represented by $\mathbf{r}_n = (0, -n, 0, 0, 0, 0)$. Thus we can obtain any vector $\mathbf{r}_n$ as a linear multiple of $\mathbf{r}_1$, namely $n \times \mathbf{r}_1$. In general it is true that if $\mathbf{r}$ is a vector representing a reaction r, and $\mathbf{r}'$ is a linear multiple of r representing a reaction r', then r' conserves any quantity that r conserves, for much the same reason that $x = y$ entails that $cx = cy$. That is, the set of vectors representing reactions that conserve a set of quantities is closed under scalar multiplication. Similarly it is the case that if $\mathbf{r}_1$ is a vector representing a reaction r, and $\mathbf{r}_2$ is a vector representing a reaction r_2, then any reaction represented by the vector sum $\mathbf{r}_1 \oplus \mathbf{r}_2$ conserves all the quantities that are conserved by both r_1 and r_2. This is true for the same reason that $x = y$ and $a = b$ entails that $x + a = y + b$. For example, the reaction $p+\mu^- \to p+\pi+e^-+\nu_\mu+\overline{\nu_e}$ is represented by the vector $(0, -1, 1, -1, -1, -1)$. This vector is the sum of $(0, -1, 0, 0, 0, 0)$ and $(0, 0, 1, -1, -1, -1)$, which represent the reactions $p + p \to p + p + \pi$ and $\mu^- \to e^- + \nu_\mu + \overline{\nu_e}$, respectively. So we may deduce that if $p + p \to p + p + \pi$ and $\mu^- \to e^- + \nu_\mu + \overline{\nu_e}$ each conserve a set of quantities, then so does $p + \mu^- \rightarrow p + \pi + e^- + \nu_\mu + \overline{\nu_e}$.

Hence if a particle world is such that a conservation theory is empirically adequate for it, the set of reactions possible in it must be closed under linear combinations. More precisely, if $\mathbf{r}_1, \mathbf{r}_2, ..., \mathbf{r}_k$ is a set of vectors representing possible reactions, and $a_1, a_2, ..., a_k$ are scalars, then all reactions represented by the vector $\sum_{i=1}^{k} a_i \mathbf{r}_i$ conserve whatever quantities the reactions represented by $\mathbf{r}_1, \mathbf{r}_2, ..., \mathbf{r}_k$ conserve. It turns out that this condition is not only necessary but also sufficient to characterize particle worlds ruled by conservation principles. That is, if we have a set of possible reactions that is closed under linear combinations, then there is always some set of additive conservation principles that rules in exactly the reactions inside the set, and rules out exactly those outside the set.

Before the formal statement of this result, let us take care of one more detail. Conservation principles are only one way of deciding whether a given reaction is possible or not; a physicist may have other knowledge that rules out certain reactions even if they are linear combinations of possible reactions. For example, the collision $p + p \rightarrow p + p + \pi$ occurs, and as we have seen, it follows that the decay $p \rightarrow p + \pi$ must be consistent with conservation principles also. But this decay violates conservation of energy.[8] So even though empirically derived conservation principles would permit the decay $p \rightarrow p + \pi$, another part of the physicist's knowledge rules it out.

The physicist's background knowledge about particle reactions may have a complex structure and a variety of sources, but for our purposes it is sufficient to represent it simply by a set K of reactions that are possible for all that the physicists knows. Then we can refine our result about what conservation particle worlds are like as follows: They are those particle worlds in which any linear combination of possible reactions is possible, except for those ruled out by background knowledge.

Proposition 2 *Let K be a set of reactions representing background knowledge, and let ε be an infinite sequence of observed reactions. Then there is an empirically adequate conservation theory for $\varepsilon \iff$ all reactions are observed that are linear combinations of the observed reactions and consistent with K.*

The proof is in Section 13. Proposition 2 gives a complete characterization of what the particle world must be like if there is an empirically adequate conservation theory. We may think of this result as spelling out the import of the assumption that some set of selection rules is empirically adequate.

How can a theorist go about finding a set of correct conservation principles assuming that one exists? Given some initial set of observed reactions, it follows from Proposition 2 that the theorist may choose a set Q of conservation principles that will rule in exactly the linear closure of these reactions. If all future observations are within the linear closure of these reactions, then this

[8] This follows easily from conservation of energy and momentum applied in what would be the rest frame of the decaying proton. See also (Neeman and Kirsch [1983], p.139).

conservation theory is empirically adequate: It never rules out an observed reaction, and eventually any reaction within the linear closure of the initial ones will be observed (except for those ruled out by other background knowledge K), since we are assuming that some conservation theory is empirically adequate. So a theorist may conjecture Q until a reaction is observed that is ruled out by Q. Such a reaction would be outside the linear closure of reactions observed so far—that is, the reaction would be linearly independent of the observations. Now in an n-dimensional vector space such as R^n it is a basic theorem of linear algebra that there can be at most n linearly independent reactions. Hence if at each stage, the theorist conjectures that exactly the linear closure of all observed reactions is possible, there are at most n times that the theorist may have to change his conservation theory before settling on an empirically adequate one. If indeed n linear independent reactions were found to occur among n particles, basic linear algebra entails that the only conservation principle consistent with all observations would be the trivial one that assigns 0 to every particle. Presumably at this stage scientists would abandon the search for conservation principles and look for other laws of particle interactions instead. In this sense physicists cannot go on forever inventing new conservation principles no matter how the evidence turns out.

There is another interesting wrinkle. If we allow a theorist to posit *undetected* particles, the scope of the phenomena that conservation principles can accommodate is considerably enlarged. Just how much more expressive power this possibility adds to conservation theories is the topic of the next section.

7 Reliable Inquiry Into Conservation Principles With Hidden Particles

Consider two particles, say K and μ. Suppose that the transitions $K \to \mu$ and $K + K \to K + K + \mu + \mu$ were observed. Taking the K particle as the first particle, these transitions are represented by the vectors $(1, -1)$ and $(0, -2)$ respectively. Since these two vectors are linearly independent, they span the entire space R^2, and it follows from Proposition 2 that any set of selection rules that rules in the two observed transitions must be consistent with any transition whatsoever involving the K and μ particle. A direct way to see is this to note that if the process $K + K \to K + K + \mu + \mu$ conserves a quantity, then the μ particle must carry 0 of that quantity. But then the K particle must carry 0 of the quantity too since $K \to \mu$ conserves the quantity.

Now let's hypothesize that during the transition $K + K \to K + K + \mu + \mu$ a neutrino ν was present, and the reaction that actually took place was $K + K \to K + K + \mu + \mu + \nu$. 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 selection rule that is consistent with the observed transitions $K \to \mu$ and $K + K \to K + K + \mu + \mu$ but not with all other processes involving the K and μ particle. For example, the selection rule is inconsistent with observing the transition $K + K \to \mu$: there is no way to add neutrinos to the products to make the **q**-count balance out.

Thus hidden particles allow a theorist to introduce selection rules that save the phenomena and yet rule out certain reactions that are linear combinations of observed phenomena. Hidden particles do not in general allow the theorist to rule out any transition whatsoever that is in the linear closure of the observed reactions. In the example above, introducing the neutrino is sufficient to rule out the unobserved transition $K + K \rightarrow \mu$, but not the unobserved transition $K + K \rightarrow \mu + \mu$. How much expressive power exactly do undetected particles add to selection rules? The answer lies with distinguishing between linear combinations that involve *fractional* coefficients and those that involve *integer* coefficients. Conservation principles can rule out the former with undetected particles, but not the latter. In the examples above, $K + K \rightarrow \mu + \mu$ is represented by the vector $(2, -2) = 2(1, -1)$ which is an *integral* multiple of the vector representing the observed transition $K \rightarrow \mu$. By contrast, $K + K \rightarrow \mu$ is represented by $(2, -1) = 2(1, -1) \oplus -\frac{1}{2}(0, -2)$; writing $(2, -1)$ as a linear combination of the observed transition vectors requires proper *fractions*. Say that the **integral span** of a set $E = e_1, e_2, ..., e_k$ of observed transitions contains exactly those vectors $\mathbf{v}$ that can be obtained as a linear sum $\sum_{i=1}^{k} n_i \mathbf{e}_i = \mathbf{v}$, for suitable *integers* n_i. Among the linear combinations of $\mathbf{e}_1, \mathbf{e}_2, ..., \mathbf{e}_k$ there will be some vectors that represent reactions (basically, those with integral components since elementary particles do not come in fractions), and among those there will be some that are not in the integral span of $\mathbf{e}_1, \mathbf{e}_2, ..., \mathbf{e}_k$. According to the following proposition, this is the group of transitions that can be ruled out with undetected particles but not without them.

