

# Boolean Coherence: Does it make sense?

W. T. Harwood, J. A. Clark, J. L. Jacob

University of York  
Department of Computer Science

**Abstract.** We continually face the problem of making sense of the world by resolving conflicting reports from multiple sources of information. This is particularly so if we attempt to formulate Qualitative Safety<sup>1</sup> Arguments. Traditional logic offers little to assist in this process.

In every day reasoning we usually assume that, without information to the contrary, we should use all information from all sources and that “information to the contrary” is the presence of an inconsistency between the sources. In order to resolve these conflicts we must make use of additional information which gives preference to one source of information over another when conflict arises between sources.

The suggestion put forward in this note is that it is possible to reach a ‘best’ conclusion by taking the most coherent theory that respects the preference ordering on sources and that this theory is the maximal consistent theory with respect to the ordering. This process is parameterized by an underlying notion of logic that provides the notions of consistency and entailment. This notion is straightforward when applied to the standard notion of a strict preference ordering but is a little more involved when we consider partially ordered preferences.

## 1 Introduction

A safety case documents the argument that links evidential claims to a safety claim that they support. The evidential claims are propositions whose truth is supported by evidence. The safety claim is a direct or indirect logical consequence of the evidential claims. This is essentially the view set out by e.g. Bishop and Bloomfield [3] or Wilson, Kelly and McDermid [14] (although both include additions and extensions to this basic view). If this was all there was to consider, the world would be simple and standard logical reasoning would allow one to take the evidential claims and prove (or not) the safety claim. Indeed, Rushby considers how a large system of claims might be formalized in a way that could be handled by existing theorem proving tools [11]. The problem, however, is that the real world is not so simple. Evidential claims need to be selected from a collection of potentially conflicting claims that arise from different sources of information and that are supported by varying degrees of

---

<sup>1</sup> Here, and throughout, safety refers to the the real world safety of a system, as opposed to the technical notions of safety (and liveness) introduced into program verification by Lamport [8].

evidence. Moreover the notion of degree, or strength, of evidence is far from straightforward. In general we are neither neutral to what is asserted as an evidential claim nor neutral about the source from which the claim arises. We rate the plausibility of a claim by such assessments as the degree to which it accords with our own experience, the degree of bias we might think is being expressed, the reproducibility of results, the experience of an individual making a claim, the methodology (or lack of it) that lead to the claim, etc.

A safety case will simply be that some *suitably selected* set of the evidential claims entails the safety claim. The issue is what does *suitably selected* actually mean?

The approach taken here is to at least partially formalize these notions in a framework we call Boolean Coherence<sup>2</sup>. It is related to Rescher and Manor's paraconsistent logic [10] and Default Logic [9] in the way that it deals with the presence of inconsistencies, and is strongly related to Prioritized Default Logic (see, for example, [2, 5, 6]) in its use of a preference relation to decide which assertions should be considered active during reasoning. Unlike these logics however, our concern is with selecting a consistent subset of the evidentially supported claims rather than in defining a different notion of entailment. The approach is parameterized by an underlying notion of logic which can be varied, but for this paper we will assume the underlying logic to be classical propositional logic. Entailment comes from the underlying logic and is used to show that the claims either hold, or fail to hold, with respect to this selected subset. This has the effect of making our notion of entailment more conservative than those used either in systems of paraconsistent or default logic.

To select one claim over another requires the use of additional information about the claims. We require each claim to be labeled by a *source* and that a *preference ordering* is imposed over sources.

A *source* is a formal notion reflecting whatever it is that supports the veracity of the claim (e.g. an individual making the claim, a process that produces a result, etc.), and the *preference ordering* over sources expresses the idea that some sources reflect a higher strength of evidence than others. The selection rule is that we start with the strongest sources and progressively include information from weaker sources. When contradictions arise between sources ordered by the preference relation then we preserve stronger sources in favour of weaker sources. When contradictions arise between claims not strictly ordered by the preference relation we remove the *weakest set* of sources that removes the contradiction, treating both sides of the contradiction symmetrically.

