Kinds of Inconsistency*†

Gregory R. Wheeler
Department of Philosophy
University of Rochester
Rochester, NY 14627 U.S.A.

April 21, 2001

Abstract

Typically paraconsistency is motivated by one of two types of arguments: either an ontological argument is given whereby it is claimed that inconsistent objects or events of some kind demand a logic able to reason about such items, or an epistemological argument is provided that denies that language or the world is infected with inconsistency but claims instead that the problem is all in a reasoning agent's head. Approaches to designing paraconsistent logics have thus far tended to heed this distinction. Stronger, dialetheic logics are called on to handle contradictory objects or events in a manner that does not trivialize inference, while these logics are out of favor with epistemological approaches which favor preserving as much of classical logic as possible. This latter approach is typically viewed as one of introducing a set of techniques for how to manipulate an agent's inconsistent yet rational corpus to preserve a familiar form of logical consequence. This essay argues against the practice of linking epistemological concerns with weak paraconsistent logics. It introduces an epistemological motivation for adopting a strong paraconsistent logic by considering the results of measuring physical objects and how we might go about reasoning with these results. Measurement behaves in a way as to be a source of inconsistency that is neither best understood as a problem between agents nor the result of either a paradoxical property of language or the world.

1 Introduction

There are distinct kinds of inconsistency. My interest is to introduce a kind that occurs in languages whose domains include physical objects. In other words, I am interested in a kind of inconsistency that is fundamental to talk about

*This research was supported in part by grant SES 990-6128 from the National Science Foundation and generous support from the National Research Council of Brazil.

†Thanks to Eva Cadavid, Henry Kyburg, Graham Priest, Catherine McKeen, Gabriel Uzquiano and WCP2000 referees for their comments.

domains like our very own –that is, domains consisting of things we observe, measure, and about which we reason.

The claim is not that there are many kinds of inconsistent theories. This claim can be easily substantiated by citing a pair of axiomatic systems or corpora of beliefs that, when represented in a classical first-order language, each entail a sentence satisfying the schema $\ulcorner \phi \wedge \neg\phi \urcorner$. For instance, naïve set theory and the book of *Genesis* are at least *functionally* different kinds of theories: one and only one of the pair is a mathematical theory. But both are inconsistent. Whatever interest there is for this kind of classification, it will not be ours.

Instead, our interest concerns a source of inconsistency that is peculiar to physical theories. The resulting kind of inconsistency is existent in both folk and scientific theories, but its structure is better revealed in the empirical sciences. The source of the inconsistency is a probabilistic acceptance rule that is fundamental to experimental methods and measurement procedures common in scientific practice.[1] This paper raises the specter of a distinct 'kind' of inconsistency and presents a pair of probabilistically motivated examples to serve two ends. The first objective is to put to an end the practice of identifying certain philosophical motivations for paraconsistency with particular classes of paraconsistent logics. The tendency to view epistemological issues as solely treatable by 'weak' paraconsistent logics, for instance, is a mistake that impedes progress both in philosophy and logic. The second goal is of greater importance: an understanding of scientific inference. Paraconsistent logicians have often turned to the sciences for examples to apply paraconsistent systems. Below I introduce a structure within scientific inference that is both commonplace and ready made for paraconsistent logic.

2 Revising Theories, Revising Logic

Why concern ourselves with an inconsistent theory of *any* kind? After all, ours is an intellectual tradition that treats inconsistency as a pox rather than a provision. W. V. O. Quine [29, p. 81] refers to so-called logics that tolerate inconsistencies as a "change of subject" rather than logic at all. As for inconsistent theories, Karl Popper's well known polemics against the very idea rest on what he takes to be the fundamental requirement for theoryhood:

> The requirement of consistency plays a special role among the various requirements which a theoretical system, or an axiomatic system, must satisfy. It can be regarded as the first of the requirements to be satisfied by *every* theoretical system, be it empirical or non-empirical. [24, p. 91f.]

[1] Some recent work in the philosophy of science and computer science on probabilistic acceptance rules and the distinction between an acceptance view and so-called subjective probability views are Deborah Mayo [22] and Choh Man Teng [33], the latter being technical work that builds upon Kyburg [15]. See also Kyburg [17] for a defense of acceptance and replies in the same volume.