Proposition 3 *Let E be a set of observed transitions $e_1, e_2, ..., e_k$. Then*

1. *every conservation theory, whether it involves undetected particles or not, is consistent with linear combinations of $\mathbf{e}_1, \mathbf{e}_2, ..., \mathbf{e}_k$ that represent transitions and have integral coefficients only.*

2. *there is a conservation theory, possibly involving undetected particles, that is consistent with exactly those linear combinations of $\mathbf{e}_1, \mathbf{e}_2, ..., \mathbf{e}_k$ that represent transitions and have integral coefficients only.*

The proof is in Section 13. As in Proposition 2, we can accommodate background knowledge K by considering only integral linear combinations of the observed transitions that are consistent with K.

At first glance, Proposition 3 might seem surprising: Did Hempel not show that any theory positing hidden entities can be reformulated without hidden entities but with the same observational consequences ([1965], Secs. 8–9)? Why does the 'theoretician's dilemma' not arise with particle reactions? For the answer we must take a close look at Hempel's argument. Hempel relied on a theorem of Craig's which roughly says that, under certain general conditions, for a theory T formulated in first-order logic there is *some* other first-order theory T' with no hidden objects and the same observational consequences as T. Craig's theorem leaves open the possibility that the observationally equivalent theory

T' may fail to meet further criteria satisfied by the original theory T. Thus if a scientist focuses on a certain subclass of first-order theories, that class may not contain a theory T' without hidden entities that is observationally equivalent to the theory T featuring theoretical entities. This is the case when the subatomic phenomena are such that they permit no empirically adequate conservation theory without hidden particles: there is indeed some other theory of particle reactions that is empirically adequate and introduces no hidden particles—but this theory demonstrably cannot be expressed with selection rules.

Proposition 3 delineates precisely how much expressive power hidden particles add to selection rules. This extra expressive power enlarges the range of alternative empirical phenomena that a conservation theorist can adequately describe. At the same time, the theorist has to sift through more alternative theories so as to find an empirically adequate one—the set of alternative particle worlds under consideration has grown, and with it the inductive complexity of the inquirer's task. Indeed, if we allow that there may be infinitely many hidden particles, then there is no inductive method that is guaranteed to settle on an empirically adequate theory.[9] Reliable inquiry becomes possible if we assume that the particle world is such that a finite number of observable and of hidden particles generate all observed transitions. In that case a theorist who predicts that exactly the transitions in the integral span of the observed processes are possible is guaranteed to eventually arrive at an empirically adequate theory of particle reactions. (This is not obvious but I will not take the time to prove it here.) There is, however, no bound on the number of times that this method might change its mind, unless we assume an a priori bound on the number of hidden particles. Otherwise we cannot put a limit on the number of times that a theorist might revise her theory, nor a limit on the number of hidden particles that she will introduce—but we do know that *eventually* she will stop and settle on a theory that makes the correct predictions about what observable transitions are and are not possible.

Our investigation so far shows the impact of various assumptions about the particle world on the difficulty of empirical inquiry. Without assumptions, every inductive method fails on infinitely many possible ways in which the evidence about particle reactions might arrive. If we assume that there is an empirically adequate conservation theory that posits only a finite number of hidden particles, then there is an inductive method that reliably settles on an empirically adequate conservation theory. If we assume further that there is an empirically adequate conservation theory without hidden particles, then there is a reliable inference procedure for finding an empirically adequate conservation theory that changes its mind at most n times, where n is the number of known particles. Thus the strength of different methodological assumptions is measured by the

[9] Outline of diagonal argument: Start with a transition of the form $(0, -2)$, such as $p+p \rightarrow p+p+\pi+\pi$. Show nothing else until the theorist produces a conservation theory whose empirical content is exactly the integral span of $\{(0, -2)\}$. Choose a prime number b and present $(0, -b)$. Since b is prime, $(0, -b)$ is a linear but not integral multiple of $(0, -2)$. Show nothing else until the theorist produces a conservation theory whose empirical content is exactly the integral span of $\{(0, -2), (0, -b)\}$. Choose another prime number b' and continue.

extent to which they make different standards of empirical success feasible. The different gradations of inductive complexity that we see in the particle reaction problem apply generally to inductive problems: They are part of a complexity hierarchy that ranks inductive problems according to what epistemic aims are feasible in them (Schulte [1999a]). The complexity gradations that define this complexity hierarchy correspond exactly to a topological measure of the complexity of the hypotheses under investigation (Schulte [1999b]; Kelly [1996], Ch.4).

So far we have discussed the intrinsic difficulty of empirical questions, relative to applicable background assumptions. Whenever reliable inquiry is feasible, there are infinitely many inductive methods that are guaranteed to settle on an empirically adequate theory. Among these methods we can discriminate further by applying other standards of empirical success in addition to logical reliability. I now turn to the investigation of reliable *optimal* methods for making inferences about particle reactions.

8 Reliable and Efficient Inquiry Into Conservation Principles

Epistemologists have proposed a number of aims of inquiry, such as convergence to the truth, finding simple theories, providing informative content, avoiding errors, minimizing belief changes, and others. Such aims sanction hypothetical imperatives: employ the epistemic means that promote your epistemic aims. We can systematically investigate which inference methods for particle theories attain these aims. In this paper, I focus on three aims that turn out to be particularly interesting and generally applicable (Schulte [1999a]): reliable, fast and steady convergence to an empirically correct hypothesis. Let's begin with fast convergence.

We may measure the time-to-truth of a reliable inductive method by the number of observations the method makes until it settles on a correct hypothesis. That is, the convergence time of an inductive method δ on a data stream ε is the earliest time t by which δ converges to an empirically adequate theory on ε. A minimal criterion of adequacy for the speed of an inductive method δ is that there should be no other reliable method δ' such that (1) δ' converges at least as early as δ on every data stream ε (consistent with background knowledge), and (2) there is some data stream consistent with background knowledge on which δ' converges earlier than δ does. Following Kelly, I refer to methods that meet this criterion as **data-minimal** (Kelly [1996], Ch.4.8).[10] What inference rules are data-minimal in our conservation theory problem? The following general

[10] Data-minimality is the result of applying the decision-theoretic principle of weak dominance to convergence time (Schulte [1999a]). E. Gold first introduced the criterion into learning theory ([1967], Sec.10). The properties of data-minimal methods are investigated in (Kelly [1996], Ch.4.8; Juhl [1994]; Schulte [1999a]; Martin and Osherson [1998], Secs. 2.3, 3.4). For stricter requirements on the speed of methods see (Martin and Osherson [1998]; Schulte [1999a], Sec.9).

characterization of data-minimal methods makes it easy to answer this question.