Consider a simple hypothetical example: A company intends to place a radio transmitter mast in a location near a school. Legislation permits the company to operate in one of two bands, Band X and Band Y. Recent medical opinion from extensive experimentation on mice is that Band X is unsafe. An expert biophysicist believes that if Band X is unsafe then the same is true of Band Y, although he has no experimental evidence to back this up. A technical expert is willing to testify that the company operates in Band Y. We might

---

<sup>2</sup> In contradistinction to Bayesian Coherence as discussed in, e.g. [4]

consider the ordering of plausibility of this evidence as: legislation is the most definite fact, the medical expertise and the technical expertise are on a par and that the biophysicist is expressing an opinion which is less well founded i.e.  $\text{legislation} > \text{medicalExpertise} = \text{technicalExpertise} > \text{biophysicist}$ . The conclusion we arrive at is that it is unsafe to place the transmitter mast given the best evidence available. If now, however, additional evidence from a medical expert is received to the effect that the same experiments that were performed on mice using Band X were performed on mice using Band Y and there were no ill effects, then we would revise our conclusion even if the biophysicist still holds the same opinion about the relation between Band X and Band Y. The reason we revise our assessment is that the new medical claim is stronger than the biophysicist opinion because it is backed by experimental evidence. In essence the overall safety case can be expressed as the hypothesis 'safe' follows from the most plausible theory we can construct given the available sources of evidence and our relative evaluations of the plausibility, reliability or trustworthiness of the evidence. Using a fairly self evident formalism the overall safety case can be expressed as:

```

legislation >
medicalExpertise = technicalExpertise >
biophysicist |

legislation: (BandX & ~BandY) + (~BandX & BandY),
medicalExpertise: BandX => ~safe,
medicalExpertise: BandY => safe,
biophysicist: (BandX => ~safe) => (BandY => ~safe),
technicalExpertise: ~BandX

hypothesis
safe

```

The ordering gives the reason that we reject the evidence of the biophysicist, i.e. that, in the context of the other claims, it contradicts 'stronger' claims.

Our goal is to formalize such reasoning to enable the capture of all of the claims, whether used or rejected, in constructing the safety case.

## 2 Formalisation

The reasoning process can be viewed as finding a maximal consistent set of sources that is also a maximal set in an ordering obtained when the source ordering is extended to an ordering on sets of sources. This extension to sets of sources should obey two simple conditions:

- Given two sets of sources,  $A$  and  $B$ , any sources they have in common cannot help decide between them.

- Given two sets of sources,  $A$  and  $B$ , with no sources in common,  $A$  is greater than  $B$  if for every element  $b$  of  $B$ ,  $A$  has some element greater than  $b$

That is, when we ignore the elements that  $A$  and  $B$  have in common,  $A$  is bigger than  $B$  if, whichever element of  $B$  we look at, we can always find a bigger element of  $A$ .

We formalize this as a relation over non-empty sets, by: Let  $\succ$  be the ordering on sources, which is a strict partial order (i.e. irreflexive and transitive), possibly obtained as the strict part of a non-strict partial order (i.e. reflexive, transitive and anti-symmetric), then, for non-empty, distinct, sets  $A$  and  $B$ <sup>3</sup>,

$$A \sqsupset B \equiv \forall b \in B \setminus A. \exists a \in A \setminus B. a \succ b$$

and otherwise false.

If  $\succ$  is a strict total order then there is only one maximally consistent set that is also maximal in this ordering  $\sqsupset$ . If  $\succ$  is a strict partial order then this is no longer the case and there may be none, one or many maximally consistent sets that are ‘largest’. That is, we have a collection of sets that are incomparable under the ordering  $\sqsupset$  and are inconsistent with one another. In this case we opt to take the common elements of all the maximally consistent sets that are also maximally in  $\sqsupset$ .

Let  $P$  be a propositional language and let  $L$  be a set of labels denoting sources. The set of pairs  $L \times P$  is the labeled propositional language generated by  $L$  and  $P$  and we write elements of  $L \times P$  as  $l : p$  where  $l \in L$  and  $p \in P$ . We define the projections on sets of pairs  $prop(S) = \{p \mid \exists l. l : p \in S\}$ ,  $lab(S) = \{l \mid \exists p. l : p \in S\}$  and the selective projection, for a set of labeled propositions  $S$  and a set of labels  $l$ ,  $S \circ l = \{p \mid x : p \in S \wedge x \in l\}$ .