So, an *accepted* inconsistent theory is either by Quine's lights a non-logical program or, by Popper's, not something deserving of the name 'theory' at all. Pox indeed.

And the cure? Revision. A contradiction deduced from a set of axioms tells against that set: at least one of the axioms must be revised in order to restore consistency. How to go about this, though, is a matter of considerable controversy. Quine, for one, considers all statements fair game. Theories are partially-ordered sets of statements –a.k.a a 'web of belief'– whose ordering corresponds to each statement's proximity to observation. In most cases it will turn out cheaper and easier to revise the statements closer to the observational end of this scheme, but nothing precludes tinkering with the axioms of a theory to restore consistency. Popper, on the other hand, accepts a sharp analytic/synthetic distinction, holds axioms of logic and mathematics fixed and certain, and demands severe testing of all 'corroborated' empirical statements. He does so while coaching us not to *accept* any non-mathematical or non-logical truth, no matter how rigorously tested. Despite considerable disagreement surrounding how to characterize revision, and even disagreement over how to regard theoretical structures in general, the discussion about revision is carried out with the assumption that inconsistency is intolerably problematic. The question, tradition tells us, is not *whether* to revise but *how*.[2]

Enter paraconsistent logic. We might view paraconsistent logic, and paraconsistency in general, as a proposal to make 'whether to revise' a reasonable question. Rather than view the problem of revising an inconsistent theory as simply one of setting a theory straight, paraconsistency suggests that we consider the costs involved in doing so. Determining just what those expenses are, however, has divided the field. On the one hand we find the reform-minded dialetheists. These logicians recommend adopting a new logic to accommodate *bona-fide* inconsistent objects. Their favored strategy is to argue for cases where revising an inconsistent theory would be an outright mistake. On the other hand we find neo-classicists who take logic and the world to be consistent, even if rather complicated. The motivation behind this movement is the recognition that fallible agents often reason with useful but imperfect knowledge bases. Given these circumstances, it may prove cognitively expensive to revise an inconsistent theory or the choices available to an agent may take him farther away from the truth.

These two motivations for paraconsistency have settled into a peculiar form of partisanship.[3] Epistemic motivations are thought to be tied to conserva-

---

[2]The philosophical problem of how to revise an inconsistent empirical theory has its first explicit formulation in Pierre Duhem's *The Aim and Structure of Physical Theory*. For recent work on Duhem's problem, or the problem of how to revise a theory given disconfirming observational evidence that is contrary to a theory's predicted value for that observation, see Worrall [35], Mayo [22] and [23], Howson [11], and Wheeler [34]. For theory revision in computer science, see Alchourrón et. al. [1], the collection edited by Gärdenfors [10] and Kyburg and Teng [19].

[3]Exceptions hold, of course. Two to note are Brandon and Rescher's logic of inconsistency [5] and the recent debate concerning logical pluralism [30]. Brandon and Rescher's inconsistent logic is an ontologically motivated project which nevertheless retains classical logical

tive, neo-classical approaches to paraconsistency, while ontological concerns are thought to demand bolder, dialetheic approaches. This habit is unfortunate. One problem with conflating weak paraconsistent logics with epistemic concerns, for instance, is an oversimplification of the issues at hand. Consider the probabilistic examples cited in the epistemic case for paraconsistency. These examples most often share the same underlying structure of Henry Kyburg's lottery paradox [15], [18]. An inconsistency can be generated by an n-member collection of accepted statements such that the chance of mistakenly rejecting each statement of the collection is no greater than ϵ, but that the chance of mistakenly rejecting the n-member conjunction when joined by adjunction is greater than ϵ. Most proposed solutions to the lottery paradox turn on some non-adjunctive strategy to contain the loss of probability mass that occurs by using the rule of adjunction. But, as we shall see, there is a problem which has its roots in probabilistic acceptance rules yet is overlooked by this approach. The problem's structure suggests a stronger paraconsistent treatment, one which would otherwise be dismissed out of hand. I return to this discussion in the next section.