Consider a theorist who proposes a theory of particle reactions after a sequence of reactions e has been observed. The theorist would be in a strange epistemic state if she were sure to eventually abandon her theory *no matter what* evidence is obtained in the future course of inquiry. If she is not in this state, then there must be *some* way for the future evidence to come in so that she maintains her theory. In this case I say that the theorist *projects* her theory. Formally, an inductive method δ **projects its theory along reaction stream** ε after observing evidence e, given background knowledge K, iff (1) the reaction stream ε is consistent with K and the observations e, and (2) δ converges to its current theory $\delta(e)$ along ε by the stage at which the observations e have been made. It can be shown that the reliable data-minimal methods are exactly those that project their theories at each stage of inquiry.

Proposition 4 *A reliable inductive method δ is data-minimal given background knowledge $K \iff$ at each stage of inquiry e, the method δ projects its theory $\delta(e)$ along some reaction stream ε given K.*

The proof is in (Schulte [1999a]); see also (Martin and Osherson [1998]). Thus the aim of speedy convergence—data-minimality—rules out, for example, methods that 'wait and see' until further evidence is obtained, since those methods do not propose a theory, much less project one, while they are waiting.

Now let us consider steady convergence. Since Plato's *Meno* contrasted the stability of genuine knowledge with the fickleness of mere opinion, philosophers have been familiar with the idea that stable belief is better than an oscillating mindset. Sklar says that 'stability of belief is itself a desirable state of affairs and an end to be sought' ([1975]). Kuhn points to the cost of retraining scientists in a new paradigm as a factor that opposes revolutionary changes ([1970]). Learning theorists have considered methods that minimize the number of belief changes before settling on the right answer (Putnam [1965]; Case and Smith [1983]; Schulte [1999a], [1999b]; Kelly [1996]). Let us say that an inductive method δ changes its mind on a reaction stream ε at time $t+1$ just in case the conjecture of δ after observing the evidence $\varepsilon|t+1$ is different from its conjecture at time t. An inductive method δ changes its mind at most n times given background knowledge K just in case δ changes its mind at most n times on every reaction stream ε consistent with K. In that case the background knowledge entails a priori a bound on the number of times that the method will change its mind.[11]

We encountered an example of a method with a bounded number of mind changes before: The method from Section 6 that reliably finds an empirically adequate set of conservation principles (if no hidden particles are involved) changes its mind at most n times, where n is the number of known particles. For an another illustration, let's consider the relatively simple problem of determining

[11] In decision-theoretic terms, the success standard of convergence to the truth with bounded mind changes is the result of applying the minimax choice criterion to the aim of minimizing mind changes. It turns out that, in contrast with data-minimality, applying weak dominance to mind changes leads to an unfeasibly stringent standard of empirical success (Schulte [1999a]).

whether a given process is possible or not. To be definite, let us take $\mu^+ \to e^+ + \gamma$ as the process under investigation. There is a natural method—call it δ—for reliably arriving at the right answer as to whether $\mu^+ \to e^+ + \gamma$ is possible or not with at most one mind change. The natural method conjectures that the reaction is impossible until it has been produced experimentally. If the reaction is indeed impossible, it will never be produced, and the method correctly settles on the right belief without ever changing its mind. If the reaction is possible, eventually it will be observed; at that point, δ changes its mind once to conclude that the reaction is possible. Furthermore, it is easy to see that this method always projects its current conjecture, and hence by Proposition 4 minimizes time-to-truth. Hence in this problem, the maxim 'assume that the past is like the future' yields a reliable data-minimal method that succeeds with at most one mind change.[12]

What is more surprising is that the method δ is the *only* method that achieves this performance. To see why this is so, let any other inductive method δ' be given. We may assume that δ' conjectures that the process $\mu^+ \to e^+ + \gamma$ is possible as soon as it is observed (indeed, data-minimality requires this). What happens if the process fails to be observed in a repeated number of experiments? A reliable method may wait for some number of failures without drawing a conclusion, but to arrive at the right answer, it must eventually issue a conjecture. What if δ' first conjectures that the process is possible after some number of failures to produce it? Then reliability requires that δ' eventually change its mind and conjecture that the process $\mu^+ \to e^+ + \gamma$ is impossible. Suppose that this first mind change occurs after t failures to observe the process. Then if the $t+1$-st experiment realizes the transition, δ' has to change its mind *for the second time*. Hence in the worst case, δ' changes its conjecture twice, whereas we saw that the natural method δ changes conjectures at most once. This shows that if the process $\mu^+ \to e^+ + \gamma$ is never observed, the first conjecture of every method that minimizes the number of belief changes must be that the process is impossible. And if the method minimizes convergence time, it follows from Proposition 4 that δ' does not wait to issue a conjecture, but agrees with the natural method δ in conjecturing that the process is impossible until it has been observed.

Thus in the particle reaction problem, the aim of reliable convergence to an empirically adequate theory, while minimizing convergence time and mind changes, singles out a unique method that is optimal for that aim: at each stage of inquiry, choose the 'closest fit' to the observed reactions, that is, choose a theory that rules out all unobserved reactions. Given our guiding background assumption that there is an empirically adequate set of selection rules, the closest fit to the data is to hypothesize that the only possible reactions are the observed ones and those that the logic of conservation principles requires to be possible given the observations. As we saw in Section 7, achieving this closest fit may require introducing hidden particles.

[12] This argument relies on the assumptions from Section 4, namely that possible reactions are eventually included in the theorist's evidence and impossible reactions never are. Section 9 discusses evidence in particle physics further.

Salmon criticized Reichenbach's 'pragmatic vindication' of induction on the grounds that long-run convergence to a correct hypothesis is consistent with any crazy behaviour in the short-run. His criticism did not take into account, however, the possibility that *additional* epistemic aims may lead to short-run constraints on inductive inferences (cf. (Juhl [1994])). In the particle reaction problem, adding the goals of minimizing convergence time and theory changes to long-run reliability leads to short-run constraints that determine an essentially unique inference at each stage of inquiry, including a prescription for when exactly to posit undetected particles.

The particle reaction problem is not the only example in which the objectives of reliable, fast and steady convergence give strong guidance for what to conjecture in the short run. The same criteria (1) single out the natural projection rule in a Goodmanian Riddle of Induction—namely, project that all emeralds are green as long as that hypothesis is consistent with the evidence—and (2) vindicate a variant of Occam's Razor—namely 'do not posit the existence of entities that are unobservable in principle but that haven't been observed yet' (Schulte [1999a]). It is no accident that means-ends considerations yield similar results in these three problems: even though they may appear quite different from each other on the surface, a deeper analysis shows that they share a similar topology of hypotheses and observation sequences. My paper 'The Logic of Reliable and Efficient Inquiry' provides an explicit description of this topological structure, with a proof that it is exactly the structure that characterizes reliable inquiry with bounded mind changes (Schulte [1999b]).

The cited vindication of Occam's Razor may seem to contradict the finding of Section 7 that avoiding mind changes can require positing the existence of undetected entities. In fact, this result points to an interesting tension between different aspects of particle theories. Since particle theories make both ontological claims about what particles exist as well as predictions about what reactions are possible, a theorist may strive to avoid changes to either component of his particle theory. The aims of minimizing mind changes with respect to these components can come into conflict when the closest fit to the observed reactions requires undetected entities. Occam's Razor can also conflict with a theorist's desire to avoid abandoning previously accepted conservation principles. The next section explores these tensions in some detail.

9 Conservation Principles and Occam's Razor

Unobserved entities have aroused much debate in the philosophy of science. What is their function in scientific theories, and how can observations give reasons for inferring the existence of the unobserved? In the context of the models of inquiry into conservation principles that I developed in this paper, it is possible to give precise answers to these questions from the point of view of a means-ends epistemology.

Occam's Razor is often taken to be a fundamental principle of epistemic rationality. A commonly endorsed form of Occam's Razor is: 'do not needlessly multiply entities'. Does the method for positing hidden particles from Section 7 violate Occam's Razor?