We assume  $P$  is equipped with a consistency predicate  $CONS$   $Q$  which determines for each  $Q \subseteq P$  whether or not  $Q$  is consistent and an entailment relation  $\vdash$  which determines if a set  $Q \subseteq P$  entails a given  $p$ , element of  $P$ <sup>4</sup>. Given a set of labeled propositions,  $S$ , we define  $cons S$  as the set of all consistent subsets of  $S$  by<sup>5</sup>:

<sup>3</sup> We should note that because this relation is only used over maximally consistent sets to decide on which of two sets makes a ‘better’ choice with respect to resolving inconsistencies, there is a degree of leeway in the exact relation that could be used. We may also note that the relation we have used is closely allied to the Hoare ordering used in defining the lower, or Hoare, power domain.

<sup>4</sup> If we restrict ourselves to classical logic we may avoid having both  $CONS$  and  $\vdash$  as they are inter-definable. However, they are not necessarily so in a more general setting. Moreover, one may be tempted to think of the setup with both relations as a Scott information system [13] but this is not necessarily the case, as the underlying logic not obey the axioms of information systems. For example, the underlying logic could fail transitivity of entailment for consistent sets.

<sup>5</sup> We treat the collection of assertions with the same label as if they were a single conjunctive assertion with that label.

$$\text{cons } S = \{S \circ I \mid I \subseteq \text{lab}(S) \wedge \text{CONS}(S \circ I)\}$$

Next we define the maximal sets under subset ordering ( $\supset$ ) and the extended preference ordering ( $\sqsupset$ ):

$$\text{max}_{\supset} X = \{M \in X \mid \neg \exists N \in X. N \supset M\}$$

$$\text{max}_{\sqsupset} X = \{M \in X \mid \neg \exists N \in X. \text{lab}(N) \sqsupset \text{lab}(M)\}$$

and these are used to define the set of alternative theories each of which is maximal in the extended preference ordering:

$$\text{alternatives } S = \text{max}_{\sqsupset}(\text{max}_{\supset}(\text{cons } S))$$

Finally the most plausible theory is defined as the common elements of the alternative maximal theories:

$$\text{plausible } S = \begin{cases} \emptyset & \text{if alternatives } S = \emptyset \\ \bigcap(\text{alternatives } S) & \text{otherwise} \end{cases}$$

The effect of the first maximization, by  $\text{max}_{\supset}$ , is to take maximally consistent sets. To aid the understanding the second maximization,  $\text{max}_{\sqsupset}$ , we introduce the notion of a *conflict* within a set of propositions. A conflict in a set  $S$  is a minimal contradiction within  $S$ , i.e. a conflict is a set  $C \subseteq S$  such that  $\neg \text{CONS}(C)$  and  $\forall C' \subset C. \text{CONS}(C')$ . Let  $\text{conflicts}(S)$  be the set of all conflicts in  $S$ . We will say  $\succ$  uniquely resolves the conflicts of a set  $S$  if every  $C \in \text{conflicts}(S)$  has a minimum element under this ordering.

Given a maximally consistent set  $A$  and a conflict  $C$  in  $A$ , then at least one  $c \in C$ , is not in  $A$ . We say  $A$  excludes  $c$ . For maximally consistent sets,  $A$  and  $B$ ,  $A \sqsupset B$  means that  $A$  excludes less preferred elements of conflicts than does  $B$ . If  $\succ$  uniquely resolves the conflicts of a set  $S$  then there is a single largest (under  $\sqsupset$ ) maximally consistent subset of  $S$ . If  $\succ$  does not uniquely resolve the some conflicts of  $S$  then there are multiple maximally consistent sets under  $\sqsupset$  that differ in the elements they exclude of each conflict. As there is no preference between these resolutions we cannot decide between them and so reject all the alternative resolutions in favour of the common part of the maximally consistent sets.

