

Reasoning from Conflicting Sources

GILBERT PLUMER and KENNETH OLSON

Test Development
Law School Admission Council
662 Penn St., Box 40
Newtown, PA 18940-0040 U.S.A.
gplumer@lsac.org
kolson@lsac.org

ABSTRACT: One might ask of two or more texts—what can be inferred from them, taken together? If the texts happen to contradict each other in some respect, then the unadorned answer of standard logic is *everything*. But it seems to be a given that we often successfully reason with inconsistent information from multiple sources. The purpose of this paper will be to attempt to develop an adequate approach to accounting for this given.

KEYWORDS: defeasible reasoning, inconsistency, paraconsistency, standardized tests

We often draw information from multiple sources and reason on the basis of this information. Indeed, the knowledge base for each of us is a highly complex amalgamation of information from countless sources. So one might ask of two or more sources—what can be inferred from them, taken together? If the sources happen to contradict each other in some respect, then the answer given by standard logic is *everything*. In that case, all of our inferences would be equally worthless, since for every p we would also infer *not-p*. But surely it is a given that we often successfully reason with information from multiple sources that are not entirely consistent with one another (or even from a single source that contradicts itself). There seem to be two basic approaches to accounting for this given; they may be distinguished from one another according to whether inconsistent information (I) *is* or (II) *is not* ultimately allowed in the premise set from which one reasons. In this paper we first briefly consider features and drawbacks of approach I. Then we develop approach II by articulating strategies for expunging inconsistent information and try to motivate them by applying them in some detail to a practical situation.

APPROACH I

Paraconsistent logic

A “paraconsistent” logic is one that abandons the principle that anything follows from contradictory premises (*ex contradictione quodlibet*). Graham Priest, the leading proponent of such logics, cites a number of motivations for them (Priest and Tanaka 2004). The most controversial is his claim that there are “true contradictions,” propositions p such that p and *not-p* are both true. Priest and others who accept this claim seem to think that the main objection to it is that together with classical logic it makes

every statement true (the danger of “triviality”). Of course, if there are true contradictions, and not every statement is true, then the world does not obey classical logic.

In our opinion, Priest has gotten things backwards. The problem with contradictions is not the danger of triviality, it is that they are *contradictions*. A theory from which everything follows is bad because some of those things are false, or worse, self-contradictory. Quine made a related point years ago, adding that “not” ceases to be recognizable as negation as soon as some conjunctions of the form “ p and $not-p$ ” are regarded as true and taken as not implying everything (1970, p. 81). The proponent of true contradictions is unlikely to be moved by these considerations, which he will regard as expressions of prejudice rather than arguments (as if prejudice were always a bad thing). But we are inclined to agree with David Lewis that attempting to debate this issue is futile, since debate presupposes common ground, and once one has placed the principle of non-contradiction in doubt there is no more obvious principle on the basis of which it could be defended (2004, p.176).

What arguments can be given in *favor* of true contradictions? Priest appears to put the greatest weight on the semantic paradoxes, especially the Liar. At first glance, “This statement is false” implies its own negation and vice versa. Why not take this appearance at face value and accept that both the statement and its negation are true? It must be confessed that every solution to this paradox that has been tried comes at some cost to our naïve intuitions. To take only the most famous example, Tarski concluded that natural languages either do not contain their own truth predicates or are incoherent. Proponents of true contradictions have the unenviable task of showing that their solution comes at a lower cost than the competition, or that it confers greater benefits.

Let us call the alleged existence of true contradictions the “radical” case for paraconsistent logic. Whatever one may think of the doctrine, it must be admitted that, if we accept it, the only way we can avoid admitting every statement whatever as true is to adopt a paraconsistent logic. Other motivations for paraconsistent logic cited by Priest do not depend upon contradictions’ being literally true, but only on their being, to one degree or another, unavoidable, at least in certain contexts, where they must be kept from doing any more damage than necessary. These other motivations constitute what we shall call the “moderate” case for paraconsistent logic.