To begin with, Occam's Razor is vague to the extent that there are different notions of what constitutes a 'need' to postulate the existence of an entity. In particular, a means-ends epistemologist would hold that we should 'multiply entities' if doing so serves the aims of inquiry, and that we should refrain from doing so when it does not. The kind of situations that can arise in particle inquiry show a wealth of considerations that are relevant to whether or not positing undetected particles helps or hinders cognitive objectives. We saw in Section 7 that the goal of finding an empirically adequate theory of particle reactions can require hidden particles. Even when they are not necessary for producing an empirically adequate theory, the aim of maintaining conservation principles—a form of epistemic conservatism—can motivate introducing hidden particles. Once scientists have accepted a principle such as conservation of lepton number as a 'law of nature', we would not expect them to abandon it easily. The procedure from the proof of Proposition 3 shows how hidden particles can make existing conservation principles consistent with the observations.

However, when hidden particles are not necessary for producing the closest fit to the particle reaction data, an alternative is to abandon previous conservation principles. This will lead to a conservation theory that fits the phenomena as closely as possible, and has fewer selection rules than the theory that results from introducing hidden particles.[13] Thus in this case, parsimony opposes epistemic conservatism.

Moreover, there are issues not only about what is epistemically desirable, but also about what is cognitively feasible. The procedure that saves falsified conservation principles with hidden particles proceeds *incrementally*, in the sense that the new conservation theory Q_{t+1} is a function of the previous conservation theory Q_t and the most recent datum r_t only, rather than a function of the total available evidence. Several epistemologists have recommended that inquiry proceed incrementally. For example, Harman argues against 'foundation theories' on the grounds that remembering the entire past data is infeasible for cognitively bounded agents ([1986]). In view of the hundreds of reactions that constitute the available data for even a small number of elementary particles, it seems plausible that the limitations of human memories would lead scientists to proceed incrementally. Thus even if revising the current theory in light of the total accumulated evidence leads to a more parsimonious theory and avoids hidden particles, it may be too difficult to find this theory for inquirers with the kind of memory limitations that we humans have.

If we look at the actual historical development of neutrino theories, one more complication becomes apparent: With regard to neutrinos, the distinction between hidden and observable particles is not necessarily a permanent one.

[13] The procedure that accomplishes this is a combination of those used in the proofs of Propositions 2 and 3. First, find a maximal linearly independent subset E' of the observed transitions E such that the integral span of E' is equal to the vectors representing reactions in its span. Then apply the procedure from the proof of Proposition 3 to rule out vectors representing reactions that aren't in the integral span of E, if there are such vectors.

When Pauli proposed the existence of the neutrino in 1930, he did so to reconcile the evidence with conservation of energy and angular momentum (Williams [1997], Ch.5.3). Even if he had no particular reason to expect that the neutrino could be detected with the experimental technology at the time, it was certainly possible that in the future it would be. Indeed, since then the neutrino has been 'more or less directly observed', as Cooper puts it ([1992], p.408). In practice, various ways of manipulating neutrinos also count as strong evidence for their existence, even if they fall short of 'direct observation'. An example is the way in which neutrino-quark scattering constitutes evidence for two distinct types of neutrinos, the muon neutrino ν_μ and the electron neutrino ν_e. Essentially, we posit that two sources S_1 and S_2 each produce neutrinos but that the ones produced by S_1 are different ('muon-flavoured') than those produced by S_2 ('electron-flavoured'). Experiments confirm that the neutrinos from the two sources respond differently to the same manipulation, which indicates that they are indeed distinct types of entities (Williams [1997], Ch.12.3). This sort of arguments fits well with Hacking's 'experimental realism' ([1983]).

Thus the development of observational techniques in particle physics shows that a straightforward application of Occam's Razor to neutrinos is too myopic: we ought to consider what the theorist can accomplish with an ongoing research program, not just what she should infer on the basis of the evidence here and now. A theorist who wants to introduce a hidden entity for the sake of a cognitive objective—such as saving conservation principles, one of Pauli's motivations—can take an epistemic gamble that the entity will eventually be observed ('more or less directly'). Sometimes the entity will be observed, as has often been the case in particle physics, and the gamble pays off doubly: not only was the original epistemic aim achieved (e.g., the conservation principle was saved), but science has uncovered a previously hidden part of nature. Sometimes the entity will fail to be observed; then the theorist faces a difficult choice: (1) to argue for continuing the experimental search for the entity, (2) to abandon the search and try to rework her theory without the hidden entity, or (3) to posit that the entity exists but is impossible to observe. The history of science provides many examples of all three courses of action. As for (1), particle physics has seen several successful extended searches for particles whose existence was predicted. In cosmology, the discovery of Neptune is another example of a successful extended search. In Putnam's account ([1975], Ch.16.7):

> When the predictions about the orbit of Uranus that were made on the basis of the theory of universal gravitation and the assumption that the known planets were all there were turned out to be wrong, Leverrier in France and Adams in England simultaneously predicted that there must be another planet. In fact, this planet was discovered - it was Neptune.

Just as the background assumption that a conservation theory is empirically adequate can entail the presence of an undetected particle, so the background theory of universal gravitation can entail the presence of an undetected planet.

Einstein's disposal of the ether illustrates option (2). Putnam claims that the situation (3) is prevalent.

> It may be argued that it was crucial that the new planet [Neptune] should itself be observable. But this is not so. Certain stars, for example, exhibit irregular behavior. This has been explained by postulating companions. When those companions are not visible through a telescope, this is handled by suggesting that the stars have *dark companions* - companions which cannot be seen through a telescope. The fact is that many of the assumptions made in the sciences cannot be directly tested - there are many 'dark companions' in scientific theory. ([1975], p.256, Putnam's emphasis)

In sum, in the perspective of means-ends epistemology the decision to posit undetected particles involves a complex choice among different epistemic values, constrained by what is cognitively feasible. Parsimony—seeking a smaller set of laws—and the difficulty of empirically settling the presence or absence of hidden particles can speak against undetected particles. The aim of finding an empirically adequate theory of particle reactions, memory limitations, and epistemic conservatism, support hidden particles, especially given the prospect of verifying their presence through later experimental progress. How an inquirer, or an epistemic community, resolves these conflicts will depend on the historical context and on their methodological attitudes. It would be an interesting study to see if and in what form such conflicts among epistemic values have arisen in the history of particle physics, and how physicists resolved them. The volume of relevant data makes it difficult to decide what alternative theories were consistent with the data at various points in the history of the field, but with the aid of a computer program such a study should be feasible.

As I noted above, the line between what particles are unobservable in principle rather than just unobserved so far may not be clear-cut. There is considerable flexibility in my formal models to accommodate such gradations. We could extend the analysis to models in which certain particles are observed in some processes, but only hypothesized to be present in others. Or we might consider situations in which the evidence entails that some undected entities were present, but not how many. Within the analysis that I developed in this paper, we need certainly not insist that neutrinos are hidden particles, but could simply accept what physicists tell us about when they are and are not present. This would lead to a model of particle inquiry with no hidden particles, considered in Section 6. The issue there is not whether to introduce hidden particles, but solely what conservation principles to adopt. How does the learning-theoretic procedure derived in this model compare to the practice of particle physics?