Returning to the example above and applying this definition of plausible using a simple Boolean Coherence calculator program, the plausible theories for “before” and “after” the additional medical information is obtained are:

```

medicalExpertise: BandX => ~safe
biophysicist: BandX => ~safe => BandY => ~safe
technicalExpertise: ~BandX
legislation: BandX & ~BandY + ~BandX & BandY

```

and

```

legislation: BandX & ~BandY + ~BandX & BandY
medicalExpertise: (BandX => ~safe)    &
                  (BandY => safe)
technicalExpertise: ~BandX

```

A hypothesis holds if it is entailed by the most plausible theory. In our example case this gives either  $\sim$ safe or safe respectively.

### 3 A Comparison with Some Other Logics

We briefly compare this approach to the paraconsistent logic of Rescher and Manor and prioritized default logic.

Rescher and Manor define two closure operators derivable from the set of maximally consistent sets in the presence of inconsistency, here modified to deal with labeled propositions:

$$\begin{aligned}
strong(S) &= \bigcap (Th(max_{\supset}(cons S))) \\
weak(S) &= \bigcup (Th(max_{\supset}(cons S)))
\end{aligned}$$

with  $strong(S) \subseteq weak(S)$ .

A proposition,  $p$ , is true in  $strong(S)$  if it follows from every maximally consistent subset, i.e. it may be true for different reasons in each maximally consistent set. A proposition,  $p$ , holds in  $weak(S)$  if it follows from any maximally consistent set. In this case  $p$  may be true in some maximally consistent sets and  $\neg p$  may be true in others. This does not necessarily mean that all propositions are true in  $weak(S)$  since these contradictions do not collide inside a  $Th$ -closure.

If we restrict our attention to situations where  $\succ$  uniquely resolves all conflicts in  $S$  then  $strong(S) \subseteq Th(plausible(S)) \subseteq weak(S)$ .

Prioritized default logic is an extension of Reiter's default logic. We give a brief sketch of Reiter's logic and its extension to prioritized default logic.

Reiter's default logic is constructed by extending classical logic with *default rules* of the form  $(\frac{\alpha:\beta_1,\dots,\beta_n}{\gamma})$ , where  $\alpha$  is called the prerequisite,  $\beta_i$  the justifications and  $\gamma$  the conclusion, of the default rule. A rule is *active* if  $\alpha$  is true in the current theory and each  $\beta_i$  is consistent with the current theory. If a rule is active then the current theory can be extended by  $\gamma$ . If  $T$  is a classical theory (deductively closed collection of propositions under classical inference) and  $D$  is a collection of default rules, a theory extension of the default presentation  $(D,T)$  is generated by repeatedly, whilst possible, non-deterministically selecting an active default rule and applying it to obtain a new theory  $T'$  in which all consequences of the conclusion of the rule are added to the theory. When it is no longer possible to continue because there are no more active rules to apply, if every justification of every default rule used in generating  $T'$  is consistent

with  $T'$ , then  $T'$  is a default extension of  $T$ . A conclusion  $p$  follows *skeptically* from  $(D, T)$  if  $p$  is true in all default extension of  $(D, T)$  and  $p$  follows *credulously* from  $(D, T)$  if  $p$  is true in some default extension of  $(D, T)$ . An extension to these approaches is to say a proposition  $p$  is true *preferentially* if it holds in all *best preferred extension* defined by a selection rule[7].

Restricting default rules to the form  $(\frac{:p}{p})$  gives a logic essentially the same as Rescher and Manor's with skeptical consequence corresponding to strong closure and credulous consequence corresponding to weak closure.

Prioritized default logic extends default logic by allowing the specification of priorities between default rules. A priority relation is a strict partial order. Priorities may be used in at least two distinct ways: default extensions are built using the highest priority active default rules available at each step of the non-deterministic iteration; or default extensions are built in the standard way except when a conflict arises between rules, in which case the highest priority rule is applied. In either case, if there are multiple extensions, default entailment may be taken skeptically, credulously or preferentially.