For example, a computer may store inconsistent information and be called upon to draw inferences from this information. A paraconsistent logic, it is claimed, would cut down on the number of false inferences generated, since it would not license inferring anything whatever from the contradictions. Whether paraconsistent logics have a contribution to make here is clearly an empirical, even a technological, question, possessed of no wider significance.

Priest’s discussion of set theory appears at first to be in service of the moderate case, but it seems to us in the end to depend upon the radical one. Some early set theorists (but not, apparently, Cantor) operated on the implicit assumption that every condition on sets (and perhaps other objects as well) determines a set, the set of all sets (or objects) that satisfy the condition. This is called the “(unrestricted) comprehension axiom schema.” If the condition in question is taken to be that of not belonging to itself, this assumption gives rise to Russell’s Paradox of the set of all sets that do not belong to themselves: if this set belongs to itself it doesn’t, and if it doesn’t, it does. Priest claims

REASONING FROM CONFLICTING SOURCES

that a paraconsistent approach would allow us to infer that it both does and does not without infecting the rest of the theory. The comprehension axiom schema could thus be upheld in its unrestricted form, thereby restoring naïve set theory (in one popular sense of that term).

Since paraconsistent logic is weaker than classical logic, the question naturally arises how much of set theory can be developed using it. The standard proof of Cantor's Theorem, which says that the power set of a set A has greater cardinality than A , begins by assuming that some set can be mapped onto its power set and deduces a contradiction. This proof is not available in a paraconsistent logic, and Priest acknowledges in his book *In Contradiction* (1987, p. 179) that, as far as he is aware, the theorem may not in fact be provable using such a logic. This was arguably the theorem that first put set theory on the map.

Priest seems to think that the possible loss of this and other central theorems of set theory is made up for by the recovery of naïve set theory, in which there are such things as the set of all sets. One advantage he claims for this theory is that it could provide a better model for category theory, a very general and abstract theory of mathematical structures, than does current set theory. In category theory, the category of all categories is itself a category, and it is natural to view categories as sets of certain kinds. But current set theory does not allow a set to be a member of itself. Priest's naïve set theory does; for example, the set of all sets is a member of itself. These benefits, however, such as they are, are provided, not by the use of paraconsistent logic alone, but by the acceptance of true contradictions. Someone who accepted paraconsistent logic only on the strength of what we have called the moderate case for it, i.e., as a means of keeping unavoidable contradictions from doing any more damage than necessary, would surely be reluctant to embrace a principle, like the comprehension axiom schema, that is *known* to lead to a contradiction. After all, we are assuming that such a person still regards contradictions as bad, and principles that lead to them as false. Moreover, Priest cannot be certain, even with his safeguards, that his set theory does not imply everything. The paradise that he believes he has regained may turn out to be a fool's paradise.

Other instances of inconsistent theories may be seen as more clearly bolstering the moderate case. In this connection, Priest cites Bohr's theory of the atom, according to which an electron orbits the nucleus without radiating energy, contrary to Maxwell's equations, which were assumed by the theory. Bohr's theory was thus inconsistent. "Yet" notes Priest, "patently, not everything concerning the behavior of electrons was inferred from it. Hence, whatever inference mechanism it was that underlay it, this must have been paraconsistent" (Priest and Tanaka 2004, p. 1). This is a non sequitur. By classical logic, everything *follows* from Bohr's theory, assuming that Priest's description of it is accurate. That does not mean that everything was in fact *inferred* from it. The information we are given concerning Bohr's theory underdetermines the logic that was in fact employed to derive inferences from it. We are not aware of any false inferences Frege drew from the axioms of the *Grundgesetze*, but we know that Russell's paradox followed from one of them. Does this show that the logic that underlay the *Grundgesetze* must have been paraconsistent? It is clear that Frege regarded this logic, which, after all, he stated explicitly, as classical.

Defeasibility logic

The moderate case for paraconsistent logic may equally be taken to lend support to a somewhat different variety of alternative logic. Lord Keynes, when criticized for altering his opinions, replied “When the facts change, I change my mind—what do you do, sir?” “Defeasibility logic” emphasizes the dynamic aspect of knowledge and belief. Since we modify our beliefs on the basis of new information, defeasibility logicians hold that we employ “nonmonotonic” reasoning, in which the addition of further premises can invalidate conclusions previously drawn (cf., for example, Kyberg and Teng 2001, p. 117).