10 Conservation Principles in Practice

The procedures from Sections 6 and 7 produce the closest fit to a given list of possible reactions, which as we saw in Section 8 serves the aims of reliable, fast and steady convergence to an empirically adequate set of selection rules. Practitioners endorse the general principle of choosing the closest fit (Omnes [1971], Ch.2); sometimes the importance of ruling out unobserved processes is expressed by saying that 'what does not occur is as important as what does occur' (see the quotation from Cooper in Section 2). Thus there is good reason to suppose that if we were to apply the procedure from Proposition 2 on the current data, with neutrinos included where physicists have placed them, then the procedure would produce selection rules that rule in the same processes as the ones currently accepted by physicists. It is not difficult to see that there are a number of empirically equivalent selection rules (choosing a set of quantities to be conserved essentially amounts to choosing a basis for a linear space), so an inference procedure would not necessarily produce exactly the set of conserved quantities that physicists have introduced. Which of the empirically equivalent sets of selection rules are adopted in practice is a complicated matter that has to do with, among other things, considerations of particle systematics such as those provided by the Standard Model based on quarks. The reasoning that guides physicists in finding systematic, unified classifications of particles is an interesting topic that is beyond the scope of this paper. Valdés-Pérez and Żytkow have provided a sophisticated analysis of particle systematics that has led to a computer program which reproduces several of the physicist's predictions and classification schemes (Valdés-Pérez and Żytkow [1996]). Their program posits one of the undetected entities (the Ω^-) that physicists inferred based on the quark model and subsequently discovered.

There is another reason why we can expect the procedure from Section 6 to yield results that are empirically equivalent to the current set of selection rules. It turns out that selection rules are highly constrained: Given conservation of energy and momentum, the number of irredundant selection rules that are valid for a set of reactions is bounded by the number of particles that are stable in the set of reactions. This is a strong constraint because the number of stable particles is relatively small.

Theorem 5 *Let Q be a conservation theory with process-invariant quantities. Let D be a set of decays of distinct particles such that all decays in D conserve energy, momentum and all the quantities posited by Q. Then the number of independent quantities posited by Q is no greater than the number of particles without decays in D.*

The proof is beyond the scope of this paper. To illustrate the theorem, let us consider the particle situation as we currently know it. Of the over 100 known elementary particles, 11 are stable as far as the evidence indicates[14] $\gamma, e^-, e^+, p, \overline{p}, v_e, \overline{v_e}, v_\mu, \overline{v_\mu}, v_\tau, \overline{v_\tau}$. Thus if we choose one decay for each particle p_i that has a known decay, we obtain a set D of decays of distinct particles that conserve energy and momentum. Theorem 5 entails that if Q is a set of quantities that are conserved in all decays in the set D, the number of independent quantities in Q is no greater than the number of stable particles. We

[14] As of 1998. I neglect quarks.

find that physicists currently take 5 independent process-invariant quantities to be conserved in all reactions (the baryon, lepton, muon, and tau numbers, as well as electric charge) (Williams [1997], Ch. 13.1; Wolfenstein and Trippe [1998]).[15] This leaves a gap of 11-5 = 6 possible selection rules between the number of accepted universally conserved process-invariant quantities and the number allowed by Theorem 5. However, it is possible to strengthen Theorem 5 to show that we do not need to count antiparticles in the bound on conservation principles. More precisely, if we assume that whenever a particle carries x units of a conserved quantity $\mathbf{q}$, then its antiparticle carries $-x$ units of $\mathbf{q}$, as it is the case with the 6 selection rules listed above (cf. Table 1), then we need count each member of a particle-antiparticle pair only once. Thus the list of 11 stable particles reduces to 6, and it follows that conservation of energy and momentum allow 6 quantities to be conserved. Thus the procedure from Proposition 2 would add at most one more selection rule to the ones that physicists have adopted. We need to further examine the data to determine whether an additional selection rule is required to produce the closest fit to the reaction data, or whether the currently accepted rules already achieve the closest fit.[16]

11 Conclusion

This paper has examined a complex inductive problem from the point of view of a means-ends epistemology: finding an empirically adequate theory for describing reactions among elementary particles. Without any background assumptions, there is no reliable investigative method for solving this problem, even in the limit of inquiry. I showed how various background assumptions reduce the complexity of this problem: the stronger the background knowledge, the higher the standard of empirical success that we can expect inquiry to achieve. The higher the standard of empirical success that is feasible, the less complex a given inductive problem is. In the particle reaction problem, the levels of inductive complexity that correspond to various common background assumptions are instances of a general hierarchy of inductive complexity. For example, we find the problem of settling on the right answer as to whether a given reaction is possible or not is at the same level of complexity as a Goodmanian Riddle of Induction (Schulte [1999a]).

One of the central assumptions that allow reliable inquiry into particle reactions is that an empirically adequate theory of particle reactions can be found

[15] Quark number (vector: $Q\#$) is also conserved but is not independent of baryon number (vector: $B\#$): A reaction conserves quark number iff it conserves baryon number. That is because the principle that hadrons are colour singlets implies that $Q\# = 3B\#$ (Williams [1997], Ch. 10.3, 10.7).

[16] Theorem 5 gives an upper bound on the number of possible selection rules. Valdés-Pérez and Erdman provide a theorem concerning the number of required selection rules ([1994]): They show that if we seek to rule out only a *finite* set of unobserved reactions, a *single* selection rule will suffice. The optimal procedures considered in this paper, however, seek conservation principles that rule out infinitely many unobserved reactions, so that their theorem does not apply.

within the traditional framework of conservation principles. I showed that this approach makes it feasible to reliably find an empirically adequate theory of particle reactions, whereas this performance is impossible to attain without such assumptions. The fact that conservation principles sharply reduce inductive complexity provides a means-ends interpretation of their central methodological role in particle physics.

To examine the methodological consequences of various background assumptions, we must carefully study the conceptual structure of the scientific domain in question. For this reason, means-ends analysis and the logic of a scientific domain often illuminate each other. One example in this paper is the result that conservation of energy and momentum entail that the number of valid, logically independent selection rules (additive conservation principles) cannot be greater than the number of stable, nondecaying particles. Another example is the fact that for some particle phenomena, selection rules can be empirically adequate only if they posit the presence of undetected particles. Even though in such situations there are other empirically adequate theories that do not introduce hidden entities, none of these theories can be cast in the form of conservation principles. This is an instance of a general rationale for positing undetected entities: Given contextual constraints on the theoretical alternatives—constraints which in the case of conservation principles play an indispensable role for reliable inquiry in particle physics—observations can logically entail the presence of unobserved entities.

A long-standing objection against the aim of converging to a correct theory in the long run is that long-run convergence does not entail any constraints on conjectures in the short run. In my paper 'Means-Ends Epistemology' I showed that this is no longer true if we consider additional epistemic aims besides long-run convergence to a correct hypothesis. It turns out that fast and steady convergence to a correct theory are especially interesting theoretical objectives, in that they are feasible in a wide range of inductive problems and provide strong guidance for what to conjecture in the short run. In 'Means-Ends Epistemology' I showed that fast and steady convergence to a correct hypothesis vindicates natural inferences in several inductive problems, including a Goodmanian Riddle of Induction. The analysis of particle reaction theories provides another illustration of the strong constraints that long-run reliability entails when combined with the criteria of speedy and steady convergence: These considerations require a theorist to choose the selection rules with the 'closest fit' to the data, which predict that reactions not observed in the past will not occur in the future, except for those reactions that the logic of conservation principles requires to be possible given the evidence. Obtaining the closest fit to the data can, in precisely specified circumstances, require the introduction of undetected particles.

Inductive norms, like 'inductive logic', typically aim to sort inferences into a 'justified' vs. 'unjustified' dichotomy, in analogy with the notion of a valid argument in deductive logic. Means-ends analysis, on the other hand, asks whether a given inference rule helps or hinders the aims of inquiry. An inquirer may have several cognitive objectives at the same time, which can stand in tension with each other. How a rational inquirer balances different epistemic values depends on a number of factors, including the weight she assigns to various considerations, her cognitive capacities, and the general historical and disciplinary context of her research. Means-ends methodology does not aim to determine a priori what constitutes the uniquely 'justified' balance among different epistemic considerations. Instead, the goal is to systematically map out the terrain of epistemic values, determine when a theoretical objective is feasible, what methods attain it when feasible, in what situations exactly conflicts among epistemic values arise, how great such conflicts can be, and what the options for resolving them are.