Our logic closely corresponds to a prioritized default logic with rules restricted to the form  $(\frac{:p}{p})$  and partially ordered with the ordering used to resolve conflicts and the semantics taken as skeptical entailment. It differs in the entailment when the order relation does not fully resolve all conflicts in a set of propositions because we take a more conservative entailment. A proposition  $p$  is only entailed if it can be entailed from the same premises in all default extensions i.e. it follows from the common part of the default extensions.

## 4 Conclusion

The logic presented here captures a notion of entailment from the most plausible theory given some notion of ordering of the sources of information. Considering the most plausible theory drawn from all the available evidential claims causes us to be explicit in rejecting information from sources. Information is only rejected if there is more plausible, that is more preferred, information. Rejection of information therefore becomes a matter of providing explicitly the contrary case together with the assertion that the contrary case is more plausible than the rejected information thereby making the overall safety case more explicit.

The example discussed here uses classical propositional logic but it should be clear from the formalization that any logic that offers a notion of consistency and entailment can be used in its place. Practical calculation proceeds by using the ordering to decide how to attempt to extend consistent sets of information and backtracking when inconsistencies are encountered. The ability to perform the calculation is limited by the complexity of the satisfaction problem, which in the current case is NP. This problem notwithstanding, it is interesting to consider the use of logics, such as a conditional logic or relevance logic, to better reflect the information relation between assertions. One may also consider how much *effort* is put into finding the most plausible theory. We may, for

example, consider limiting consistency checking by some computational limit, reflecting the idea that there is only bounded foresight when considering how claims interact. Such changes have the potential for formally capturing more of the practical reasoning of real world safety cases.

## 5 Reviewer Acknowledgments

We thank the anonymous reviewers for their insights, comments and questions. We have attempted to address most of them without substantially changing the flow of the paper. However, one observation has gone unaddressed: there is a connection between the approach taken here and the field of multi-valued model checking (c.f. [12]). This is an interesting observation as there is a significant connection between paraconsistent and relevance logics, and multi-valued logics, particularly the 4-valued logic of Belnap [1]. The approach in [12] seems to extend Belnap's framework to multiple participants suggesting an approach to multiple sources we had not considered. We hope to investigate this suggestion in a later paper.

## 6 Sponsorship Acknowledgement

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

## References

1. Alan Ross Anderson and Nuel Belnap. *Entailment: v. 2: Logic of Relevance and Necessity*. Princeton University Press, 1992.
2. Franz Baader and Bernhard Hollunder. Priorities on defaults with prerequisites, and their application in treating specificity in terminological default logic. *J. Autom. Reasoning*, 15(1):41–68, 1995.
3. Peter Bishop and Robin Bloomfield. A methodology for safety case development. In *Safety-Critical Systems Symposium*. Springer-Verlag, 1998.
4. Luc Bovens and Stephan Hartmann. *Bayesian Epistemology*. Oxford University Press, 2003.
5. John Horty. Defaults with priorities. *Journal of Philosophical Logic*, 36(4):367–413, 2007.
6. Yarden Katz and Jennifer Golbeck. Social network-based trust in prioritized default logic. In *AAAI*. AAAI Press, 2006.

7. Sarit Kraus, Daniel J. Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *CoRR*, cs.AI/0202021, 2002.
8. Leslie Lamport. Proving the correctness of multiprocess programs. *IEEE Trans. Software Eng.*, 3(2):125–143, 1977.
9. Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
10. N. Rescher and R. Manor. On inference from inconsistent premises. *Theory and Decision*, pages 179–217, 1970.
11. John Rushby. Formalism in safety cases. In *Making Systems Safer*, pages 3–17. Springer, 2010.
12. Mehrdad Sabetzadeh and Steve Easterbrook. Analysis of inconsistency in graph-based viewpoints: A category-theoretic approach. *Automated Software Engineering, International Conference on*, 0:12, 2003.
13. Dana S. Scott. Domains for denotational semantics. In Mogens Nielsen and Erik Meineche Schmidt, editors, *ICALP*, volume 140 of *Lecture Notes in Computer Science*, pages 577–613. Springer, 1982.
14. S P Wilson, T P Kelly, and J A McDermid. Safety case development: Current practice, future prospects. In *1st ENCRESS/12th CSR Workshop*. Springer, 1995.