What is interesting about paraconsistency and the family of logics going by the same name is the possibility of using them to understand how what David Israel [12] has called "real, honest-to-goodness inference" works. Necessary to understanding such inference practices is an understanding of how working theories are revised with accepted but fallible evidence. For this latter problem, scientific inference presents an ideal test case. Assume that revisions to physical theories are due neither to capricious Quinean fancy nor merely to trial-and-error Popperian hypothesis testing, but are instead evidence-driven events. If we understand scientific inference as an evidence driven enterprise, what we must first come to terms with is that the domain of such reasoning isn't the world *simpliciter*, but a model of the world. In the best case, this model is built in part from experimental observations and measurement practices that each come along with known error probabilities. Discussing paraconsistency in light of scientific inference raises a third question to add to the pair cited above. In addition to the questions of *how* to revise a theory and *whether* to revise a theory, we might add to the discussion the question of exactly *what* is to be revised. I turn to this question next.

3 What what?

What are the objects with which science reasons? Furthermore, what role, if any, should paraconsistent logic play in understanding scientific reasoning? Let's begin with the second question. A review of the paraconsistency literature turns up two examples from science that are often proposed as candidates

consequence. They wind up handling inconsistent objects by introducing a semantics for inconsistent worlds instead. The second exception is to how reform-minded a dialethist must be. Logical pluralists argue that accepting a paraconsistent logic needn't entail the rejection of classical logic *per se*. See Priest [27] for a reply to Restall and Beall.

for paraconsistent treatment. The first is the so-called dual nature of quantum physical states.[4] But this is a problem that has little to do with scientific inference, at least as far as paraconsistency is concerned, since Quantum mechanics is typically given a statistical interpretation, such as Born's. The idea here is to concede that Quantum theory doesn't predict what the measured value of X will be but, instead, provides a probability distribution over a set of possible measured values. The kind of inference involved then is statistical, not deductive. Hence, worries about paraconsistency are beside the point. The second example is the problem of reconciling the mathematical assumptions of quantum theory with those of relativity theory.[5] This problem, however, is primarily the problem of how best to unify two distinct yet incompatible *theories*. The problem of unification should be distinguished from the problem of revision. A serious answer to the unification problem is just to give up on the largely philosophical project of unifying science.[6] Notice that whatever the merits of this proposal, a similar reply to the problem of theory revision isn't an option. The revision problem has it that there is recalcitrant evidence within the domain described by a theory. Giving up on squaring theory with evidence is simply giving up doing science.

If we wish to understand scientific inference we'll need to focus on evidence. The temperatures of solutions, tensile strengths of compounds, and the masses of slabs function as evidential bedrock for evaluating scientific hypotheses and, in turn, scientific theories. Since the bulk of scientific reasoning concerns measured magnitudes of physical objects or physical events, it's here that we should look for a connection with paraconsistent logic.

Consider how the magnitude *length* is measured. The measurement of length is paradigmatic of all fundamental additive measurements in the physical sciences, such as mass, velocity, and the like. Imagine that a pole S is exactly 4 meters long. What, on the basis of measurement, does one conclude: that the length of S is 4 meters? No, not exactly. If careful, one concludes that the length of S is plus-or-minus d units of r, where r is the recorded value of the measurement. If measured carefully, the true value, 4, will very likely fall within

[4] The textbooks tell us that the physical state of any isolated system behaves deterministically in accordance with Schrödinger's equation until a measurement of some physical magnitude (*e.g.*, position, energy, spin) is made. A pre-measured physical state, modeled by a state-vector Φ, is not a linear combination of vectors $\langle \phi_1, \phi_2, ..., \phi_n \rangle$ that represent particular values $\langle x_1, x_2, ..., x_n \rangle : x \in \mathbb{R}$ of the physical magnitude X in question, but rather measurement of some other observable not commensurate with X. Thus, $\Phi = a_1\phi_1 + a_2\phi_2 + ... + a_n\phi_n$, where each a_i is a complex scalar, many of which in fact are, typically, non-zero. The result then is that Φ does not have a definite value X. But on measurement of X, the state Φ changes, or "collapses", taking one of the measured values $x_1, x_2, ..., x_n$ of X. See Krips [14] for discussion.

[5] Relativity theory concerns macroscopic objects and treats collision as a primitive, or would if one assumes that the colliding particles are infinitely small. Quantum theory, in turn, concerns microscopic objects and cannot be said to treat collision as primitive since collision need not be isolated in space-time. This conflict may be traced to the mathematics underpinning of each theory: Relativity theory relies on a 4-dimensional Riemann space, whereas Quantum theory operates on an infinite-dimensional Hilbert space.