The contrast between Occam's Razor, taken as a norm for inductive inference along the lines of 'do not needlessly multiply entities' and my analysis of the function of undetected entities in inquiry illustrates the spirit of means-ends analysis. I asked whether 'multiplying entities' serves cognitive objectives. We encountered a wealth of considerations that can potentially pull a theorist in different directions: Epistemic conservatism and cognitive feasibility towards undetected entities, parsimony and the risk of falsely positing hidden entities away from them.

The aim of this paper was to illustrate in a specific problem, close to scientific practice, how means-ends methodology can point us to some of the pivotal methodological issues: what are the crucial assumptions that make the aims of inquiry attainable? How do various elements of scientific theories serve the aims of inquiry? The models in this paper display a wealth of epistemological structure in particle physics that highlights the complex methodological challenges that arise in scientific inquiry.

12 Acknowledgments

Clark Glymour pointed me to Valdés-Pérez' work on the methodology of selection rules. I am grateful to Raúl Valdés-Pérez for many helpful explanations and suggestions. Jeff Schanding pointed out to me the role of undetected particles in achieving the closest fit to the observed reactions. Jesse Hughes suggested Proposition 3 and an outline of its proof. Kevin Kelly contributed many methodological insights. Daniel Osherson and Alexandre Korolev provided helpful comments and suggestions on a draft of this paper. Robert Coleman and Manuella Vincter gave useful guidance to the relevant physics. For the relationship between conservation of energy and momentum, on the one hand, and the number of selection rules on the other (Theorem 5) I am indebted to exchanges with Alexander Rueger, David Sharp, Robert Coleman, and Paul Teller. This research was supported by a grant from the Natural Sciences and Engineering Research Council of Canada.

Department of Philosophy
University of Alberta
4-89 Humanities Centre
Edmonton, AB T6G 2E5

Canada
oliver.schulte@ualberta.ca

13 Proofs

For a set of vectors V, I denote the set of its linear combinations by $span(V) = \{\mathbf{v} : \mathbf{v} = \sum a_i \mathbf{v}_i$, where $\mathbf{v}_i \in V$ and a_i is a scalar$\}$.

Proposition 2 *Let K be a set of reactions representing background knowledge, and let ε be an infinite sequence of observed reactions. Then there is an empirically adequate conservation theory for $\varepsilon \iff$ all reactions are observed that are linear combinations of the observed reactions and consistent with K.*

Proof. Let n be the number of known particles, so that we are considering the vector space R^n. Let Q be an $m \times n$ conservation matrix with m quantum properties. Multiplying Q by an n-ary vector v yields an m-ary vector. The set of all vectors that are consistent with Q is the set of all vectors v that are orthogonal to each row in Q, that is, $Qv = \mathbf{0}$; this set is called the kernel of Q and denoted by $\ker(Q)$. It is easy to see that the kernel of Q is a linear subspace, i.e., closed under linear combinations. It follows that a conservation theory Q is empirically adequate on a reaction stream ε given background knowledge K iff $\ker(Q) \cap K = range(\varepsilon)$. Thus if a conservation theory Q is empirically adequate for ε given K, then all reactions that are (1) linear combinations of observed reactions and (2) consistent with K must be eventually observed on ε. So the proof of the proposition is complete if we establish that whenever $range(\varepsilon) = span(range(\varepsilon)) \cap K$, then there is a conservation theory Q such that $range(\varepsilon) = \ker(Q) \cap K$. I will show that (*) for every set of vectors E representing observed transitions, there is a conservation matrix Q such that $\ker(Q) = span(E)$. The proposition follows from (*) by taking $E = range(\varepsilon)$.

The orthogonal complement of a set of vectors V is denoted by $V^{\perp}$, defined as $V^{\perp} = \{v' : v \cdot v' = 0 \text{ for all } v \in V\}$. That is, $V^{\perp}$ is the set of all vectors that are orthogonal to *every* vector in V. It is a standard fact that if V is a linear space, then $(V^{\perp})^{\perp} = V$. Since $span(E)$ is a subspace of R^n, it follows from a familiar theorem of linear algebra that R^n is the direct sum of $span(E)$ and $[span(E)]^{\perp}$, that is, every vector $\mathbf{r} \in R^n$ can be uniquely written as $\mathbf{r_E} \oplus \mathbf{r_{E^{\perp}}}$, where $\mathbf{r_E} \in span(E)$ and $\mathbf{r_{E^{\perp}}} \in [span(E)]^{\perp}$. Now choose a basis $\{\mathbf{b_1}, \mathbf{b_2}, ..., \mathbf{b_m}\}$ for $span(E)$, where $m \leq n$. By another standard fact, the orthogonal complement of a set of vectors is a subspace, so we may choose a basis $\{\mathbf{q_{m+1}}, ..., \mathbf{q_n}\}$ for $[span(E)]^{\perp}$ such that $\{\mathbf{b_1}, \mathbf{b_2}, ..., \mathbf{b_m}\} \cup \{\mathbf{q_{m+1}}, ..., \mathbf{q_n}\}$ is a basis for R^n. I observe that (**) $V^{\perp} = [span(V)]^{\perp}$. It is immediate that $[span(V)]^{\perp} \subseteq V^{\perp}$ because every vector that is orthogonal to all vectors in the span of V is orthogonal to those in V. Conversely, let $\mathbf{v}^{\perp}$ be a vector in $V^{\perp}$, and let $\mathbf{v}$ be a vector in $span(V)$. Then $\mathbf{v} = \sum a_i \mathbf{v_i}$ for scalars a_i and vectors $\mathbf{v_i} \in V$. So $\mathbf{v}^{\perp} \cdot \mathbf{v} = \mathbf{v}^{\perp} \cdot (\sum a_i \mathbf{v_i}) = \sum a_i (\mathbf{v}^{\perp} \cdot \mathbf{v_i}) = 0$ since $\mathbf{v}^{\perp}$ is orthogonal to each vector in $\mathbf{v_i}$. Hence $V^{\perp} \subseteq [span(V)]^{\perp}$, which establishes (**).

Let Q be the matrix whose rows are $\mathbf{q_{m+1}}, ..., \mathbf{q_n}$. The kernel of Q is the set of all vectors that are orthogonal to each row in Q; so $\ker(Q) = \{\mathbf{q_{m+1}}, ..., \mathbf{q_n}\}^\perp$. By (**), $\{\mathbf{q_{m+1}}, ..., \mathbf{q_n}\}^\perp = [span(\{\mathbf{q_{m+1}}, ..., \mathbf{q_n}\})]^\perp$. Since $\{\mathbf{q_{m+1}}, ..., \mathbf{q_n}\}$ is a basis for $[span(E)]^\perp$, we have that $span(\{\mathbf{q_{m+1}}, ..., \mathbf{q_n}\}) = [span(E)]^\perp$. Hence $\ker(Q) = (span(E)^\perp)^\perp = span(E)$, as required. $\square$

In a conservation theory Q, each quantum property $\mathbf{q}$ assigns a value to one of a set P of particles. I refer to P as the set of particles in Q. Among the particles P, we distinguish a subset $D \subseteq P$ of detectable particles. For a given vector $\mathbf{v} \in R^{|P|}$, the *observable part* of $\mathbf{v}$ is the restriction of $\mathbf{v}$ to the particles in D; formally, $\mathbf{v}|D$ is the vector $\mathbf{v'}$ defined by $\mathbf{v'}(i) = \mathbf{v}(i)$ if $p_i \in D$, and $\mathbf{v'}(i) = 0$ otherwise. For a set of vectors V, I define $V|D = \{\mathbf{v'} : \mathbf{v'} = \mathbf{v}|D$ for some vector $v \in V\}$. For a set of particles P, let $R(P) = \mathbf{N}^{|P|}$ be the set of vectors that could represent reactions among the particles in P, that is, those vectors with integral components. Finally, let $int - span(V)$ be the set of linear combinations of vectors in V with integral coefficients.