It is well-known that inductive reasoning is nonmonotonic. But defeasibility logic has also been proposed as an alternative to classical deductive logic for dealing with inconsistent information. A favorite example in the literature goes as follows. Initially the information base includes the proposition that *Tweety is a bird*. In the absence of any contrary information in the base, it can be inferred that *Tweety can fly*. Here, the classical notion of logical consequence is replaced by a mere requirement of consistency governing inferences: if x has the property A (or properties A_1 – A_n), and if it is consistent with what is known to assume that x also has the property B , then it can be inferred that x has the property B . So, if further propositions containing contrary information were added to the information base, for example, that *Tweety is a penguin* and *penguins cannot fly*, then the original conclusion would have to be “retracted” and the contradictory conclusion can be inferred, viz., that *Tweety cannot fly* (e.g., Reiter 1980, p. 81ff.; Kyberg and Teng 2001, pp. 119-120; Prakken and Vreeswijk 2002, p. 219ff.; Ferguson 2003, pp. 337-338).

This brief synopsis and illustration of the approach taken by defeasibility logics is instructive in that it indicates the kinds of problems that fuller accounts must wrestle with. Surely the most glaring is, as Ferguson puts it, that, “it is simply not intuitive that $B(x)$ follows from $A(x)$ merely because the two statements are *consistent!*” (2003, p. 339). After all, that standard would often leave us free to “infer” either p or $not-p$ as we see fit. What kind of inference is that?

There are also difficulties in applying such a logic. Granted that we are permitted to retract conclusions on the basis of new information, we still face

the problem of which conclusions we are to retract and how we are going to accomplish it [which] is by no means an easy one. A conclusion can be supported by multiple monotonic or nonmonotonic derivations, and it can also be supporting the derivation of other conclusions (Kyberg and Teng 2001, p. 120n, cf. p. 131; also see Oaksford and Hahn 2006, p. 247)

In their discussion of an example that is basically parallel to the Tweety case above, Prakken and Vreeswijk (2002) find “two conflicting arguments,” (corresponding to the argument that *Tweety can fly* and the argument that *Tweety cannot fly* above.) But they acknowledge that from “the point of view of FOL [first order logic],” one should “have refused to accept one of the premises” (p. 221) or effected “a reinterpretation of a premise” (p. 222). The premise in question corresponds to *birds can fly*, which is only implicit in the defeasible reasoning, and the “reinterpretation” is the proposition that *some birds can fly*, which leads to the conclusion that *possibly, Tweety can fly*. From this point

REASONING FROM CONFLICTING SOURCES

of view, there is no conclusion that needs to be retracted even given the addition of *Tweety is a penguin* and *penguins cannot fly* to the information base.

Has anyone in the real world ever been tempted to reason even only implicitly from a premise like *birds can fly* without first attempting to disambiguate it? Does it mean that some birds can fly? Maybe Tweety is one who can't. That all birds can? Then Tweety certainly can too, but who would accept the premise? In the context given, it seems to mean that any bird that fits a certain stereotype (which does not include penguins) can fly. But if you put it that way, you will not be tempted to infer that Tweety can fly unless you have reason to believe that Tweety fits the stereotype. If we appear to belabor a single unfortunately chosen example, we reply that the literature on these logics is full of such sentences. They seem to constitute its prime motivation. Other examples of such premises include *academics are unfit* and *runners are fit* (Oaksford and Hahn 2006, p. 247).

A more technical drawback is that such logics are not “semidecidable” (Kyberg and Teng 2001, p. 147). This is just another way of saying that the consequences of a set of premises are not recursively enumerable. In contrast, classical first order logic is semidecidable in that for any formula within it, to determine whether the formula is provable, there is an algorithm that terminates if the formula is provable (this is the decidability); but there is no assurance of a terminating algorithm if the formula is not provable (this is the undecidability). Undecidability particularly besets defeasibility logics because in performing the consistency check, all the exceptions inherent in the situation (such as Tweety's being a penguin or Tweety's being an ostrich in the example above, as initially described) must be not provable. But there no effective way of showing this.