[6] The philosopher Nancy Cartwright and physicist Richard Feynman are notable critics of the unification of science thesis.

the interval $r - d < r < r + d$, where $d \ll r$. The justification for this conclusion is the belief that it is simply incredible, given our knowledge of the methods and technology used to measure S, that the actual value which r is intended to represent (in this case, 4) does in fact fall outside the interval $r - d < r < r + d$. In standard practice we set the value of d only as high as our needs demand and as low as our instruments warrant to insure our incredulity. Once all this is in place we may *accept* that S is $r \pm d$ units long and then go about using or manipulating S, confident (although not certain) that its actual length is within specification.

Turning now to reasoning, consider a simple theory whose object language contains the sentence

(P) '$\texttt{Length\_of}(S)$ $\texttt{is}$ 3.95 ± 0.10 $\texttt{meters}$',

where S is the name of a flagpole, and r takes the value 3.95 and d, 0.10. If P is accepted, that is if P is assigned the value 1, we assume that there is a class of models which satisfy this sentence. Classically speaking we say that P is true in just the models making up this class, false otherwise. Standard declarative semantics has it that the actual length of S is an element we are firmly in command of; setting aside differences over how to represent S's 4-meter-longishness, S's *being* 4 meters long is generally taken to be the truth-maker for models of P.[7] But underpinning this abstraction is an assumption that, perhaps, we now are in a position to improve upon. The operative assumption behind explaining what counts as classical models of P is the assumption that physical properties and relations are readily accessible. Poles may be tall and not tall, but they aren't both 4 meters and not.[8] Likewise, the proposition expressed by (P) is either true or false. Yet the one thing conspicuously missing from any realistic method of *discovering* the truth value of statements like (P) is S's actual length. If S's length were directly accessible, we wouldn't need to measure S and science would be considerably easier than it in fact is. Instead, what stands between us and S's length is a fallible measurement procedure. So when considering arguments involving evidence statements like (P), rather than working with the propositions expressed by such statements we have instead just a collection of accepted statements –statements that give a reliable report about the physical state in question. So, while scientists reason *about* the length of a pole, very often they reason *with* accepted statements about the length of a pole. This collection of accepted statements is simply the brute data that can be shared or generated by researchers. Let's call this collection of statements a *domain of discourse*.[9] In the best of circumstances we know quite a bit

[7] Or, S being r meters long is taken as a truth-maker for P when $r \in [3.85, 4.05]$.

[8] In other words, the issue of vagueness is orthogonal to the topic at hand.

[9] I prefer 'domain of discourse' to 'data reports' or Deborah Mayo's 'models of data' only because these latter terms leave open the question of just where the statements stand with respect to data analysis –namely, pre- or post-analysis. I intend 'domain of discourse' to concern statements about pre-processed, raw observational data. 'Observation statement' might do as well so long as it is not taken to denote beliefs, a tempting move for those charmed by subjective interpretations of probability. At any rate, data analysis turns out to

about how frequently such reports are mistaken and how such mistakes occur. Returning to poles, this knowledge is simply the know-how of measuring poles with a tape.[10]

Falling into error simply amounts to accepting a false report within the domain of discourse. A false report is just an accepted statement that either records a value for r (on the basis of a reliable measurement procedure) when in fact the value of r is outside this interval, or fails to record a value for r as falling within the acceptance interval that is in fact within that interval. Though erroneous, an accepted but false report is warranted. No report *enters* the domain of discourse by mistake. A reader troubled by the uncertainty introduced by measurement procedures might recommend that to decrease the chances of falling into error, simply increase the value for d. Increasing the value for d would increase the margin of error and thus, assuming that the recorded values for r are normally distributed, would reduce the number of false reports appearing in the domain of discourse. Though mathematically sound, this would be bad advice to follow. Remember that the final goal isn't error elimination but finding the length of poles. Our interest in the truth presses us to accept a minimal value for d whose associated frequency of error is known. The goal of finding the true length of poles is compromised by increasing the interval to stamp out error. Besides, error is eliminated completely only at the price of triviality.[11]

What is interesting about a domain of discourse –that is, a set of statements interpreted as accepted reports rather than expressing true propositions– is that it can be inconsistent. Consider just what good measurement practices deliver: reliable reports about *bona fide* magnitudes. These reports are approximations at best, which is the reason they are interval-valued. However, these interval-valued reports are also fallible. This is the point behind accepted statements: a measurement may report of S that $r \in [x, y]$ when in fact $r \notin [x, y]$. It is important to remember that this situation describes what I prefer to call an error, not an inconsistency. The interesting case is when we happen to encode such errors by measurement, that is when we have an acceptable report of S that $r \in [x, y]$ and an acceptable report of S that $r \notin [x, y]$. In this case we have a pair of statements that ascribe inconsistent properties to S, namely that the value r is and is not in some interval x,y.