Proposition 3 *Let E be a set of observed transitions with empirical parts $e_1, e_2, ..., e_k$. Then*

1. *every conservation theory, whether it involves undetected particles or not, is consistent with linear combinations of $\mathbf{e_1}, \mathbf{e_2}, ..., \mathbf{e_k}$ that represent transitions and have integral coefficients only.*

2. *there is a conservation theory, possibly involving undetected particles, that is consistent with exactly those linear combinations of $\mathbf{e_1}, \mathbf{e_2}, ..., \mathbf{e_k}$ that represent transitions and have integral coefficients only.*

Proof. Let $E = \{\mathbf{e_1}, \mathbf{e_2}, ..., \mathbf{e_k}\}$ be a set of vectors representing observed transitions, let $r \in span(E)$, let Q be a conservation theory, and let P be the set of particles in Q.

Part 1. For each $\mathbf{e_i}$, choose a reaction r_i involving only the particles in P such that the visible part of r_i is e_i, i.e. $r_i|D = e_i$, and r_i is physically possible according to Q, that is, $\mathbf{r_i} \in \ker(Q)$. Let t be a transition in the integral span of E, such that $\mathbf{t} = \sum z_i \mathbf{e_i}$ for integers $z_1, ..., z_{|E|}$. Then $[Q](\sum z_i \mathbf{r_i}) = \sum [Q](z_i \mathbf{r_i}) = \mathbf{0}$, so $\mathbf{r} = \sum z_i \mathbf{r_i}$ is physically possible according to Q. But $\mathbf{r}|D = \sum z_i \mathbf{r_i}|D = \sum z_i \mathbf{e_i} = \mathbf{t}$. So the empirical content of Q includes the observation of t, as required.

Part 2. I show how to find the required conservation matrix Q by induction on $|E|$.

Base Case: $E = \emptyset$, so $span(E) = \{\mathbf{0}\}$. Let Q_0 be the $|D| \times |D|$ identity matrix. Clearly $\ker(Q_0) = \{\mathbf{0}\}$. Inductive Step: Let P be the set of particles in Q_{k-1} (so Q_{k-1} is an $m \times |P|$ matrix). Let $\mathbf{e}_1, .., \mathbf{e}_k$ be the vectors of dimension $|P|$ that represent the observed transitions $e_1, ..., e_k$. Suppose that $(\ker(Q_{k-1}) \cap R(P))|D = int - span(\{\mathbf{e_1}, \mathbf{e_2}, ..., \mathbf{e_{k-1}}\})$. If $int - span(\{\mathbf{e_1}, \mathbf{e_2}, ..., \mathbf{e_{k-1}}\}) = int - span(\{\mathbf{e_1}, \mathbf{e_2}, ..., \mathbf{e_k}\})$, we may just take $Q_k = Q_{k-1}$ with the same set of particles P. Otherwise, reinterpret the transition

$e_k = agents(e_k) \rightarrow products(e_k)$ by adding a hidden particle h among the products to arrive at the 'actual' reaction $e_{k,h} = agents(e_k) \rightarrow products(e_k) + h$. Formally, let $h = p_{n+1}$, and set $e_{k,h}(n+1) := -1$. Note that $e_{k,h}|D = e_k$ since all particles reported in e_k are detectable. Extend the matrix of quantum properties Q_{k-1} with $|P|$ columns to a matrix Q_k with $|P|+1$ columns by modifying each quantum property (row) $\mathbf{q}$ in Q_{k-1} as follows: $\mathbf{q}'|P = \mathbf{q}|P$, and $\mathbf{q}'(n+1) = \mathbf{q} \cdot \mathbf{e_k}$. That is, consider each quantum property $\mathbf{q}$, and assign to h the value $\mathbf{q}'(n+1)$ that is required to restore the balance of $\mathbf{q}$ in e_k, if there is an imbalance. The introduction of the new particle adds a dimension to the vector space. The component of each quantum property in the new dimension is determined by choosing the component just so that it balances the count of e_k for the quantum property.

One notational difficulty that arises now is that the vector representation of a reaction depends on the set of all particles under consideration. I shall continue to write $\mathbf{r}$ for a vector of dimension $|P|$ that represents the reaction r with respect to the particles P, and I write $\mathbf{r}'$ for the vector of dimension $|P|+1$ that represents the same reaction with respect to the particles $P \cup \{h\}$. Note that $\mathbf{r}|D = \mathbf{r}'|D$ since $h \notin D$, and that $\mathbf{r}|P = \mathbf{r}'|P$ since $h \notin P$.

I argue that the empirical content of Q_k is exactly the integral span of $\{\mathbf{e'_1}, ..., \mathbf{e'_k}\}$, that is, $(\ker(Q_k) \cap R(P \cup \{h\})|D = int-span(\{\mathbf{e'_1}, ..., \mathbf{e'_k}\})$.

($\supseteq$) By inductive hypothesis, we may choose for each $i < k$ a reaction r_i among only the particles in P that is physically possible according to Q_{k-1} and whose visible part is e_i. Since $r_i(n+1) = 0$, for each extended quantum property $\mathbf{q}'$, we have that $\mathbf{q}' \cdot \mathbf{r'_i} = \mathbf{q} \cdot \mathbf{r_i} + \mathbf{q}' \cdot r_i(n+1) = \mathbf{q} \cdot \mathbf{r_i} = 0$ by the assumption that r_i is physically possible according to Q_{k-1}. In other words, r_i conserves all of the extended quantum properties in Q_k because r_i conserves the unextended quantum properties in Q_k and does not feature the hidden particle h. Therefore r_i is physically possible according to Q_k. Also, we have that $\mathbf{q}' \cdot \mathbf{e'_{k,h}} = \mathbf{q} \cdot \mathbf{e_k} + e_{k,h}(n+1)\mathbf{q}'(n+1) = \mathbf{q} \cdot \mathbf{e_k} - \mathbf{q} \cdot \mathbf{e_k} = 0$. So $e_{k,h}$ is physically possible according to Q_k, and the visible component of $e_{k,h}$ is e_k.

It follows by Part 1 that the visible parts of the physically possible transitions according to Q_{k+1} include all transitions in the integral span of $\{\mathbf{e'_1}, ..., \mathbf{e'_k}\}$.

($\subseteq$) Note that Q_k contains as many linearly independent quantum properties as Q_{k-1}, that is, $rank(Q_k) = rank(Q_{k-1})$. The dimension of the kernel of a matrix Q is called the *nullity* of Q, written $null(Q)$. By a standard theorem of linear algebra, the nullity of Q_{k-1} is the number of particles in P minus the rank of Q_{k-1}, and the nullity of Q_k is $|P| + 1 - rank(Q_k) = null(Q_{k-1})+1$. Let $\{\mathbf{b'_1}, ..., \mathbf{b'_{null(Q_{k-1})}}\} \subset R(P \cup \{h\})$ be a basis for $\ker(Q_{k-1})$ comprising possible reactions involving *only* the particles in P. (We may choose $\{\mathbf{b'_1}, ..., \mathbf{b'_{null(Q_{k-1})}}\}$ as a subset of $\{\mathbf{e'_1}, ..., \mathbf{e'_k}\}$.) Then since the reactions encoded by $\mathbf{b'_1}, ..., \mathbf{b'_{null(Q_{k-1})}}$ do not involve h, we have that $\{\mathbf{b'_1}, ..., \mathbf{b'_{null(Q_{k-1})}}\} \subset \ker(Q_k)$, and that $\mathbf{e'_{k,h}}$ is independent of $\{\mathbf{b'_1}, ..., \mathbf{b'_{null(Q_{k-1})}}\}$ because $e_{k,h}(h) = -1$. So $\{\mathbf{b'_1}, ..., \mathbf{b'_{null(Q_{k-1})}}, \mathbf{e'_{k,h}}\}$ is a basis for $\ker(Q_k)$.