A good case can be made that logic is a priori and unrevisable. But even if one is open in principle to revising one's logic, as Quine famously was, it should surely be contemplated only as a last resort, after everything else has been tried. This was the upshot of Quine's “web of belief” metaphor. In the face of an inadequacy or mismatch between one's beliefs relative to the world, one should attempt to revise beliefs or theories at the empirical periphery of the web and work one's way inward only as necessary. Powerful justifications would be required for altering the centermost region where logic resides. Imposing novel logics as a way of handling inconsistent beliefs is like breaking a butterfly upon a wheel.

APPROACH II

Such theoretical considerations incline us to take approach II in accounting for reasoning with potentially inconsistent information, whereby inconsistent information is *not* ultimately allowed in the premise set. A curious fact is that in at least one area where logic is practically applied, we feel that one is *forced* (whether theoretically inclined to it or not) to take approach II. This seems to provide evidence that approach I—whereby inconsistent information *is* ultimately allowed in the premise set from which one reasons—may be inherently implausible.

The area of practical application is that of high-stakes and highly visible standardized testing. We will focus on the United States-Canadian (and soon to be,

Australian) Law School Admission Test (LSAT),¹ but our remarks pertain essentially equally well to reading and reasoning questions on such tests as the Graduate Record Examinations (GRE), the Medical College Admission Test (MCAT), the Graduate Management Admission Test (GMAT), and the SAT.

The problem of reasoning with potentially inconsistent information arises most vividly on the LSAT in a question type called Comparative Reading, which is being introduced on the test in June 2007 (the question type is similar to “paired passages” on the SAT). In a Comparative Reading set there are two reading passages followed by five to eight questions, most of which concern how the two passages relate to one another. In Comparative Reading there tends to be a premium placed on passages that oppose one another. This is useful for generating test questions, since it lets us ask about their points of agreement and disagreement. It also has construct and face validity for the study of law: In law there is often conflicting material to deal with, such as conflicting judicial opinions or inconsistent witness accounts of the same event.

So for such a pair of opposing passages, suppose that a question simply asks:

(1) Which one of the following statements can be inferred from the two passages taken together?

Then we could be in serious trouble if the two passages contradict each other in some respect, because an examinee could challenge the test question on the grounds that it has no single correct answer: every answer choice is correct, i.e., can be inferred. The examinee’s challenge would be presupposing classical or standard logic whether or not the examinee knew it (compare the riposte that takes the form “if absurdity x is the case, then I’m a monkey’s uncle”). But the telling point is, we could not defend the test question against this challenge by any sort of appeal to paraconsistent or defeasibility logic as *more accurate* logic than classical logic. For one thing, it would be unfair to the examinee to expect him or her to have and to apply any particular training in logic; such preparation is not a stated prerequisite for any of the standardized tests mentioned above. Moreover, if the challenge went to an outside panel of experts (such panels often include informal logicians) for adjudication, we think that we would *lose* and thereby suffer the embarrassment and expense of having to rescore the test. The underlying reason for all of this seems to be that ordinary thinking deeply embodies classical logic, whether or not it is ever consciously attended to, and therefore mature reasoners can generally be held to account (by such standardized tests) for reasoning by its principles, not some more exotic set of principles.

There appear to be certain strategies employed in everyday reasoning with potentially inconsistent information from multiple sources, strategies that are designed to expunge any inconsistent information. These strategies include:²

¹ The Central European University also uses the LSAT. Countries such as Japan and Romania have been allowed to “copy the blueprint” for the LSAT.

² Some of these strategies have affinities with some of those discussed by Woods (2003, esp. pp. 99-200). We focus on seemingly common strategies that are also useful in developing Comparative Reading for the LSAT, unlike, for example, the strategy where a person is “faced with a (noticed) contradiction in his belief-stock,” and he “quickly forget[s] that it is there if it seems harmless and inessential to any of his cognitive tasks at hand” (Woods 2003, p. 99).