This last point is obscured by common, epistemological arguments for paraconsistency. An example should make this point clear. Consider a set of poles, each thought equal in length, yet nevertheless the cause of controversy when installed outside the UN. Though accepted as a set of poles of roughly equal length, seeing a nation's flag atop one of the new poles flying noticeably below those

be one of the important tools needed to do revision –namely the one that does most of the clean-up work on some initial set of statements.

[10] There is an important point about measurement that needs to be brought out here. The choice of technology is irrelevant. A tape and the frustrations of using one are hopefully common enough to most readers to grasp the general point: measurements don't provide direct, certain access to physical properties.

[11] At extremes one could say that the length of r is a real number. But that tells you only what a magnitude is and not, say, where to cut the stock.

atop the others has the makings for a small scandal. Likewise, sawing boards and assembling bookcases presents a similar, perhaps even familiar, problem. The reader can well imagine a situation where someone accepts that each of the boards cut is a certain length but is flummoxed to learn that they don't all fit together to make a bookcase. If we're interested in representing such problems formally, the argument typically goes, then we are discussing a situation that is, even if temporarily, inconsistent.

That it is easy to see around these two problems has more to do with how the stories are told than about how the agents in each story found their troubles. The poles and planks aren't themselves in and out of specification, they are just out of specification. Furthermore, we know they are out of specification. In fact, the stories work the way they do only because we know so. What is important to consider is that we haven't always the benefit of knowing this last piece of information: when faced with an irregularity, it is a luxury to know immediately its source. In fact, science is simply the art of reasoning without this luxury.

Ignoring this point, it is tempting to think that all epistemological troubles rest with individual agents. Neo-classicists think so. The epistemological problem seems to them to be one epitomized by joining the craftsman's corpus (i.e., in particular, that the poles are equal in length) with the assembler's (i.e., they aren't). The structure of the problem, this view has it, lends itself to solution either by tossing out the obviously incorrect belief or, in tough cases, partitioning the joined sets into mutually exclusive but self-consistent sets (See Brown and Schotch [6], and Jennings and Schotch [13]).

But notice that this approach isn't as compelling when we aren't provided with the privileged perspective of knowing whom is to blame. As I've told these stories, we know that the poles and planks aren't all in spec and that their not being so is the key to solving the problems raised in each tale. But suppose now we develop the stories, adding to either one the presence of irregular mounting brackets, warped stock, optical illusions, paranoid diplomats–take your pick. Notice that changing a story in such a manner amounts to changing genre: rather than a morality tale, we have now a mystery. This is a critical point. In such cases we don't know what the cause of the problem is but only that there is one. The methods used to land us in such a fix –those which, at root, rely on a host of measurements reporting on relations and properties of the objects in play– are precisely those we must rely on to get us out.

The question then is how best to represent the problem of what to do in the face of wayward evidence. Following the neo-classical line of treating problems like this as essentially involving disagreement between reasoning agents passes over the state of affairs *any* agent confronts when reasoning about physical objects. All of us start with the same kind of evidence –a more-or-less reliable, inter-subjective representation of the properties and relations of the physical objects we're interested in manipulating. But this evidence can be –and as a practical matter *is*– inconsistent. This isn't to say that the world is inconsistent. Nor either is it to say that the problem is always one of mediating a dispute between agents, each of whom provides self-consistent accounts of some shared event. The problem, in short, is a dispute between us and the world: at our best we only get it almost right, nearly all of the time.