Now let $\mathbf{r}' \in \ker(Q_k)$ encode a physically possible reaction r according to Q_k. Then $\mathbf{r}' = \sum_{i=1}^{null(Q_{k-1})} a_i \mathbf{b'_i} + a\mathbf{e'_{k,h}}$ for scalars $a, a_{1,...}, a_{|P|}$. Since none of the reactions encoded by $\mathbf{b'_i}$ involve h, and $e^h_k(h) = -1$, it follows that a must be an integer if r is a possible reaction. Let $\mathbf{r'_{old}} = \sum_{i=1}^{null(Q_{k-1})} a_i \mathbf{b'_i}$ and $\mathbf{r'_{new}} = a\mathbf{e'_{k,h}}$. Then r_{old} is a possible reaction according to Q_{k-1}, that is, $\mathbf{r_{old}} \in \ker(Q_{k-1}) \cap R(P)$. So by inductive hypothesis, the visible component of r_{old} is in the integral span of $\{\mathbf{e_1}, ..., \mathbf{e_{k-1}}\}$, and hence in the integral span of $\{\mathbf{e'_1}, ..., \mathbf{e'_{k-1}}\}$; let $\mathbf{r'_{old}}|D = \sum_{i=1}^{k-1} z_i(\mathbf{e'_i}|D)$ for integers z_i. Thus the visible component of r is in the integral span of $\{\mathbf{e'_1}, ..., \mathbf{e'_k}\}$, since $\mathbf{r'}|D = \mathbf{r'_{old}}|D + \mathbf{r'_{new}}|D = \sum_{i=1}^{k-1}(z_i\mathbf{e'_i}|D) + a(\mathbf{e'_{k,h}}|D)$, where a is an integer. This shows that the empirical content of all reactions that are physically possible according to Q_{k+1} is the integral span of $\{\mathbf{e'_1}, ..., \mathbf{e'_k}\}$, which completes the inductive step. $\square$

References

[Aris 1969] Aris, R. [1969]: *Elementary Chemical Reactor Analysis*, Englewood Cliffs, N.J.: Prentice Hall.

[Case and Smith 1983] Case, J. and Smith, C. [1983]: 'Comparison of Identification Criteria for Machine Inductive Inference', *Theoretical Computer Science*, **25**, pp.193–220.

[Cooper 1992] Cooper, L. [1992]: *Physics: Structure and Meaning*, Hanover: University Press of New England.

[Earman 1992] Earman, J. [1992]: *Bayes or Bust?*, Cambridge, Mass.: MIT Press.

[Feynman 1965] Feynman, R. [1965; 19th ed. 1990], *The Character of Physical Law*, Cambridge, Mass.: MIT Press.

[Franklin 1990] Franklin, A. [1990]: *Experiment, right or wrong*, Cambridge: Cambridge University Press.

[Galison 1997] Galison, P.L. [1997]: *Image and logic : a material culture of microphysics*, Chicago : University of Chicago Press.

[Gold 1965] Gold, E.M. [1965]: 'Limiting Recursion', *Journal of Symbolic Logic*, **30**[1], pp. 28–48.

[Gold 1967] Gold, E. [1967]: 'Language Identification in the Limit', *Information and Control*, **10,** pp. 447–474.

[Goldman 1986] Goldman, A. [1986], *Epistemology and Cognition*, Cambridge, Mass.: Harvard University Press.

[Hacking 1983] Hacking, I. [1983]: *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*, Cambridge: Cambridge University Press.

[Harman 1986] Harman, G. [1986]: *Change in View: Principles of Reasoning*, Cambridge, Mass.: MIT Press.

[Hempel 1965] Hempel, C.G. [1965]: *Aspects of Scientific Explanation*, New York: Free Press.

[Juhl 1994] Juhl, C. [1994]: 'The Speed-Optimality of Reichenbach's Straight Rule of Induction', *British Journal for the Philosophy of Science*, **45**, pp. 857–863.

[Juhl 1995] Juhl, C. [1995]: 'Is Gold-Putnam Diagonalization Complete?', *Journal of Philosophical Logic*, **24,** pp. 117–138.

[Juhl 1997] Juhl, C. [1997]: 'Objectively Reliable Subjective Probabilities', *Synthese*, **109**, pp. 293-309.

[Kelly 1996] Kelly, K. [1996]: *The Logic of Reliable Inquiry*, Oxford: Oxford University Press.

[Kelly *et al.* 1997] Kelly, K., Schulte, O. and Juhl, C. [1997]: 'Learning Theory and the Philosophy of Science', *Philosophy of Science*, **64**, pp. 245–267.

[Kuhn 1970] Kuhn, T. [1970]: *The Structure of Scientific Revolutions.* Chicago: University of Chicago Press.

[Martin and Osherson 1998] Martin, E. and Osherson, D. [1998]: *Elements of Scientific Inquiry*, Cambridge, Mass.: MIT Press.

[Neeman and Kirsch 1983] *The Particle Hunters*, Cambridge: Cambridge University Press.

[Omnes 1971] Omnes, R. [1971]: *Introduction to Particle Physics*, London, New York: Wiley Interscience.

[Putnam 1963] Putnam, H. [1963]: "Degree of Confirmation' and Inductive Logic', in *The Philosophy of Rudolf Carnap*, ed. A. Schilpp, La Salle, Ill: Open Court.

[Putnam 1965] Putnam, H. [1965]: 'Trial and Error Predicates and a Solution to a Problem of Mostowski', *Journal of Symbolic Logic*, **30**, pp. 49–57.

[Putnam 1975] Putnam, H. [1975]: 'Probability and Confirmation', in *Mathematics, Matter and Method, Philosophical Papers Vol. 1*, London: Cambridge University Press.

[Reichenbach 1949] Reichenbach, H. [1949]: *The Theory of Probability*, London: Cambridge University Press.

[Salmon 1991] Salmon, W. [1991]: 'Hans Reichenbach's Vindication of Induction', *Erkenntnis*, pp. **35,** pp. 99–122.

[Schulte 1999a] Schulte, O. [1999]: 'Means-Ends Epistemology', *The British Journal for the Philosophy of Science*, **50**, pp.1–31.

[Schulte 1999b] Schulte, O. [1999]: 'The Logic of Reliable and Efficient Inquiry', *The Journal of Philosophical Logic*, **28**, pp. 399–438.

[Sextus Empiricus 1985] Sextus Empiricus [1985]: *Selections from the Major Writings on Skepticism, Man and God,* ed. P. Hallie, trans. S. Etheridge, Indianapolis: Hackett.

[Sklar 1975] Sklar, L. [1975]: 'Methodological Conservatism', *Philosophical Review* **LXXXIV**, pp. 374–400.

[Valdés-Pérez and Erdmann 1994] Valdés-Pérez, R. and Erdmann, M. [1994]: 'Systematic Induction and Parsimony of Phenomenological Conservation Laws', *Computer Physics Communications*, **83,** pp. 171–180.

[Valdés-Pérez and Żytkow 1996] Valdés-Pérez, R. and Żytkow, J. [1996]: 'Systematic Generation of Constituent Models of Particle Families', *Physical Review E*, **54,** pp. 2102–2110.

[Van Fraassen 1980] Van Fraassen, B. [1980]: *The Scientific Image*, Oxford: Clarendon Press.

[Williams 1997] Williams, W.S.C. [1997]: *Nuclear and Particle Physics*, New York: Oxford University Press.

[Wolfenstein and Trippe 1998] L. Wolfenstein and T.G. Trippe [1998]: 'Tests of Conservation Laws', *The European Physical Journal C*, **3**[1–4], pp. 62–67 .