REASONING FROM CONFLICTING SOURCES

(A) The “reinterpretation of a premise” as in the Tweety case discussed above, that is, revising specific propositions in the premise set as necessary to eliminate any contradiction.

(B) Considering the largest intersection of consistent statements expressed by the different sources.

(C) Recognizing that our primary data in dealing with multiple sources of information take the form of attributed or propositional attitude statements like ‘Source x says that p ’, not simply the proposition or statement p . This enables taking sources distributively rather than collectively.

(D) Narrowing the sets of statements from which the inference is supposed to be drawn from whole books or articles or passages to statements relevant to a particular topic.

(E) Treating as inductive phenomena (involving both confirming and disconfirming evidence) situations that otherwise might be treated as involving contradictions or that paraconsistent logic treats as deductive.³

A variation of strategies (A) and (B) in developing Comparative Reading for the LSAT is carefully checking the two passages in their entirety and perhaps making some revisions to ensure that they do not contradict each other in any respect. In doing this one must not go overboard and see contradictions where there are none. For example, in situations where inductive logic applies, there can be both confirming and disconfirming evidence with respect to the same thing—without there being any contradiction, e.g., in drug trials⁴ (this invokes strategy (E)). The fact is that in reading and comparing multiple

³ In discussing Priest’s 1987, Smith (1991, p. 382) says: “Priest supplies further examples of ‘positive facts’ which supposedly force us to accept the truth and falsity of a statement. These examples always involve two different pieces of evidence alternatively confirming and disconfirming that statement. For example, he points out that two different means of measuring temperature can supply evidence for inconsistent statements about the temperature of an object. . . But only one type of confirming evidence can be taken as proof of the truth of a statement: deductive inference from premises of which we are apodeictically certain. . . If inductive confirmation is not proof that a statement is true, *a fortiori*, confirmation of both a sentence and its negation by *different* pieces of evidence cannot be accepted as proof that a statement can be both true and false. . . In fact, the tension between confirmation and consistency that arises in such cases is precisely what makes them vital to scientific discovery. For example, inconsistent temperature reading obtained by what are believed to be reliable methods would rightly be taken by the scientists involved as warranting a renewed investigation into the physics of temperature [strategy A]. Priest’s theory would counsel our removing this tension in certain cases by accepting the mutually inconsistent statements as true and using his logic to explore the content of the theory containing them.”

⁴ Again, Priest seems not to take adequate account of this possibility. He says, “suppose I am playing cards and have some reason to believe the dealer to be honest. Despite this, at a crucial time, he is dealt a perfect hand. The two hypotheses are now that a) the dealer cheated: my belief that he was honest was wrong; and b) the dealer is honest and the improbable has happened by sheer chance. How one might go about deciding between these hypotheses, though in practice a sensitive issue, is in principle clear enough. One may in the end, have to accept that one was just unlucky. Similarly, if in the end there is no satisfactory way the

texts it may be critical for understanding to see how they compliment each other—to see how they are more informative or persuasive considered jointly than they are considered separately. Moreover, sometimes a larger set of statements simply will allow more challenging or sophisticated inferences than a smaller set. So the kind of task evoked by question stem (1) is a valid reading comprehension task, and the attempt is made to accommodate it in Comparative Reading on the LSAT.

Alternatively, this kind of question stem could be used in Comparative Reading:

(2) Which one of the following statements can be inferred from each of the passages?

This question means that one is supposed to consider each passage independently of the other and determine whether on its own it supports a given statement. Since one is not supposed to combine in any way the statements of one passage with those of the other, it does not matter if the two passages contradicted each other in some respect. In contrast, question stem (1) explicitly tells one to take the two passages together. In other words, question stem (2) reads distributively where question stem (1) reads collectively (strategy (C)). For example, if one passage argues that global warming is human-caused and the other passage argues that it is not human-caused, an inference as per (2) from each passage might be that the arctic icecap is shrinking (given what the passages say otherwise).

Still imagining that the passage pair is about global warming, another kind of question is exemplified by (3):

(3) The statements in the passages about worldwide temperature records most strongly support which one of the following generalizations?