Once we recognize that domains of discourse about physical objects are likely inconsistent, then a proposal to partition the language into consistent cells misses the point. Of course it is preferable to work with a consistent set of statements. Hence, it is desirable to work toward removing errors from a domain of discourse. But what is important to recognize is that this work does not precede making that domain available for reasoning. In practice it may not even be possible to completely remove errors from the domain of discourse. And, even if it were, we still would need to deploy a logic to sort out this set of sentences. So, if a paraconsistent logic is simply a logic whose consequence operation is non-trivial when closing an inconsistent set of statements, then the set of sentences modeled true in a domain of discourse is a prime candidate for paraconsistent logic. Furthermore, since the source of inconsistency is not inter-agent but systemic –it is simply a side-effect of how we interact with our entire surroundings, not simply with each other– it is a mistake to imagine partitioning schemes to be the best approach to handling a collection of such sets of statements.

The idea of an inconsistent domain of discourse is likely to meet strong resistance. It might be thought, for instance, that the advice to measure twice and cut once holds at least the promise of eliminating error and, hence, a reduction of the problem under discussion to one already covered by neo-classicists. Notice that what this suggestion to measure twice/cut once amounts to is simply to run an experiment. With multiple measurement 'trials' we can expect to catch the very kind of errors under discussion. Unless one is a systematic incompetent, one could discover false reports by repeating the measurement procedure and tossing stray values out. The hope is, then, that we can dismiss this talk of accepted statements and holdout for honest-to-goodness *propositions*. Unfortunately, while it is true that you can reduce the occurrence of error with this approach, you can't eliminate it.

To see why this is so, consider a standard experimental practice found in sciences as disparate as psychology, chemistry, and medicine. In each science, experiments are designed to test a null 'no-effect' hypothesis, h_0, by choosing a region of rejection within a well-defined sample space. If evidence lies in this region of rejection, then h_0 is rejected. The region is selected such that if the appropriate experimental assumptions of randomness, independence and their kin hold, then there is only a small chance, ϵ, that given the supposition that h_0 is true, evidence falling in the rejection region will be collected. Another way to put it is to say that if h_0 is true, the probability of mistakenly rejecting it is less than the specified value of ϵ, where ϵ is made as small as one likes. Put in practice, we sample, check that the experimental assumptions hold, and then, should the sample obtained fall in the rejection region, we reject the null h_0 which states that the controlled variable has no effect. Note that the rejection of h_0 isn't hedged, but full-out; for instance, we *reject* the hypothesis that cigarette smoking has no effect on cancer rates in mice and men. The *grounds* for rejecting h_0 rest on the statistical –and *ipso facto* uncertain– claim that there is only a small, preassigned chance that we shall do so mistakenly.

Notice that what we've spelled out are the very same rules that measurement follows.[12] The proposal to buy certainty at the cost of taking additional measurements fails because our best experimental methods are themselves fallible. There are three points to notice about this result. First, error cannot be eliminated but only, even under the best of circumstances, controlled. So, for those interested in understanding scientific inference, holding out for honest-to-goodness propositions is an exercise in wishful thinking.[13] Second, controlling error is expensive. We get the best (but not certain) results when we are keenly aware of what kind of errors we're liable to commit and design experiments or conduct measurements in a manner that reduces those risks. Not knowing all the ways one can go wrong contributes, in part, to the difficulty of the problem of revision: we're constantly discovering new ways to err. What can make a case a tough one is figuring out whether one has stumbled upon a new way of bungling or is in the position of needing to reject part of a well confirmed theory. Finally, a comment on domains of discourse. What is compelling about accepted statements for recording the results of measurement and experiment alike is that they reflect the uncertain but likely results of a probabilistic acceptance rule that lies at the heart of these seemingly disparate activities. A domain of discourse so described is very reliable; it just isn't certain. Nearly all of what the reports state about the world is actually true. Taken together, these statements give a well confirmed yet likely inconsistent representation of the world. This structure is what science reasons with when turning to the world for physical evidence. So, any language capturing such a structure will likely contain double-talk that isn't a function of a particular agent's limited knowledge-base. The revision problem thus breaks down to at least two problems: the problem involving a noisy world and how to correct our immediate representations of it, and then the problem of squaring this imperfect representation with a given theory. Given the nature of the evidence used to revise a theory, a strong paraconsistent logic may present itself as part of a successful approach.