Question stem (3) narrows the sets of statements from which the inference is supposed to be drawn from whole passages to statements relevant to a particular topic—worldwide temperature records. Even more specifically, perhaps, it narrows the sets to those statements that support some generalization about this particular topic. This narrowing or focusing makes it less likely that the sets will contain statements that contradict each other (strategy (D)). Relatively speaking, question stem (1) is completely unfocused.

In addition, the fact that for (3) the inference is a matter of *generalizing* and the fact that the evidence is characterized as *most strongly supporting* the generalization indicate that the inference is supposed to be inductive (or nondeductive). For (3) the task is to go from scattered temperature data to a general statement, such as that the global average temperature has risen over the past century. Since the problem of *ex contradictione quodlibet* pertains only to deductive inference, it would be uncooperative for the examinee to read question stem (3) as allowing anything to follow if the relevant passage material turned out to harbor a contradiction. Nevertheless, since technically a full degree of support (valid deduction) is still support (support at the high limit), this strategy, like the narrowing one, does not necessarily get the test developer out of the woods.

Variations on preceding kinds of question stems include the following:

contraction can be avoided, one may one may have to accept that the domain in question contains dialetheias” (1987, p. 132).

REASONING FROM CONFLICTING SOURCES

- (4) The authors of the two passages would be most likely to disagree over whether
- (5) Which one of the following is most strongly supported by information common to both passages?

For (4) the examinee is supposed to consider the particular claim in the answer choice and determine whether one author would think that it is true, and the other author think that it is false. This (distributive task) involves inferring the view of each author on a particular topic based on what is said in the respective passage. Needless to say, there is no problem of *ex contradictione quodlibet* in different people holding contradictory views (“*p* and *not-p*” is the form of a contradiction, not “*X* believes that *p* and *Y* believes that *not-p*”). The examinee is not by any means supposed to combine the two views into one (inconsistent) view. Of course one or more of the other strategies would have to be brought to bear if one happened to accept both sources as authoritative.

For (5) the examinee is supposed to consider which particular claim among the answer choices is most strongly supported by information that is shared between both passages (the information intersection). This narrows the information to be considered from all the information conveyed by either passage (the information union), and so makes it less likely that there will be an internal conflict in the information to be considered. Finally, note that the terminology of “most likely” in (4) and “most strongly supported” in (5) suggest that the inference is supposed to be inductive.

[link to commentary](#)

REFERENCES

- Ferguson, K.G. (2003). Monotonicity in practical reasoning. *Argumentation*, 17, 335-346.
- Kyberg, H.E. & Teng, C.M. (2001). *Uncertain Inference*. Cambridge: Cambridge University Press.
- Lewis, D. (2004). Letters to Beall and Priest. In G. Priest, J.C. Beall, & B. Armour-Garb (Eds.), *The Law of Non-Contradiction: New Philosophical Essays* (pp. 176-177, Ch. 10), Oxford: Oxford University Press.
- Oaksford, M. & Hahn, U. (2006). Non-monotonicity and informal reasoning: Comment on Ferguson. *Argumentation*, 20, 245-251.
- Prakken, H. & Vreeswijk, G. (2002). Logics for defeasible argumentation. In D.M. Gabbay & F. Guenther (Eds.), *Handbook of Philosophical Logic*, 2nd Ed., Vol 4 (pp. 219-318). Dordrecht: Kluwer.
- Priest, G. (1987). *In Contradiction: A Study of the Transconsistent*. Dordrecht: Martinus Nijhoff Publishers.
- Priest, G. & Tanaka, K. (2004). Paraconsistent logic. In: E.N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Winter 2004 Edition)*, URL = <http://plato.stanford.edu/archives/win2004/entries/logic-paraconsistent/>.
- Quine, W.V. (1970). *Philosophy of Logic*. Englewood Cliffs: Prentice-Hall.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81-132.
- Smith, J.M. (1991). Review of G. Priest (1987). *Noûs*, 25, 380-383.
- Woods, J. (2003). *Paradox and Paraconsistency: Conflict Resolution in the Abstract Sciences*. Cambridge: Cambridge University Press.