The moral to draw about revision is this. Tradition has it that the scientific enterprise comes full-stop when confronted with an inconsistency, which is clearly not the case. Some defenders of paraconsistency have suggested that there are cases where revising an inconsistent theory is mistaken because there are actual inconsistent objects–a controversial position, indeed. The main alternative to this view is that the best application of paraconsistency is to resolve disagreements between agents. What the proposal in this essay amounts to is to recognize that the basic relations and properties of physical objects are the bedrock of scientific reasoning and that, *qua* objects of reasoning, they are best represented by accepted statements that are the result of a reliable measure-

[12] Applied to measurement, we would propose some factor as responsible for getting an unreliable measurement. So in our example we might say that the null hypothesis is that having me measure S is *not* positively correlated with recording readings outside of the margin of error using a tape rule of kind K. Then we test to see whether or not my performance *is* outstanding and, if so, reject the null.

[13] We could chase certainty by starting a regimen of infinite trials, of course.

ment procedure. Taken together these statements make for a well-confirmed yet likely-inconsistent representation of the actual physical environment. Our world isn't inconsistent; just our evidence is. So while we strive to correct errors and mistakes when we can, we must face that this amounts to an ongoing project and is not one which we can pass off to inductive logicians while we wait for them to deliver *bona fide* properties and relations.

4 Conclusion

Hume put the matter this way: while we may know what a bad egg *is*, nothing *looks* so alike as two eggs. Traditionally, logicians think about domains containing bad eggs and good ones and reasoning about each, scoring arguments on truth preservation. But when we stop to worry about how to train this practice on our physical environment, Hume's comment becomes wise counsel. We haven't the luxury to crack every egg for inspection. This is what makes science so difficult.

What is desired is an account of how agents reason successfully about physical objects in worlds like ours with limitations very much like those we face. From a logician's point of view, we'd especially like to know the preferred way of doing this. Keeping this in mind, it may turn out that the ideal should not be constrained by representations in classical, first-order languages. I've suggested a reason here why it indeed shouldn't.[14]

Finally, a comment about honest-to-goodness inference. Another intriguing parallel between the study outlined here and 'real' inference concerns perception. Perception may be viewed as simply a measurement procedure. Within the field of computer vision, object recognition is treated as a stochastic process [32], and recent trends in computer speech recognition are toward modeling language understanding with probabilistic models [2],[21]. The point is that for perfectly epistemic reasons we may well have to work with formal representations that ascribe inconsistent properties to objects we wish to reason about. Moreover, the repair work on this domain may occur at any point during its useful life-span, not prior to its becoming available for reasoning.

On the view being developed, a well-behaved, classical, model-theoretic semantics is an afterthought. It's a tool brought to bear on a localized problem where the spread of error can be safely ignored. When dealing with empirical objects, one cannot ignore the threat of error. Indeed, the very best experimen-

[14] Graham Priest has commented that even if one grants all that I've argued for, an issue remains whether there are in fact inconsistent objects. Hence, there remains the dispute which I labored at the outset to set aside. However, I understand the best dialetheic arguments only to establish the possibility of inconsistent objects. Establishing that there are actual inconsistent objects involves adopting a logic in which such a question could be posed. The dialetheist strategy, I mentioned, is to find objects to justify the logic. On the view proposed here, this has matters turned around. So, while the objects in this essay don't satisfy the dialethiests, they are a good reason to adopt the very kind of logic which is necessary to settle such a question.

tal methods catalogue the ways to err and the chances an agent stands of falling prey to each. Our logic should take this into account.

References

[1] Alchourrön, C. E.; Gärdenfors, P.; Makinson, D. 1985. "On the logic of theory change: Partial meet contraction and revision functions", *Journal of Symbolic Logic*, 50:510-530.

[2] Allen, J. 1995. *Natural Language Understanding, 2nd ed.* Redwood City, CA: Benjamin Cummings Publishing.

[3] Benferhat, S; Dubois, D.; Prade, H. 1997. "Some Syntactic Approaches to the Handling of Inconsistent Knowledge Bases: A Comparative Study", *Stud. Logic* 58:17-45.

[4] Boole, G. 1847. *The Mathematical Analysis of Logic*. Cambridge; Reprinted 1948. Oxford: Oxford Press.

[5] Brandon, R. and Rescher, N. 1979. *The Logic of Inconsistency: a study in non-standard possible-world semantics and ontology.* Totowa, NJ: Rowman and Littlefield.

[6] Brown, B. and Schotch, P. 1999. "Logic and Aggregation", *Journal of Philosophical Logic* 28:265-287.

[7] Da Costa, N. C. A. 1982. "The Philosophical Import of Paraconsistent Logic", *Journal of Non-Classical Logic* 1:1-19.

[8] Da Costa, N. C. A. 1997. "Overclassical Logic", *Logique et Analyse* 40(157):31-44.

[9] Da Costa, N. C. A.; Bueno, O.; French, S. 1998. "The Logic of Pragmatic Truth", *Journal of Philosophical Logic* 27:603-620.

[10] Gärdenfors, P. [ed.] 1992. *Belief Revision.* Cambridge: Cambridge University Press.

[11] Howson, C. 1997. "A Logic of Induction", *Philosophy of Science* 64:268-290.

[12] Israel, D. 1980. "What's Wrong with Non-Monotonic Logic?", *Proceedings of the First Congress on Artificial Intelligence, Stanford University, August 18-21, 1980.* Palo Alto, CA: AAAI Press.

[13] Jennings, R. and Schotch, P. 1989. "On Detonating", appearing in Priest, G., and Routley, R. [ed.] *Paraconsistent Logic: Essays on the Inconsistent.* Hamden [Conn.]: Philosophia.

[14] Krips, H. 1999. "Measurement in Quantum Theory", *Stanford Encyclopedia of Philosophy.* <http://plato.stanford.edu/entries/qt-measurement/>.

[15] Kyburg, H. E., Jr. 1961. *Probability and the Logic of Rational Belief.* Middletown, CT: Wesleyan.

[16] Kyburg, H. E., Jr. 1984. *Theory and Measurement.* Cambridge: Cambridge University Press.

[17] Kyburg, H. E., Jr. 1994. "Believing on the Basis of Evidence", *Computational Intelligence* 10:3-22.

[18] Kyburg, H. E., Jr. 1997. "The Rule of Adjunction and Reasonable Inference", *Journal of Philosophy* 94:109-125.

[19] Kyburg, H. E., Jr., and Teng, C. 2000. *Uncertain Inference.* Forthcoming, Cambridge: Cambridge University Press.

[20] Laudan, L. 1997. "How about Bust? Factoring Explanatory Power Back into Theory Evaluation", *Philosophy of Science* 64:306-316.

[21] Manning, C. and Shütze, H. 1999. *Foundations of Statistical Natural Language Processing.* Cambridge: MIT Press.

[22] Mayo, D. 1996. *Error Statistics and the Growth of Experimental Knowledge.* Chicago: University of Chicago Press.

[23] Mayo, D. 1997. "Severe Tests, Arguing from Error, and Methodological Underdetermination", *Philosophical Studies* 86:243-66.

[24] Popper, K. 1959. *The Logic of Scientific Discovery.* New York: Routledge.

[25] Priest, G. 1979. "The Logic of Paradox", *Journal of Philosophical Logic* 8:219-241.

[26] Priest, G., and Routley, R. 1989. *Paraconsistent Logic: Essays on the Inconsistent.* Hamden [Conn.]: Philosophia.

[27] Priest, G. 1999. "Logic: One or Many?" Unpublished Manuscript.

[28] Quine, W. V. O. 1969. *Ontological Relativity and Other Essays.* New York: Columbia University Press.

[29] Quine, W. V. O. 1986. *Philosophy of Logic.* 2nd ed. Cambridge: Harvard Press.

[30] Restall, G. and Beall, J. C. 1999. "Logical Pluralism", Unpublished Manuscript.

[31] Routley, R., Meyer, R., Plumwood, V., and Brady, R. 1982. *Relevant logics and their rivals 1.* Atascadero, CA: Ridgeview.

[32] Sonka, M., Hlavac, V., and Boyle, R. 1999. *Image Processing, Analysis, and Machine Vision.* 2nd. ed. Pacific Grove, CA: Brooks/Cole Publishing.

[33] Teng, C. 1998. *Non-monotonic inference: characterization and combination.* Ph.D. thesis. Department of Computer Science, University of Rochester.

[34] Wheeler, G. 2000. “Error Statistics and Duhem’s Problem”, *Philosophy of Science* 67(3):410-420.

[35] Worrall, J. 1993. “Falsification, Rationality, and the Duhem Problem”, appearing in *Philosophical Problems of the Internal and External Worlds: Essays on the Philosophy of Adolf Grünbaum*, edited by J. Earman, A. Janis, G. Massey, and N. Rescher. Pittsburg: University of Pittsburg Press.